



A Structuralist Formal Account of Sensorimotor Contingencies in Perception

Jean-Merwan Godon

► To cite this version:

Jean-Merwan Godon. A Structuralist Formal Account of Sensorimotor Contingencies in Perception. Automatic. Sorbonne Université, 2022. English. ⟨NNT : ⟩. ⟨tel-04000131⟩

HAL Id: tel-04000131

<https://hal.science/tel-04000131v1>

Submitted on 22 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

SORBONNE UNIVERSITÉ

DOCTORAL THESIS

A Structuralist Formal Account of Sensorimotor Contingencies in Perception

Author:

Jean-Merwan GODON

Supervisor:

Dr. Sylvain ARGENTIERI
and Pr. Bruno GAS

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy*

in the

Institute for Intelligent Systems and Robotics - ISIR
Engineering Department

October 18, 2021

This PhD thesis has been defended on January 12 2022, with the jury:

Philippe SOUERES, LAAS-CNRS, rapporteur

Olivier GAPENNE, Costch-UTC, rapporteur

Angelo ARLEO, Institut de la vision, Sorbonne Université

Kévin O'Regan, INCC, Université de Paris

Bruno GAS, ISIR, Sorbonne Université

Sylvain ARGENTIERI, ISIR, Sorbonne Université

Declaration of Authorship

I, Jean-Merwan GODON, declare that this thesis titled, “A Structuralist Formal Account of Sensorimotor Contingencies in Perception” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

"Thanks to my solid academic training, today I can write hundreds of words on virtually any topic without possessing a shred of information, which is how I got a good job in journalism."

Dave Barry

"Les sciences sont des lunettes pour grossir les problèmes."

Louis Scutenaire

To my own.

SORBONNE UNIVERSITÉ

Abstract

Faculty of Sciences and Engineering
Engineering Department

Doctor of Philosophy

**A Structuralist Formal Account of Sensorimotor Contingencies in
Perception**

by Jean-Merwan GODON

The classical approach to perception revolves about a bottom up paradigm where it is the result of applying feature extraction and information processing techniques on the raw flow supplied by sensors. It is quite famously illustrated by the “Sense, Plan, Act” triptych in the field of robotics. Seeing as one generally has precise knowledge of the implementation of their robots, this approach leverages *a priori* modelisation (of the body of the agent, of its interaction with its environment...) by which to interpret and process the sensory flow. However, the self-evident aspect of this knowledge is naturally questioned when one chooses to focus on the development process of autonomous agents. In this particular context, one must instead search for perception (and where applicable for the models themselves) in the raw sensory data. One such theory of perception, introduced by O’Regan and Noë in 2001, is that of *Sensorimotor Contingencies* (SMCT). It posits that perception is precisely the skillful exercise of certain discovered *regularities* (the aforementioned contingencies) present in the *sensorimotor flow*. Importantly, it imparts crucial importance on the role of voluntary motor action in the emergence of perception; it therefore falls under the paradigms of (Inter)active perception. Since, several works have set out to test and exploit some of the insights it provides in robotic settings.

This thesis aims to expand these and propose a formalism suitable to the study of SMCT in robotic contexts. We therefore study how a “naive” agent can be made to discover structural contingencies regulating its sensorimotor flow, and how this discovery mediates “perception”. Perhaps as importantly, we carefully examine what the previous “naive” and “structural” qualifiers denote on a formal level. The goal doing so is to provide rigorous grounds to SMCT in robotics, especially in its focus on bootstrapping perceptual processes.

It starts with the exposition of our new formalism, making explicit how it relates to and generalizes previous works. We then employ it to study the structure of shifts of receptive fields that occur during motion of the agent. In particular we show that they can be leveraged by a naive agent *via* sensory prediction, and doing so they allow for discovering the structure of sequences of displacements; we also show that the algebraic language of (semi)groups provides a natural setting for such study. We then turn to addressing the issue of subjective continuity in sensorimotor experiences, that is investigating where this feeling of continuity can arise from. Developing insight from topology, which we show is unable to capture such phenomena,

we show how metric geometry can be made to characterize “typical” sensorimotor events. This allows our agent to assess the regularity of its experience in a quantitative sense, subsequently paving the way for much desired robustness and scalability properties.

Acknowledgements

We all have heard some variation on the tale of these dwarves that, following most peculiar happenings¹, found themselves perched atop the shoulders of giants. So it is honesty that compels me to concede this work had me realize how they are so much more than pretty words.

Its original form, it seems, focuses on those “great thinkers” we are privileged enough to base our works on. Accordingly, there is quite a number of them I should mention; I am however of the impression that I should refrain from naming them here, lest I postpone the dissertation itself several more pages. They are therefore, for the convenience of the reader², collated at the end of this paper in a most remarkable section titled “Bibliography”. Not all names there have proven equally as influent in this development, and regrettably neither are they all equally revered; nevertheless they all contribute their own ways to what I am meant to claim as *my contribution*. One such mention will however be found absent where I reckon it should appear in all recent theses, that of Alexandra Elbakyan: I am (and really we are) much obliged.

But such would hardly be all the work that requires to be acknowledged. Because it seems to me that science is foremost the labour of scientists, equally important is the contribution of those who helped shape who (and how) we are. So must I heartily thank the ones who taught me that questions *demand*³ asking, and that one should not yield ere they had been dealt with in the most minute detail. Friends and family first come to mind, and they should know who they are: they would find children have no business seeking imaginary numbers on a piece of paper. I must also mention professors that were instrumental in me pursuing this endeavour altogether. Many thanks to you, M. Sainte-Colombe, Goisque, and of course you Mme Rambaud and M. Boutémy⁴. It is in no small part your achievement that I forsook a sure blunder of a literary career, where I now know the thrill of spending the darkest hours of the night in the throes of anguish that is the search for a good representation of (sadly non E-unitary) inverse semigroups. Here alas I forget so many more; pray, forgive the discourtesy of me being this remiss,

¹which I hear colleagues down at the Humanities Department have taken to calling “History”

²and in conformity with established practice

³in fact, that questions demanded *questioning*

⁴Although in the case of the latter, it remains to be seen whether I shall take more from his excellent mathematical practice or his superb one man show performances.

be assured that I am grateful nonetheless and that I will right all such wrongs as they come to my attention.

Finally I address my thanks to the people at ISIR where I could make this work a reality. This is, of course, my advisors Sylvain Argentieri and Bruno Gas to which I am immensely indebted for allowing me to embark on this project. This is also the staff at the lab, who I must commend on their unwavering resolve in ensuring oblivious doctoral students go about their projects without setting too unfortunate events in motion. Yes, that is you Adela, Awatef, Sylvie, or Michèle. This is at last these fellow oblivious doctoral students whose inputs and conversation was the source of many a fruitful idea and otherwise made the daily search for *the Truth* a merrier affair. Thank you Corentin for the many questions, you Ghani for the algebraic Twister seance and you Valentin for the simulation express challenges !

Last but not least, I must precise that my thanking Sylvain Argentieri doubles as congratulations. Indeed if you are reading this it means that despite my *very best* efforts he has had me finish this work.

And if you are required to read it, I promise the writing gets markedly less *extra*.

Contents

Declaration of Authorship	iii
Abstract	viii
Acknowledgements	xi
1 The sensorimotor account of perception	1
1.1 Some highly symbolic history	2
1.1.1 The classical sandwich model	2
1.1.2 Limitations	3
1.2 Modelling the internal experience	4
1.2.1 Sensorimotor grounding	5
1.2.2 Learning to perceive	5
1.2.3 The necessary role of action	7
1.3 The Sensorimotor Contingencies Theory	9
1.3.1 Perception as sensorimotor mastery	9
1.3.2 Further applications	11
2 A formal account of sensorimotor flows under motor actions: a revised formalism	13
2.1 Out with the old, in with the new	15
2.2 State of the framework and standing limitations	16
2.2.1 A framework for describing sensorimotor experiences	16
2.2.1.1 The sensorimotor diagram	16
2.2.1.2 Interpretation in concrete cases	17
2.2.2 A problem of bootstrapping: internal data and <i>a priori</i> knowledge	18
2.2.2.1 Mathematical structure as a formal correlate of knowledge	18
2.2.2.2 From my point of view...	20
2.2.2.3 The nature of the motor set	22
2.2.3 Limitations	23

2.2.3.1	Locomotive motives	23
2.2.3.2	Stateful descriptions make for static	24
2.2.3.3	The tradeoff of numericity	25
2.3	Towards a formalism of sensorimotor flow	26
2.3.1	Reparameterizing the description	27
2.3.1.1	The <i>where</i> and the <i>what</i> : environmental concerns in space	27
2.3.1.2	Motor data accounting for locomotion	32
2.3.2	The naive sensorimotor experience	36
2.3.2.1	<i>Moteur...</i> action!	36
2.3.2.2	Sensels, sensors, receptive fields	43
2.3.2.3	How can it be so naive: a model of knowledge	51
2.4	Conclusion	58
3	Discovery of the motor algebra by prediction of sensory shifts	61
3.1	A zero-th level of <i>spatial</i> sensorimotor contingencies	63
3.2	A formal study of learning motor structure <i>via</i> sensory prediction	64
3.2.1	Introducing conservative actions	64
3.2.1.1	Special spatial exchanges in sensation	64
3.2.1.2	An algebraic conversation of conservation	67
3.2.1.3	Partial successor map and quasiconservative actions	71
3.2.2	Perceptive properties of sensory prediction	74
3.2.2.1	Using predictable actions as motor proxies	74
3.2.2.2	Prediction in the case of conservation	79
3.2.2.3	Degeneracy in the semigroup case	81
3.2.3	Conclusion	86
3.3	Experiment: a simulated account of inferring action structure by naively assessing conservativity	88
3.3.1	Description of the setup	88
3.3.2	Simulation 1: searching the sensorimotor flow for conservativity	92
3.3.2.1	Representing conservative actions	92
3.3.2.2	Progressive discovery of the conservative permutations	94
3.3.2.3	Results	101

3.3.3	Simulation 2: investigating the dynamics in the discovery of sensel permutations	102
3.3.3.1	Assessing progress in the convergence of the empirical matrices	102
3.3.3.2	Sources of variance in the dynamics of discovery	103
3.3.3.3	Experimental setup	104
3.3.3.4	Results	105
3.3.4	Simulation 3: internally building the conservative sub-algebra	108
3.3.4.1	Generative Dijkstra-like algorithm for computing the group of sensory predictions	108
3.3.4.2	Managing loss of information in nonconservative actions	112
3.3.4.3	Force directed methods for plotting the Cayley graphs	115
3.3.4.4	Results	116
3.3.5	Simulation 4: reparameterizing basis actions	123
3.3.5.1	Dependence of the graph with respect to the starting action set	123
3.4	Conclusion	126
4	Towards a naive continuity of sensorimotor experiences	129
4.1	Speaking from experience: the need for continuity	131
4.2	Finding smoothness in statistical regularities	133
4.2.1	Characterizing continuity: the topological approach	133
4.2.1.1	Let's get real: the initial model case	133
4.2.1.2	Abstracting topology out of the real example	135
4.2.1.3	Application to the sensorimotor experience	138
4.2.1.4	Discrete spaces: so close yet so far...	141
4.2.2	A naive process for building quantified continuity	143
4.2.2.1	Ascribing experiential meaning to metric values	143
4.2.2.2	Building the sensory metric	144
4.2.3	Some properties of the sensory metric	147
4.2.3.1	Fitting spatial and sensory dynamics in the exploration	147
4.2.3.2	Using the metric in regularity assessment	150
4.3	Experiment: building the continuity map under perturbation	151

4.3.1	Description of the experimental setup	151
4.3.2	Computing the sensory distance	154
4.3.3	Spatially representing the sensory metric	156
4.3.3.1	The use of dimensionality reduction	156
4.3.3.2	ISOMAP	158
4.3.4	Results	161
4.3.4.1	The sensory metric	161
4.3.4.2	Crawling in my film: Effects of amplitude in the agent's movements on the discovered reg- ularities	165
4.3.4.3	Introducing noise in the sensorimotor interac- tion	169
4.4	Conclusion	174
	Conclusion	179
	Bibliography	185

List of Figures

2.1	Representation of the model interaction used throughout the development	15
2.2	Fitting the mobile camera in the MXS scenario	18
2.3	Distinct motor configurations related by locomotion.	23
2.4	A bijection $\mathbb{N}^2 \cong \mathbb{N}$ illustrated.	29
2.5	Modification of the motor parameterization to circumvent the issue of locomotion.	32
2.6	Visualisation of the effects of motor actions as vectors fields for $\mathcal{B} = \mathbb{R}^2$	37
2.7	Decomposition of arbitrary motor trajectories over a repertoire of primitive actions.	40
2.8	Sensels, sensors, sensory apparatus.	44
2.9	Undefined receptive field in pathological environment systems.	50
2.10	Sensation as arrays of (strongly) uninterpreted symbols.	52
3.1	Illustration of the way receptive fields shifts under motor displacements of the agent	65
3.2	<i>Quasiconservativeness</i> in finite receptive fields	71
3.3	Partial maps generalizing the sensel permutations of conservative actions.	73
3.4	Redundant motor configurations and pose predictability.	75
3.5	Partial maps generalizing the sensel permutations of conservative actions (<i>repeated</i>).	82
3.6	Composing inverse partial mappings only results in a partial identity.	83
3.7	Accumulated loss of information in compositions of partial bijections	84
3.8	Experimental setup for the prediction experiment	88
3.9	Empirical permutation matrices discovered	100
3.10	Evolution of the certainty in sensel permutations during exploration	105

3.11	Statistical survey of the dynamics in the internal discovery of senses permutations.	106
3.12	Directed graph of empirical permutation matrices M_{a_k}	116
3.14	Visualisation of compatible predictions resulting from distinct actions.	121
3.15	Erroneous virtual prediction associated to spurious compati- bilities	122
3.16	Effects on the Cayley graph of various random draws of start- ing actions.	125
4.1	Illustration of metric continuity: the case of maps $\mathbb{R} \rightarrow \mathbb{R}$. . .	134
4.2	Illustration of the final topology in the case of the toy model. .	138
4.3	Illustration of the limited capabilities of topology applied to discrete spaces.	142
4.4	Compression of the input color values by the camera sensor .	153
4.5	Illustration of the internal code of color sensations as numeric labels	153
4.6	Comparison of the ambient and intrinsic metrics on a Rieman- nian manifold.	157
4.7	Sensory metric matrices resulting from the sensorimotor inter- action of the agent in the base case.	162
4.8	Evaluation of the symmetric component of the quasimetrics empirically discovered	163
4.9	Comparison of experienced local spatial coherence in the en- vironment for various amplitudes of motor displacement. . .	167
4.10	Effectiveness of the sensory metric as a typicality benchmark for distinct choices of movement amplitude.	168
4.11	Influence of the amount of noise on the discovered sensory metric.	170
4.12	Evolution of the distributions of sensory distances with envi- ronmental noise.	172
4.13	Evolution of the variance in sensory distances with environ- mental noise.	173

Chapter 1

The sensorimotor account of perception

Robotic agents, much as living organisms, exist in and interact with the physical world. This presumes that they are endowed with some capabilities with which they can conduct this interaction: namely they can use their sensors to gather information about what lies out there in their environment, while their actuators allow them to enact whichever *interaction* they are partaking in. In particular, these capabilities can support some behavioral adaptation to the wealth of situations and environments an sensorimotor agent may be faced with in realistic settings. As such they are of particular importance to the operation of *autonomous* robots which, by design, should confront various settings using only their own interactive capabilities.

At the heart of this issue lies the notion of perception. Generally, one can observe that the sensors of any such agent yield readings that are not directly suitable for proper exploitation by an “intelligent” agent¹. Instead they provide very rich but low level information, (hopefully) flexible enough to capture the details of the outside environment but that lacks structure the agent can quickly and clearly apprehend. In fact the same observation can be made for effectors: for complex kinematic chains, the relation between actuation of some joints and resulting movement of the whole body is not always obvious. Perception, then, is precisely the process by which agents get *meaning* from these readings; in that it informs them with an understanding of their interaction fit for their sensorimotor capabilities, it further enables them to adapt their behavior to the situation they are currently facing.

¹or, as William James famously put it, a “blooming buzzing confusion”.

1.1 Some highly symbolic history

1.1.1 The classical sandwich model

Many models of the sensorimotor interaction, both in robotics and in cognitive psychology (Burr, 2017; S. L. Hurley, 2002), share a common classical approach based on the “Sense, Plan, Act” (SPA) internal paradigm as per

$$\text{Sense} \longrightarrow \text{Plan} \longrightarrow \text{Act}$$

It aims to delineate the sensorimotor interaction as the succession of these three processes, in which

1. The agent accesses the readings its sensors are reporting, which it then processes and interprets to further understand the conditions it is interacting with. For instance in a visual scene it may attempt to discern and separate the objects in its field of view, determining where and what they are.
2. From the understanding of the scene it has acquired in the first stage it can plan its next step, further determining its behavior. Continuing the same example, it can attend to some object of interest it has detected in the scene and choose to go grasp it.
3. Finally it uses its actuators to enact the behavior it has devised, which involves both motor planning and motor control. Thus, performing the grasping task requires the agent to determine which effectors it should actuate (and how) as well as ensure its actual motor trajectory follows its specifications.

This organization, for example illustrated in (Russell and Norvig, 1995), transparently supposes a flow going from sensors to effectors². Accordingly, works which adhere to this paradigm usually employ terms such as “sensory *inputs*” and “motor *outputs*”. Moreover it attempts to separate the central cognitive functions, which modelling and planning capabilities form the meat of the sandwich, from surrounding sensorimotor abilities of the agent.

²In fact, the reference asserts in its preface that one of its five principal features is “Intelligent agent design” of which it says “The unifying theme of the book is the concept of an intelligent agent. In this view, the problem of AI is to describe and build agents that receive percepts from the environment and perform actions. Each such agent is implemented by a function that maps percepts to actions, [...]”.

As a result, it has been dubbed a cognitivist or computational account of the sensorimotor experience (Brooks, 1999; Rescorla, 2020).

As the details of its three processes show, the SPA diagram has perception occur at the later stages of the “Sense” process which combines both querying of the sensations and their interpretation as perceptual items. This is classically realized by applying signal processing techniques to the data stream supplied by the agent’s sensors, which extract features and combine them into a representation of the interaction it can readily use in its later planning stages. This especially entails that the agent must *a priori* be capable of conducting the corresponding algorithms, which rely on carefully crafted models of its body (Mutambara and Litt, 1998), its sensors or its interaction with its environment (Marconi, Naldi, and Gentili, 2011).

Finally, although rigid adherence to the SPA paradigm is generally more prevalent in older works it still largely permeates current sensorimotor models (Oddi et al., 2020; Siegel, 2003; Srivastava, 2019). In particular, many modern control architectures such as e.g. that found in (Bouman et al., 2020) are built by refining the basic SPA scheme and introducing specialized sub-systems; as such, they all exhibit the previously described traits.

1.1.2 Limitations

Because it so clearly decouples higher cognitive processing from the actual sensorimotor interaction, the Sense, Plan, Act architecture directly allows designers to tap into symbolic AI techniques to implement their agents’ executive functions. As such the early promises researchers sought in these methods (McCorduck, 2004; Newell and Simon, 1963) redounded on SPA, hinting that it may prove a very efficient framework for modelling physical robots. This enthusiastic outlook depended on the working assumption that the core “intelligence” agents had to replicate lay in higher cognitive functions. Furthermore, it came out on the hypothesis that these functions could generally be readily formulated and fulfilled using the symbol semantics of our conscious understanding of our experience; for example this corresponds to the idea that since we know how to physically reach and grasp some object we see, so must we (*consciously*) be able to know how it is that we perform this endeavour³.

³For instance, Rodney Brooks recounts in (Brooks, 2008) that “in 1966, some AI pioneers at MIT thought it would take three months—basically an undergraduate student working during the summer—to completely solve the problem of object recognition”.

This approach, however, proved much harder to proceed with than predicted⁴. Evidence not only in the field of AI (Dreyfus, 1965, 1986; Lighthill, 1973; Lungarella, Iida, et al., 2007), but also that of robotics (Brooks, 1990, 1991a,b) and neurophysiology (Cisek, 2007; Cisek and Pastor-Bernier, 2014) shows that the serial and symbolic focus of this architecture is sometimes ill suited to reproducing natural “intelligent” behavior. Moreover its framework offloads the burden of establishing *how* an artificial agent should determine its behavior and further enact it onto its designer. Indeed, it is the choice of models and algorithms (e.g. that of feature extraction, decision making, motor planning...) they program into the agent that dictates its behavior down the road; as a result it has been argued that the intelligence of these artificial agents is in fact that of their builders, and that they neither perceive nor cognize themselves (Grabiner, 1986).

This last point entails further limitations relative to the breadth of agents and sensorimotor interactions the architecture can efficiently model. Indeed it relies on the agent leveraging externally provided models and methods to conduct its own sensorimotor experience, thus requiring that the model faithfully characterize the actual experience of the agent. This is especially troublesome since external models are at times inaccurate (C. Lee et al., 2017) and incomplete (Nguyen, Sreedharan, and Kambhampati, 2017). As a result, both the specific models (e.g. of features to extract from the sensory flow) and algorithms (e.g. for motor path planning) chosen beforehand by its designer may end up suboptimal in the particular conditions an artificial agent experiences. Finally, in as much as one views AI as a model of intelligence (Morse et al., 2011) then the symbolic SPA paradigm raises the issue of the knowledge agents *a priori* possess. Artificial agents designed along this architecture perceive and act in accordance with finely crafted models provided by some designer; on the other hand it does not itself account for how natural intelligence can emerge without such external input.

1.2 Modelling the internal experience

⁴Of course, Brooks adds, “the student failed”.

1.2.1 Sensorimotor grounding

Critics of the previous approach (Brooks, 1990; Dreyfus, 1986) have argued that symbolic perspectives are fundamentally flawed for modelling the cognitive functions involved in sensorimotor interactions. According to them, the brunt of said cognition pertains to *subsymbolic* thinking, and crucially it is carried out by subconscious processes in human cognition. Consequently our conscious reasoning about our own sensorimotor interaction need not faithfully and effectively model our experiences, let alone that of artificial agents we design. By contrast, this subsymbolic cognition has the added benefit of being *physically grounded* (Harnad, 1990) in the actual sensorimotor experience of the agent; in particular the (sub)“symbolic” units the agent perceives and acts upon should involve no representation, instead directly appearing in its sensorimotor interaction with its environment.

This reasoning relies on there being some very rich structure in any one robotic agent’s sensorimotor flow, imposed by its embodiment (Pfeifer and Bongard, 2006; Varela, Rosch, and Thompson, 1993). Indeed it posits that the particular interactions and dynamics thereof any one agent experiences largely depend on its specific physical body. According to this perspective of *situated* cognition, experiences not only shape the contents of the agent’s sensorimotor knowledge but also the terms wherewith it can be formulated (Maturana and Varela, 1980). Thus the level at which an agent’s perception and cognition can adequately be conducted is to be determined internally to the agent itself, as opposed to the external representations of the symbolic AI paradigms. This means that the internal *workings* of an agent’s perception may not seem intuitive to us, especially if its embodiment (either its sensory apparatus or motor capabilities) is largely distinct from ours; nevertheless this helps ensure that what perception they get from their sensorimotor interaction is actually relevant to their own capabilities, which in turn enables them to produce adapted *behavior*.

1.2.2 Learning to perceive

Some of the most severe limitations of the SPA architecture we previously outlined followed from the models provided to artificial agents not being optimal in their respective use cases. In particular, they then lacked the capabilities to adapt their *a priori* knowledge to the actual conditions they could be confronted with. One way to tackle this issue is to instead have the agent

learn its models in some capacity. For instance, robotic agents can learn models of their sensory apparatus (Censi and Murray, 2012) or of their spatial interaction with their environment (Jonschkowski and Brock, 2015). In fact the practice has since become common enough that it corresponds to a specific domain, that of developmental robotics (Asada et al., 2009; Lungarella, Metta, et al., 2003). In this approach, we consider the cognitive capabilities of the agent (as well as the knowledge it possesses) as something it incrementally constructs (Ziemke, 2001) as it pursues its experience. Indeed another essential characteristic of the sensorimotor interaction is that it is something the agent *partakes in* rather than is subjected to; thus the particular way in which it enacts its exploration of its environment gradually and actively shapes its sensorimotor experience (A. Noë, 2004; Stewart, Gapenne, and Di Paolo, 2010).

Taking these considerations to their extreme leads us to the bootstrapping problem (B. J. Kuipers et al., 2006). It raises the question of determining what “minimal” knowledge of its interaction the agent must possess so as to progressively build higher cognitive capabilities (including perceptive capabilities), as well as how this building can be achieved. This reflects the observations that such an incremental development of sensorimotor cognitive capabilities occurs at the intra-individual level in infants (Piaget, 1937; L. B. Smith and Thelen, 1994); among others they have been shown to internally develop notions as basic as space and time, further advancing to objects and then causality (Sexton, 1983). Besides, B. Kuipers’s works on the Spatial Semantic Hierarchy (B. Kuipers, 2000, 2008; B. Kuipers and Byun, 1991; Remolina and B. Kuipers, 2004) have shown how robotic agents could iteratively construct spatial knowledge of their interaction, before using said internal knowledge e.g. in spatial navigation.

The elaboration of these developmental methods therefore answers several previous issues with the cognitivist approach. First, it provides the agents with some degree of adaptability relative to their actual conditions of experience. One clear benefit is that an artificial developmental agent may alter what models its designer endowed it with to better match its own interaction as opposed to the ideal modelled case. Moreover, it also potentially enables the agent to adapt on the fly should the conditions change during its experience; for example this situation could happen as its actuators wear or even break down, thus increasing its autonomy. Second, it begins addressing the considerations relative to *a priori* knowledge of an agent. In particular

the bootstrapping scenario doubles as an examination of the ways in which “natural” intelligence can develop from the most bare set of assumptions⁵.

1.2.3 The necessary role of action

The SPA architecture posits that the cognitive processes involved in the sensorimotor experience flow from sensory readings to the actuation of an agent’s joints. As a result, it has perception emerge entirely as a product of the sensory flow. This however runs contrary to experimental evidence observed in the biological setting, which tends to show that perception crucially depends on the agent actively exercising its motor capabilities. A first experiment is famously found in (Held and Hein, 1963). Its experimental setup involves pairs of kittens placed in a carousel environment acting as a visual scene. Importantly, the kittens are linked by a harness contraption which ensures that they follow identical displacements (up to central symmetry) at all times. In each kitten pair one –the *active* kitten– is free to actively move, whereas the other, firmly encased in a gondola, is *passively* subjected to the same movement. The kittens have been chosen as to first see the world under these conditions, having been reared in darkness until then. After some time, the authors’ tested their perceptual abilities as visually guided behavioural responses. Importantly, they found that all active kittens displayed traits coherent with that of cats that had had normal visual experiences; on the other hand passive kittens displayed visually impaired behaviors, for instance failing to respond to fast approaching objects. Importantly the mechanical contraption of the setup guarantees that, as far as their sensory flows were concerned, the two kittens of any one pair exactly shared their learning experiences. This result therefore shows how *active* motor exploration is necessary for the emergence of some perceptive capabilities.

The second experiment is that of Bach-y-Rita’s sensory substitutions (P. Bach-y-Rita, 1972; P. Bach-y-Rita et al., 1969; Paul Bach-y-Rita, 2005). In these experiments blind participants were given a device which performed tactile vision sensory substitution (TVSS), rerouting visual sensory information over the tactile modality. Concretely the substitution device acquired a visual image via a camera sensor, which it converted into black and white. Then, it reproduced the corresponding stimulus on 2-d arrays of tactile stimulators

⁵Although in a biological setting, the incremental “development” stages may well occur at the inter-individual evolutionary level (which is transparently mirrored in robotic settings *via* evolutionary robotics (Doncieux et al., 2015)) as well as at the intra-individual psychological one.

in contact with the skin of the user, with the activation level of a stimulator denoting the luminance of the corresponding visual point. The participants were then asked to discriminate between objects which were being filmed by the camera on the basis of the reproduced tactile stimulus. First observations determined that they did not perform better than chance level whenever they only passively experienced the tactile stimulus and they had no feedback as to the correctness of their answers (Sampaio, 1994). On the contrary, when they were allowed to actively manipulate the camera to change the way (e.g. distance, angle, inclination) in which it filmed the target object then the blind subjects developed remarkable object recognition capabilities. First, they became able to tell lines and volumes from one another. Then they learned to tell concrete objects (among others a chair, a cup and a telephone) apart; finally, they succeeded in discriminating between faces, postures and movements of other individuals. As in Held and Hein's experiment before, this indicates that action plays a pivotal role in the development of one's perception: without actively leveraging their interaction capabilities with the TVSS device, participants could not "get" meaning (that is, perception) from the sole sensory feed they were provided by the tactile stimulators. The result, however, extends further. Indeed subjects learned to describe their experiences using *a priori* visual properties. Among them they reported positional relations between recognized objects when several were concurrently filmed (i.e. whether –and which– one was "on the left" of the other); additionally they reported "perspective, parallax, looming and zooming, and depth judgments" (Paul Bach-y-Rita et al., 2005). But where all these relations are expected to occur in visual perception from the way the physics of light mediate our visual sensory experiences, they are instead usually absent from tactile perception which generally lacks a gradual concept of "closeness" as that implied by perspective. Therefore this shows that these perceptual relations are not as much properties of a specific sensory modality (i.e. of the physical ontology of the signal captured) as that of the interaction the agent is conducting.

Similar observations about an agent's perception being facilitated by *active* interaction with its environment have been made in robotics. They have shown that the constraints under which physical agents can move in their environments –and the resulting constraints on the sensorimotor interaction as a whole– are among the key regularities imposed by their embodiment. On

the opposite side disembodied methods which cannot or choose not to leverage this added structure, such as (especially Deep Learning based) Computer Vision, instead need to rely on much larger amounts of training data to statistically recover these regular constraints. As a result, a number of approaches have been designed with the intent to base robots' perception not only on their sensory flow but on the whole sensorimotor one; they have generally been subsumed under the field of Interactive Perception (Bohg et al., 2017).

1.3 The Sensorimotor Contingencies Theory

1.3.1 Perception as sensorimotor mastery

One particular proposed theory of perception displaying all three previous traits is the Sensorimotor Contingencies Theory (SMCT), introduced in (J K O'Regan and Noë, 2001) and further developed in (J. O'Regan, 2012). It puts its focus on the regular laws by which a physical agent's sensations change as it interacts with its environment, that is precisely the sensorimotor *contingencies*. As in the previous argument it emphasizes that these sensorimotor experiences are largely constrained by, among others, the specific embodiment of the agent and physical properties of its environment. Then it asserts that perception is precisely the interactive state of agents that have recognized and mastered what structure on their sensorimotor flow emerges from these constraints. Thus, to an agent, perceiving some quality is recognizing that their current sensorimotor experience exhibits the contingencies characteristic of that quality. Let us borrow an example from (J. O'Regan, 2012) and consider a situation where we hold a sponge in our hand. By the SMCT account, our perceiving the *flexibility* of the sponge is precisely the internal sensorimotor assessment that we would feel our fingers press *into* it should we close our fist. This can of course be compared to the case where we would hold a solid cube: we would perceive it to be stiffer, corresponding to the sensation of our fingers abruptly stopping at its surface.

Interestingly, the Sensorimotor Contingencies Theory somewhat reverses the intuitive notion of what perception is. We introduced it as a process by which the agent "gets" meaning out of its experienced sensation; but here SMCT has it be instead the felt experience of an agent which, having recognized the contingencies structuring its interaction, exercises the corresponding mastery. Because of this, its proponents claim that it provides an

account of the sensory *qualia*⁶ which eschews the pitfalls other theories encounter. Indeed, it follows an argument (made for instance in (J K O'Regan and Noë, 2001)) that theories which appeal to objective qualities of the supporting material processes (e.g. the “neural correlates”) to explain subjective qualities of the corresponding experiences inevitably result in an explanatory gap (Levine, 1983). This is thus a particular instance of the hard problem of consciousness (Chalmers, 1995; Nagel, 1974), which states that no material theory of mental processes can successfully reductively account for the subjective conscious experience associated to them.

The previous discussion makes apparent that the SMCT approach dismisses the need for any internal representations, such as those of classical symbolic paradigms. It posits that all required information is to be found in the sensorimotor flow of the agent, that is precisely what is available from its interactive capabilities; moreover it argues that it should not necessitate any internal “re”presentation of the world, instead probing its actual environment whenever it needs to access its actual state (Melnik et al., 2018; J. K. O'Regan, 1992). This entails in particular that it only requires very basic initial knowledge from the agent, by which it also addresses the bootstrapping problem. More precisely, assuming that the agent can discover the contingencies in its experienced sensorimotor flow equates to it learning the corresponding perceptual capabilities. As such, the SMCT is also a theory of perceptual development. Finally, its very definition of perception presupposes there is some action to which the recognized sensorimotor contingencies correspond. Therefore it is part of the approaches of active perception we previously described. This is particularly clear in the case of tactile perception, such as in the sponge example discussed earlier; nevertheless the same argument holds for other modalities, albeit less transparently⁷. By this account, not only does perception require some specific motor actions but it is even a particular way of acting, one that leverages what mastery of sensorimotor contingencies the agent has acquired. Accordingly, it asserts that perceiving is the agent *making* sense of its experience rather than “getting” meaning from it.

⁶that is the *subjective* qualities of the sensory experience, or “what it’s like” (for the agent) to have this experience.

⁷In particular, (A. Noë, 2004) discusses at length how vision displays such “touch-like” qualities.

1.3.2 Further applications

A number of works have since been conducted in the theoretical setting supplied by the SMCT, pursuing either of two complementary goals. First they may attempt to follow its principles and observe what sort of “intelligent” behavior or agent one may obtain from them, which furthers its role as study and explanation of (natural) intelligence. On the other hand, its adhering to the traits of 1.2 lends it desirable properties as a framework for implementing robotic perception. Thus, its developmental focus and lack of representational requirement should make it result in highly generic solutions; ideally we would hope to need supply it only with the capabilities to recognize certain (types of) sensorimotor contingencies, and let it further determine how they actually manifest in its own specific sensorimotor flow.

Several early works have been devoted to showing how sensorimotor contingencies allowed *naive*⁸ agents to discover intrinsic properties of their interaction. One such property is the dimension of the ambient space in which both the agent and its environment lie (Philipona, J K O'Regan, and Nadal, 2003; A. V. Terekhov and J. K. O'Regan, 2013), which can be discovered using the contingencies of compensable transformations introduced in (Poincaré, 1902). This illustrates the two distinct sorts of results. Indeed it shows that monitoring the contingencies in one's sensorimotor flow allows determining this important value, which therefore can be assumed to be unknown to the agent at first. Moreover it also demonstrates that the corresponding cognitive architecture is generic enough to handle sensorimotor agents which may experience settings of any dimensions, using a unique method for processing its sensorimotor flow across all possible values. Another such result is that of (D. L. Philipona and J. Kevin O'Regan, 2006), which found that the (*a priori* subjective) property of colors to be “pure” can be found as an intrinsic sensorimotor contingency of their reflectance. On the motor side, it has correspondingly be shown that basic sensorimotor contingencies allow robotic agents to discover an internal image akin to a body schema (Hoffmann, 2014; Roschin and Frolov, 2011). As before, this importantly shows that these structures need not be known *a priori* by the agent.

More recent works have investigated how SMCT could support the emergence of higher level perceptual processes, associated to certain “extended”

⁸While mostly informal at this point this term denotes an agent which has very limited sensorimotor “knowledge”, that is the kind of agents considered in the bootstrapping problem. The notion and how it can be made precise is examined in much greater detail later in this work (see e.g. 2.2.2.1 or 2.3.2.3).

sensorimotor contingencies. In particular they have been shown to readily yield capabilities of (visual) object perception (Le Hir, Sigaud, and Alban Laflaquière, 2018; Maye and Andreas K. Engel, 2011), motor planning and action generation (Maye and Andreas K Engel, 2013). Besides, experiments have successfully tested some of its predictions in human subjects: performance in saccadic eye fixation tasks supports the hypothesis that visual perception does not rely on internal representation (Melnik et al., 2018), whereas (König et al., 2016) shows that new contingencies (here corresponding to a “magnetic” sense) can be learned to yield new percepts.

Getting formal Although some authors have tried to put forward formal frameworks for dealing with SMCT statements (see e.g. (Buhrmann, Di Paolo, and Barandiaran, 2013; Seth, 2014)), distinct works falling under the setting are still vastly different in their approaches. This can be attributed to the very diverse disciplinary backgrounds authors who draw from the theory may have, as well as it being kept informal in its original formulation (J K O’Regan and Noë, 2001). However this entails that expressions (such as e.g. “sensorimotor contingency” or “naive agent”) may not always denote the same idea between different bodies of works. But in turn this means that statements which may hold under one acceptation may be false for another. As a result, works such as (Buhrmann, Di Paolo, and Barandiaran, 2013; Alban Laflaquière, J. Kevin O’Regan, et al., 2015; Valentin Marcel, 2020) put great emphasis on explicitly highlighting the formal extent of their statements and the conditions under which they are examined. Moreover, this endeavour also makes stronger the point about genericity of the proposed approach⁹.

We aim to further develop this approach in this thesis. Accordingly in the following most arguments are simultaneously followed both formally and informally, corresponding to the definition and development of mathematical statements to mirror the SMCT discussion. In this aspect this work can be seen as a continuation of (Valentin Marcel, 2020), which itself largely drew from (Alban Laflaquière, 2013); we detail them in the next section.

⁹(as can be seen in (Censi, 2012) for another –non-SMCT – contribution devoted to the bootstrapping problem)

Chapter 2

A formal account of sensorimotor flows under motor actions: a revised formalism

Contents

2.1	Out with the old, in with the new	15
2.2	State of the framework and standing limitations	16
2.2.1	A framework for describing sensorimotor experiences	16
2.2.1.1	The sensorimotor diagram	16
2.2.1.2	Interpretation in concrete cases	17
2.2.2	A problem of bootstrapping: internal data and <i>a priori</i> knowledge	18
2.2.2.1	Mathematical structure as a formal correlate of knowledge	18
2.2.2.2	From my point of view...	20
2.2.2.3	The nature of the motor set	22
2.2.3	Limitations	23
2.2.3.1	Locomotive motives	23
2.2.3.2	Stateful descriptions make for static	24
2.2.3.3	The tradeoff of numericity	25
2.3	Towards a formalism of sensorimotor flow	26
2.3.1	Reparameterizing the description	27
2.3.1.1	The <i>where</i> and the <i>what</i> : environmental concerns in space	27
2.3.1.2	Motor data accounting for locomotion	32

2.3.2	The naive sensorimotor experience	36
2.3.2.1	<i>Moteur...</i> action!	36
2.3.2.2	Sensels, sensors, receptive fields	43
2.3.2.3	How can it be so naive: a model of knowl- edge	51
2.4	Conclusion	58

2.1 Out with the old, in with the new

As outlined in the previous part, the SMCT approach sets out to explain perceptive capabilities as a collection of properties a *naïve* agent can observe in its *sensorimotor flow*. As such, it challenges some otherwise accepted assumptions about said “naïveness”. In particular, what body of *a priori* knowledge about itself and its interaction with the environment may it possess? This chapter aims at proposing a formal framework suitable to the study of such problems.

It starts with a detailed examination of formal characteristics in previous proposals. We discuss their stance relative to prior knowledge of the agent as well as how it is formally represented by assumed mathematical structure, before highlighting some systematic limitations inherent in their representational choices. Then we develop in a second part our formal proposal. We show how it tackles previous limitations and expand on the duality of “points of view”, carefully distinguishing between description of sensorimotor interaction and that of sensorimotor experience. Throughout we carefully show how it allows for formulations equivalent to that found in of the work of V. Marcel (V. Marcel, Argentieri, and Gas, 2019) relative to the discovery of one’s workspace through refinement, thereby showing how it generalizes such approaches.




FIGURE 2.1: Representation of the model interaction used throughout the development

An illustrative model system A lasting aspect of the SMCT approach, as its name implies, is the *contingent* nature of the structural sensorimotor regularities. It therefore posits that the mechanisms supporting the emergence of perception do not depend on precise structure of the agent or even that of “space” as we may usually take for granted. Accordingly, the theoretical exposition we follow intendedly uses abstract formal concepts. To better

capture the concrete meaning discussed in typical cases, we will always circle back to (a variant of) the toy system illustrated in Fig. 2.1. It features an agent which motor capabilities are enacted through either the wheels supporting its frame or its robotic arm. The former allow it to perform locomotion in its environment, which varies along the horizontal direction. On the other hand the latter is a linkage of rigid sections, endowed with rotational joints: it can have the sections move about their joints (which axes of rotation lie perpendicular to the plane of the figure), accordingly altering the geometric configuration of the body of the agent. Furthermore, the agent is equipped with a camera on the terminal joint of its robotic arm. This camera allows it to sample its environment, which is depicted here as the horizontal colored strip at the top of the figure. The content of this strip, including the positions, the dimensions and the color of its blocks, is poised to change during the sensorimotor experience. The particular pattern of colors depicted therefore corresponds to one particular state of the environment, much as the precise placement of the agent in Fig. 2.1 is but one state at which it may find itself.

2.2 State of the framework and standing limitations

2.2.1 A framework for describing sensorimotor experiences

2.2.1.1 The sensorimotor diagram

Following classical robotics descriptions, the discussed formalisms take the sensorimotor interaction underlying an agent's experience to be a diagram of the MXS form, that is

$$\mathcal{M} \xrightarrow{f} \mathcal{X} \xrightarrow{\varphi_\epsilon} \mathcal{S} \quad (2.1)$$

where the $\mathcal{M}$, $\mathcal{X}$ and $\mathcal{S}$ denote sets and f , φ maps between these sets. More specifically, the *motor configuration space* $\mathcal{M}$ corresponds to the space of all possible configurations of the agent's actuators, i.e. the C-space/joint space. It is typically realized as the product space of the respective configuration spaces of each individual actuator (Farber, Tabachnikov, and Yuzvinsky, 2003). On the contrary, the *sensory configuration space* $\mathcal{S}$ mirrors that for the agent's sensors; it can similarly be represented as the tuple of all sensory states, which can range over a large number of diverse sensors for complex agents.

Between these spaces lies the representational *pose space* $\mathcal{X}$ which represents the set of all reachable poses of the agent's sensors, i.e. the workspace/task space. As such it is usually represented as a set of configurations in the ambient space containing the agent.

The sets $\mathcal{M}$ and $\mathcal{X}$ are related by the forward geometric model f , which maps motor configurations to the resulting end-effector poses. It is therefore encoding information about the body and motor structures of the agent, classically illustrated by the study of forward models of linkages. On the other side, $\mathcal{X}$ and $\mathcal{S}$ are related by the forward sensory function φ_ε which, for given environment state ε , maps a sensor pose $\mathbf{x} \in \mathcal{X}$ to the resulting sensory output $\mathbf{s}$. Correspondingly, φ encodes constraints imposed on the sensorimotor flow caused by the physical interaction between the agent and its environment.

It must be noted that for f and φ_ε to be functional relations imposes further constraints. For example, although one appeals to spatial intuition for $\mathcal{X}$ it first and foremost must act as a sufficient parameter space for φ_ε ; as such, $\mathcal{X}$ may be required to encode non spatial information. Consider the case of a camera sensor which actuators can adjust the aperture setting. For two configurations states $\mathbf{m}$ and $\mathbf{m}'$ which only differ in their aperture, we ought to observe that

$$(\varphi_\varepsilon \circ f)(\mathbf{m}') = \varphi_\varepsilon(f(\mathbf{m}')) = \mathbf{s}' \neq \mathbf{s} = \varphi_\varepsilon(f(\mathbf{m})) = (\varphi_\varepsilon \circ f)(\mathbf{m}) \quad (2.2)$$

for at least some environmental state ε . But in turn this implies that $\mathbf{x}' = f(\mathbf{m}') \neq f(\mathbf{m}) = \mathbf{x}$ since, being a proper function, φ_ε cannot take a single $\mathbf{x} \in \mathcal{X}$ to two distinct images $\varphi_\varepsilon(\mathbf{x}) \in \mathcal{S}$. In other words, states $\mathbf{x} \in \mathcal{X}$ necessarily contain aperture information. This illustrates the fact that $\mathcal{X}$ is primarily characterized by its “representational” property, that of condensing the complex and often highly redundant motor data of $\mathcal{M}$ into easily understandable structure; it *just so happens* to tightly correspond to spatial data in typical cases (Alban Laflaquière and Ortiz, 2019; V. Marcel, Argentieri, and Gas, 2019)

2.2.1.2 Interpretation in concrete cases

Fig. 2.2 shows how this representational choice accounts for our illustrative example. The state of each actuator is recorded as a component m_i (with $1 \leq i \leq 5$ since there are 5 distinct actuators), the vector of which forms




FIGURE 2.2: Fitting the mobile camera in the MXS scenario

the agent's motor configuration $\mathbf{m} = (m_1 \dots m_5)$. The pose $\mathbf{x}$ of the agent, then, is that of the camera at the end of its robotic arm. It is therefore parameterized by the planar coordinates (x, y) of the camera as well as its current orientation θ . The forward geometric model f is thus the mapping, determined by the specific embodiment of the agent¹, which takes the m_i to these 3 spatial parameters. Finally, the sensory output $\mathbf{s}$ of the agent is the image in its photoreceptive array of the color strip it is currently sampling. The figure makes apparent that for a fixed environment state ε , corresponding to a particular color content, the sensory output $\mathbf{s}$ depends only on $\mathbf{x} = (x, y, \theta)$ irrespective of the precise actuator configuration $\mathbf{m}$ that results in $\mathbf{x}$. This reflects the formal property that the forward sensory map φ_ε takes its input arguments in $\mathcal{X}$ instead of in $\mathcal{M}$.

2.2.2 A problem of bootstrapping: internal data and *a priori* knowledge

2.2.2.1 Mathematical structure as a formal correlate of knowledge

Since SMCT intends to ensure the minimality of the body of assumptions underlying our formal framework, one should certainly devise a bushel to

¹Here, for example, it depends on the lengths of the different sections of the linkage, the radius of the wheels and the dimensions of the frame of the robot.

characterize and order various collections of assumptions. This is made especially relevant by the theory questioning the essential and objective nature of formal structures such as internal representation (J K O'Regan and Noë, 2001) and external space (Alexander V. Terekhov and J. K. O'Regan, 2016). Of particular interest are the assumptions, often implicit, that come from selecting a particular mathematical structure to leverage. This is the subject matter of structuralist mathematics (Reck and Schiemer, 2020), which we turn to.

To illustrate our point, let us consider $\mathbb{R}^n$ the collection of n -tuples of real numbers, which is both very familiar and rich with various structures:

- Let us first consider its set structure, that is the structure that can be distinguished within the confines of set theory. In this context, the only basic “logical brick” one may employ is comparing any two elements of $\mathbb{R}^n$ for equality; equivalently, this is the structure that is preserved under bijections (i.e. set isomorphisms). It may be useful to think of it as labeling each element of $\mathbb{R}^n$ with an abstract symbol, with differing symbols having no discernible relation whatsoever.
- It can be endowed with a metric structure provided by a metric d (e.g. the Euclidean norm distance), where one may inquire as to the value $d(x, y)$ for any two $x, y \in \mathbb{R}^n$. This allows for distinguishing the underlying elements *up to isometry* (i.e. metric isomorphisms); this structure therefore defines a unique rigid geometry.
- Closely related are the cases of smooth structure (i.e. that of diffeomorphisms) and topological structure (that of continuous isomorphisms). These allow us to illustrate dual emerging hierarchies. On the one hand, the metric structure can be said to be *richer* than the smooth one which, in turn, is richer than the topological structure. Indeed, the metric structure of d corresponds to a unique smooth structure, which itself corresponds to a unique topological structure. Furthermore, this comparison is strict since the converse does not hold: it is known that the same topological structure can emerge from (infinitely) many distinct smooth structures (Taubes, 1987). This results in a classification of geometric hierarchies, which has previously been applied in a robotic context for qualifying robustness of validity conditions in signal processing (Censi, 2012).
- A more orthogonal structure is that brought forth by considering subtraction relations, which involves algebra. With it, one may make

sense of (and use) the classical $y - x$ syntax for any two $x, y \in \mathbb{R}^n$. It allows among others to enunciate that $P : (0, 1)$ has “the same relation” to $Q : (2, 0)$ that $R : (-4, 2)$ has to $S : (-2, 1)$, in that $(\vec{PQ} =) Q - P = S - R (= \vec{RS})$. As this notational choice hints at, this essentially yields the notion of affine spaces.

- Finally, observe that $\mathbb{R}^n$ is in particular the Cartesian product $\mathbb{R} \times \dots \times \mathbb{R}$ of n copies of $\mathbb{R}$, which itself can be endowed with such structures. This leads to distinguishing two broad types of structures. In the first we may only leverage the structure as possessed by the total tuple set $\mathbb{R}^n$, while in the second we may assume access to the component level. Let us assess the difference for the basic set structure. Let $x = (x_i)_{1 \leq i \leq n}$ and $y = (y_i)_{1 \leq i \leq n}$, the total equality comparison corresponds to querying the (binary) truth value $\delta_{x,y}$ of the one assertion

$$“\forall i \in \{1, \dots, n\}, x_i = y_i” . \quad (2.3)$$

Comparatively, being able to compare x and y at a component level is equivalent to separately accessing all n truth values $\delta_{x,y}^i = \delta_{x_i,y_i}$, for $1 \leq i \leq n$. By definition, it is clear that the collection of the $\delta_{x,y}^i$ entirely determines $\delta_{x,y}$. Indeed, it suffices that one of the $\delta_{x,y}^i$ indicate difference for $\delta_{x,y}$ to do so; conversely, whenever all $\delta_{x,y}^i$ denote equality we may be certain that $x = y$. On the contrary, there is no definitive way to conclude as to the $\delta_{x,y}^i$ in all cases where we know only $\delta_{x,y}$. This happens because two (total) vectors may be different irrespective of the number of their components which coincide, as long as not all of them do. This phenomenon is a direct consequence of $\delta_{x,y}$ being the conjunction of the $\delta_{x,y}^i$ for $1 \leq i \leq 2$. It shows that allowing component level comparison leads to a richer structure, which conversely represents more restrictive assumptions².

2.2.2.2 From my point of view...

One crucial point of the SMCT approach lies in investigating and characterizing which data, when operated upon with a given set of computations, leads

²While the comparison between the structure of a product (of sets) and the corresponding product structure has been made solely for set structures (i.e. for equality), it can be further developed and generalized to other contexts using the language of *category theory* (see (Mac Lane, 1978) for a reference) precisely *via* the notion of categorical product. In particular, all other structures mentioned here have such products in their internal logic.

to the discovery of relevant structure (i.e. perceptive capabilities) in the sensorimotor flow. It is therefore no surprise that a great deal of attention has been devoted to the choice of these starting sets, corresponding to notions of what it precisely means for an agent to be *naive*. As such, a distinction is usually made between two *points of view*, respectively external and internal. The former is that of the observer (e.g. the designer of the robot) and involves any object or structure which is relevant to the description or the understanding of the sensorimotor experience. The latter, instead, aims at emulating the naive agent. As such, two different sorts of considerations emerge: adding assumptions about external objects restricts the scope of *physical systems* to which the framework applies, while assumptions that apply to internal ones instead narrow that of *naive agents*. Furthermore, this leads to dual sorts of results and their proofs. Those leveraging (both internal and) external objects demonstrate *existence*; those who can be deduced from purely internal propositions also demonstrate *discoverability* (or computability) by a naive agent.

The MXS diagram outlined in 2.2.1.1 illustrates the two sorts of objects. Generally, both the motor and sensory sets $\mathcal{M}$, $\mathcal{S}$ are assumed to be known *a priori* to the agent, on the respective basis of proprioceptive capabilities and that accessibility is precisely what makes sensation. This knowledge is at least that of (distinct) set structure, that is the agent may compare either any two motor states ($\mathbf{m}$ and $\mathbf{m}'$) or sensory ones ($\mathbf{s}$ and $\mathbf{s}'$) for equality. Additionally, the agent may generally issue motor commands, in particular to reach any single motor state $\mathbf{m}' \in \mathcal{M}$. On the contrary, owing to its intrinsic relation to ambient space the pose space $\mathcal{X}$ (initially) comes under the external point of view, as well as its relating maps f and φ_ϵ . This distinction motivates the search for internal representations of $\mathcal{X}$ such as found in (Alban Laflaquière and Ortiz, 2019; V. Marcel, Argentieri, and Gas, 2019).

In this aspect, one can observe significant divergences between lines of work. A. Laflaquière typically embraces the specific robotic implementation details: sensory and motor states are routinely defined as tuples of $\mathbb{R}^N$, with N the number of relevant components. This is an important internal assumption in as much as their agents have the *a priori* capability to leverage its *metric structure* to perform interpolation and clustering (Laflaquière et al., 2015; Alban Laflaquière, 2013, 2017) or its *algebraic* one in predictive neural networks (Alban Laflaquière and Hafner, 2019; Alban Laflaquière and Ortiz, 2019). Instead V. Marcel opts for the bare minimum with agents only able to

compare motor (or sensory) states for a given equivalence relation, typically equality (Valentin Marcel, 2020); this generally does not even include access to the component level of states, which falls under product structure.

2.2.2.3 The nature of the motor set

The motor set $\mathcal{M}$ is a core object of the formal results, owing to it being the starting point of the “motor” part of the sensorimotor interaction. As previously discussed, it is introduced as the familiar C-space of joint actuator configurations. However, in the SMCT context it fulfills three important functions:

- The entire sensorimotor interaction is defined in the $\mathcal{M} \rightarrow \mathcal{X} \rightarrow \mathcal{S}$ diagram, which intrinsically requires the motor set to be a *parameter set* of the sensorimotor interaction. Because there is only for hidden parameters ε in $\varphi_\varepsilon : \mathcal{X} \rightarrow \mathcal{S}$, let us consider a sensorimotor experience with fixed environment. In this case, any two points of the sensorimotor interactions resulting in distinct sensory outputs $\mathbf{s}' \neq \mathbf{s} \in \mathcal{S}$ must originate from distinct $\mathbf{m}' \neq \mathbf{m} \in \mathcal{M}$. This was previously formulated as the property that the states of $\mathcal{M}$ entirely define the interaction the agent has with its environment.
- As its name implies, the motor set $\mathcal{M}$ is taken to be that which can be *acted upon* by the agent. It supports the motor capabilities of the agent, which motor commands are represented as changes in state $\mathbf{m} \in \mathcal{M}$. Nearly all previous experiments assumed the corresponding equivalence: the sequence of environment states ε is entirely independent of the motor commands of the agent, whereas the sequence of motor states is entirely determined by the same commands.
- Finally, the motor state $\mathbf{m} \in \mathcal{M}$ is assumed to be available to the internal knowledge of the agent as its *proprioceptive input*. This entails two noteworthy properties. First, because of the first point it requires the proprioceptive feedback to be *complete*: the agent is assumed to internally distinguish any two distinct motor configurations. Second, the exteroceptive and proprioceptive feeds are assumed to be separated *a priori*. This provides grounds to dual *sensory* and *motor* flows, and from that to the approach of structuring the former using the latter.




FIGURE 2.3: Distinct motor configurations related by locomotion.

2.2.3 Limitations

Having discussed in finer detail the previous formal frameworks, we now expose several of their important limitations which we try to address going further. The first two, concerning the indescribability of locomotive agents and the staticity of the description, have previously been recognized and discussed (Valentin Marcel, 2020).

2.2.3.1 Locomotive motives

The several functions ascribed to the motor space $\mathcal{M}$ have the consequence of eliminating from the scope of the formalism any agent that performs *locomotion*, that is agents which can end situated at several distinct places in spaces while having the same motor configuration. This closely follows from $\mathcal{M}$ being both the internal motor states of the agent (as reported by its proprioception) as well as it the *parameter set*. Let us get back to our toy system and consider the particular situation depicted in Fig. 2.3 The same agent is pictured at two different “states” of its sensorimotor interaction with a fixed environment. Importantly these states are obtained from one another *via* locomotion, corresponding to them sharing the same *internal* motor state $\mathbf{m} \in \mathcal{M}$. Instead, they differ only by the places in space x and x' at which the agent is located. It is clear that whichever the means, we expect to find

that these situations result in distinct sensory outputs $\mathbf{s}' \neq \mathbf{s} \in \mathcal{S}$ since different parts of the environment fall in the field of view of the agent. However, this raises further issues with the previous definitions: if we accepted that these two states correspond to the same motor configuration $\mathbf{m}' = \mathbf{m} \in \mathcal{M}$, then so must we do that they end up mapped to the same sensory output $\mathbf{s}' = \varphi_\varepsilon(f(\mathbf{m}')) = \varphi_\varepsilon(f(\mathbf{m})) = \mathbf{s}$. By contraposition, starting from the assertion that $\mathbf{s}' \neq \mathbf{s}$ yields that the motor states $\mathbf{m}$ must differ between situations; but this then runs at odds with the interpretation of $\mathbf{m}$ as proprioceptive feedback, which importantly motivates $\mathbf{m}$ being an *internal* object.

As has previously been argued, we can try and salvage the description using the equivalent description of compensation. It corresponds to taking the point of view of these situations occurring *as if* the agent had not moved at all, and the environment had performed the inverse displacement. In effect, the figure should then depict side by side ε and ε' instead of the unique environment state ε ; each would correspond to one of the two (formally) distinct environment states the agent is interacting with at “fixed” configuration $x' = x$. However, even putting aside the apparent clumsiness brought forth, the argument only paints a partial picture of a solution. Indeed, while the compensability point of view holds for describing how motor, environment, and sensory states are related at any point of the sensorimotor interaction, it fails to provide an account of the agent’s experience by *a priori* erasing the distinction between exterior displacements and those caused by the agent. This formal trick also causes the second function of $\mathcal{M}$ to cease to hold: motor commands of the agent may now modify the environment state ε so that it is much more difficult to talk of fixed environment and separate causes of change. Finally, further complications arise when we consider multiple agents: take the same schema, except interpret it now as describing two identical agents simultaneously experiencing in a shared environment ε , and apply the argument presented in the previous paragraph. We find that either their motor states must differ –but that is precisely what the compensation operation set out to avoid– or we must distinguish the environment state ε with which each agent is interacting, contrary to the important idea that it is a shared “outside”.

2.2.3.2 Stateful descriptions make for static

A second significant constraint on $\mathcal{M}$ is imposed by the conjunction of it parameterizing the interaction and it being precisely which the agent acts upon.

It being the set of parameters makes it necessary that it contain information that is “positional” in nature –as opposed to kinematic or dynamic. On the previous illustrated case, this corresponds to the motor state $\mathbf{m} \in \mathcal{M}$ being a tuple of joint angles rather than angular velocities; and indeed, excluding the case of accelerometers³ we would not expect to be able to map a velocity (or torque for that matter) motor state vector to a resulting sensation⁴. Instead, the sensory output would change during the movement generated by the dynamic motor state, begging the question of the choice of a corresponding $\mathbf{s} \in \mathcal{S}$.

Consequently, describing the system in this fashion implicitly assumes that the agent can issue commands directly in terms of angular configurations. In the cases where its implementation details require a robotic agent to be velocity or torque controlled, integration must therefore be preemptively performed to yield motor state $\mathbf{m}$, raising the question of the *a priori* capabilities to compute such integrations.

This property also makes any assertion about two displacements being “similar” cumbersome. Indeed if one only considers the configurations (either spatial or motor) at which the sensorimotor experience takes place, then there is no special relationship between e.g. the coordinate values x_1 and $x'_1 = x_1 + \Delta x$ in our illustrative example. However the experience shows that observable properties in the sensorimotor flow are instead directly linked to this Δx quantity, such as the retinal smears found in (Alban Laflaquière, 2017; Le Hir, Sigaud, and Alban Laflaquière, 2018; Montone, J. Kevin O'Regan, and Alexander V. Terekhov, 2015). But then, following this sort of parameterization one lacks the way of efficiently characterizing which other pairs x_2, x_3 have the same relationship so that $x_3 = x_2 + \Delta x$. In particular other works –such as those previously mentioned– had the agent leverage this structure without making explicit how and to which extent it had access to it.

2.2.3.3 The tradeoff of numericity

The endeavour of limiting as best as possible to which structure the agent has *a priori* access has the predictable downside of reducing the breadth of valid problem-solving approaches. This is generally made manifest with the choice to assume that states are coded as *numeric* values (or tuples thereof),

³which incidentally could be argued to belong to proprioception as much as to exteroception

⁴Note that, even for event cameras, one needs to know where it currently *is* in addition to how it is moving to determine whether an event is happening.

and of special influence with that of whether to use natural structures of $\mathbb{R}^N$ in $\mathcal{M}$ and $\mathcal{S}$. Indeed, most developments which try to achieve robustness and scalability do so *via* extrapolation and clustering (Alban Laflaquière, 2017; Le Hir, Sigaud, and Alban Laflaquière, 2018; V. Marcel, Argentieri, and Gas, 2017). Nevertheless these techniques require referring to preexisting external metrics, which constitutes assumption about *a priori* knowledge as discussed in 2.2.2.1. These experiments importantly show that although the results may formally hold with less restrictive structure, they may not be tractable *in practice* because the spaces surveyed are too complex. Furthermore, this approach is made more debatable in the SMCT context because it operates on the underlying assumption that the closeness of perceptual items (which e.g. clustering strives to achieve) is faithfully represented by the closeness of representational code (as denoted by the metrics on $\mathcal{M}$ and $\mathcal{S}$). This forms in essence a claim that the code used by the agent is naturally fitting to the processes it is intended to represent, a claim which is debated at great length from the onset of the theory (J K O'Regan and Noë, 2001; J. K. O'Regan and Block, 2012).

On the other hand, contributions which observe stricter restrictions, while more compliant in this aspect, generally suffer from greatly reduced applicability. Foregoing otherwise innocuous structure leads to systems which do not scale (Philipona, J K O'Regan, and Nadal, 2003), are not robust to noise (V. Marcel, Garcia, et al., 2015) or cannot evaluate the relevancy of context (Maye and Andreas K Engel, 2013, 2011). These shortcomings consequently narrow these contributions to be proofs of concept. In each of these cases, contrast with the typical numeric case makes clear what is “missing”. However, it remains to be determined how we should approach bridging this particular gap; that is how we should derive structures (both metric and algebraic) akin to that of numeric sets, structures which would fulfill the same functions while being *internally* computable.

2.3 Towards a formalism of sensorimotor flow

We now turn to introducing and detailing the formal framework we propose. Our exposition is mostly twofold, corresponding to the dual points of view previously outlined. First, we describe the external picture, corresponding to the parameterization of the sensorimotor interaction; it is naturally followed by the corresponding internal account of the sensorimotor experience. Both

parts are further subdivided between description of the interior (or motor) objects and that of exterior (or sensory) ones. It is concluded with a detailed discussion about what we choose to consider part of the *internal* point of view (see 2.2.2.2), in which we finish showing how our framework proposal can be made to generalize the one presented before.

2.3.1 Reparameterizing the description

We previously discussed how the definition $\mathcal{M}$ confounding system parameterization and bodily state of the agent causes issues in the expressiveness of the formalism. Therefore we turn to a “supermotor” set acting as parameter set, one which encodes not only the postural state of the agent but also accounts for the localization of its embodiment in space. This in turn is made possible by the introduction of *ambient space*, which we explicitly define in a departure from previous frameworks.

2.3.1.1 The *where* and the *what*: environmental concerns in space

Giving some space to the interaction The notion of exterior space has been a staple of the SMCT approach, with many of the first proofs of concept attending to the naive determination of its properties (e.g. dimensionality, isotropy...). This makes somewhat surprising the observation that the frameworks used do not in fact characterize said space, instead relying on some environment state similar to the ε variable. The substitution finds its ground in that it is merely the environment the agent is interacting with, whereas space itself is intangible. This perspective culminates in a radical antirealist undercurrent of the SMCT that challenges the objective nature of external space (Alexander V. Terekhov and J. K. O'Regan, 2016).

Notwithstanding these provocative considerations, our most basic intuition runs so: the sensorimotor experience is a process in which an *embodied* agent meets with its physical environment by way of its sensorimotor apparatus, a process further occurring *in* an embedding immutable structure which is precisely the “ambient space” of the system. This embedding manifests *via* several structure constraints. First, motor commands of the agent act on the localization of its embodiment in space, modifying both its posture and its position. Second, the laws governing its sensory apparatus are spatially mediated, so what we can talk of *where* an agent is directing its sensors. Third, spatiality not only mandates that things be localized somewhere but

also distinguishes certain “regular” laws of motion, which we internally experience as rigid. Importantly, these constraints all involve the same shared space, that is a *common medium* in which the agent’s motor and sensors operate. This commonality, while seeming trivial, then supports some contingencies a naive agent is supposed to find in its sensorimotor flow.

It remains to be determined how we should formally characterize this space. Indeed an important argument can be made that, contrary to the intuition conveyed by the language used, spaces cannot be reduced to their constituent “points”. Let us illustrate this point on (real) n -dimensional spaces which cases with $0 \leq n \leq 3$ certainly feel most familiar. In a straightforward way the set of points corresponding to n -dimensional space is taken to be $\mathbb{R}^n$, that is the collection of all n -tuples of reals $(x_1, \dots, x_n)$. This correspondence is remarkably robust in that it intuitively carries geometrical insight. For example, each of the coordinates x_i of the tuple uniquely corresponds to a dimension of n -d space; moreover, many geometric constructions (such as length, distance, angles...) can be defined in terms of these coordinates. However none of these notions can actually follow from the bare $\mathbb{R}^n$ set itself. A classical result, illustrated in Fig. 2.4, shows that $\mathbb{N}^2$ and $\mathbb{N}$ share the same structure as sets. More precisely it shows the existence of

$$f : n \in \mathbb{N} \mapsto f_n = (k_n, l_n) \in \mathbb{N}^2$$

that is an isomorphism of sets, that is a bijection. In this case, bijectivity is identical to f “enumerating” tuples of whole numbers such that each tuple $(k, l) \in \mathbb{N}^2$ appears once and exactly once in the sequence of the f_n . It then follows that $\mathbb{N}^3$ also shares its set structure with $\mathbb{N}$: one can bijectively “fold” three natural numbers into one by way of

$$(k, l, m) \in \mathbb{N}^3 \mapsto (f^{-1}(k, l), m) \in \mathbb{N}^2 \mapsto f^{-1}(f^{-1}(k, l), m) \in \mathbb{N}.$$

This can clearly be generalized further to $\mathbb{N}^k$ for any $k \in \mathbb{N}$, which are all equivalent as sets to $\mathbb{N}$ as a result. On the other hand a similar argument, albeit slightly more technical can be made to show that this property holds when we replace $\mathbb{N}$ with $\mathbb{Q}$ or even $\mathbb{R}$ (Nicolay and Simons, 2014)⁵. Therefore in every n -d space there exists a way to faithfully denote each point of space

⁵Otherwise, set theory (specifically cardinal set arithmetic) also shows in much greater generality that for infinite sets X , there are bijections between X^n and X for every whole number $n \geq 1$ (Berger, 2009; Morgan, 2006). However, the corresponding construction is far less illustrative than that made above.




FIGURE 2.4: Illustration of a way a bijective mapping $\mathbb{N} \xrightarrow{\sim} \mathbb{N}^2$ can be constructed. Starting at $(0,0)$ for $f(0)$, it follows the red arrows and enumerates the points of the grid (which correspond to tuples $(m,n) \in \mathbb{N}^2$) as it encounters them. Because of the way they diagonally map out $\mathbb{N}^2$ we are assured that f will miss no possible output value.

(otherwise usually denoted $(x_1, \dots, x_n) \in \mathbb{R}^n$) by a single numeric label $\lambda \in \mathbb{R}$. Nevertheless, this runs very much unlike our intuition of what “space” should be. Indeed, for distinct values of n , n -dimensional spaces should not be “equivalent” if only on the basis that they have (to us) perceivably distinct dimensions. Therefore, we cannot hope to find the structure of spaces in the mere collection of their points.

An alternative viewpoint is to characterize spaces *via* collections of mappings defined over them. One such approach was developed in Klein’s Erlangen Program (Klein, 1893) (and further extended *via* Cartan geometry (Sharpe, 2000)) and has since proved very fruitful. It shows how a correspondence can be made between “geometries” and their corresponding transformations. Roughly speaking it gives a unified account of how various geometrical structures (and the distinct “spaces” that can carry these structures) result in different algebraic properties of their geometric motions. As a formal object, we therefore choose to define the *external space* in which the experience takes

place as some set $\mathcal{X}$ together with its geometry automorphism group $\mathcal{G}(\mathcal{X})$. The elements τ of $\mathcal{G}(\mathcal{X})$ are *structure preserving transformations* $\mathcal{X} \rightarrow \mathcal{X}$, whichever the choice of structure considered (see 2.2.2.1 for some such geometrical structures and corresponding isomorphisms). Our experience therefore corresponds to the case where $\mathcal{X}$ is the 3-d real space $\mathbb{R}^3$ endowed with its natural metric structure. Its corresponding automorphism group $\mathcal{G}(\mathcal{X})$ is that of transformations $\mathbb{R}^3 \rightarrow \mathbb{R}^3$ that preserve distances, which precisely characterizes the Special Euclidean group $SE_3(\mathbb{R})$ of rigid transformations. Importantly, for a given space as underlying set of points there can be several candidate geometrical structures, each corresponding to its transformations $\mathcal{G}(\mathcal{X})$; we previously illustrated this point when discussing hierarchies of structures in 2.2.2.1.

It should be noted that this viewpoint is remarkably reminiscent of that previously taken by H. Poincaré (see (Poincaré, 1895) and especially Chapter 4 in (Poincaré, 1902)), if maybe more explicit. It is central to the role played by *compensable transformations*. Indeed, the informal argument Poincaré makes identifies compensability as a property of a certain distinguished class of spatial transformations. In particular, he equates compensable transformations of *our* environment with rigid motions of solid objects, from which he draws his conclusion “If there were no solid bodies, there would be no geometry”. On the other hand, the contrasting thought experiment in the “dilat-ing” sphere⁶ finds another collection of transformations to be compensable, including certain particular deformations. In both cases, the argument characterizes compensability as a property of “displacements”⁷, that is precisely those geometry preserving transformations of space. Consequently, in this approach it is as much transformations of $\mathcal{G}(\mathcal{X})$ as it is $\mathcal{X}$ itself that one should strive to learn about by way of their experience⁸.

Highlighting the environmental factors of experience Once space provides the notion of “where”, the environment is then roughly the “what”

⁶which corresponds to the *Poincaré disk model*, a particular instance of (hyperbolic) non-euclidean geometry. Interestingly, the study of such spaces (and the comparison with previously known geometries) was precisely the primary object of the Erlangen Program; as a result, it easily fits under the assumptions presented here despite its apparent strangeness.

⁷This context, in turn, gives all its weight to his statement “Thus is defined [...] a particular class of phenomena we will call displacements. *It is the laws of these phenomena which is the object of geometry.*” (in (Poincaré, 1902), translation ours)

⁸It also explains *a posteriori* why certain unintuitive results are in fact no inaccuracy, such as the set of compensable transformations (in rigid 3-d space) being 6-dimensional (instead of 3). They should be recognized as a property of “displacements”, by which the argument made here lends epistemic value to their discovery.

that happens to lie in the various regions of space. Accordingly, we choose to model states of the environment after functions $p : \mathcal{X} \rightarrow \mathcal{P}$, that is valuations that to each point in space $x \in \mathcal{X}$ ascribe a physical state $p(x) \in \mathcal{P}$. It should be noted that by the same curious account we made for external space, this physical state can be considered an abstraction limited to external understanding. The agent never directly probes these states; instead, they always only become accessible after transduction by its sensory apparatus. However, the interesting structure is that states of the environment can be *localized* to regions of space, corresponding to the usual restriction $\cdot|_{\cdot}$ of functions. These properties essentially hint at us defining environment states as a sheaf $\mathcal{E}$ on $\mathcal{X}$ ⁹, where more precisely any single environment state ε is a global section of $\mathcal{E}\mathcal{X}$. The context of sheaves then explicitly provides restriction operations $\mathcal{E}\mathcal{X}' \rightarrow \mathcal{E}\mathcal{X}''$ whenever applicable, mirroring the properties of the usual restriction of functions. Maintaining consistency, we will denote $\mathcal{E}\mathcal{X}'$ the collection of all restrictions of environmental states to subspace $\mathcal{X}' \subset \mathcal{X}$. Note that although this formulation has the benefits of increased generality and avoiding reference to an arbitrary $\mathcal{P}$, in all practical cases we will here present one may equivalently think of the usual functional case $\varepsilon : \mathcal{X} \rightarrow \mathcal{P}$.

This separation between space and environment makes apparent another structure caused by the motions of space $\mathcal{G}(\mathcal{X})$. From the previous definitions, we can now for any state of the environment ε talk of the state that is obtained “from ε following motion τ ”. Indeed, because the spatial component is free we can now formulate the resulting motion of environment as the mapping

$$\begin{aligned} \tau_{\mathcal{E}} : \mathcal{E}\mathcal{X} &\rightarrow \mathcal{E}\mathcal{X} \\ \varepsilon &\mapsto \tau\varepsilon \end{aligned} \tag{2.4}$$

where the $\tau\varepsilon$ environment state is defined by $\varepsilon \circ \tau$, or at a pointwise level $\forall x \in \mathcal{X}, \tau\varepsilon(x) = \varepsilon(\tau x)$.

⁹The reader can find useful introductory material on the subject of sheaves in (Tennison, 1975) or (Urbanik, n.d.). They should note however that most expositions are deeply interleaved with that of further applications (such as theoretical logic or algebraic geometry), making them notoriously arduous. In particular we will only make superficial use of the tool as a convenient internal axiomatic framework, and do not require any degree of familiarity with the results of sheaf theory.




FIGURE 2.5: Modification of the motor parameterization to circumvent the issue of locomotion.

2.3.1.2 Motor data accounting for locomotion

The previous discussion about the functions previously attributed to the motor set (see 2.2.2.3) and the corresponding limitation to fixed-base agents (2.2.3.1) calls for reparameterizing the motor description of the formalism. In particular it points to the need for accounting for the localization of the agent in space, localization which allows distinguishing configurations that are related to one another by locomotion.

By this account, consider the set $\mathcal{B}$ of all configurations the embodiment of the agent may assume in its ambient space $\mathcal{X}$. Each of these configurations $\mathbf{b} \in \mathcal{B}$ corresponds to a posture of the agent, mirroring the previous motor configuration $\mathbf{m}$; however, it is now enriched by the data of *how* this posture is embedded in $\mathcal{X}$, that is “where” the agent stands. The resulting object $\mathbf{b}$ can therefore be sketched as a tuple $\mathbf{b} = (\mathbf{m}, \tau)$, where τ is the *pose* (i.e. position and orientation) of the configuration. Subsequently, we refer to $\mathbf{m}$ as *internal* motor configuration and $\mathbf{b}$ as *absolute* motor configuration.

This decomposition is in line with Fig. 2.5 where the added external reference is denoted by the τ position variable. Three distinct (absolute) configurations $\mathbf{b}, \mathbf{b}', \mathbf{b}''$ are depicted. As before $\mathbf{b}'$ can be obtained from $\mathbf{b}$ by locomotion, while $\mathbf{b}''$ corresponds to a distinct internal motor configuration $\mathbf{m}''$. But now it clearly follows from the previous examinations that the addition of this position data allows for a functional relation, should one replace $\mathcal{M}$ by $\mathcal{B}$ in the $\mathcal{M} \rightarrow \mathcal{X} \rightarrow \mathcal{S}$ diagram. Indeed, the two configurations which previously corresponded to a same internal motor state $\mathbf{m}$ are now denoted by $(\mathbf{m}, \tau') = \mathbf{b}' \neq \mathbf{b} = (\mathbf{m}, \tau)$, so that each can be mapped to its corresponding output $\mathbf{s}'$ and $\mathbf{s}$ by a proper function.

Solving the problem of external references It should be mentioned that the aforementioned construction is largely arbitrary. It is especially cumbersome to try and provide satisfactory reasoning for determining the τ in $\mathbf{b}$: it is not clear from the previous objects what exactly constitutes a pose; moreover, our intuition signals that we may only hope to give such a τ only up to a choice of reference frame. To illustrate this point, consider again Fig. 2.5 and imagine that we replace the measurement of position x with that of the left wheel of the robot: clearly this is a different but equivalent parameterization, and thus it should be equally as valid. The same process can be made more involved if we instead label the position of the agent by the coordinate of its middle joint (corresponding to m_2): because this point of the body is not rigidly linked to its physical frame, changes in posture can also result in changes of the associated reference coordinate.

As for environment states, the construction can however be made more intrinsic to eschew this issue. Indeed, on the basis that absolute configurations $\mathbf{b}$ are configurations of the embodiment in $\mathcal{X}$ we assume that some may be related by its motions $\tau \in \mathcal{G}(\mathcal{X})$. This relation on $\mathcal{B}$, denoted as $\mathbf{b}' = \tau'\mathbf{b}$, corresponds to the case where the transformation of embodiment of the agent from $\mathbf{b}$ to $\mathbf{b}'$ is equivalent to its spatially moving along τ' . Should one assume the previous decomposition available, it could then be concretely defined as $\tau'\mathbf{b} = \tau'(\mathbf{m}, \tau) = (\mathbf{m}, \tau'\tau)$. In the familiar case of rigid geometry, it therefore expresses that two configurations are related to one another by locomotion, that is by a combination of translation and/or rotation. This is also pictured on Fig. 2.5, where $\mathbf{b}'$ is obtained from $\mathbf{b}$ by sliding from τ to τ' . By definition, two absolute configurations $\mathbf{b}$ and $\mathbf{b}'$ are thusly related by any space automorphism $\tau \in \mathcal{G}(\mathcal{X})$ if and only if one can be obtained from another by locomotion. Therefore, consider the relation $\sim_{\mathcal{G}(\mathcal{X})}^{\mathcal{B}}$ defined by

$$\forall \mathbf{b}, \mathbf{b}' \in \mathcal{B}, \mathbf{b} \sim_{\mathcal{G}(\mathcal{X})}^{\mathcal{B}} \mathbf{b}' \Leftrightarrow \exists \tau \in \mathcal{G}(\mathcal{X}) \text{ such that } \mathbf{b}' = \tau\mathbf{b}. \quad (2.5)$$

We can show that it inherits strong properties from the group structure of the relating transformations $\tau \in \mathcal{G}(\mathcal{X})$. First there is some particular transformation $\text{id}_{\mathcal{G}(\mathcal{X})} \in \mathcal{G}(\mathcal{X})$ under which all of space is invariant, so that

$$\forall \mathbf{b} \in \mathcal{B}, \mathbf{b} = \text{id}_{\mathcal{G}(\mathcal{X})}\mathbf{b}.$$

In particular this implies that $\mathbf{b} \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}$ for any absolute configuration $\mathbf{b}$, by

which we say $\sim_{\mathcal{G}(\mathcal{X})}$ is *reflexive*. Additionally, the reversibility of transformations $\tau \in \mathcal{G}(\mathcal{X})$ yields

$$\forall \mathbf{b}, \mathbf{b}' \in \mathcal{B}, \forall \tau \in \mathcal{G}(\mathcal{X}), \mathbf{b} = \tau \mathbf{b}' \Leftrightarrow \mathbf{b}' = \tau^{-1} \mathbf{b}.$$

Accordingly, whenever we have $\mathbf{b} \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}'$ then it must necessarily be true that $\mathbf{b}' \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}$ (and *vice versa*). This property is that of relation $\sim_{\mathcal{G}(\mathcal{X})}$ being *symmetric*. Finally, the relation can be chained along the composition of transformations of $\mathcal{G}(\mathcal{X})$ as per

$$\forall \mathbf{b}, \mathbf{b}', \mathbf{b}'' \in \mathcal{B}, \forall \tau, \tau' \in \mathcal{G}(\mathcal{X}), \mathbf{b}'' = \tau' \mathbf{b}' \text{ and } \mathbf{b}' = \tau \mathbf{b} \Rightarrow \mathbf{b}'' = (\tau' \tau) \mathbf{b}.$$

In terms of relations, this means that whenever $\mathbf{b} \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}'$ and $\mathbf{b}' \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}''$ it must hold that $\mathbf{b} \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}''$; we call this particular property the *transitivity* of $\sim_{\mathcal{G}(\mathcal{X})}$. These three properties make $\sim_{\mathcal{G}(\mathcal{X})}$ into an *equivalence relation*¹⁰. In this context a particular object of interest is that of *equivalence classes*, that is collections of elements that are related to one another. Indeed, for any particular absolute configuration $\mathbf{b}_0 \in \mathcal{B}$ consider the collection of related configurations $\mathbf{b} \in \mathcal{B}$ defined by

$$[\mathbf{b}_0]_{\mathcal{G}(\mathcal{X})} = \left\{ \mathbf{b} \in \mathcal{B} \text{ such that } \mathbf{b} \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}_0 \right\}.$$

From the three aforementioned properties of $\sim_{\mathcal{G}(\mathcal{X})}$, it can be shown that any two configurations $\mathbf{b}, \mathbf{b}'$ in this subset are not only related to $\mathbf{b}_0$ *via* $\sim_{\mathcal{G}(\mathcal{X})}$ but also to one another. Moreover, any given configuration $\mathbf{b} \in \mathcal{B}$ must appear in such a $[\mathbf{b}_0]$ for some $\mathbf{b}_0$: in fact, it must precisely appear in $[\mathbf{b}]_{\mathcal{G}(\mathcal{X})}$. Let us denote $\mathcal{B}/\mathcal{G}(\mathcal{X})$ the collection of equivalence classes in $\mathcal{B}$ relative to the relation $\sim_{\mathcal{G}(\mathcal{X})}$, which we call the *quotient of $\mathcal{B}$ by $\sim_{\mathcal{G}(\mathcal{X})}$* . What we have just seen is that

- the equivalence classes *partition* the total set $\mathcal{B}$: $\bigcup_{A \in \mathcal{B}/\mathcal{G}(\mathcal{X})} A = \mathcal{B}$ and $\forall A, B \in \mathcal{B}/\mathcal{G}(\mathcal{X})$, either $A = B$ or $A \cap B = \emptyset$;
- for any two absolute configurations $\mathbf{b}, \mathbf{b}' \in \mathcal{B}$, $\mathbf{b}' \in [\mathbf{b}]$ if and only if $\mathbf{b}' \sim_{\mathcal{G}(\mathcal{X})} \mathbf{b}$, in which case we even have $[\mathbf{b}'] = [\mathbf{b}]$.

¹⁰This exposition is motivated by the fact that we will make regular use of equivalence relations in this work; the reader should find this enumeration of their properties as well as the construction of equivalence classes sufficient to follow developments later. Should they search for a more thorough reference nonetheless they would find it in most set theory textbooks such as (Hrbacek and Jech, 1999) which also provides concrete illustrations of the notion.

Finally, we will say that configuration $\mathbf{b} \in \mathcal{B}$ is a *representative* of class $A \in \mathcal{B}/\mathcal{G}(\mathcal{X})$ whenever $[\mathbf{b}] = A$. In particular an equivalence class may admit any number of representatives, which are necessarily all related by $\sim_{\mathcal{G}(\mathcal{X})}$ to one another.

Importantly, in our specific case Eq 2.5 shows that the corresponding equivalence class of any configuration $\mathbf{b} \in \mathcal{B}$ is

$$[\mathbf{b}]_{\mathcal{G}(\mathcal{X})} = \{\mathbf{b}' \in \mathcal{B} \text{ such that } \mathbf{b}' = \tau\mathbf{b}, \tau \in \mathcal{G}(\mathcal{X})\}.$$

This enunciates how equivalence classes precisely group together all configurations related by general locomotions; in intuitive (if external) terms, the class $[\mathbf{b}]_{\mathcal{G}(\mathcal{X})}$ is the set of all configurations which share the same posture with $\mathbf{b}$. But we can therefore adopt the converse reasoning provides internal logic to characterize “posture” and “pose”. Indeed, consider the “posture” $\mathbf{m}$ associated to an absolute configuration $\mathbf{b} \in \mathcal{B}$ to be its equivalence class for this locomotion relation. Then, for each of these classes $\mathbf{m} \in \mathcal{B}/\mathcal{G}(\mathcal{X})$ choose one particular representative $\mathbf{b}_\mathbf{m}^0 \in \mathcal{B}$, that is a configuration that verifies $[\mathbf{b}_\mathbf{m}^0]_{\mathcal{G}(\mathcal{X})} = \mathbf{m}$; each of the $\mathbf{b}' \in \mathcal{B}$ that share posture with $\mathbf{b}$ are by definition related to $\mathbf{b}$ *via* a locomotion $\mathbf{b}' = \tau\mathbf{b}$ for some $\tau \in \mathcal{G}(\mathcal{X})$. This entails that $\mathbf{b}'$ can be described as the tuple $(\mathbf{m}, \tau)$. Importantly this identifies the ambiguous pose parameter with transformations of space $\tau \in \mathcal{G}(\mathcal{X})$, and that of reference positions with those $\mathbf{b}_\mathbf{m}^0 = (\mathbf{m}, \text{id}_{\mathcal{G}(\mathcal{X})})$ representatives of equivalence class.

Revisiting the functions of $\mathcal{B}$ This $\mathcal{B}$ space is explicitly characterized by its parameter function, and as such we intend it to replace $\mathcal{M}$ in the $\mathcal{M} \rightarrow \mathcal{X} \rightarrow \mathcal{S}$ diagram. The discussion provided in 2.2.3.2 therefore shows that it is intrinsically required to provide positional parameters, that is static snapshots of the way the agent is embodied in ambient space. Furthermore, by separately defining $\mathcal{B}$ and the space of environment states $\mathcal{E}\mathcal{X}$ we tacitly assume that these states should be considered independent. This especially lends itself to descriptions where any absolute configuration $\mathbf{b}$ can exist in any environment state ε . One should therefore note that this is at odds with certain interactions of the agent-environment system: taking the example of rigid collisions, this allows for cases where $\mathbf{b}$ indicates that the agent should be located where ε instead mandates that there should be a solid obstacle. This is actually a situation which was *de facto* avoided in previous experiments and which this framework will *not* improve.

We can already make another observation relative to the functions of $\mathcal{B}$: because it is adapted from $\mathcal{M}$ by enriching it with the unarguably *external* pose variable, we cannot continue assuming that it corresponds to proprioceptive input. Consequently, we cannot reasonably assume that the agent be *a priori* capable of accessing its absolute configuration $\mathbf{b} \in \mathcal{B}$.

2.3.2 The naive sensorimotor experience

The previous considerations were largely external, as the prominence of ambient space in their discussion may hint at. We now instead attend to the definition of those structures related to the internal picture of the sensorimotor experience we will assume. In particular, the previous observation about $\mathcal{B}$ being intrinsically external indicates that we need to characterize how motor commands of the agent manifest in our formalism. We also circle back to our starting intuition about spatially mediated sensation and formally particularise the sensory function to make the notion of “looking somewhere” amenable.

2.3.2.1 Moteur... action!

One of the important functions of $\mathcal{M}$ as previously defined was that it represented the motor capabilities of the agent *via* the active changes of state $\mathbf{m} \in \mathcal{M}$ its motor commands entailed. Our illustrative example shows that this point can also be made of $\mathcal{B}$: the pictured agent’s commands can alter both its internal configuration, as well as its pose should it action its wheels. Therefore, although we cannot hope for the agent to be naively able to tell where it “is” in $\mathcal{B}$ we nonetheless characterize its motor capabilities as a collection of transformations $\mathcal{B} \rightarrow \mathcal{B}$ of absolute configurations. Such transformations, which we henceforth refer to as *motor actions* are therefore concretely functions

$$\begin{aligned} a: \mathcal{B} &\rightarrow \mathcal{B} \\ \mathbf{b} &\mapsto a\mathbf{b} = a(\mathbf{b}) \end{aligned} \tag{2.6}$$

Additionally, we will now denote the relation $\mathbf{b}' = a\mathbf{b}$ by $\mathbf{b} \xrightarrow{a} \mathbf{b}'$. The previous discussion about $\mathcal{B}$ being a set of static configurations makes motor actions a into differences of static configurations; reiterating our point from 2.2.3.2 they cannot denote differential quantities such as velocity or acceleration of motor state.




FIGURE 2.6: Visualisation of the effects of motor actions as vector fields for $\mathcal{B} = \mathbb{R}^2$.

This approach therefore places precedence on the differences between configurations $(\mathbf{b}, \mathbf{b}')$ rather than on any individual configuration $\mathbf{b}$. In particular it makes these differences somewhat *homogeneous*: it naturally formulates the fact that pairs $(\mathbf{b}_1, \mathbf{b}_2)$ and $(\mathbf{b}'_1, \mathbf{b}'_2)$, seen as starting/ending configurations, are “equally related” whenever $\mathbf{b}_1 \xrightarrow{a} \mathbf{b}_2$ and $\mathbf{b}'_1 \xrightarrow{a} \mathbf{b}'_2$ simultaneously hold. This is illustrated in Fig. 2.6, in which the space undergoing transformations is restricted to $\mathbb{R}^2$ for visualisation purposes. Each action a is then represented as a vector field, which attaches to any starting point $X \in \mathbb{R}^2$ the affine vector $\vec{u}$ such that aX is its endpoint $X + \vec{u}$.

The algebraic structure of actions We however mainly attend to the algebraic structure the set $\mathcal{A}$ of such actions exhibits by way of their compositionality. Indeed, we can start by observing that “staying still” is represented by the particular action $e \in \mathcal{A}$ that verifies

$$\forall \mathbf{b} \in \mathcal{B}, \mathbf{b} \xrightarrow{e} \mathbf{b} \quad (2.7)$$

This makes it into an *identity* action, which maps all absolute configurations $\mathbf{b}$ to themselves and corresponds to the null vector field. Furthermore, successive performance of two motor actions $a, a' \in \mathcal{A}$ can also be formally represented by the motor action $a'' = a'a$, which for concrete function corresponds to the classical composition $a' \circ a$ (denoting “ a then a' ”) and verifies

$$\forall \mathbf{b}, \mathbf{b}', \mathbf{b}'' \in \mathcal{B}, \quad \begin{array}{ccc} \mathbf{b} & \xrightarrow{a''} & \mathbf{b}'' \\ & \searrow a \quad \nearrow a' & \\ & \mathbf{b}' & \end{array} \quad (2.8)$$

This relation functionally defines a composition operation on $\mathcal{A}$

$$\begin{aligned} \cdot_{\mathcal{A}}: \mathcal{A} \times \mathcal{A} &\rightarrow \mathcal{A} \\ (a, a') &\mapsto a'a \end{aligned} \quad (2.9)$$

which is associative, that is $(aa')a'' = a(a'a'')$ for any three motor actions a, a', a'' of $\mathcal{A}$. As for illustrating, this leads to the field $\vec{F_2 F_1}$ illustrated in Fig. 2.6c: starting from any point $X \in \mathbb{R}^2$, we first follow $\vec{F_1}$ to determine $X' = \vec{F_1}X$. We then follow the vector supplied by $\vec{F_2}$ at X' to its endpoint $X'' = \vec{F_2}(\vec{F_1}(X)) = \left(\vec{F_2 F_1} \right) X$.

In algebraic terms, these two properties specify that $\mathcal{A}$ is a *monoid* for its composition operation. We will frequently require the additional property that $\mathcal{A}$ be a *group*, which in effect mandates that for any action $a \in \mathcal{A}$ there should exist an inverse action $a^{-1} \in \mathcal{A}$ in the sense that

$$\forall \mathbf{b} \in \mathcal{B}, \quad \mathbf{b} \xrightleftharpoons[a^{-1}]{a} \mathbf{b}'. \quad (2.10)$$

It requires in particular that a be invertible as a function, that is a bijection; this is in fact also a sufficient condition. As a consequence, this introduces additional restrictions on which transformations of configurations are admissible. However, one can alternatively frame this requirement as temporarily focusing on those monoidal actions that are themselves invertible; indeed, the subset of any monoid made of its invertible elements is a group in itself. In terms of the illustrative visualisation, a vector field represents an invertible transformation if and only if one (and only one) arrow ends at each point of $\mathbb{R}^2$, corresponding to the unique pair (X, Y) related by the bijection. In this case, the inverse vector field is obtained by keeping the shafts of the arrows in place and reversing the direction they point to.

Compositional reduction and generated algebras The algebraic structure introduced by the “successive performance” operation involves the compositionality of motor actions; it naturally evokes a prototypic notion of *complex* actions being a combination of simple actions reminiscent of the modular multiple model approach present both in biological systems (Botvinick, 2008; Flanagan et al., 1999; Wolpert and Kawato, 1998) and in robotic ones (Billing, Lowe, and Sandamirskaya, 2015; Nicolescu and Matarić, 2002). Our intuition also informs us that it may provide a way to reduce the combinatorial complexity an agent encounters when it needs to explore its motor configuration space.

To illustrate this point, consider the range of movement provided by actuating the joints of the robotic arm depicted in Fig. 2.7a. Since they are equipped with step motors, up to a choice of reference angle their state space can be labelled by an integer set $[0, \dots, N - 1]$. By this correspondence N is the number of angular steps the motors takes to achieve a full rotation, so that integer k denotes the $\frac{2k\pi}{N} = k \cdot \frac{2\pi}{N}$ radians angle. It immediately follows that the total reachable configuration space is that of all such integer pairs (j, k) (each taken mod. N), for a total number of N^2 configurations. Now consider the two special actions $a_i = \Delta q_i$ corresponding to adding a step increment




FIGURE 2.7: Decomposition of arbitrary motor trajectories over a repertoire of primitive actions.

to the state of actuator m_i . As functions of configurations this respectively corresponds to

$$\Delta q_1(j, k) = (j + 1, k)$$

and

$$\Delta q_2(j, k) = (j, k + 1)$$

applying the modulo N operation where necessary. It is therefore clear that starting from any configuration (j_0, k_0) of its actuator, the agent may reach any target configuration (j_t, k_t) by finite successive operation of its two Δq_i actions. Namely, it can first perform $\Delta q_2 (k_t - k_0)$ times before performing $\Delta q_1 (j_t - j_0)$ times, since $\Delta q_1^{j_t - j_0} \Delta q_2^{k_t - k_0} (j_0, k_0) = (j_t, k_t)$. Fig. 2.7b illustrates this particular situation with the highlighted chain of arrows; it also shows that the path need not be unique, even if one chooses to minimize its length. Nevertheless, this example shows how the exploration of a state space of size N^2 (for arbitrary N) can thus be reduced to that of sequences of 2 elementary actions.

As the illustration suggests, the relation between the entire motor capability set and the 2 “simple” actions is analogous to that between vector spaces and their bases; furthermore, this observation can be made formal in the full generality of the composition structure. Begin with any set of “simple” actions $\mathcal{A}_0 = \{a_i, i \in I\}$. Then by definition any action that can be obtained by finite successive composition of these simple actions is precisely of the form

$$a = a_{i_N} \dots a_{i_1} = \prod_{k=1}^N a_{i_k} \text{ where } N \in \mathbb{N} \text{ and } i_k \in I \forall 1 \leq k \leq N. \quad (2.11)$$

As with the usual case, the empty product (which always occurs for $N = 0$) is taken to be the identity, here $e \in \mathcal{A}$. As a consequence, the set $\langle \mathcal{A}_0 \rangle$ of all actions a that can be generated as per Eq.(2.11) is a monoid of actions itself: it contains the identity action, while the composition of two combinations of the a_i (which can be seen as the concatenation of their strings) yields another of these combinations. It is in fact minimal in the sense that *any* monoid of actions which contains all of the a_i must contain the entire $\langle \mathcal{A}_0 \rangle$ one; this follows from the requirement that it must in particular be stable under composition, that is contain all combinations of the a_i described in Eq. (2.11). Because of this property, we call $\langle \mathcal{A}_0 \rangle = \langle a_i \rangle$ the *(sub)monoid generated by the a_i* .¹¹

The approach to reducing the complexity in configuration space we illustrated is therefore represented as the search for *generating sets* of actions. More precisely, to any action monoid $\mathcal{A}$ corresponding to given motor capabilities of the agent it raises the question of determining which repertoire of simple actions $\{a_i\}$ are sufficient for compositionally generating the whole range of actions as per $\langle a_i \rangle = \mathcal{A}$. For reduction purposes, the same question naturally appears for finite, and then minimal (as per the cardinality of $\{a_i\}$) repertoires. This can be further refined, starting from a repertoire $\{a_i\}$, to finding another (perhaps more desirable) set of actions $\{a_j\}$ that allow for the same motor capabilities; this formally corresponds to the search of particular action sets $\{a_j\} \subset \langle a_i \rangle$, with the generation constraint $\langle a_j \rangle = \langle a_i \rangle$ that the generated structures be equal.

¹¹The same argument can be made *mutatis mutandis* for groups involving the combinations both of the a_i and their inverses; the resulting structure is also a (sub)group, and the minimal group containing all the a_i at that. We therefore predictably call it the *(sub)group generated by the a_i* . The two structures need not necessarily coincide, with the monoid always being a subset of the group; however, this holds in several cases which most of the practical examples we will use fall into.

Continuous transformations to account for dynamical constraints The formal definitions we assume ascribe to $\mathcal{B}$ both the function of initial parameter set as well as supporting motor commands. From the point made in 2.2.3.2, this implies that motor actions are assumed to be differences of position in the configuration state $\mathbf{b}$ of the agent. It would then seem that this formulation still fails to account for agents which are velocity or force controlled. However, while we generally operate on decidedly discrete groups for simplicity reasons we actually make no requirement of this point. On the contrary, the formalism allows for agents which continuous motor capabilities are described by Lie groups, that is precisely the groups that also possess an intrinsic differentiable structure (Belinfante, Kolman, and H. A. Smith, 1966).

If we assume that the interaction follows a continuous time process (without which the consideration is mostly made irrelevant), the requirement of velocity (or force) control can then be formally matched. Indeed, represent a velocity command by a trajectory $\delta_a : \mathbb{R}_+ \rightarrow \mathfrak{a}$ with values tangent vectors in the Lie algebra $\mathfrak{a}$ of $\mathcal{A}$. In turn, this δ_a trajectory defines a unique primitive $\gamma_a : \mathbb{R}_+ \rightarrow \mathcal{A}$ for Darboux derivation starting from the identity $\gamma_a(0) = e_{\mathcal{A}}$. In other words, given any starting absolute configuration $\mathbf{b}_0 \in \mathcal{B}$ the function

$$\begin{aligned} \gamma_{\mathbf{b}} : \mathbb{R}_+ &\rightarrow \mathcal{B} \\ t &\mapsto \gamma_a(t)\mathbf{b}_0 \end{aligned} \tag{2.12}$$

represents the trajectory of absolute configurations $\mathbf{b} \in \mathcal{B}$ matching initial conditions and velocity command provided. The construction can be extended further to force command (resp. higher order derivatives) by first integrating the corresponding trajectory $\mathbb{R} \rightarrow \mathfrak{a}$ as a vector valued function once (resp. a corresponding number of times); this yields a velocity trajectory from which the argument can be followed.

This development therefore involves notably more technical machinery. It is nevertheless conceptually valuable in as much as it presents a natural framework for dynamical constraints imposed by the physicality of the agent's embodiment. For example, the mechanistic requirement that all joint accelerations be continuous can now be formulated as requiring effective force trajectories to be continuous functions $\mathbb{R} \rightarrow \mathfrak{a}$; as expected, this is equivalent to restricting the configuration trajectory γ_b to be C^2 as a map of manifolds. The approach also somewhat addresses the objection, first made in 2.2.3.2, that these complex computations are certainly not accessible to a naive agent. Indeed, we do not assume that the parameter configuration

$\mathbf{b} \in \mathcal{B}$ is internally available; therefore this integration relation only pertains to the external point of view, which makes the objection void. In any case, we do not detail nor exploit further the Lie case. Instead, all proposed applications are presented in discrete setups. We hypothesize that the development associated with continuous transformations should be amenable to a direct, if lengthy, extension of this work.

2.3.2.2 Sensels, sensors, receptive fields

During its sensorimotor interaction with its environment, the agent must in particular be able to obtain some input that depends on the exterior state; recalling the discussion of 2.2.1.1 this is essentially the function attributed to its sensory capabilities by way of the forward sensorimotor function φ_ε , which importantly is characterized by the *a priori* internal availability of its output.

Sensel structure and forward sensory map The sensory apparatuses of distinct sensory agents may cover a wide range of complexity, from the simplistic Braitenberg vehicles (Braitenberg, 1986) to the plethora of instruments equipping NASA’s Curiosity rover (NASA, n.d.). The classical view, informed by our knowledge as designers, is to partition the sensory output of complex agents into partial outputs corresponding to each instrument or otherwise relevant criterion (e.g. modality for multimodal sensors). This is frequently further extended to characterizing the output of a single sensor as a vector ranging over its sensory elements. A textbook example of such an identification is that of a visual CCD sensor, which output is thought of as the collection of the numeric activation values of its photosensitive elements. In this perspective, complexity in the implementation of sensors corresponds to high dimensionality of their output space. This clearly hints at the necessity of providing efficient methods for processing their signals, starting from *a priori* separation of the output corresponding to distinct sensors.

We instead opt for an intermediate approach to preserve genericity: while we assume some preexistent partitioning, we refrain from requiring it to be defined in terms of the physical implementation of the robot. This is intended to allow formal sensory output to directly represent data *which the agent directly accesses*; amusingly, both cognitive viewpoints that “vision occurs in the brain” (rather than in the eye) (P. Bach-y-Rita, 1972) as well as the opposite




FIGURE 2.8: Illustration of the receptive fields for a sensor made of two rigidly linked cameras at configuration $\mathbf{b}$. Each pixel c_i of either camera produces a sensory value s_{c_i} in the overall sensory array $\mathbf{s}$ explained only by a small subset of space $F_{c_i}(\mathbf{b})$. The same applies for both cameras, thus explaining how a sensation for the agent can be explained by the perception of a subset of space.

radical anti-localizationism (Morgese, Lombardo, and De Pascalis, 2017) support maintaining the distinction. This is made formal by having each such *sensory element* (*sensel* for short) of the partition correspond to one sensorimotor mapping $\psi_c : \mathcal{B} \times \mathcal{E}\mathcal{X} \rightarrow \mathcal{S}_c$. The new c subscript, then, serves to explicitly outline that the mapping and the set of output values are that of one particular sensel c .

If we denote by $\mathcal{C}$ the set indexing the partitioning of the agent's sensory apparatus –that is, each $c \in \mathcal{C}$ labels a distinct sensel– then the entire sensorimotor mapping of the agent is realized as the product map

$$\begin{aligned} \psi_{\mathcal{C}} : \mathcal{B} \times \mathcal{E}\mathcal{X} &\rightarrow \mathcal{S} = \prod_{c \in \mathcal{C}} \mathcal{S}_c \\ (\mathbf{b}, \varepsilon) &\mapsto \mathbf{s} = (s_c)_{c \in \mathcal{C}} = (\psi_c(\mathbf{b}, \varepsilon))_{c \in \mathcal{C}} \end{aligned} \quad (2.13)$$

These definitions are illustrated in Fig. 2.8, where the example agent is physically equipped with two cameras. Its sensory array at this particular point of its interaction with the environment is depicted as the row of colored blocks to the right of the figure. Each of these blocks corresponds to one particular sensel of the agent, which may represent down to the physical unit of a pixel. Each physical camera provides a part of the total sensory output, as depicted by the partitioning of the array into its $\mathcal{C}_1$ and $\mathcal{C}_2$ segments. One should note here that the correspondence between physical sensors and their respective subarrays is not *a priori* as self-evident as might seem from the figure: indeed, the previous definition of sensel carries no notion of type or contiguity. Nevertheless, it shows that physical sensors are formally realized as (particular) subsets $\mathcal{C}'$ of the total collection of sensels $\mathcal{C}$.

The spatially mediated perception of receptive fields The embodied nature of the agent implies that it is physically interacting with its environment. Its spatial aspect, i.e. the localization of said embodiment in space, further imposes locality structure on the constraints of this interaction: at any given point of its experience the agent is in fact interacting with the part of the environment which lies in a certain part of space, a part which is poised to change as the agent moves. This locality constraint especially implies that the determination of environment state in this particular region of space in turn entirely determinates the current interaction; admittedly, such a property mimics the psychological concept of *receptive field* (Hartline, 1938; Jeffries, Killian, and Pezaris, 2013) which we intend to formally capture.

By this account, the receptive field of a sensor $\mathcal{C}' \subset \mathcal{C}$ ought to be formally represented as some region of space $\mathcal{X}' \subset \mathcal{X}$ dependent on the motor configuration $\mathbf{b} \in \mathcal{B}$ under which it becomes precisely the region that is being sampled by $\mathcal{C}'$. Accordingly, we elect to denote it as a mapping

$$\begin{aligned} F_{\mathcal{C}'}: \mathcal{B} &\rightarrow \mathfrak{P}(\mathcal{X}) \\ \mathbf{b} &\mapsto F_{\mathcal{C}'}(\mathbf{b}) \end{aligned} \tag{2.14}$$

where $\mathfrak{P}(\mathcal{X})$ is the collection of subsets of ambient space $\mathcal{X}$. Although this may not be clearly enlightening yet, it allows for a formulation of the locality constraint previously discussed. Indeed, it hints at characterizing receptive fields by the property

$$\begin{aligned} \forall \varepsilon, \varepsilon' \in \mathcal{E}, \forall \mathbf{b} \in \mathcal{B}, \\ \varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})} = \varepsilon'|_{F_{\mathcal{C}'}(\mathbf{b})} \Rightarrow \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon) = \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon'). \end{aligned} \tag{2.15}$$

The construction is also illustrated in Fig. 2.8, which represents receptive fields of its agent at several levels. The first, $F_{\mathcal{C}_i}(\mathbf{b})$, is that of a single pixel of the camera. Also depicted are $F_{\mathcal{C}_1}(\mathbf{b})$ and $F_{\mathcal{C}_2}(\mathbf{b})$, the respective receptive fields of its two cameras; they spatially correspond to the union of the receptive fields of their individual pixels. Finally, the total receptive field $F_{\mathcal{C}}(\mathbf{b})$ of the agent covers its entire field of view. As before, it coincides with the union of the receptive fields of its cameras. Moreover, Fig. 2.8 also depicts two distinct environment states ε_1 and ε_2 the agent may interact with. They are visibly distinct from one another since for example the leftmost part of ε_1 is a shade of dark green whereas it is pink in ε_2 . However, we can see that they *locally coincide* on the whole region $F_{\mathcal{C}}(\mathbf{b})$ covered by the gaze of

the agent. Therefore we expect that its sensory apparatus $\mathcal{C}$ report the same output $\mathbf{s} = \psi_{\mathcal{C}}(\mathbf{b}, \varepsilon_i)$ whether $i = 1$ or 2 . This property, then, is precisely which is summarized in Eq. (2.15).

We can deduce several structure constraints of receptive fields from the property of Eq. (2.15) alone. First, our formal viewpoint of sensors $\mathcal{C}' \subset \mathcal{C}$ being unions of their constituting sensels $c \in \mathcal{C}$ is directly mirrored by their receptive fields as per

Proposition 1. *Let $\mathcal{C}' = \{c\} \subset \mathcal{C}$ be a collection of sensels, with sensorimotor maps ψ_c . Given a corresponding collection of mappings $F_c : \mathcal{B} \rightarrow \mathfrak{P}(\mathcal{X})$ verifying the property of Eq. (2.15) each with respect to their sensels, then the mapping*

$$\begin{aligned} F_{\cup \mathcal{C}'} : \mathcal{B} &\rightarrow \mathfrak{P}(\mathcal{X}) \\ \mathbf{b} &\mapsto \bigcup_{c \in \mathcal{C}'} F_c(\mathbf{b}) \end{aligned} \quad (2.16)$$

verifies the same property relative to the resulting sensor $\mathcal{C}'$.

Proof. Let $\varepsilon, \varepsilon' \in \mathcal{E}\mathcal{X}$, and $\mathbf{b} \in \mathcal{B}$. Because $\forall c \in \mathcal{C}', F_c(\mathbf{b}) \subset \bigcup_{c \in \mathcal{C}'} F_c(\mathbf{b})$, we have

$$\varepsilon|_{\bigcup F_c(\mathbf{b})} = \varepsilon'|_{\bigcup F_c(\mathbf{b})} \Rightarrow \forall c \in \mathcal{C}', \varepsilon|_{F_c(\mathbf{b})} = \varepsilon'|_{F_c(\mathbf{b})}.$$

But since we assume each of the F_c possess the locality property of Eq. (2.16), it follows that

$$\forall c \in \mathcal{C}', \psi_c(\mathbf{b}, \varepsilon) = \psi_c(\mathbf{b}, \varepsilon').$$

Finally this entails that $\psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon)$ and $\psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon')$ are equal by equality of their corresponding components. $\square$

We could, if perhaps maliciously, term this result *the union of receptive fields is a receptive field of the union*. More importantly, it shows that the spatial intuition the illustration of Fig. 2.8 elicits is not fortuitous but rather a result we can systematically expect.

Moreover, the isolation of the spatial support of sensation *via* the $F_{\mathcal{C}'}$ receptive field naturally extends to the formal equivalent of a “physical sensitivity” function. While the former encodes the spatial localization that is constraining the sensory flow at any point of the experience, the latter instead dictates which of the observable physical properties a given sensor is sensitive to and how they relate to its output. It is all summarized as per

Proposition 2. Let $\mathcal{C}' \subset \mathcal{C}$ be a sensor with corresponding receptive field $F_{\mathcal{C}'}$. Then for every $\mathbf{b} \in \mathcal{B}$, there exists a sensitivity function mapping

$$\begin{aligned} f_{\mathcal{C}'}(\mathbf{b}) : \mathcal{E}F_{\mathcal{C}'}(\mathbf{b}) &\rightarrow \mathcal{S}_{\mathcal{C}'} \\ \varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})} &\mapsto (f_{\mathcal{C}'}(\mathbf{b}))(\varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})}) \end{aligned} \quad (2.17)$$

such that the sensorimotor map of sensor $\mathcal{C}'$ admits the factorization

$$\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon) = (f_{\mathcal{C}'}(\mathbf{b}))(\varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})}). \quad (2.18)$$

Proof. Let $\mathcal{C}'$ be a sensor and $F_{\mathcal{C}'}$ be its receptive field. For any $\mathbf{b} \in \mathcal{B}$ consider the relation $\sim_{F_{\mathcal{C}'}(\mathbf{b})}$ on environment states that collapses those coinciding on $F_{\mathcal{C}'}(\mathbf{b})$, i.e.

$$\forall \varepsilon, \varepsilon' \in \mathcal{E}\mathcal{X}, \varepsilon \sim_{F_{\mathcal{C}'}(\mathbf{b})} \varepsilon' \Leftrightarrow \varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})} = \varepsilon'|_{F_{\mathcal{C}'}(\mathbf{b})}$$

Since it is the kernel relation of the restriction mapping $\cdot|_{F_{\mathcal{C}'}(\mathbf{b})} : \varepsilon \in \mathcal{E}\mathcal{X} \mapsto \varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})} \in \mathcal{E}F_{\mathcal{C}'}(\mathbf{b})$, it is an equivalence relation.

Additionally, the characteristic property of Eq. (2.15) specifically indicates that for any two environment states ε and ε'

$$\forall \mathbf{b} \in \mathcal{B}, \varepsilon \sim_{F_{\mathcal{C}'}(\mathbf{b})} \varepsilon' \Rightarrow \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon) = \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon'),$$

so that all states of a given equivalence class are mapped to the same sensory output in configuration $\mathbf{b} \in \mathcal{B}$. Therefore, for given $\mathbf{b} \in \mathcal{B}$ consider the mapping $f_{\mathcal{C}'}(\mathbf{b})$ which takes the (equivalence class of the) restriction $\varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})} \in \mathcal{E}F_{\mathcal{C}'}(\mathbf{b})$ to the sensory output $\psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon^0) \in \mathcal{S}_{\mathcal{C}'}$ at configuration $\mathbf{b}$ for any representative ε^0 of $\varepsilon|_{F_{\mathcal{C}'}(\mathbf{b})}$. It therefore ensures commutativity of

$$\begin{array}{ccc} \mathcal{E}\mathcal{X} & \xrightarrow{\psi_{\mathcal{C}'}(\mathbf{b}, \cdot)} & \mathcal{S}_{\mathcal{C}'} \\ \downarrow \cdot|_{F_{\mathcal{C}'}(\mathbf{b})} & \nearrow f_{\mathcal{C}'}(\mathbf{b}) & \\ \mathcal{E}F_{\mathcal{C}'}(\mathbf{b}) & & \end{array} \quad (2.19)$$

which concludes the argument. $\square$

An illustration of the point of the distinction can be made using the language of (Schwartz) distributions. Indeed, consider the prototypic case of the distribution I_f associated to some smooth density function $f : \mathbb{R} \rightarrow \mathbb{R}$,

which applied to test function φ yields output

$$I_f(\varphi) = \langle I_f, \varphi \rangle = \int_{\text{Supp}(I_f)} f \varphi \, d\lambda.$$

In this expression, the support $\text{Supp}(I_f)$ of the distribution characterizes which points of the underlying space of φ actually bear on the value of I_f : any particular behavior of functions φ outside of this subset of $\mathbb{R}$ is discarded by the distribution. On the contrary f specifies how the selected values are processed by the distribution, acting as generalized weights. Therefore, this analogy associates the sensorimotor map $\psi_{\mathcal{C}'}$ of a sensor with a family of distributions, each corresponding to one configuration $\mathbf{b}$ of the agent. For any such distribution, its support corresponds to the spatial receptive field at the same configuration $F_{\mathcal{C}'}(\mathbf{b})$; similarly, the integrand f of the distribution corresponds to the physical sensitivity function $f_{\mathcal{C}'}(\mathbf{b})$ (although it is applied as a function rather than a product for lack of a linear structure). It also makes the reason for the seemingly cumbersome dependency of sensitivity on $\mathbf{b}$ clearer, as it illustrates how distinct configurations of the agent can correspond to differences in how its sensors process the same part of the environment. For example, in the case of a numeric camera sensor, this dependency may represent different possible active tuning of gain and aperture values.

An intrinsic definition for receptive fields The preceding discussion makes points of what properties and structure receptive fields possess; however, at no point did it actually define them. While the notion is intuitive enough that we can agree on what they are for most illustrative agents and practical cases, a formal definition in terms of the previous objects should still be available per our approach.

The “defining” property of receptive fields in the preliminary points is that of Eq.(2.15). To make the discussion more fluid, in the following for given sensor $\mathcal{C}' \subset \mathcal{C}$ and configuration $\mathbf{b} \in \mathcal{B}$ we will say that $F_{\mathcal{C}'}$ (and respectively $F_{\mathcal{C}'}(\mathbf{b})$) is a receptive field (resp. a receptive field at b) of $\mathcal{C}'$ if they verify the conditions of Eq.(2.15). As reflected by the indefinite article, the condition does not define unique receptive fields since we have instead

Proposition 3. *Let $\mathcal{C}'$ be a sensor, let $\mathbf{b} \in \mathcal{B}$, then*

- $\mathcal{X}$ is a receptive field of $\mathcal{C}'$ at b

- if $\mathcal{X}'$ is a receptive field of $\mathcal{C}'$ and $\mathcal{X}' \subset \mathcal{X}'' \subset \mathcal{X}$, then so is $\mathcal{X}''$

Proof. • This point is trivial: by Eq. (2.15) it corresponds to the tautology

$$\forall \varepsilon \in \mathcal{E} \mathcal{X}, \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon) = \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon)$$

- Let $\mathcal{X}' \subset \mathcal{X}$ be a receptive field of $\mathcal{C}'$ at $\mathbf{b}$, and let $\mathcal{X}'' \subset \mathcal{X}$ such that $\mathcal{X}' \subset \mathcal{X}''$. Then for any two environment states $\varepsilon, \varepsilon'$

$$\varepsilon|_{\mathcal{X}''} = \varepsilon'|_{\mathcal{X}''} \Rightarrow \varepsilon|_{\mathcal{X}'} = \varepsilon'|_{\mathcal{X}'}$$

which, chained with the implication of Eq. (2.15) relative to receptive field $\mathcal{X}'$ yields

$$\varepsilon|_{\mathcal{X}''} = \varepsilon'|_{\mathcal{X}''} \Rightarrow \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon) = \psi_{\mathcal{C}'}(\mathbf{b}, \varepsilon')$$

corresponding to $\mathcal{X}''$ being a receptive field of $\mathcal{C}'$ at $\mathbf{b}$.

□

This result at fixed configuration $\mathbf{b}$ is in turn reflected at the level of the receptive field mapping by

Corollary 1. *Let $\mathcal{C}'$ be a sensor, then*

- $\mathcal{C}'$ admits a receptive field, given by $\mathbf{b} \in \mathcal{B} \mapsto \mathcal{X} \subset \mathcal{X}$
- If F is a receptive field of $\mathcal{C}'$ and $F' : \mathcal{B} \rightarrow \mathfrak{P}(\mathcal{X})$ verifies

$$\forall \mathbf{b} \in \mathcal{B}, F(\mathbf{b}) \subset F'(\mathbf{b})$$

then F' is a receptive field of $\mathcal{C}'$.

The result states in essence that whenever some region of space is a valid receptive field, going up only yields other receptive fields. The sensible idea then, corroborating the intuition developed in concrete examples, is to define *the* receptive field as the smallest of such regions. Let us consider a given sensor $\mathcal{C}'$ and configuration $\mathbf{b}$, and temporarily denote X the collection of all $\mathcal{X}' \subset \mathcal{X}$ which are a receptive field of $\mathcal{C}'$ at $\mathbf{b}$. From the previous result we know that X is not empty, since $\mathcal{X} \subset X$. Furthermore, the classical process for taking the smallest such region is to consider $\bigcap_{\mathcal{X}' \in X} \mathcal{X}'$: indeed, if there




FIGURE 2.9: Illustration of a system of environments where no (unique) receptive field can be determined. The ambient space $\mathcal{X}$ contains only three points a, b, c . Environment states $\varepsilon \in \mathcal{E}\mathcal{X}$ are restricted so as to be entirely determined by the two colors in the translucent bars, resulting in the two points of space b and c always being attributed the same value.

exists one $\mathcal{X}_- \in X$ minimal for inclusion between all receptive fields, we must have

$$\forall \mathcal{X}' \in X, \mathcal{X}_- \subset \mathcal{X}'$$

by minimality so that

$$\mathcal{X}_- \subset \bigcap_{\mathcal{X}' \in X} \mathcal{X}'.$$

But since $\mathcal{X}_- \in X$, it also appears in the right hand side intersection, by which it results that $\mathcal{X}_- = \bigcap_{\mathcal{X}' \in X} \mathcal{X}'$.

The previous definitions, however, do not guarantee the validity of this construction; it instead hinges at least on additional considerations relative to the diversity of possible environment states. Fig. 2.9 illustrates this point; despite its symbolic setting, it falls squarely in the formal definitions. It represents a space $\mathcal{X}$ composed of only three points $\{a, b, c\}$ where crucially the values of environment state at different points are constrained: here, *all* environment states $\varepsilon \in \mathcal{E}\mathcal{X}$ necessarily verify $\varepsilon(b) = \varepsilon(c)$. In particular, this intuitively implies that one need only query the value at any one of these points to “know” the environment state at both. Accordingly, for any sensor $\mathcal{C}'$ we will find that $F_b = \{a, b\}$ and $F_c = \{a, c\}$ are valid receptive fields for $\mathcal{C}'$. However, their intersection is the singleton $\{a\}$ which may not be a receptive field itself. The previous construction fails because the collection of receptive fields here is

$$X = \{\{a, b\}, \{a, c\}, \{a, b, c\}\}$$

which is not stable under intersection, and as such does not have a smallest element. As a consequence, there is no natural way of defining *the* receptive field in such cases. This issue, however, may be resolved by a number of additional assumptions, one of which should in all but the strangest cases: first, if ambient space $\mathcal{X}$ is assumed to be (at least) topological, the same development can be made by restriction to closed subsets. Else, by contrast with this pathological example, it is valid whenever the collection of environment space contains *all* functions of the form $\mathcal{X} \rightarrow \mathcal{P}$ (*a fortiori* without constraints on their values). In any case, all the settings developed further in illustrations or experiments fall under these “nice” cases, therefore we can consider the construction valid during their examination.

2.3.2.3 How can it be so naive: a model of knowledge

The previous development laid the notions we use in describing the sensorimotor interaction of an agent with its environment. However it did not definitively delimit the internal picture of this interaction, that is the subjective basis of the internal sensorimotor *experience*. We previously mentioned how this distinction is crucial to SMCT : in as much as it is a theory of learning, the breadth of internal cases considered directly mirrors its scope and conditions of validity. We therefore turn our attention to addressing this point and provide a formal characterization of the dual points of view of 2.2.2.2.

The previous framework (see 2.2) operated under the assumption that the agent could access its internal motor states $\mathbf{m} \in \mathcal{M}$ and sensory states $\mathbf{s} \in \mathcal{S}$. Besides it could leverage its motor capabilities by *actively* modifying $\mathbf{m} \in \mathcal{M}$, via motor commands realizable as functions $f_{\mathcal{M}} : \mathcal{M} \rightarrow \mathcal{M}$. By “actively” we mean that the sequence of motor states parameterizing the sensorimotor interaction is effectively defined by the sequence of commands; importantly, the agent also has some degree of knowledge¹² about which command it issues so that it can at least reiterate the same command. This awareness of its own action is crucial in that it allows to internally distinguish between external and self-induced changes should ambiguity arise (e.g. in the case of compensations).

¹²We would informally want to qualify this knowledge as *conscious* were it not for the multiple problems associated with the notion. What matters, instead, is perhaps that these commands be volitional (Maye and A. Engel, 2016).




FIGURE 2.10: Illustration of the knowledge the agent *a priori* possesses of its sensory output. Each line corresponds to a part of its sensory array $\mathbf{s}_{\mathcal{C}}$ at two different states of its sensorimotor experience.

Sensory knowledge: the Array of Babel Following the argument exposed in 2.3.2.2 our description makes use of a straightforward stand-in for sensory state *via* the output of the total sensorimotor map $\psi_{\mathcal{C}}$. We continue to assume its *a priori* accessibility on the basis that this property is precisely which characterize sensation. However, the discussion made in 2.2.2.1 about interpreting mathematical structure as *a priori* assumptions of the agent shows that we must make precise the ways in which the naive agent can access its sensory readings. Indeed there is a number of “levels” of knowledge, ranging from that of bare set structure of $\mathcal{S}_{\mathcal{C}}$ to the wealth of expressiveness afforded by its identification to $\mathbb{R}^N$ ¹³. In the following we will assume that the agent knows about the product set structure of $\mathcal{S}_{\mathcal{C}} = \prod_{c \in \mathcal{C}} \mathcal{S}_c$. This means that the agent may *a priori* access its sensory readings on a *per component* basis (where “components” correspond to the sensel decomposition); however, the agent may only further access the set structure of its components $\mathbf{s}_c \in \mathcal{S}_c$.

Fig. 2.10 gives an illustration of the situation *from the point of view of the agent*. The horizontal arrays represent the output $\mathbf{s} \in \mathcal{S}_{\mathcal{C}}$ of the total sensory apparatus $\mathcal{C}$ at two (distinct) stages of the sensorimotor experience. Accordingly, each square in an array corresponds to a fixed sensel $c \in \mathcal{C}$. The capability to discern each of the squares and query their content is precisely that afforded by the knowledge of the product structure $\prod_{c \in \mathcal{C}} \mathcal{S}_c$ as opposed to that of set $\mathcal{S}_{\mathcal{C}}$. Then, each sensel square reports its corresponding state

¹³(see 2.2.2.1 for specific examples)

by the symbol¹⁴ it displays. Importantly, distinct sensels do not share their collection of output symbols: this corresponds to each square having its own “alphabet” as well as its color. This illustrates the minimal semantics we will assume for sensory symbols, by which we say the agent processes “uninterpreted” sensory signals. The meaning ascribed here is “stronger” than that usually found in the bootstrapping literature (see (Censi, 2012; B. J. Kuipers et al., 2006; David Philipona et al., 2004; Pierce and B. J. Kuipers, 1997)). There, the form assumed by the signal (and the expected transformations thereof) is usually known and leveraged by the agent; what it ignores instead is how these signals relate to the sensorimotor interaction. In our case, however, there are no “natural” metrics nor algebraic operations on the symbols depicted in Fig. 2.10, contrary to our understanding of the usual numeric values. We can instead only compare them for equality: that is, provided two distinct symbols in the alphabet of any given sensel (corresponding to any two possible values s_c and s'_c of any *single* $\mathcal{S}_c$) we can definitively determine whether these symbols are equal. Crucially the comparison cannot be extended to values that occur in distinct squares due to them not sharing their alphabets of symbols: the agent cannot assess whether s_c and $s_{c'}$ are equal when $c \neq c'$ as sensels. It is very much possible that the values of distinct sensels are actually related, for instance because they represent the same external physical sampling; however, this relation is not a (*de facto*) *a priori* consequence of the equality of their codes of symbols. Instead, it should be learned by observation whenever it occurs.

We can describe this approach using the formal terms previously defined. Indeed, the capability to compare values of any given sensel c for equality corresponds to the internal computability of functions

$$\delta^c : s, s' \in \mathcal{S}_c \mapsto \delta_{s,s'}^c = \begin{cases} \top & \text{if } s = s' \\ \perp & \text{if } s \neq s'. \end{cases} \quad 15$$

for all sensels $c \in \mathcal{C}$. Accordingly access to the total sensory output $\mathbf{s} \in \mathcal{S}_{\mathcal{C}}$ is given by the sensorimotor map $\psi_{\mathcal{C}}$, while the capability to access any component s_c is provided by projections

$$\pi_c : \mathbf{s} = (s_{c'})_{c' \in \mathcal{C}} \in \mathcal{S}_{\mathcal{C}} \mapsto s_c \in \mathcal{S}_c$$

¹⁴(the more *bizarre* the better)

¹⁵One can think of $\top$ as “true” and $\perp$ as “false”, as that is their intended semantics. It only matters that the agent be able to perform these binary tests, not the particular symbol they use.

with c ranging over $\mathcal{C}$. Therefore, we assume that these (collections of) mappings are all *a priori* internally computable by the naive agent. The discussion of hierarchies of structures made in 2.2.2.1 showed how product structure provides more capabilities than set structure. Therefore, this choice of assumption is possibly more restrictive than that made in V. Marcel’s works (see (V. Marcel, Argentieri, and Gas, 2019) and parts of (Valentin Marcel, 2020)) where the agent is only capable of comparing its *total* sensory outputs $\mathbf{s}$ for equality. It can however be identically realized under our proposal if we choose to adopt a sensel decomposition $\mathcal{C} = \{c\}$ with a unique sensel c . Indeed the previous functions thereby become somewhat trivialized: the δ^c become the (unique) comparison of total sensory values while the π_c is the projection of $\mathbf{s}$ onto its unique component, that is the identity $\text{id}_{\mathcal{S}_{\mathcal{C}}}$. It should be noted that these works already are parcimonious with respect to their *a priori* assumptions relative to other practices, which shows how our framework does not require imposing additional restrictions in this aspect.

No aesthetics for kinesthetics The picture concerning the agent’s *a priori* knowledge of its motor structure is starkly different. Indeed, our description of sensorimotor experiences has $\mathcal{B}$ the set of absolute configurations replace $\mathcal{M}$ that of internal motor states. Previously, internal motor states could be assumed to be accessible to the agent based on proprioceptive feedback, providing it with a kinesthetic sense. However, this argument cannot hold when one considers locomotion capabilities as depicted in 2.3.1.2. Instead, several absolute configurations $\mathbf{b}$ related to one another by locomotion should result in the same internal kinesthetic output. The internal construction we provided to identify $\mathbf{b}$ with tuples $(\mathbf{m}, \tau) \in (\mathcal{B}/\mathcal{G}(\mathcal{X})) \times \mathcal{G}(\mathcal{X})$ supports this distinction: the “postural” data $\mathbf{m} = [\mathbf{b}]_{\mathcal{G}(\mathcal{X})}$ corresponds to the previous internal motor state, whereas τ is intrinsically external. We could therefore adapt the previous argument and assume the *a priori* knowledge of this $\mathbf{m}$ internal value, for instance by computability of $\mathbf{b} \mapsto [\mathbf{b}]_{\mathcal{G}(\mathcal{X})}$. Nevertheless we choose to assume no *a priori* distinguished account of kinesthetic sense along such a mapping, and instead have proprioceptive feedback appear as sensels $c \in \mathcal{C}$. Of course, this implies in particular that sensory output $\mathbf{s}$ does not specifically represent the exteroceptive signal.

The availability of proprioception is known to entail many capabilities

(see (Gapenne, 2014), (Bhanpuri, Okamura, and Bastian, 2013) or (González-Grandón, Falcón-Cortés, and Ramos-Fernández, 2021)), as the fact that previous contributions extensively leveraged the kinesthetic sense afforded by internal motor states $\mathbf{m}$ would hint at. Consequently, this assumption seems considerably weaker than that of *a priori* distinguished sources. In fact, it corresponds to the converse of the process that separates the total sensory output $\mathbf{s} \in \mathcal{S}_{\mathcal{C}}$ into the array of its sensel decomposition $(s_c)_{c \in \mathcal{C}}$. Moreover, previous approaches leveraged the asymmetrical roles of proprioceptive and exteroceptive outputs in that they used the former to structure the latter, and not *vice versa*. Instead, we hypothesize that this separation should be learned by the agent. More precisely, we posit that for a sensel the property that it carries internal proprioceptive data should be ascertainable from the sensorimotor flow under our relaxed assumptions; among other examples, it has been shown how information processing can lead to such discovery (Schmidt et al., 2013).

The other motor function in the internal experience is that supporting motor command, concerning which we previously mentioned how it can be implicitly leveraged as “active” changes in internal motor states. We can in fact readily provide an account of this function using motor actions $a \in \mathcal{A}$. The similarity appears at a very formal level since motor actions of $\mathcal{A}$ can be realized as functions $\mathcal{B} \rightarrow \mathcal{B}$ ¹⁶, that is as changes in (absolute) motor configurations $\mathbf{b}$. Therefore, *a priori* knowledge (and, in this case, performativity) of its motor actions endow the agent with the capabilities described above for recognizing the changes induced by its actions. We can hence understand the internal knowledge of $a \in \mathcal{A}$ as the *effference copy* (Sperry, 1950; von Holst and Mittelstaedt, 1950) associated to the induced motor command $\mathcal{B} \rightarrow \mathcal{B}$. Incidentally this also allows for relaxing the previous constraint that the effference copy be entirely identical to its corresponding motor command, which is in line with several experimental observations showing these correspondences are non trivial (Bridgeman, 2007; Imamizu, 2010; Li, Zhu, and Tian, 2020; Niziolek, Nagarajan, and Houde, 2013; Pynn and DeSouza, 2012).

In contrast with the case of sensory outputs, we make a distinction here between (externally) existing structure and that (internally) known to the agent: we do not assume the naive agent knows of the algebraic structure of $\mathcal{A}$ but only of its set one. As before this implies that it can compare actions

¹⁶In fact, they were preliminarily introduced as such in 2.3.2.1

$a, a' \in \mathcal{A}$ for equality, and corresponds to the internal computability of

$$\delta^{\mathcal{A}} : a, a' \in \mathcal{A} \mapsto \delta_{a,a'}^{\mathcal{A}} = \begin{cases} \top & \text{if } a = a' \\ \perp & \text{if } a \neq a' \end{cases}$$

However actions have several algebraic properties –that is related to their compositionality– the agent cannot *a priori* probe. For example, it may not tell whether two actions are inverse of one another. In fact, it may not even determine whether any one action is the identity since these properties are realized as that of maps in $\mathcal{B}$ the agent has no direct access to. To reiterate, the distinction on structures of $\mathcal{A}$ illustrates that of points of view on a subtler level. On the one hand, the algebraic structure of $\mathcal{A}$ (whether of monoid or group) is integral to the description of the interaction between agent and environment; in particular, we show later (see Chapter 3) some ways in which it necessarily induces constraints on the effectively observable sensorimotor sequences. On the other hand, there is no reason –and indeed no need– that the agent actually be aware of it. As far as it is concerned, in fact, it may suffice to observe the aforementioned constraints in its naive experience. This paves the way for an anti-realist interpretation parallel to that of geometrical space, where one claims space is not needed as much as the regularity it imposes on the sensorimotor flow (A. V. Terekhov and J. K. O'Regan, 2013).

The compositional structure of knowledge along experience The previous discussion determined which objects are *a priori* internal (and which are not). However it does not provide yet a definitive account of the notions a naive agent may successively discover as its sensorimotor experience unfolds. Accordingly it hinges on the experience occurring along a “duration” which supports this unfolding. Of course, the corresponding “points in time” are classically taken to be elements either of $\mathbb{N}$ for discrete time systems or of $\mathbb{R}$ for continuous time cases. In general however, we are only concerned with the incremental aspect of learning for the naive agent; therefore, we only ask for an *ordered* set $\mathcal{T}$ to index the unfolding of the sensorimotor experience. This property specifically mandates that there exist a “prece-
 dence” relation $\leq_{\mathcal{T}}$ between elements of $\mathcal{T}$ such that $t \leq_{\mathcal{T}} t'$ whenever t

occurs before t' ¹⁷. In particular it provides a natural basis for characterizing past and future experience at any point of the sensorimotor interaction. Indeed, suppose that it occurs at (possibly “continuous”) timesteps denoted by $\mathcal{T}$. At any particular step t_0 we can consider $\mathcal{T}_{t_0}^- = \{t \in \mathcal{T} \mid t \leq_{\mathcal{T}} t_0\}$ and $\mathcal{T}_{t_0}^+ = \{t \in \mathcal{T} \mid t_0 \leq_{\mathcal{T}} t\}$: they respectively denote the collection of timesteps occurring before and after t_0 . We naturally expect the available knowledge of the agent to grow over time: if K_t denotes the collection of internally available mappings at timestep $t \in \mathcal{T}$ then we have $K_t \subset K_{t'}$ whenever $t \leq_{\mathcal{T}} t'$, by which $(K_t)_{t \in \mathcal{T}}$ is a filtration. In this context, *a priori* knowledge of the agent is characterized by $K_0 = \bigcap_{t \in \mathcal{T}} K_t$ which contains exactly the mappings that have always been internally available¹⁸. The growth of K_t follows from the agent gradually acquiring new samples during its sensorimotor experience: denoting a_t the actions it performed and $\mathbf{s}_t$ the sensory outputs it experienced, at any timestep $t_0 \in \mathcal{T}$ of the interaction the agent only knows of the *past* interaction, that is all $a_{t'}$ and $\mathbf{s}_{t'}$ with $t' \leq_{\mathcal{T}} t_0$. K_t is thus made to grow larger as does the collection of such $a_{t'}$ and $\mathbf{s}_{t'}$, that is precisely the “evidence” it can take as basis for further assessing the properties of its sensorimotor flow.

Moreover we can sketch an induction scheme according to which the agent incrementally learns the regularities of its sensorimotor flow from one another. First we will assume that the agent can chain internal computations, that is whenever $O_1, \dots, O_N$ and O are all internally computable mappings then so is

$$(x_1, \dots, x_n) \mapsto O(O_1(x_1), \dots, O_N(x_n))$$

provided the composition of the right hand side makes sense¹⁹. Further consider a –somewhat contrived– example where our toy camera can additionally adjust its sensor gain *via* some particular motor actions; in particular, let

¹⁷It comes with a number of conditions akin to that of equivalence relations: it is necessarily true that

1. $t \leq_{\mathcal{T}} t$: every instant occurs before itself (in the same way that $5 \leq 5$ as number)
2. $t \leq_{\mathcal{T}} t'$ and $t' \leq_{\mathcal{T}} t \Rightarrow t = t'$: an instant cannot simultaneously occur before and after any *other* instant
3. $t \leq_{\mathcal{T}} t'$ and $t' \leq_{\mathcal{T}} t'' \Rightarrow t \leq_{\mathcal{T}} t''$: precedence can be chained

¹⁸Of course, should there be a “starting” timestep $t_0 \in \mathcal{T}$ (such that it precedes all other timesteps $t \in \mathcal{T}$ as per $t_0 \leq_{\mathcal{T}} t$) then we would also have $K_0 = K_{t_0}$.

¹⁹This requirement mirrors that of (*general*) *recursive functions* being closed under composition in computability theory (Kleene, 1936). The argument for it being a reasonable assumption is also similar, since this class of functions has been studied as representing which is *effectively* (or algorithmically) computable (Turing, 1937).

a particular action set it to null that all sensels subsequently report no activity. Then the agent can *a priori* ascertain a certain pattern there since it can isolate the specific action a considered and ψ is computable: it can entirely assess that performing that specific action *always* leads to the same resulting sensory output. This observable property corresponds to a mapping that assigns this particular output to the action that causes it, arising as the conjunction of two other observations –namely the fact that the action performed *is equal to a* and the fact that the resulting output *is equal to* the null array. Although this simplistic case is of limited interest, it provides the prototype for a powerful induction rule: indeed, we now assume that the conjunction of any collection of internally computable properties is itself internally computable. Formally this means that whenever any $N \in \mathbb{N}$ boolean mappings $B_1, \dots, B_N$ with $B_i : X_i \rightarrow \{\top, \perp\}$ are internally computable, then specifically the (boolean) mapping

$$\begin{aligned} \wedge_i B_i : \prod_{1 \leq i \leq N} X_i &\rightarrow \{\top, \perp\} \\ (x_i)_{1 \leq i \leq N} &\mapsto \begin{cases} \top & \text{iff } \forall 1 \leq i \leq N, B_i(x_i) = \top \\ \perp & \text{else} \end{cases} \end{aligned} \quad (2.20)$$

is also computable. Similarly we assume that the disjunction and the negation of internally computable boolean mappings are themselves internally computable; then we know from propositional calculus that this especially entails that *any* boolean formula of the B_i observations is itself internally computable (Gries and Schneider, 1993).

2.4 Conclusion

The sensorimotor contingencies theory makes a point of questioning not only what it is to perceive, but also what it precisely is that a sensorimotor agent perceives as well as how it does so. As such, formal SMCT accounts must take great care in their definitions not to squash these nuances by overspecifying their settings along our learned intuition about our experience. Indeed, we have discussed how even just taking usual data under their familiar forms –for instance, denoting a robotic agent’s motor configuration as a tuple of real numbers– may constitute a sizeable implicit assumption on top of that of sensorimotor contingencies.

In this chapter, we detailed how the frameworks used in A. Laflaquière and V. Marcel’s works approached this issue. They generally revolve around a “MXS” diagram drawing from classical robot control descriptions. We then showed how this modelling does not entirely account for both internal and external pictures of the sensorimotor experience. We especially examined how the choices in mathematical structures and semantic functions ascribed to the various objects of the diagram correspond to implicit but potentially strong assumptions about the system. This argument led us to demonstrate that several limitations of these frameworks, such as being unable to account for locomotion, necessarily result from the formal choices made in their definitions.

We then developed our proposal for a revised formal framework. We chose to shift the focus of our definitions on some preexisting “spatial” structure hosting the sensorimotor interaction²⁰ as well as how the agent can interact with this structure by way of its motor actions. The construction was explicitly made to address the previously discussed limitations, a point we assessed on our illustrative system. We further showed how these definitions subsume previous constructions. For example, it allows for expressing formal correlates of internal motor configurations (i.e. posture of the agent), sensors, and their receptive fields. Throughout, we dedicated great care to ensure that all arguments be entirely expressed in terms of previously defined objects, without relying on our external *ad hoc* foreknowledge of the system. In particular, we provided a lengthy discussion about the distinction between the external and internal points of view (relative to the interacting agent). It culminated in 2.3.2.3 where we proposed a (partial) formal characterization of internal knowledge of the agent. Among others, it enabled us to assess that our assumptions about *a priori* knowledge of the agent are not more restrictive than that of many other contributions, in line with our intention to account for the bootstrapping function of SMCT .

²⁰(which formally appears *via* $\mathcal{X}$, $\mathcal{G}(\mathcal{X})$, and $\mathcal{B}$)

Chapter 3

Discovery of the motor algebra by prediction of sensory shifts

Contents

3.1	A zero-th level of <i>spatial</i> sensorimotor contingencies . . .	63
3.2	A formal study of learning motor structure <i>via</i> sensory prediction	64
3.2.1	Introducing conservative actions	64
3.2.1.1	Special spatial exchanges in sensation	64
3.2.1.2	An algebraic conversation of conservation	67
3.2.1.3	Partial successor map and quasiconservative actions	71
3.2.2	Perceptive properties of sensory prediction	74
3.2.2.1	Using predictable actions as motor proxies	74
3.2.2.2	Prediction in the case of conservation	79
3.2.2.3	Degeneracy in the semigroup case	81
3.2.3	Conclusion	86
3.3	Experiment: a simulated account of inferring action structure by naively assessing conservativity	88
3.3.1	Description of the setup	88
3.3.2	Simulation 1: searching the sensorimotor flow for conservativity	92
3.3.2.1	Representing conservative actions	92
3.3.2.2	Progressive discovery of the conservative permutations	94
3.3.2.3	Results	101

3.3.3	Simulation 2: investigating the dynamics in the discovery of sensel permutations	102
3.3.3.1	Assessing progress in the convergence of the empirical matrices	102
3.3.3.2	Sources of variance in the dynamics of discovery	103
3.3.3.3	Experimental setup	104
3.3.3.4	Results	105
3.3.4	Simulation 3: internally building the conservative subalgebra	108
3.3.4.1	Generative Dijkstra-like algorithm for computing the group of sensory predictions	108
3.3.4.2	Managing loss of information in nonconservative actions	112
3.3.4.3	Force directed methods for plotting the Cayley graphs	115
3.3.4.4	Results	116
3.3.5	Simulation 4: reparameterizing basis actions	123
3.3.5.1	Dependence of the graph with respect to the starting action set	123
3.4	Conclusion	126

3.1 A zero-th level of *spatial* sensorimotor contingencies

The SMCT approach as originally outlined by O'Regan and Noë (J K O'Regan and Noë, 2001) introduces a distinction between two essential types of sensorimotor contingencies. Contingencies of the first type are related to the particular organisation in the implementation of a sensory apparatus, in as much as they are assumed to be regularities of the way by which the corresponding apparatus is brought to interact with its environment. Modality-specific contingencies (such as those of vision or tactile perception) generally fall under this category. On the other hand the second class is that of contingencies characteristic of a perceptual object, closer to the ecological perception of Gibson (Gibson, 1979). It therefore naturally appears that regularities in the first class are somewhat more general than those of the second; another tentative distinction is that the former are contingencies relevant to *any* environment for a specific embodied agent (since they only rely on the fixed laws of the sensorimotor embodiment) whereas the latter are only relevant in the specific environments in which figure their associated objects.

The following part is then devoted to showing how the sensorimotor interaction being spatially mediated supports a “zeroth” type of contingencies, that is contingencies that fundamentally only depend on the constraint the localization of an embodied agent imposes in its relation to ambient space. In particular, we investigate the existence and the discoverability of a function that would allow an agent to predict the sensory outcome of its motor actions. First, we leverage the formalism we previously introduced to theoretically characterize the shifts of receptive fields that occur as the agent moves in its environment. We show that these shifts possess an internal algebraic structure mirroring that of the motor actions that cause them. Crucially, we show that this sensory image of the motor structure is internally discoverable under our assumptions of *a priori* knowledge of the agent. It therefore provides the agent with a proxy for probing its motor capabilities.

The part is concluded by a series of interrelated simulations illustrating the scope of such a correspondence. It is in essence a direct application of the theoretical development: a naive agent is set to determine the sensory predictions associated with its displacements, before internally discovering their underlying algebraic structure. Although it follows in a straightforward from the preceding exposition, it sheds some light on the relevancy of

details that otherwise seem purely formal. In particular, great care is taken in making explicit how the processes the simulated agent leverage fall in the confines of the *internal* point of view developed in 2.3.2.3.

3.2 A formal study of learning motor structure *via* sensory prediction

3.2.1 Introducing conservative actions

3.2.1.1 Special spatial exchanges in sensation

We showed in the preliminary exposition about our formal description for sensorimotor interaction how receptive fields ensure the distinction between spatial and physical content of the interaction. The former is that of where the part of the environment that is relevant to the interaction at hand lies in space, whereas the latter deals with how said part of the environment constrains the sensory image of the interaction. Recall the relevant property, that is the expression of the sensorimotor map in terms of both these factors as per

$$\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \psi_c(\mathbf{b}, \varepsilon) = f_c(\varepsilon|_{F_c(\mathbf{b})})$$

where we restrict sensors to sensels and drop the notational dependence of the sensitivity function f_c on $\mathbf{b}$ for simplicity.

Consider now the relation any two sensels c_i, c_j have when, for a given motor action $a \in \mathcal{A}$, sensel c_i systematically experiences *after* a the region of space c_j was experiencing immediately before. From the previous definitions, this situation corresponds to the assertion

$$\forall \mathbf{b} \in \mathcal{B}, F_{c_i}(a\mathbf{b}) = F_{c_j}(\mathbf{b}). \quad (3.1)$$

In the context of visual modality this corresponds to the phenomenon of *retinal smear* (Festinger and Holtzman, 1978), which was previously exploited for model free calibration of camera sensors in robots (Montone, J. Kevin O'Regan, and Alexander V. Terekhov, 2015). Since it is formulated as an equality of spaces, it seems to firmly fall under the external point of view. However, if it also holds that sensors c_i and c_j share the same sensitivity




FIGURE 3.1: Illustration of the way receptive fields shifts under motor displacements of the agent

function $f_{c_i} = f_{c_j}$, it further entails the *sensory equality*

$$\begin{aligned} \forall \mathbf{b} \in \mathcal{B}, \quad \forall \varepsilon \in \mathcal{E}\mathcal{X}, \\ \psi_{c_i}(a\mathbf{b}, \varepsilon) &= f_{c_i}(\varepsilon|_{F_{c_i}(a\mathbf{b})}) \\ &= f_{c_j}(\varepsilon|_{F_{c_j}(\mathbf{b})}) = \psi_{c_j}(\mathbf{b}, \varepsilon). \end{aligned} \quad (3.2)$$

The important distinction here appears in the internal point of view: assuming that the agent is capable of comparing the outputs s_i and s_j of its sensels c_i and c_j , it can therefore assess for any action $a \in \mathcal{A}$ whether property (3.2) holds.

Figure 3.1 depicts this particular situation. Each sensory array corresponds to a particular absolute configuration of the agent as it moves along the axis, so that sensels and the parts of environment they sample are suitably aligned. The particular shift in receptive fields discussed occurs between sensels c_i and c_j : indeed, the amplitude of the displacement from $\mathbf{b}$ and $\mathbf{b}' = a\mathbf{b}$ matches the separation between their receptive fields. Accordingly, the region c_i samples at $\mathbf{b}'$ (that is $F_{c_i}(a\mathbf{b})$) is precisely the one c_j was sampling before at $\mathbf{b}$ (i.e. $F_{c_j}(\mathbf{b})$). Further, the corresponding color contents in the respective sensory arrays must match as a result of the assumption that sensitivity functions f_{c_i} and f_{c_j} be equal.

The precision about needing the capability of comparing the outputs of distinct sensels is significant, since we assumed in 2.3.2.3 that the actual values sensel outputs take specifically have no relation between separate sensels. Notwithstanding, it corresponds to the assumption that c_i and c_j share their sensitivity function. Indeed, it stands to reason that the spatial

relation denoted by Eq. (3.1) would have no bearing on the resulting sensory value if, say, c_i and c_j were respectively sensitive to audio frequency and local temperature. On the contrary, whenever the two sensels have exact same sensitivity, Eq. (3.2) show that sensory outputs directly mirror the spatial relation of Eq. (3.1). Further, an intermediate criterion for existence of a *one to one* correspondence of sensory outputs can easily be obtained as follows

Proposition 4. *Let c_i and c_j be two sensels with respective sensitivity functions f_{c_i} , f_{c_j} and output sets $\mathcal{S}_i$, $\mathcal{S}_j$.*

Then there exists a one to one transduction mapping $t_{i,j} : \mathcal{S}_j \rightarrow \mathcal{S}_i$ such that

$$t_{i,j} \circ f_{c_j} = f_{c_i}$$

if and only if f_{c_i} and f_{c_j} share the same kernel sets, i.e.

$$\forall s_j \in \mathcal{S}_j, \exists s_i \in \mathcal{S}_i \text{ such that } f_{c_j}^{-1}(s_j) = f_{c_i}^{-1}(s_i)$$

Additionally, this mapping verifies that whenever Proposition (3.1) holds for a given action $a \in \mathcal{A}$, then

$$\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E} \mathcal{X}, \psi_{c_i}(a\mathbf{b}, \varepsilon) = t_{i,j}(\psi_{c_j}(\mathbf{b}, \varepsilon)). \quad (3.3)$$

Proof. • Let $t_{i,j}$ be such a bijective mapping. For any sensory value s_i the sensel c_i can take, its preimage by f_{c_i} is precisely the associated kernel set of f_{c_i} . Moreover, its preimage by $t_{i,j}$ is a singleton $\{s_j\}$ of $\mathcal{S}_j$ since it is the image of the bijective inverse of $t_{i,j}$. But therefore this implies that $(t_{i,j} \circ f_{c_j})^{-1}(s_i) = f_{c_j}^{-1}(t_{i,j}^{-1}(s_i)) = f_{c_j}^{-1}(s_j)$ which is in turn the kernel set of f_{c_j} associated to its value s_j .

- Assume that f_{c_i} and f_{c_j} share their kernel sets, and let s_j be any output value of c_j . Lift s_j to its kernel $f_{c_j}^{-1}(s_j)$: by our assumption, this set is also a kernel set of f_{c_i} for an arbitrary output value s_i of c_i . Accordingly define $t_{i,j}$ by $t_{i,j}(s_j) = s_i$: by construction, it satisfies $t_{i,j} \circ f_{c_j} = f_{c_i}$.

The corollary then follows as a reiteration of Prop. (3.2): assuming such a mapping and under the conditions of Prop. (3.1),

$$\begin{aligned} \forall \mathbf{b} \in \mathcal{B}, \quad \forall \varepsilon \in \mathcal{E} \mathcal{X}, \\ \psi_{c_i}(a\mathbf{b}, \varepsilon) &= f_{c_i}(\varepsilon|_{F_{c_i}(a\mathbf{b})}) \\ &= t_{i,j} \circ f_{c_j}(\varepsilon|_{F_{c_j}(\mathbf{b})}) = t_{i,j}(\psi_{c_j}(\mathbf{b}, \varepsilon)). \end{aligned}$$

□

Therefore, the conditions for a perfect sensory image of the retinal smear existing up to transduction are rather permissive, or at least as much as we could hope for. However, the problem with dissimilar sensitivity function is not one of existence but one of practical exploitability. Indeed, the sensory property of Eq. (3.2) is *readily* available in the identical sensibility case thanks to the equality comparison. On the contrary, the distinct sensitivity case seems to imply that the agent should also learn the transduction mapping in some sort in order to observe Prop. (3.3). And while there are only $|\mathcal{C}|^2$ pairs of sensels possibly so related, there are as many as $|S_i|!$ such transduction mappings if we only restrict them to bijections: the space of possible candidate grows exponentially with the number of output values. In fact, the works that exploit this structure in sensory space either implicitly use the identical sensitivity assumption (Le Hir, Sigaud, and Alban Laflaquière, 2018; Montone, J. Kevin O'Regan, and Alexander V. Terekhov, 2015) or make use of clustering techniques for complexity reduction (Alban Laflaquière, 2017). However, our assumptions (see 2.3.2.3) do not allow for the agent *a priori* performing such computation. Instead, our hypothesis is that dealing with this added complexity should become effectively possible at a later stage where the agent has learned approximation capabilities with which to perform sensory clustering.

3.2.1.2 An algebraic conversation of conservation

The discussed property of sensation substitution discussed relates the output states of separate sensels at the distinct steps of the sensorimotor experience the agent encounters by performing its motor actions. It therefore seems plausible that chaining motor actions exhibiting this behavior should result in sequences of sensels successively transferring their stimuli to one another. In the event of such combinations, the “retinal smears” follow the same sequencing order as that of the displacements that caused them; importantly, though, these smears are naively available to the naive agent. As a consequence, a naive agent capable of recognizing the sensory signature of sequences of smears might infer from it the structure of its originating motor actions.

In the following, for any motor action $a \in \mathcal{A}$ and any two sensels $c_i, c_j \in \mathcal{S}$ let us denote $c_i \xrightarrow{a} c_j$ the previously discussed relation

$$\forall \mathbf{b} \in \mathcal{B}, F_{c_i}(a\mathbf{b}) = F_{c_j}(\mathbf{b}) \quad (3.4)$$

which we also state by saying that c_j is a *successor* of c_i by a . The natural framework for assessing the iterated relation corresponding to sequences of smear is that of actions that maintain such relations across most (if not all) possible pairs of sensels. Accordingly, let us consider those motor actions a of $\mathcal{A}$ which verify

$$\forall c_j \in \mathcal{C}, \exists c_i \in \mathcal{C} \text{ such that } c_i \xrightarrow{a} c_j. \quad (3.5)$$

Following the previous discussion, these actions are precisely those such that the agent does not interact with any new region of space after it has performed them, by which we thus refer to as *conservative actions*.

In fact, the previously exploited Fig. 3.1 illustrates such a conservative action: to each sensel clearly corresponds a successor in the same way as that highlighted in the previous part. Conservativeness of action, however, is a significantly restrictive property: not only is it most suited to locomotion as illustrated by the example, but also a locomotion action may fail to be conservative purely on the basis that the agent happens to “lack” successor sensels.

An interesting result is that some light formal work shows that such motor actions admit an alternative interpretation as permutations of sensels along

Proposition 5. • *Let $a \in \mathcal{A}$ be a conservative action, then there exists a unique map*

$$\begin{aligned} \sigma_a: \mathcal{C} &\rightarrow \mathcal{C} \\ c &\mapsto c' \end{aligned} \quad (3.6)$$

such that

$$\sigma_a(c) = c' \Leftrightarrow c \xrightarrow{a} c'. \quad (3.7)$$

- *If there is a finite number of sensels, then for any conservative action $a \in \mathcal{A}$ its successor map σ_a is bijective.*

Proof. • Let $a \in \mathcal{A}$ conservative and $c \in \mathcal{C}$. By conservativity $\exists c' \in \mathcal{C}$ such that $c \xrightarrow{a} c'$. Let $c'' \in \mathcal{C}$ such that $c \xrightarrow{a} c''$, then

$$\forall \mathbf{b} \in \mathcal{B}, F_c(\mathbf{b}) = F_{c'}(a\mathbf{b}) \text{ and } F_c(\mathbf{b}) = F_{c''}(a\mathbf{b})$$

so that

$$\forall \mathbf{b} \in \mathcal{B}, F_{c'}(a\mathbf{b}) = F_{c''}(a\mathbf{b}).$$

But $a : \mathcal{B} \rightarrow \mathcal{B}$ must be surjective because it is bijective, so that *all* $\mathbf{b} \in \mathcal{B}$ can be written $a\mathbf{b}'$ for some $\mathbf{b}' \in \mathcal{B}$. Therefore $F_{c'} = F_{c''}$, from which $c' = c''$ ¹: successor sensels are necessarily unique. We therefore declare σ_a to be the map that takes each sensel $c \in \mathcal{C}$ to its *unique* successor sensel.

- Let a be a conservative action, and let $c, c' \in \mathcal{C}$ be sensels such that $\sigma_a(c) = \sigma_a(c')$. From this it follows that

$$\forall \mathbf{b} \in \mathcal{B}, F_c(\mathbf{b}) = F_{c''}(a\mathbf{b}) \text{ and } F_{c'}(\mathbf{b}) = F_{c''}(a\mathbf{b})$$

for some common successor $c'' \in \mathcal{C}$. But it entails in particular

$$\forall \mathbf{b} \in \mathcal{B}, F_c(\mathbf{b}) = F_{c'}(\mathbf{b})$$

that is $F_c = F_{c'}$, which further yields $c = c'$: σ_a is injective.

From injectivity of σ_a , it follows that $|\sigma_a(\mathcal{C})| = |\mathcal{C}|$. But since $\mathcal{C}$ is finite it in turns follows that $\sigma_a(\mathcal{C}) = \mathcal{C}$, i.e. that σ_a is also surjective.

□

The manner in which this viewpoint is equivalent is made apparent by considering the algebraic structure of permutations of $\mathcal{C}$: by definition, it is the textbook example of a group of (bijective) transformations $\mathcal{C} \rightarrow \mathcal{C}$ with internal operation supplied by function composition. Remarkably, it can be shown that the algebraic structure of conservative motor actions and the combinatorial one of their sensel permutations coincide:

¹*Stricto sensu*, the definition allows for two distinct sensels $c' \neq c''$ to have the same receptive field function F . However, this describes the case where whatever the state of the interaction with its environment, these two sensors always survey the same region of space. Considering that they are already assumed to have equal sensitivity, they are entire duplicate of one another as far as the sensory experience is concerned. This case can therefore be managed by quotienting over the equivalence classes of such redundant sensels; we prefer the notational simplicity offered by the implicit assumption that they simply don't occur.

Proposition 6. • Let $\mathcal{A}_{\mathcal{C}} \subset \mathcal{A}$ be the subset of all conservative actions. Then $\mathcal{A}_{\mathcal{C}}$ is in fact a subgroup of $\mathcal{A}$.

- Let $\sigma : a \in \mathcal{A}_{\mathcal{C}} \mapsto \sigma_a \in \mathfrak{S}_{\mathcal{C}}$ map conservative actions to their associated sensel permutation, then σ defines a group isomorphism $\mathcal{A}_{\mathcal{C}} \cong \sigma(\mathcal{A}_{\mathcal{C}})$.

Proof. • $\mathcal{A}_{\mathcal{C}} \subset \mathcal{A}$ by its very definition, therefore we need only prove it is actually a group.

- $\forall c \in \mathcal{C}, c \xrightarrow{e} c$ with e the identity action: e is conservative.
- Let a and a' be conservative actions, and $c \in \mathcal{C}$: since $a \in \mathcal{A}_{\mathcal{C}}$, $\exists c' \in \mathcal{C}$ such that $c \xrightarrow{a} c'$. But since $a' \in \mathcal{A}_{\mathcal{C}}$ too, there also exists $c'' \in \mathcal{C}$ verifying $c' \xrightarrow{a'} c''$, so that finally

$$\forall c \in \mathcal{C}, \forall a, a' \in \mathcal{A}_{\mathcal{C}}, \exists c'' \text{ such that } c \xrightarrow{a'a} c''$$

that is $a'a$ is conservative itself.

- Let $a \in \mathcal{A}_{\mathcal{C}}$ and let σ_a be its successor map $\mathcal{C} \rightarrow \mathcal{C}$. For any $c \in \mathcal{C}$, since σ_a is surjective, we have $c = \sigma_a(c')$ for some $c' \in \mathcal{C}$, or equivalently

$$\forall c \in \mathcal{C}, \exists c' \text{ such that } c' \xrightarrow{a} c.$$

Finally, since $c' \xrightarrow{a} c \Leftrightarrow c \xrightarrow{a^{-1}} c'$ it follows that a^{-1} is conservative too.

- The final point of the previous proof shows that for any conservative action $a \in \mathcal{A}_{\mathcal{C}}$, we have $\sigma_{a^{-1}} = (\sigma_a)^{-1}$. Let a and a' be two such actions of $\mathcal{A}_{\mathcal{C}}$, and c be any sensel of $\mathcal{C}$. Then, let $c' = \sigma_a(c)$ and $c'' = \sigma_{a'}(c')$ be the corresponding successors, by definition we have $c'' = (\sigma_{a'} \circ \sigma_a)(c)$. Alternatively, as a displacement of receptive fields we can observe that

$$\forall \mathbf{b} \in \mathcal{B}, F_c((a'a)\mathbf{b}) = F_c(a'(a\mathbf{b})) = F_{\sigma_a(c)}(a'\mathbf{b}) = F_{\sigma_{a'} \circ \sigma_a(c)}(\mathbf{b}) = F_{c''}(\mathbf{b}).$$

Since $a'a$ is conservative by account of $\mathcal{A}_{\mathcal{C}}$ being a group, this therefore identifies sensel c'' as the successor of c by $a'a$ and shows that σ is a group homomorphism. Then, because two formal $a, a' \in \mathcal{A}_{\mathcal{C}}$ that differ




FIGURE 3.2: Adaptation of the previous conservative setup to an agent with bounded receptive field. As opposed to the ideal situation previously discussed, bordering sensels may fail to have suitable successors

in their motor effect induce distinct permutations, this homomorphism is even an isomorphism onto its image $\sigma(\mathcal{A}_{\mathcal{C}})$.²

□

3.2.1.3 Partial successor map and quasiconservative actions

The previous arguments all occurred in the particular case of conservative actions, where moves of the agent caused successor relations on its total sensor array. Although it is useful in the insight it gives into the combinatorial structure of sensel shifts, it is a highly idealistic case which restrictive conditions prevent from occurring in a practical measure. The structure may however be generalized to account for some leeway while still preserving the argument. To illustrate this process, consider the agent of Fig. 3.2 following the template of previous systems. Its camera sensor is discretized into N identical sensels, which we label by the corresponding integer set $\{1, \dots, N\}$ (or alternatively by $\{c_1, \dots, c_N\}$ where ambiguity may form). As actions, it is able to stay still via its identity e as well as move left (respectively right) by 1 “unit” via action a_l (resp. a_r). The unit considered here is so that it corresponds to shifting by (the width of) 1 receptive field of sensels.

The setup is therefore almost identical to that which illustrated proper conservativity in Fig. 3.1. However, the required property is not verified here: for example the leftmost sensel c_1 does not have a successor for the “move left” a_l action, and the same can be said of c_N relative to a_r . This

²According to the mathematical definition, it remains to be shown that it is *injective*. As before, this necessitates a purely formal precaution of quotienting out the pathological case of several “distinct” actions which are nonetheless indiscernible for the sensorimotor experience.

boundary specific effect makes clear that actions are not conservative *per se*, but relative to specific sensory capabilities (here, it is “lacking” sensels): had the agent had instead its receptive fields extending infinitely, it would have corresponded to the previous ideal conservative case. We can still, however, formulate the sensel successor functions σ . of actions a_l and a_r at the cost of taking them to be partial functions. Consider $\overline{\mathcal{C}} = \mathcal{C} \cup \{*\}$ where the added $*$ acts as a null symbol, then the partial successor functions can be written as (proper) functions

$$\sigma_a : \overline{\mathcal{C}} = \mathcal{C} \cup \{*\} \rightarrow \overline{\mathcal{C}}$$

under the condition that $\sigma_a(*) = *$. Using this notation, $\sigma_a(c) = c'$ retains the previous meaning while $\sigma_a(c) = *$ denotes that sensel c has no successor for action a . Therefore, this construction is entirely general: any action a has a valid such successor map, including $\sigma_a : c \mapsto *$ should it not cause *any* shift between sensels.

Importantly, the structure of these generalized successor maps extends that of the permutations previously discussed: because the argument made in Prop. 5 for injectivity still holds, the generalized successor maps are *partial bijections*. As functions of $\overline{\mathcal{C}}$ this means that their restriction to sensels which have a successor

$$\sigma_a^{-1}(\mathcal{C} \subset \mathcal{C} \cup \{*\}) \rightarrow \overline{\mathcal{C}}$$

is injective, and therefore a bijection from its domain (i.e. predecessor sensels) onto its image (successor ones). In the illustrative example, this is figured by the maps

$$\begin{aligned} \sigma_{a_l} : i \in \{1, \dots, N\} &\mapsto \begin{cases} * & \text{for } i = 1 \\ i - 1 & \text{for } 2 \leq i \leq N \end{cases} \\ \sigma_{a_r} : i \in \{1, \dots, N\} &\mapsto \begin{cases} * & \text{for } i = N \\ i + 1 & \text{for } 1 \leq i \leq N - 1 \end{cases} \end{aligned}$$

which induce bijections (denoted by the expressions on the second lines) between subsets $\{1, \dots, N - 1\}$ and $\{2, \dots, N\}$. These bijections are in fact mutual reciprocal inverses, mirroring the inverse relation between a_l and a_r . However, this statement cannot be extended to the entire partial transformations in the usual way: indeed, we have

$$\sigma_{a_l} \circ \sigma_{a_r} : i \in \{1, \dots, N\} \mapsto \begin{cases} i & \text{for } i \in \{1, \dots, N - 1\} \\ * & \text{else.} \end{cases}$$




FIGURE 3.3: Partial maps generalizing the sensor permutations of conservative actions for the case where some sensors have no successor. The red arrows denote sensors which are mapped to the special ‘*’ element.

That is, composing generalized successor maps of inverse motor actions yields somewhat *partial identities* as opposed to the true identity of permutations in the conservative case. This particular trait is intrinsically characterized by the relation

$$\sigma_a \circ \sigma_{a^{-1}} \circ \sigma_a = \sigma_a$$

which generalizes the nice property of inverses found in groups without requiring reference to a particular identity function³. It also characterizes the structure of generalized successor maps as that of an *inverse semigroup* (Mark V Lawson, 2018; Preston, 1954), showing that it is the natural framework of sensor shifts for cases including loss on information at the boundaries. The previous results can then be reformulated as

Proposition 7. *Let $a \in \mathcal{A}$ be any motor action*

- *It defines a generalized successor map $\sigma_a : \overline{\mathcal{C}} = \mathcal{C} \cup \{*\} \rightarrow \overline{\mathcal{C}}$ via*

$$\sigma_a(c) = \begin{cases} c' & \text{if and only if } c \xrightarrow{a} c' \\ * & \text{else.} \end{cases}$$

- *The collection $I_{\mathcal{A}} = \{\sigma_a, a \in \mathcal{A}\}$ of such generalized successor maps is an inverse semigroup, corresponding to one of partial transformations of $\mathcal{C}$*

³Similarly, partial identities ι are instead internally characterized by their *idempotence*, that is they verify $\iota \circ \iota = \iota$.

Conclusion

The previous discussion has shown how the formalism we introduced of 2 enables us to account for how displacements of the agent affect the localization of its interaction with its environment. We detailed how it could be explicitated as the shifts of its receptive fields, and that for conservative actions these shifts actually share the algebraic structure of their corresponding actions. Furthermore, a similar if weaker correspondence holds in the more general nonconservative case, with permutations becoming partial because of the loss of information. In the next part, we detail how these shifts manifest in the sensorimotor flow of the agent. We show that they cause predictable sensory sequences, and that learning these patterns endows the naive agent with the (*a priori* hidden) knowledge of its motor action algebra.

3.2.2 Perceptive properties of sensory prediction

The idea that perceptive processes involve active inference of the content of one's sensorimotor experience is ubiquitous in our understanding of their mechanisms. It implies an (at least partly) top-down view of perception, where one is continually verifying and refuting hypotheses to build explanations of their sensory experience. Therefore, this comes as a rebuttal of classical "sandwich" models –Sense, Plan, Act in robotics; Perception, Cognition, Action in cognitive psychology– where perception is entirely mediated by the bottom-up extraction of information from sensory samples (Burr, 2017). In particular, a large body of works in perceptual psychology (Gibson, 1966; Helmholtz, 2013) have sought to explain the subjective phenomenon of perception as a set of (mainly unconscious) cognitive inference rules.

3.2.2.1 Using predictable actions as motor proxies

The predictive processing approach generally relies on some internal model of system states. This model, taken as the agent's understanding of its interaction, is precisely that on which it is inferring the dynamics of its experience. In the particular case of the sensorimotor experience, the agent should therefore be able to tell what transitions between states in the model are induced by its motor capabilities. This model should also accurately and efficiently capture information relevant to its objectives. Under these conditions, the agent should then be able to skillfully exercise its motor capabilities to navigate this state space.




FIGURE 3.4: Redundant configurations that result in the same end effector pose x need not remain so when affected by the same action a . This situation results in the impossibility for a mapping predicting the outcome of a purely in poses to exist.

An important detail is that this state space should be *internal*. Therefore, the naiveness assumptions outlined in 2.3.2.3 require it be discovered from the sensorimotor flow. This viewpoint is essentially that of *State Representation Learning* (Lesort et al., 2018) which we will adopt. Let us consider an arbitrary state space X with mapping $\phi_X : \mathcal{S}_{\mathcal{C}} \rightarrow X$; this mapping, taking naive observation to their representational state, is to be internally computed by the agent. This is, in particular, the case of the discovery of (a representation of) the pose space $\mathcal{X}$ and forward geometric model f found in (V. Marcel, Argentieri, and Gas, 2019). It illustrates how the equivalent identification of end effector poses as the result of quotienting out spatial redundancies shows that it is indeed a compressed representation of the sensorimotor interaction.

The algebraic formalism of motor actions provides an account of state transitions occurring during the sensorimotor experience. Indeed, action $a \in \mathcal{A}$ induces deterministic transitions in states $x \in X$ precisely through a mapping $\Pi_{X,a} : X \rightarrow X$ defined by

$$\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \phi_X(\psi(a\mathbf{b}, \varepsilon)) = \Pi_{X,a}(\phi_X(\psi(\mathbf{b}, \varepsilon))). \quad (3.8)$$

Learned by the agent, this mapping encodes the information as to the outcome of performing a in terms of states $x \in X$. We therefore say that such actions are *predictable over X* .

Figure 3.4 illustrates how the example of pose space shows that all actions need not be predictable in this way. It pictures the case where $a \in \mathcal{A}$ adds a

given, fixed, angular increment to one joint on the robotic arm. Importantly, it makes it so that configurations $\mathbf{b}, \mathbf{b}' \in \mathcal{B}$ sharing a same pose $\mathbf{x} = f(\mathbf{b}) = f(\mathbf{b}') \in \mathcal{X}$ yield distinct poses $\mathbf{x}' = f(a\mathbf{b}) \neq f(a\mathbf{b}') = \mathbf{x}''$ upon performing a ; and indeed the impossibility of determining a mapping $\Pi_{\mathcal{X},a}$ precisely comes from this ambiguity. Since by design the prediction mapping $\Pi_{\mathcal{X},a}$ needs to operate on poses $\mathbf{x} \in \mathcal{X}$, whenever $\mathbf{b}$ and $\mathbf{b}'$ share their pose along

$$\phi_{\mathcal{X}}(\psi(\mathbf{b}, \varepsilon)) = f(\mathbf{b}) = f(\mathbf{b}') = \phi_{\mathcal{X}}(\psi(\mathbf{b}', \varepsilon))$$

it should also be true that

$$\begin{aligned} f(a\mathbf{b}) &= \phi_{\mathcal{X}}(\psi(a\mathbf{b}, \varepsilon)) = \Pi_{\mathcal{X},a}(f(\mathbf{b})) \\ &= \Pi_{\mathcal{X},a}(f(\mathbf{b}')) = \phi_{\mathcal{X}}(\psi(a\mathbf{b}', \varepsilon)) = f(a\mathbf{b}'). \end{aligned} \quad (3.9)$$

This implication can clearly be generalized in terms of kernel sets, with them replacing the distinct configurations $\mathbf{b}, \mathbf{b}'$ that share their corresponding end-effector pose. Indeed, for arbitrary state spaces X we have

Proposition 8. *Let X be a state space with corresponding mapping $\phi_X : \mathcal{S}_{\mathcal{C}} \rightarrow X$, and let $a \in \mathcal{A}$ be a motor action.*

Then a is predictable over X if and only if for any $\varepsilon \in \mathcal{C}$ it maps kernel sets of

$$(\phi_X \circ \psi)_{\varepsilon} : \mathbf{b} \in \mathcal{B} \mapsto \phi_X(\psi(\mathbf{b}, \varepsilon))$$

to themselves, that is if and only if

$$\forall \mathbf{b}, \mathbf{b}' \in \mathcal{B}, (\phi_X \circ \psi)_{\varepsilon}(\mathbf{b}') = (\phi_X \circ \psi)_{\varepsilon}(\mathbf{b}) \Rightarrow (\phi_X \circ \psi)_{\varepsilon}(a\mathbf{b}) = (\phi_X \circ \psi)_{\varepsilon}(a\mathbf{b}').$$

Interestingly, this result accounts for the fact that distinct state spaces can result in various sets of predictable actions. For example, the condition adapted to the pose state $(\mathcal{X}, f)$ makes it so that pose predictable actions are precisely those that induce displacement of the end effector. Indeed, Eq. (3.9) shows that the pose $f(a\mathbf{b})$ obtained as a result of performing (pose predictable) action a depends only on its previous value $f(\mathbf{b})$. On the other hand, the same condition for sensory state $\mathbf{s} \in \mathcal{S}$ is not trivially verified of displacements. We can observe instead that modulating the output of a sensory element (say, by modifying its numerical gain) is naturally predictable over $\mathcal{S}$.

As done previously in the case of conservative actions, we dwell on the

relations *between* functions $\Pi_{X,a}$ for varying $a \in \mathcal{A}$. Following our framing these mappings as “predictions”, we could reasonably expect to be able to chain them in a meaningful way compatible with the corresponding sequences of actions. Quite familiarly now, we have

Proposition 9. • *The identity action $e \in \mathcal{A}$ is predictable over X , with associated mapping $\Pi_{X,e} = \text{id}_X$*

- *Let a and a' be predictable over X , then $a'a$ is predictable over X and $\Pi_{X,a'a} = \Pi_{X,a'} \circ \Pi_{X,a}$*
- *If any invertible action a is predictable over X and its associated mapping $\Pi_{X,a}$ is invertible, then a^{-1} is predictable over X . Moreover, whenever a and a^{-1} are predictable we have $\Pi_{X,a^{-1}} = \Pi_{X,a}^{-1}$.*

Proof. • This case trivially follows from Eq. (3.8) boiling down to the tautological

$$\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \phi_X(\psi(\mathbf{b}, \varepsilon)) = \text{id}_X(\phi_X(\psi(\mathbf{b}, \varepsilon))).$$

- Let a, a' be two such predictable actions. Then, for any $\mathbf{b} \in \mathcal{B}$ and any $\varepsilon \in \mathcal{E}\mathcal{X}$ we have

$$\begin{aligned} \phi_X(\psi(a'a\mathbf{b}, \varepsilon)) &= \phi_X(\psi(a'(a\mathbf{b}), \varepsilon)) \\ &= \Pi_{X,a'}(\psi(a\mathbf{b}, \varepsilon)) && \text{by Eq. (3.8) applied to } a' \\ &= \Pi_{X,a'}(\Pi_{X,a}(\psi(\mathbf{b}, \varepsilon))) && \text{by Eq. (3.8) applied to } a \\ &= (\Pi_{X,a'} \circ \Pi_{X,a})(\mathbf{b}, \varepsilon). \end{aligned}$$

- Let a be an invertible predictable action with $\Pi_{X,a}$ invertible as a mapping $X \rightarrow X$, then for any $\mathbf{b} \in \mathcal{B}$ and $\varepsilon \in \mathcal{E}\mathcal{X}$

$$\begin{aligned} \phi_X(\psi(a\mathbf{b}, \varepsilon)) &= \phi_X(\psi(aa^{-1}\mathbf{b}', \varepsilon)) && \text{for some } \mathbf{b}' \in \mathcal{B} \\ &= \Pi_{X,a}(\phi_X(\psi(a^{-1}\mathbf{b}', \varepsilon))) && \text{because } a \text{ is predictable} \\ &= \phi_X(\psi(\mathbf{b}', \varepsilon)) && \text{because } aa^{-1} = e. \end{aligned}$$

But left composing the last two lines with $\Pi_{X,a}^{-1}$ yields

$$\forall \mathbf{b}' \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \phi_X(\psi(\mathbf{b}', \varepsilon)) = \Pi_{X,a}^{-1}(\phi_X(\psi(\mathbf{b}', \varepsilon)))$$

which precisely shows that $\Pi_{X,a}^{-1}$ verifies Eq. (3.8) for action a^{-1} .

□

As usual, this result is more succinctly formulated as a relation of structures:

Corollary 2. *Let $\mathcal{A}_{\Pi_X} \subset \mathcal{A}$ be the set of all actions that are predictable over X , then*

- $\mathcal{A}_{\Pi_X}$ is a submonoid of $\mathcal{A}$
- The mapping

$$\begin{aligned} \Pi_X: \mathcal{A}_{\Pi_X} &\rightarrow \mathcal{F}(X, X) \\ a &\mapsto \Pi_{X,a} \end{aligned} \tag{3.10}$$

taking predictable actions to their state transition mappings is a monoid homomorphism

- *For each subgroup $A \subset \mathcal{A}_{\Pi_X}$ of predictable actions, the restriction of Π_X to A is a group homomorphism.*

This result provides a formal basis for the *ideomotor principle*, which states that actions (corresponding to the $a \in \mathcal{A}$ here) are somewhat equivalent to their effects (i.e. the $\Pi_{X,a}$) (A. Stock and C. Stock, 2004). It shows that this desired equivalence is necessarily true of the compositionality captured by the algebraic structure of $\mathcal{A}$. Importantly, it does so by providing an *isomorphism* between the composition and action and that of their effects instead of an *identity*. To reiterate, it does not ask for any common domain for the material implementation of motor actions and that of their perceptual effects. It is thus more parsimonious than approaches taking the Common Coding route (as found e.g. in (Prinz, 1997)) in this aspect.

An important consequence that follows from Prop. (2) is that *the outcome of a sequence of predictable actions is equal to the sequence of the outcomes of the corresponding action factors*, or algebraically

$$\forall a_i, 1 \leq i \leq n, \Pi_{X,a_n \dots a_1} = \Pi_{X,a_n} \circ \dots \circ \Pi_{X,a_1}.$$

This provides a way for agents which have learned their action-effects correspondences Π_{X,a_i} to guide action selection in terms of intended outcomes. Therefore, a logical next step is to determine under which conditions and by which means these correspondences are effectively discoverable.

3.2.2.2 Prediction in the case of conservation

The previous examination was somewhat generic in its discussion of the state space X on which the prediction was built. As outlined before, settling on a suitable state space and *internally* computing it is in fact a central part of the problem in the bootstrapping context of SMCT. Furthermore, the theory encourages a staunch antirepresentationalist perspective in that it dismisses the need for an internal “re”presentation of the world (J K O’Regan and Noë, 2001). Accordingly we turn to the most straightforward case where the state space X is directly sensory space $\mathcal{S}_{\mathcal{C}}$, with trivial “representational” mapping $\phi_{\mathcal{S}_{\mathcal{C}}} = \text{id}_{\mathcal{S}_{\mathcal{C}}}$.

Let us consider any conservative action $a \in \mathcal{A}_{\mathcal{C}}$ as defined: we showed in 3.2.1.2 that a is equivalent to a permutation σ_a of sensels, i.e. a bijective mapping from $\mathcal{C}$ onto itself. But following the same exposition as that made in 3.2.1.1, in the case of equal sensitivity functions

$$\forall c, c' \in \mathcal{C}, f_c = f_{c'}$$

the permutation of receptive fields under action a becomes one of sensory values. Let us denote $\mathbf{s} = \psi_{\mathcal{C}}(\mathbf{b}, \varepsilon) \in \mathcal{S}_{\mathcal{C}}$ the sensory output at configuration given by $\mathbf{b} \in \mathcal{B}, \varepsilon \in \mathcal{E}\mathcal{X}$. Then, performing action a yields the new sensory value $\mathbf{s}'$ defined by

$$\begin{aligned} \mathbf{s}' &= (\mathbf{s}'_c)_{c \in \mathcal{C}} \\ &= (\psi_c(a\mathbf{b}, \varepsilon))_{c \in \mathcal{C}} \\ &= (f_c(\varepsilon|_{F_c(a\mathbf{b})}))_{c \in \mathcal{C}} \\ &= (f_{\sigma_a(c)}(\varepsilon|_{F_{\sigma_a(c)}(\mathbf{b})}))_{c \in \mathcal{C}} \\ &= (\psi_{\sigma_a(c)}(\mathbf{b}, \varepsilon))_{c \in \mathcal{C}} \\ &= (\mathbf{s}_{\sigma_a(c)})_{c \in \mathcal{C}} = \mathbf{s} \circ \sigma_a. \end{aligned} \tag{3.11}$$

But this deterministic relation is precisely that of action a being *predictable over sensory space*. It can even be made to accomodate the more general case where transduction is required between the output of distinct sensels, resulting in

Proposition 10. Assume that sensels $c, c' \in \mathcal{C}$ all follow the condition of Prop. (4) so that there exist transduction mappings $t_{c',c} : \mathcal{S}_c \rightarrow \mathcal{S}_{c'}$ for all such pairs.

Then, every conservative action $a \in \mathcal{A}_{\mathcal{C}}$ is predictable over $\mathcal{S}_{\mathcal{C}}$ with associated

sensory prediction

$$\begin{aligned} \Pi_{S,a}: \mathcal{S}_{\mathcal{C}} &\rightarrow \mathcal{S}_{\mathcal{C}} \\ (s_c)_{c \in \mathcal{C}} &\mapsto \left(t_{c, \sigma_a(c)}(s_{\sigma_a(c)}) \right)_{c \in \mathcal{C}} \end{aligned} \quad (3.12)$$

In particular, in the identical sensitivity case –where all $t_{c',c}$ are equal to the identity $\text{id}_{\mathcal{S}_c}$ – the sensory prediction is given by the permutation of components

$$\begin{aligned} \Pi_{S,a}: \mathcal{S}_{\mathcal{C}} &\rightarrow \mathcal{S}_{\mathcal{C}} \\ s &\mapsto s \circ \sigma_a \end{aligned} \quad (3.13)$$

Corollary 3. • The conservative subgroup $\mathcal{A}_{\mathcal{C}}$ is a subgroup of the sensory predictable actions $\mathcal{A}_{\Pi_{S,\mathcal{C}}}$

- The prediction mapping of conservative actions

$$\begin{aligned} \Pi_S: \mathcal{A}_{\mathcal{C}} &\rightarrow \mathfrak{S}_{\mathcal{S}} \\ a &\mapsto \Pi_{S,a} \end{aligned} \quad (3.14)$$

is a group homomorphism.

This result can be made *vastly* stronger under the assumption that for any two motor actions a and a'

$$[\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E} \mathcal{X}, \psi_{\mathcal{C}}(a' \mathbf{b}, \varepsilon) = \psi_{\mathcal{C}}(a \mathbf{b}, \varepsilon)] \Rightarrow a = a'. \quad (3.15)$$

This condition, which we will now explicitly assume, asks that no two actions a, a' verify both that

- a and a' are “symbolically” distinct as per $a \neq a'$
- a and a' are indiscernible up to the sensorimotor interaction, i.e. at any point $\mathbf{b}, \varepsilon$ of the interaction performing either a or a' will yield the exact same sensory results.

This assumption is therefore mostly technical since we take the sensorimotor interaction to be the unsurpassable “ground truth” from which other symbols derive their meaning; it seems reasonable to argue that noncompliant cases are fringe abuses of the possibilities given by the formalism. Nevertheless, we can also take the viewpoint that such *distinct but indiscernible* actions might exist. In this case, then, observe that the assertion of Eq. (3.15) defines

an equivalence *via*

$$a \sim_\psi a' \Leftrightarrow [\forall \mathbf{b} \in \mathcal{B}, \forall \varepsilon \in \mathcal{E}\mathcal{X}, \psi_{\mathcal{C}}(a'\mathbf{b}, \varepsilon) = \psi_{\mathcal{C}}(a\mathbf{b}, \varepsilon)].$$

One can then quotient out all these strongly redundant cases all the same and reformulate Eq. (3.15) by replacing actions a and a' with their respective classes for $\sim_\psi$. Crucially, in the case of predictable actions this additional consideration is equivalent to the assertion that no two distinct (predictable) actions a can share a common sensory prediction mapping $\Pi_{\mathcal{S},a}$. This is therefore entirely identical to the injectivity of the $\Pi_{\mathcal{S}}$ mapping, from which we can conclude

Theorem 1. *The sensory prediction mapping $\Pi_{\mathcal{S}}$ gives a natural group isomorphism*

$$\mathcal{A}_{\mathcal{C}} \cong \Pi_{\mathcal{S}}(\mathcal{A}_{\mathcal{C}})$$

between conservative actions and their associated sensory predictions

In comparison to the previous result of Corollary (2) obtained for arbitrary state spaces, this one crucially makes the construction internal. Indeed the definition of *predictability over X* only guarantees that the prediction mappings $\Pi_{X,a}$ exist as formal functions; it does not ask for them to be discoverable by the agent. However the particular case of predictability induced by conservativity, as described in Prop. (10), is one where both the state space and the representational mapping are defined in sensory terms, which are *a priori* internal objects as per the assumptions of 2.3.2.3. We further detail in Section 3.3 a simulation that shows how the naive agent can actually compute the sensory predictions $\Pi_{\mathcal{S}}(a)$ using only its available rules.

3.2.2.3 Degeneracy in the semigroup case

To conclude the discussion, we examine how the ideal isomorphism we just isolated appears in non conservative cases. We again refer to the situation of Fig. 3.2 of a mobile camera moving along a 1-d axis using its translations a_l and a_r . We have shown in 3.2.1.3 that finiteness of the camera, which sensel set $\mathcal{C} = [|1, \dots, N|]$ comprises $N \in \mathbb{N}$ pixels, entails non conservativity of the translations. Correspondingly, the (generalized) sensel successor maps become *partial* bijections $\sigma_a : \overline{\mathcal{C}} \rightarrow \overline{\mathcal{C}}$. Let us also consider a virtual extended agent defined thus: it possesses the same action set $\mathcal{A} = \langle a_l, a_r \rangle$ but its sensel set $\hat{\mathcal{C}} = \mathbb{Z}$ infinitely extends $\mathcal{C}$ in both directions. Importantly actions a_l




FIGURE 3.5: Partial maps generalizing the sensel permutations of conservative actions (*repeated*).

and a_r are conservative for this extended agent, so that its maps $\hat{\sigma}_a$ are true permutations of $\mathcal{C}$.

We have shown in Prop (6) that in the conservative case, successor maps faithfully account for algebraic combinations of actions. In particular, this means that $\hat{\sigma}_{a_l} \circ \hat{\sigma}_{a_r} = \hat{\sigma}_{a_l a_r} = \hat{\sigma}_e = \text{id}_{\mathcal{C}}$ for the extended agent. These relations do not hold for the non conservative case. Recall, as previously mentioned in 3.2.1.3 that the partial bijections are given by

$$\sigma_{a_l} : i \in \{1, \dots, N\} \mapsto \begin{cases} * & \text{for } i = 1 \\ i - 1 & \text{for } 2 \leq i \leq N \end{cases}$$

$$\sigma_{a_r} : i \in \{1, \dots, N\} \mapsto \begin{cases} * & \text{for } i = N \\ i + 1 & \text{for } 1 \leq i \leq N - 1 \end{cases}$$

We still have $a_l = a_r^{-1}$ as motor actions, therefore $\sigma_{a_l a_r} = \sigma_e = \text{id}_{\mathcal{C}}$ continues to hold. However, we previously saw that

$$\sigma_{a_l} \circ \sigma_{a_r} : i \in \{1, \dots, N\} \mapsto \begin{cases} i & \text{for } 1 < i \leq N - 1 \\ * & \text{else.} \end{cases}$$

so that instead of the regular identity the conservative case would yield, composing “inverse” partial bijections only results in a *partial identity*. This raises a substantial issue in our earlier argument: in the non conservative case, we cannot directly compare (sequences of) actions by performing the corresponding comparison on (sequences of) permutations. In fact, the identities are further partialized as the loss of information due to $*$ results accumulates in longer chains of compositions. Indeed, consider longer sequences




FIGURE 3.6: Composing inverse partial mappings only results in a partial identity.

of $k \in \mathbb{N}$ iterations as per

$$\sigma_{a_l}^k \circ \sigma_{a_r}^k : i \in \{1, \dots, N\} \mapsto \begin{cases} i & \text{for } i \leq N - k \\ * & \text{else.} \end{cases}$$

This shows that sequences $a_l^k a_r^k$ and which products as actions all collapse to the identity $e \in \mathcal{A}$ result, as partial bijections, in gradually smaller restrictions of the identity permutation $\text{id}_{\mathfrak{S}_{\mathcal{C}}}$. Fig. 3.7 illustrates how the support of the partial bijection becomes arbitrarily smaller when the length of the sequence increases, with the depicted case corresponding to $k = 3$.

Still, we can yet find informative structure by comparing with the case of the corresponding extended agent (pictured translucent in Fig. 3.2). Observe indeed that the extended sensel permutations are given by

$$\hat{\sigma}_{a_l} : i \in \mathbb{Z} \cup \{*\} \mapsto i - 1 \in \mathbb{Z} \cup \{*\}$$

$$\hat{\sigma}_{a_r} : i \in \mathbb{Z} \cup \{*\} \mapsto i + 1 \in \mathbb{Z} \cup \{*\}$$

We can see that the true sensel permutations extend those of the partial case in the sense that whenever a successor relation $\sigma_a(i) = j$ holds for the finite agent, it (or more precisely $\hat{\sigma}_a(i) = j$) is also true of the extended one. Conversely, partial bijections of the finite case can be obtained by applying a partial identity to the corresponding infinite permutation. Let us denote $\text{id}_{\mathbb{Z}}^I$




FIGURE 3.7: Visualisation of $\sigma_{a_l}^k \circ \sigma_{a_r}^k$ and $\sigma_{a_r}^k \circ \sigma_{a_l}^k$ for $k = 3$. We see that the loss of information owing to the unmapped border elements accumulates: it still results in a partial identity, but the part of $\llbracket 1, N \rrbracket$ that is actually mapped to itself shrinks as k increases.

the partial identity of $\mathbb{Z}$ that is “restricted” to subset $I \subset \mathbb{Z}$, that is

$$\text{id}_{\mathbb{Z}}^I : n \in \mathbb{Z} \cup \{*\} \mapsto \begin{cases} n & \text{if } n \in I \\ * & \text{else.} \end{cases}$$

With this convention, the partial identity of Fig. 3.6 corresponds to $\text{id}_{\mathbb{Z}}^{\{1, \dots, N-1\}}$ while that of Fig. 3.7 are $\text{id}_{\mathbb{Z}}^{\{1, \dots, N-3\}}$ and $\text{id}_{\mathbb{Z}}^{\{4, \dots, N\}}$.

Now it can be checked that $\sigma_{a_l} = \hat{\sigma}_{a_l} \text{id}_{\mathbb{Z}}^{\{2, \dots, N\}}$ and $\sigma_{a_r} = \hat{\sigma}_{a_r} \text{id}_{\mathbb{Z}}^{\{1, \dots, N-1\}}$. Let us denote this *partialization* relation by

$$\sigma \preceq \sigma' \Leftrightarrow \exists \text{id}_{\mathbb{Z}}^I \text{ such that } \sigma = \sigma' \text{id}_{\mathbb{Z}}^I. \quad (3.16)$$

where σ and σ' are any two partial bijections of $\mathbb{Z}$ and $\text{id}_{\mathbb{Z}}^I$ is any partial identity of $\mathbb{Z}$. Informally, this relation accounts for the case when σ and σ'

represent the same (virtual) sensel permutation with some loss of information from σ' to σ . This is made especially relevant because we have

$$\forall a, a' \in \mathcal{A}, \sigma_a \sigma_{a'} \preceq \sigma_{aa'}$$

in the non conservative case (Exel, 1998). Therefore, although distinct sequences of actions $a_p \dots a_1$ and $a'_q \dots a'_1$ which products are the same final action a do not yield the same partial bijection we must have instead

$$\sigma_{a_p} \circ \dots \circ \sigma_{a_1} \preceq \sigma_a \text{ and } \sigma_{a'_q} \circ \dots \circ \sigma_{a'_1} \preceq \sigma_a.$$

However, this relation cannot be directly used to compare sequences of actions since it relies on us knowing the hypothetical σ_a common upper bound. We can instead observe that whenever two partial bijections σ, σ' partialize a given third σ'' they should in particular be *non contradictory*; that is whenever σ_a maps i to a “meaningful” (by which we mean non $*$) value j , σ' cannot map the same i to *another* meaningful value $j \notin \{j, *\}$. This *compatibility* relation can be checked as is or reformulated by

$$\sigma \sim \sigma' \Leftrightarrow \exists \text{id}_{\mathcal{C}'}^{\mathcal{C}'} \text{ such that } \sigma \sigma'^{-1} = \text{id}_{\mathcal{C}'}^{\mathcal{C}'} \quad (3.17)$$

for σ, σ' any two partial bijections of $\mathcal{C}$ and $\mathcal{C}'$ any subset of $\mathcal{C}$. It defines an equivalence relation on generalized successor maps which verifies

$$\forall a, a' \in \mathcal{A}, \sigma_{a'} \circ \sigma_a \sim \sigma_{a'a'}$$

thereby generalizing the corresponding equality in the conservative case.

It should be noted that although the development required referring to the virtual infinite agent, the final equivalence relation is entirely formulated in terms of the real limited case. In fact, the relation between partial shifts of receptive fields and conservative actions is general: it can be proved that any such semigroup of partial bijections arising from a group action can be made to be the “partialization” of a suitable group of true bijections (Kellendonk and Mark V. Lawson, 2004). However, the general procedure carrying out the extension is far more abstract so that it is not always trivial to envision what the extended agent actually looks like.

Finally we can reformulate the final result of 3.2.2.2 for non conservative settings; indeed, the argument we presented to demonstrate how sensel permutations result in state prediction does not depend on total permutation.

Therefore, using the definitions introduced in this section we can conclude

Theorem 2. *Let $a \in \mathcal{A}$ be a motor action*

- *It defines a generalized sensory prediction*

$$\Pi_{\mathcal{S},a} : \prod_{c \in \mathcal{C}} \overline{\mathcal{S}_c} \rightarrow \prod_{c \in \mathcal{C}} \overline{\mathcal{S}_c}$$

which verifies the prediction property of Eq. (3.8) for pairs of sensels verifying $c \xrightarrow{a} c'$.

- *The collection $I_{\mathcal{S}} = \{\Pi_{\mathcal{S},a}, a \in \mathcal{A}\}$ of generalized sensory predictions is an inverse semigroup partially acting on $\mathcal{S}$.*
- *The mapping*

$$\begin{aligned} \Pi_{\mathcal{S}} : \mathcal{A} &\rightarrow I_{\mathcal{S}} \\ a &\mapsto \Pi_{\mathcal{S},a} \end{aligned} \tag{3.18}$$

is a semigroup unital premorphism⁴, i.e. it verifies

- $\Pi_{\mathcal{S}}(e) = \text{id}_{\mathcal{S}}$
- $\Pi_{\mathcal{S}}(a^{-1}) = \Pi_{\mathcal{S}}(a)^{-1}$ (the second inverse being in the sense of semigroups)
- $\Pi_{\mathcal{S}}(a') \circ \Pi_{\mathcal{S}}(a) \preceq \Pi_{\mathcal{S}}(a'a)$

In particular, it transforms identities of compositions into equivalences as per

$$(a_p \dots a_1 = a'_q \dots a'_1) \Rightarrow (\Pi_{\mathcal{S}}(a_p) \circ \dots \circ \Pi_{\mathcal{S}}(a_1) \sim \Pi_{\mathcal{S}}(a'_q) \circ \dots \circ \Pi_{\mathcal{S}}(a'_1))$$

3.2.3 Conclusion

The point made above is twofold. First, we have shown on a formal basis how the structure of motor actions is intrinsically reflected on that of the associated predictions of their consequences. This is a significant result for our naive agent. Indeed it entails that whenever it acquires the capability to discover the outcomes associated with its actions, it can *cognize* about (the compositional properties of) the predictions it learned *as if they were* its motor

⁴This notion is far less standard than that of group (homo)morphisms it replaces; an introduction to the subject can be found e.g. in (Petrich, 1984). It is not necessary to delve in its details, instead one can and should mostly think of them as stand-ins for group morphisms in limited information conditions.

actions itself. In that way, we have shown how learned action effects can be used as a proxy of their generating actions.

Then, our development about shifts of receptive fields had us isolate one such class of predictable actions. Moreover we have found that this prediction occurs at the level of sensory readings, by which it is *a priori* internally accessible to the agent. The picture the study of conservative actions draws is particularly clear, resulting in an algebraic isomorphism between actions and their sensory effects. On the other hand, the generalization to partial permutations for nonconservative actions allows for more realistic cases; however the structure involved –that of inverse semigroups– is somewhat degenerate, and certainly more involved. Nevertheless, it allows for a precise formulation of the way in which the action-effects correspondence applies in the general case.

On a conceptual level, we need to highlight another aspect of the conservativity property. As the illustrations suggest, it is first and foremost a property associated with spatial displacements of one’s embodiment. Although the relation of shifts is formulated at the level of receptive fields, its origin is naturally found in a shift of the physical location of sensels, that is of the body of the agent. Consequently, we hypothesize that a study of one’s conservative actions largely overlaps with that of one’s displacements. But importantly this is also the intuition behind the role of Poincaré’s compensability, which also characterized locomotion. This intuition, which is mirrored by our formal account of ambient space (see 2.3.1.1 and 2.3.2.1), is that the ambient geometry is precisely the data of these displacements rather than that of the points comprising space. Therefore, learning properties of one’s displacements is a way to discover the corresponding properties of external space. We posit then that conservative actions give another access to the naive agent as to these external properties. Furthermore, we posit that it is preferable in two aspects. First, it does not ask for synchronization between motor and environment transformations as compensability intrinsically does; second, the corresponding point can be made in non-ideal cases, namely that of partial shifts. These make the conservativity property a somewhat practical outlook for the study of displacements for the naive agent.




FIGURE 3.8: Experimental setup used in simulation. A holonomic agent is placed in a 2D environment which ceiling displays grayscale images. The agent can move in this environment by performing seven distinct actions a_k . A 10-by-10 camera pointed towards the ceiling is placed on top of the agent and generates a sensory array $\mathbf{s} = (s_{c_i})_i$.

3.3 Experiment: a simulated account of inferring action structure by naively assessing conservativity

The previous discussion highlights how spatial localization of the sensorimotor interaction entails the necessary structure of shifts of receptive fields in the internal experience. However it remains very general by virtue of extending its scope to all setups considered in 2.3; we might say it establishes *essential* properties of the sensorimotor experiences, as opposed to those *contingent* on any given agent's. This is instead the focus of this section, in which we experimentally illustrate the main results of 3.2.

3.3.1 Description of the setup

As depicted in Fig. 3.8, the system we consider in this experiment is a 2-dimensional analogue of that found in previous illustrations. The sensorimotor interaction occurs on a plane the agent can explore using its omnidirectional wheels. These allow it to perform translations along both spatial coordinates x and y as well as rotate about its center. Correspondingly its sensory apparatus is now a rectangular camera; its pixels are sensitive to the luminance of the stimulus provided by a grayscale image above its body.

Following the definitions of 2.3.1, we can take as formal parameterization

- the ambient space $\mathcal{X}$ is the usual plane $\mathbb{R}^2$;
- an environment state $\varepsilon \in \mathcal{E} \mathcal{X}$ is a function which takes point $(x, y) \in \mathbb{R}^2$ of the plane to the luminance value $\varepsilon(x, y) \in [0; 255]$ of the image at said point;
- the absolute configuration space $\mathcal{B}$ is $\mathbb{R}^2 \times S_1^5 \cong \mathbb{R}^2 \times]-\pi; \pi]$ (mod. 2π): it accounts for both position $(x, y) \in \mathbb{R}^2$ of the agent and its orientation $\theta \in S_1$. Choosing the $]-\pi; \pi]$ representation over S_1 is equivalent to choosing a preferred “reference” orientation corresponding to the null angle $\theta = 0$. While the former is more convenient in a digital simulation, there is no natural choice here: instead, we can arbitrarily choose the reference orientation.

As for the internal picture of the agent, its sensory is given by an array $\mathbf{s} \in [0; 255]^{W_c \times H_c} = \mathcal{S}_{\mathcal{C}}$ where W_c (resp. H_c) is the number of pixels in one row (resp. column) of the camera; in the following simulations, $W_c = H_c = 10$ so that the total array $\mathbf{s}$ has length 100. Each of its components s_{c_i} corresponding to pixel $c_i \in \{1, \dots, 100\} = \mathcal{C}$ is given by the luminance of the image at the point c_i is sampling. As with the previous case the ordering of pixels (corresponding to a choice of labelling $\mathcal{C} \cong \{1, \dots, 100\}$) is entirely arbitrary.

Additionally, we define a set A of actions a_k , $k = 1, \dots, 7$ (depicted in Fig. 3.8):

- the identity action a_1 , by which the agent maintains any absolute configuration $\mathbf{b} = (x, y, \theta)$ it previously was at;
- two rotations a_6, a_7 corresponding to both clockwise and counter-clockwise turns, each with amplitude 90° .
- four translations a_2, a_3, a_4, a_5 , with amplitude 1 pixel: one for each direction on the plane.

This definition looks “natural” to us since we know from our understanding of geometry that these seven actions efficiently span all displacements in the plane. However, it does not necessarily correspond with the motor actions the agent may *a priori* perform. Instead, it is capable of issuing a *multiset*⁶

⁵Here S_1 denotes the topological structure of the circle.

⁶That is, a collection in which elements can appear multiple times (as opposed to sets in the mathematical sense).

$\mathcal{A}_{\text{init}}$ of $n_{\mathcal{A}}$ actions drawn in the collection of combinations of actions of A ⁷. Therefore $\mathcal{A}_{\text{init}}$ may not be as convenient as A depending on how the draw. For example we should expect the agent to “miss” some natural actions a_k of A when the size $n_{\mathcal{A}}$ of its starting repertoire $\mathcal{A}_{\text{init}}$ is small, which becomes trivially true when $n_{\mathcal{A}} < 7$. On the contrary, some actions $a, a' \in \mathcal{A}_{\text{init}}$ effectively available at first are poised to be “duplicates” of one another when $n_{\mathcal{A}}$ is large. This situation occurs when a, a' are distinct occurrences (in the multiset $\mathcal{A}_{\text{init}}$) of a same underlying motor action $a_k \in A$ ⁸. Of course both the notions of actions being “missing” and “duplicates” are entirely beyond the comprehension of the naive agent, at least initially: as outlined in 2.3.2.3 it may only perform elements of $\mathcal{A}_{\text{init}}$, and its capability to compare for equality is therefore that of occurrences in $\mathcal{A}_{\text{init}}$.

Importantly the translations $a_2, \dots, a_5$ are defined in the reference frame of the agent, which is not equivalent to translations in $\mathcal{B}$. For example, consider the situation depicted in Fig. 3.8 and assume it corresponds to orientation $\theta = 0$: then we have

$$a_3(x_r, y_r, 0) = (x_r + 1, y_r, 0).$$

Consequently, it locally coincides with translation of the first component $(x, y, \theta) \mapsto (x + 1, y, \theta)$. However the agent could instead perform rotation a_6 before, ending at configuration $\mathbf{b}' = (x_r, y_r, \frac{\pi}{2})$. But since a_3 occurs in the reference frame of the agent –as if the arrows of Fig. 3.8 were rigidly linked to its body– we must have

$$a_3\left(x_r, y_r, \frac{\pi}{2}\right) = \left(x_r, y_r + 1, \frac{\pi}{2}\right).$$

This shows how a_3 now coincides with translation of the *second* component $(x, y, \theta) \mapsto (x, y + 1, \theta)$. Therefore, we see that the direction taken by the displacement rotates with θ . More generally, let us denote R_θ the 2-by-2 matrix

⁷(which 2.3.2.1 characterized as the submonoid $\langle A \rangle$)

⁸Using classical mathematics often result in clumsy accounts of this distinction (Knuth, 1981); instead, it is better illustrated by considering the difference between *equality* and *identity* in computer science. Consider a program where two variables a and b are declared and separately set to the value 2. Then it is generally true that $a == b$ (that is a is *equal* to b) because they represent the same value, whereas it is instead usually false that a and b are *identical* since they are distinct as occurrences (e.g. as objects in physical memory).

of the planar rotation of angle $\theta \in]-\pi; \pi]$, that is

$$R_\theta = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}.$$

Then, the translation a_3 is given by

$$a_3(x, y, \theta) = (x', y', \theta)$$

where

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix} + R_\theta \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Naturally we have that orientation θ is invariant under a_3 , corresponding to it being a translation. Moreover, the $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ vector corresponds to the “direction of displacement”: it is accordingly rotated by R_θ to account for current orientation before it modifies position $\begin{pmatrix} x \\ y \end{pmatrix}$. Of course each of the translations a_k , $k = 2, \dots, 5$ are of this sort, where each corresponds to a distinct displacement vector $\begin{pmatrix} x_{a_k} \\ y_{a_k} \end{pmatrix}$. It is precisely the presence of R_θ that differentiates translations in the self-referential frame from that of the ambient space; in our experience, this corresponds to the distinction between e.g. “move 1 meter to the right” and “move 1 meter eastwards”.

Finally, we can observe that none of the translation actions are strictly conservative, generalizing the illustration of Fig. 3.2. Indeed, let us get back to the example of a_3 which we temporarily take to denote “forward” translation⁹. Each of the pixels in the inner area of the camera has a successor (and a predecessor) under a_3 , namely the one that comes in front of (resp. behind of) it. However, by the same account there is no such successor for pixels in the front row of the sensor: any time the agent performs a_3 , these pixels instead sample a previously unseen part of the environment. Therefore the W_c (or N_c) boundary pixels of the front row make a_3 fail the conservativity condition (see 3.2.1). The same unpredictability phenomenon arises for all translations, each one corresponding to its particular “front” border. Therefore they all fall instead under the partial case of 3.2.2.3. However it should be noted that they are somewhat “largely” conservative since the relevant

⁹This choice only serves to provide convenient spatial vocabulary (front, back, etc.); although it is entirely arbitrary, it also comes without loss of generality.

condition only fails to hold for 10 out of all 100 pixels, which are an order of magnitude apart.

3.3.2 Simulation 1: searching the sensorimotor flow for conservativity

The first experiment is one of *motor babbling* where the naive agent actively explores its environment, collecting samples from its camera along the way. It successively runs its starting actions $a_k \in \mathcal{A}_{\text{init}}$ and monitors the sequence of its sensory outputs to try and determine their associated permutations of sensels, should they exist. The environment state ε can entirely change during the interaction so long as it does not happen simultaneously with the generation of the a_k s: that is, it can change *between actions*. In the simulation, this is achieved by changing the image presented to the agent between each action.

3.3.2.1 Representing conservative actions

We can leverage the numerical properties of this representation to naturally make sensory permutations into matrices. Indeed let $N_c = W_c \times H_c$ be the total number of sensels and $\phi : \llbracket 1, N_c \rrbracket \rightarrow \llbracket 1, N_c \rrbracket$ be any permutation, and consider the N_c -by- N_c matrix M_ϕ given by

$$M_{\phi_{i,j}} = \begin{cases} 1 & \text{iff } j = \phi(i) \\ 0 & \text{else.} \end{cases} \quad (3.19)$$

Then for any value of sensory output $\mathbf{s} \in \mathcal{S}$, the result $\mathbf{s}_\phi = (\mathbf{s}_{\phi_i})_{i \in \llbracket 1, N_c \rrbracket}$ of permuting its components following ϕ is given by

$$\mathbf{s}_\phi = M_\phi \mathbf{s}.$$

From the definition of Eq. (3.19) we can see that only N_c out of the N_c^2 matrix entries are nonzero for any given permutation of N_c sensels, that is 100 out of 10,000 in our case. Moreover this number becomes even smaller if the permutation represented is only partial, as happens for nonconservative actions. Therefore it is natural to implement these representations as *sparse* matrices for computational efficiency.

This representation is made all the more convenient by the fact that successive composition of permutation is faithfully accounted for by matrix product. Indeed consider any value $\mathbf{s}$ of the sensory array, and apply permutation σ to its component; from the previous statement we know we must get array $\mathbf{s}' = \mathbf{s}_\sigma = M_\sigma \mathbf{s}$ as a result. Should we then permute the components of this $\mathbf{s}'$ along some σ' , we would accordingly obtain $\mathbf{s}'' = \mathbf{s}'_{\sigma'} = M_{\sigma'} \mathbf{s}'$. Now it only remains to observe that

$$\mathbf{s}'' = M_{\sigma'} \mathbf{s}' = M_{\sigma'} (M_\sigma \mathbf{s}) = (M_{\sigma'} M_\sigma) \mathbf{s}$$

so that $M_{\sigma' \circ \sigma} = M_{\sigma'} M_\sigma$ ¹⁰.

Importantly this also addresses the dependency of the representation on the ordering of sensels. Indeed consider two (possibly distinct) orderings of sensels, that is two index assignments $\nu, \nu' : \llbracket 1, N_c \rrbracket \rightarrow \mathcal{C}$; each is a bijection that maps index number i to the actual sensel it labels under this particular ordering. Consider then the sensory arrays $\mathbf{s}$ and $\mathbf{s}'$ resulting from the same interaction the agent has with its environment, only differing in that they respectively assume ordering ν and ν' . Let us denote $c \in \mathcal{C}$ the particular sensel that appears first in ν , that is $c = \nu(1)$. It must necessarily appear as $\nu'(k)$ for some $1 \leq k \leq N_c$ since ν' also indexes $\mathcal{C}$, but then we must have $\mathbf{s}'_k = \mathbf{s}_1$ since arrays $\mathbf{s}$ and $\mathbf{s}'$ only differ in ordering. This shows that should we start from components of $\mathbf{s}$, following ν^{-1} to get the underlying sensels then applying ν' would precisely yield the ordered array $\mathbf{s}'$. Therefore the arrays $\mathbf{s}$ and $\mathbf{s}'$ corresponding to respective orderings ν and ν' verify

$$\mathbf{s}' = \mathbf{s}_{\nu' \circ \nu^{-1}}.$$

In fact, the picture is equally nice concerning the matrix representation of conservative actions. Indeed consider any conservative action a with corresponding permutation of sensels $\sigma_a : \mathcal{C} \rightarrow \mathcal{C}$, and consider as before two orderings ν, ν' corresponding to ordered arrays $\mathbf{s}, \mathbf{s}'$. Let us denote $M_{\sigma_a}^\nu$ the matrix that corresponds to σ_a under ordering ν ¹¹. Then keeping ordering ν

¹⁰What this denotes is that the correspondence between “abstract” permutations and their matrices is in fact a group homomorphism (which suitably restricts to an isomorphism). However we only leverage this assertion for its convenient implementation and make no claim of any deep meaning here, therefore we need only focus on its practical aspect.

¹¹That is it represents the permutation of *whole numbers* that maps $i \in \mathbb{N}$ to $j \in \mathbb{N}$ if and only if σ_a maps the sensel of index i (for ν) to that of index j (for ν).

fixed, performing action a takes the sensory array $\mathbf{s}$ to $\mathbf{s}_a = M_{\sigma_a}^\nu \mathbf{s}$. But accordingly for ordering ν' we have

$$\begin{aligned} \mathbf{s}'_a &= (\mathbf{s}_a)_{\nu' \circ \nu^{-1}} = (M_{\sigma_a}^\nu \mathbf{s})_{\nu' \circ \nu^{-1}} \\ &= [M_{\nu' \circ \nu^{-1}} M_{\sigma_a}^\nu] \mathbf{s} \\ &= [M_{\nu' \circ \nu^{-1}} M_{\sigma_a}^\nu] (M_{\nu' \circ \nu^{-1}}^{-1} \mathbf{s}') = (M_{\nu' \circ \nu^{-1}} M_{\sigma_a}^\nu M_{\nu' \circ \nu^{-1}}^{-1}) \mathbf{s}', \end{aligned}$$

where $M_{\nu' \circ \nu^{-1}}$ is the permutation matrix corresponding to $\nu' \circ \nu^{-1} : \llbracket 1, N_c \rrbracket \rightarrow \llbracket 1, N_c \rrbracket$. This proves that the representations as permutation matrices found for various orderings of sensels are closely related along an equation of the form

$$M_{\sigma_a}^{\nu'} = P^{-1} M_{\sigma_a}^\nu P \quad (3.20)$$

where P is a matrix that only depends on ν and ν' . Consequently while the precise value of the matrix (that is which sensel indices it swaps) depends on the choice of ordering, the resulting algebraic properties do not¹².

3.3.2.2 Progressive discovery of the conservative permutations

The exploration and ongoing learning occurs as a sequential process, following timestep indices t_n , $n \in \mathbb{N}$. At the beginning of the sensorimotor experience a N_c -by- N_c matrix M_a is associated to each of its starting actions $a \in \mathcal{A}_{\text{init}}$, with all coefficients initialized to 1. Then, at any timestep t_n the agent (selects and) performs an action $a_k = a[t_n]$ from its repertoire $\mathcal{A}_{\text{init}}$: this accordingly changes the absolute configuration of the interaction from $\mathbf{b}[t_n] = (x[t_n], y[t_n], \theta[t_n])$ to $\mathbf{b}[t_{n+1}] = a_k \mathbf{b}[t_n]$. In this experiment, actions are selected by a sequence of independent uniform draws. Then the sensory arrays $\mathbf{s}[t_n]$ and $\mathbf{s}[t_{n+1}]$ the agent experiences before and after a_k are used to update its matrix M_{a_k} as per:

$$(M_{a_k}[t_{n+1}])_{i,j} = \begin{cases} 1 \text{ iff } s_i[t_{n+1}] = s_j[t_n]^{13} \text{ and } (M_{a_k}[t_n])_{i,j} = 1 \\ 0 \text{ else.} \end{cases} \quad (3.21)$$

¹²This is, down to the form of the relation, extremely similar to the familiar notion of change of basis in vector spaces. Transposed to nonlinear contexts, it still corresponds to nice *conjugacy* operations: it is a classical result that the mapping that takes the $M_{\sigma_a}^\nu$ to the corresponding $M_{\sigma_a}^{\nu'}$ performs an (*inner*) group automorphism (Dummit and Foote, 2009), in particular preserving the group structure.

We can observe that the only possible way coefficients may change is going from their initial value of 1 to 0: indeed whenever a coefficient of M_{a_k} is already set to 0 at timestep t_n , the first condition automatically fails so that it must stay so. Therefore the rough dynamics imposed by this update is that after all coefficients are initialized to 1, a certain number of them are eventually switched down to 0 during exploration until matrices M_{a_k} converge to a final (possibly null) form. Furthermore, the first condition on the first line necessarily occurs during shifts of receptive fields (see Eq. (3.2)): the idea, of course, is that coefficients $(M_{a_k})_{i,j}$ that indefinitely stay equal to 1 correspond to successor sensel pairs $c_i \xrightarrow{a_k} c_j$ for a_k . Accordingly we can externally think of coefficient $(M_{a_k}[t_{n+1}])_{i,j}$ as depicting the ongoing internal assessment of whether the past experience (i.e. using sensory samples $\mathbf{s}[t_0], \dots, \mathbf{s}[t_{n+1}]$) has rejected the hypothesis that $c_i \xrightarrow{a_k} c_j$ (corresponding to a null value) or not *as of yet*. These assertions are formally amenable *via*

Lemma 1. *Let a be any starting action of $\mathcal{A}_{\text{init}}$*

For any coefficient $m_{a_{i,j}}$ of M_a , the associated sequence $(m_{a_{i,j}}[t_n])_n$ of values taken in $(M_a[t_n])_n$ during exploration is nonincreasing with values in $\{0, 1\}$.

Proof. Let us consider an arbitrary timestep t_n in the exploration. If a is not drawn at this timestep, then

$$m_{a_{i,j}}[t_{n+1}] = m_{a_{i,j}}[t_n] \leq m_{a_{i,j}}[t_n].$$

If it is instead chosen, assuming $m_{a_{i,j}}[t_n] \in \{0, 1\}$ then as per the update rule of M_a , either

- $m_{a_{i,j}}[t_n] = 0$ and then $m_{a_{i,j}}[t_{n+1}] = 0$ too,
- or $m_{a_{i,j}}[t_n] = 1$ and $m_{a_{i,j}}[t_{n+1}] \in \{0, 1\}$

therefore $m_{a_{i,j}}[t_{n+1}] \leq m_{a_{i,j}}[t_n]$ in both cases. The lemma then follows by induction on n . $\square$

Lemma 2. *Let a be any starting action of $\mathcal{A}_{\text{init}}$.*

For any nonzero coefficient $m_{a_{i,j}} = 1$ in the matrix of its (generalized) successor map M_{σ_a} , the associated sequence $(m_{a_{i,j}}[t_n])_n$ is constant with value $m_{a_{i,j}}[t_n] = 1$.

¹³Note that this condition is precisely the result of computing the sensel equality operator $\delta_{\mathbf{s}_i, \mathbf{s}_j}$ we assumed to be *a priori* available to the agent in 2.3.2.3; therefore, this update rule is itself internally computable.

Proof. At any timestep t_n of the exploration, if a is not chosen then $m_{a_{i,j}}[t_{n+1}] = m_{a_{i,j}}[t_n]$.

If it is instead drawn, then by Equation (3.2) we know that

$$s_i[t_{n+1}] = s_j[t_n]$$

because $m_{a_{i,j}} = 1$ implies that $j = \sigma_a(i)$ as per the definition of M_{σ_a} . Then by the update rule of $M_a[t_n]$, $m_{a_{i,j}}[t_{n+1}] = 1$.

The lemma then follows by induction on n . $\square$

These guarantee that the update rule of Eq. (3.21) admits no false negatives, that is it accurately detects all successor sensel pairs as such. To have the complete result, it still needs to be established that it conversely has no false positives: each of the initial $m_{a_{i,j}}$ that correspond to sensel pairs c_i, c_j that are *not* related by a should eventually be set to 0. Here there is in fact some dependency on the environment states the agent has interacted with. Consider indeed a pathological case where all the images it samples are entirely uniform in color, e.g. they are pure black everywhere (corresponding to $\varepsilon(x, y) = 0$ for all x, y). Then clearly all sensels will always report the same output at every timestep of the interaction, so that the condition $s_j[t_{n+1}] = s_i[t_n]$ is “spuriously” verified irrespective of the choices of i, j and n . In such a case the update rule of Eq. (3.21) can never branch into the null case: all coefficients will indefinitely keep their starting value of 1, including a large majority of false positives. In the general case, however, it suffices that these coincidences fail to hold if only once for the hypothesis to be rejected as per the update rule.

Let us develop further the intuition from the uniformly black environment that false positives occur from coincidental equality of environmental state at various points of space. All displacements the agent can perform in this experiment can be spatially characterized in the form

$$\begin{aligned} \forall a \in \mathcal{A}_{\text{init}}, \exists \vec{u}_a = \begin{pmatrix} x_a \\ y_a \end{pmatrix} \in \mathbb{R}^2, \exists \theta_a \in \mathbb{R} \\ \text{such that } \forall \mathbf{b} = (x, y, \theta) \in \mathcal{B}, a\mathbf{b} = (x', y', \theta') \\ \text{where } \begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix} + R_\theta \begin{pmatrix} x_a \\ y_a \end{pmatrix} \text{ and } \theta' = \theta + \theta_a. \end{aligned}$$

We showed earlier how translations can be written thus by setting $\theta_a = 0$ and $\begin{pmatrix} x_a \\ y_a \end{pmatrix}$ be the vector of displacement. The case of rotations (a_6 and a_7) is even simpler since it corresponds to the cases where $\begin{pmatrix} x_a \\ y_a \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and $\theta_a = \pm \frac{\pi}{2}$ is the angle of the rotation. Furthermore the receptive fields of the agent move rigidly with its body when it performs locomotion, by which we have

$$\forall c \in \mathcal{C}, \forall \mathbf{b} = (x, y, \theta) \in \mathcal{B}, F_c(x, y, \theta) = \begin{pmatrix} x \\ y \end{pmatrix} + R_\theta F_c(0, 0, 0). \quad (3.22)$$

This relation between absolute configurations $\mathbf{b}$ of the agent and localization of its receptive fields allows us to make the spatial characterization of actions of Eq. (3.3.2.2) into one of displacements of receptive fields. Indeed,

Lemma 3. *Let a be any starting action, c, c' be any two sensels. There exists a unique vector $\vec{d}_{a,c,c'} \in \mathbb{R}^2$ such that*

$$\forall \mathbf{b} = (x, y, \theta) \in \mathcal{B}, \overrightarrow{F_{c'}(\mathbf{b})F_c(a\mathbf{b})} = R_\theta \vec{d}_{a,c,c'}.$$

Proof. Let $\mathbf{b} = (x, y, \theta) \in \mathcal{B}$. From the previous two assertions we have

1. $F_{c'}(\mathbf{b}) = \begin{pmatrix} x \\ y \end{pmatrix} + R_\theta F_{c'}(0, 0, 0),$
2. $F_c(a\mathbf{b}) = F_c \left(\begin{pmatrix} x \\ y \end{pmatrix} + R_\theta \begin{pmatrix} x_a \\ y_a \end{pmatrix}^{14}, \theta + \theta_a \right)$
 $= \begin{pmatrix} x \\ y \end{pmatrix} + R_\theta \begin{pmatrix} x_a \\ y_a \end{pmatrix} + R_{\theta+\theta_a} F_c(0, 0, 0)$

so that

$$\overrightarrow{F_{c'}(\mathbf{b})F_c(a\mathbf{b})} = R_\theta \left(\begin{pmatrix} x_a \\ y_a \end{pmatrix} + R_{\theta_a} F_c(0)^{15} - F_{c'}(0) \right).$$

which proves taking $\vec{d}_{a,c,c'} = \left(\begin{pmatrix} x_a \\ y_a \end{pmatrix} + R_{\theta_a} F_c(0) - F_{c'}(0) \right)$ satisfies the property. □

¹⁴This is a slight abuse of notation where we use $F_c(\begin{pmatrix} x \\ y \end{pmatrix}, \theta)$ as an alias for $F_c(x, y, \theta)$.

¹⁵Here $F_c(0)$ denotes $F_c(0, 0, 0)$, the 0 of the first expression being the 3-dimensional null element.

The $\vec{d}_{a,c,c'}$ vector, which (for fixed subscripts $a \in \mathcal{A}_{\text{init}}, c, c' \in \mathcal{C}$) is invariant during the sensorimotor experience of the agent, captures some geometry of conservation. Indeed, observe that from the definition of relation $c \xrightarrow{a} c'$ and Lemma 3 we get

$$\forall a \in \mathcal{A}_{\text{init}}, \forall c, c' \in \mathcal{C}, c \xrightarrow{a} c' \Leftrightarrow \vec{d}_{a,c,c'} = \vec{0}. \quad (3.23)$$

Let us briefly circle back to the pathological uniform example: the update rule of Eq. (3.21) failed to produce good results then because for any two sensels c and c' , c *always* sampled after a the same luminance value c had been sampling before. But then by Lemma 3 this is equivalent to the $\vec{d}_{a,c,c'}$ vectors (and their images by rotations R_θ) having *all* their endpoints $\begin{pmatrix} x \\ y \end{pmatrix}$ and $\begin{pmatrix} x \\ y \end{pmatrix} + R_\theta \vec{d}_{a,c,c'}$ share their luminance values. It is in fact a necessary condition as per

Theorem 3. *Let $a \in \mathcal{A}_{\text{init}}, c_i, c_j \in \mathcal{C}$. Then for given environment configuration $\varepsilon \in \mathcal{E} \mathcal{X}$ the two statements*

1. *There exists an absolute configuration $\mathbf{b} \in \mathcal{B}$ such that if $\mathbf{b}[t_n] = \mathbf{b}$ and $a[t_n] = a$ for some $n \in \mathbb{N}$, then $m_{a_{i,j}}[t_k] = 0 \forall k \geq n + 1$*
2. *ε is not doubly periodic with periods $\vec{d}_{a,c_i,c_j}$ and $R_{\frac{\pi}{2}} \vec{d}_{a,c_i,c_j}$*

are equivalent.

Proof.

- Proceed contrapositively, assuming that ε is both $\vec{d}_{a,c_i,c_j}$ - and $R_{\frac{\pi}{2}} \vec{d}_{a,c_i,c_j}$ -periodic. Let $n \in \mathbb{N}$ be such that $a[t_n] = a$, let $\mathbf{b} = \mathbf{b}[t_n]$. We have

$$\mathbf{s}_j[t_n] = \varepsilon(F_{c_j}(\mathbf{b}[t_n]))$$

and

$$\mathbf{s}_i[t_{n+1}] = \varepsilon(F_{c_i}(\mathbf{b}[t_{n+1}])) = \varepsilon(F_{c_j}(\mathbf{b}[t_n]) + R_\theta \vec{d}_{a,c_i,c_j}).$$

But since $\theta \in \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$, $R_\theta = \pm I_2$ or $R_\theta = \pm R_{\frac{\pi}{2}}$.

Therefore by periodicity of ε we have $\mathbf{s}_i[t_{n+1}] = \mathbf{s}_j[t_n]$, from which by induction on n we get $\forall n \in \mathbb{N}, m_{a_{i,j}}[t_n] = 1: 1 \Rightarrow 2$.

- Without loss of generality, let us assume that ε is not $\vec{d}_{a,c_i,c_j}$ -periodic (if it is instead only not $R_{\frac{\pi}{2}} \vec{d}_{a,c_i,c_j}$ -periodic, the same argument follows up

to a rotation).

Let $X_0 \in \mathbb{R}^2$ such that $\varepsilon(X_0) \neq \varepsilon(X_0 + \vec{d}_{a,c_i,c_j})$, and let $\mathbf{b}_0 = (x, y, 0) \in \mathcal{B}$ be such that $F_{c_j}(\mathbf{b}_0) = X_0$.

By definition $F_{c_i}(a\mathbf{b}_0) = X_0 + \vec{d}_{a,c_i,c_j}$ so that if $\mathbf{b}[t_n] = \mathbf{b}_0$ for some $t_n \in \mathbb{N}$, $s_i[t_{n+1}] \neq s_j[t_n]$. From the update rule of M_a we then get $m_{a,i,j}[t_{n+1}] = 0$, which by Lemma 1 concludes the proof.

□

Corollary 4. *If environment state $\varepsilon : \mathbb{R}^2 \rightarrow [0, 255]$ is aperiodic, then there exists¹⁶ a sequence of draws of actions $(a[t_n])_{n \in \mathbb{N}}$ such that following the update rule (3.21) in environment ε will verify*

$$\forall a_k \in \mathcal{A}_{\text{init}}, \lim_n M_{a_k}[t_n] = M_{\sigma_{a_k}}.$$

We now have the complete picture regarding the behavior of the empirical matrices $(M_{a_k}[t_n])_{n \in \mathbb{N}}$: not only do we know that they report $c_i \xrightarrow{a_k} c_j$ pairs via coefficients with constant value 1 but also that all other cases are eventually detected and accordingly switched to 0. Consequently we also know that the update rule (3.21) allows the naive agent to internally distinguish whether actions are (partially) conservative, with Theorem 3 guaranteeing convergence to the null matrix in the purely nonconservative case.

¹⁶In fact, seeing the draw of actions $(a[t_n])_{n \in \mathbb{N}}$ as one of displacements in a random walk on $\mathbb{Z} \times \{-\frac{\pi}{2}, 0, \frac{\pi}{2}, \pi\}$, it is known e.g. from (McCrea and Whipple, 1940; Spitzer, 2013) that in 2 dimensions such a sequence of draws will even occur with probability 1.




FIGURE 3.9: Representation of the seven binary $10^2 \times 10^2$ permutation matrices M_{a_k} for various choices ν_i of ordering of the physical pixels.

3.3.2.3 Results

Fig. 3.9 presents the final steady-state of the sequences of empirical matrices $(M_{a_k}[t_n])_{n \in \mathbb{N}}$ obtained following update rule (3.21). Each row corresponds to the matrices of the a_k (with k varying from 1 to 7) for a fixed ordering v_i of the pixels of the camera, depicted on the leftmost column. We choose the ordering of the first row as reference: it is in fact (one of) the “natural” ordering to our external eye, given by enumerating the pixels of the camera in reading order. As per the previous discussion, all other orderings can then be obtained from v_1 by permutation of the indices: the matrix corresponding to the ordering v_i used on line i ¹⁷ is depicted for each line in the leftmost column. All permutations matrices, being binary, are displayed in black and white: black pixels correspond to null coefficients, whereas white ones accordingly denote 1s.

We can observe that the resulting matrices $M_{a_k}^{v_i}$ differ greatly with respect to the pixel orderings v_i . In particular, while some orderings such as the reference v_1 yield “orderly” matrices, the behavior under other orderings looks far more erratic. As external observers, we can however make sense of these matrices by referring to the spatial layout corresponding to their particular ordering. This is done on the bottommost row of the figure, wherein the column attributed to a_k an arrow is plotted from the location of sensel of index i to that of index j whenever $(M_{a_k})_{i,j} = 1$. Note that this representation requires our *external* understanding of the spatial interpretation of orderings v_i ¹⁸, and as such is not available to the agent. However it allows us to assess that the agent has learned the *same* displacement of sensels independently of the ordering of its pixels.

Crucially, this last fact separates two “levels” in the discovery the agent makes. On the one hand the *particular value* of the empirical matrix $M_{a_k}^{v_i}$ is entirely *contingent* on the specific details of the implementation of the system, up to the ordering of the pixels; on the other hand, as previously discussed in 3.2.2.2 and Eq. (3.20), the *relations* between these values and the corresponding semantics are instead *invariant* properties of the sensorimotor experience.

¹⁷which was denoted as $M_{v_1 \circ v_i^{-1}}$ –or P in Eq. (3.20) relating the matrices expressed for distinct orderings– in the previous discussion.

¹⁸(as did our arbitrary choice of reference v_1 , which only seems convenient by means of external observations)

3.3.3 Simulation 2: investigating the dynamics in the discovery of sensel permutations

In the previous experiment, we only examined the (limit) steady-state of the empirical permutation matrices $(M_{a_k}[t_n])_{n \in \mathbb{N}}$ and the conditions under which it could be reached. However the sensorimotor experience does not begin at this limit state since it is rather a gradual process, *during* which the agent is precisely poised to have agency. Therefore we now devote attention to the dynamics with which the agent learns the sensel permutations corresponding to its actions.

3.3.3.1 Assessing progress in the convergence of the empirical matrices

The matrix representation chosen in the previous experiment make it easy to declare whether the empirical matrix of action $a_k \in \mathcal{A}_{\text{init}}$ at timestep t_n (i.e. the matrix $M_{a_k}[t_n]$) denotes a (partial) permutation. Indeed since by Lemma 1 its coefficients are always 0s and 1s, $M_{a_k}[t_n]$ is the matrix of a (partial) permutation *if and only if* each of its rows *and* its columns contain *at most* one 1. While this property can be internally ascertained by the naive agent to gauge whether it has entirely learned the sensel successor laws of its actions, it only distinguishes binary states of convergence of empirical matrices. Instead we will now refine this observation by considering an entropy-like criterion, evaluated on all possible states of the empirical matrices $M_{a_k}[t_n]$ and that hinges upon the interpretation of their coefficients as the plausibility of hypotheses $c_i \xrightarrow{a_k} c_j$ based on the past sensorimotor experience. This criterion is numerically defined as

$$\begin{aligned} C(M) &= 1 - \frac{1}{N_c \log_2(N_c)} \sum_{i=1}^{N_c} H_i, \\ \text{where } H_i &= - \sum_{j=1}^{N_c} \frac{M_{i,j}}{\mu_i} \log_2\left(\frac{M_{i,j}}{\mu_i}\right), \\ \text{and } \mu_i &= \frac{1}{\max(1, \sum_{j=1}^{N_c} M_{i,j})}. \end{aligned} \tag{3.24}$$

Here, H_i is the entropy of the post-action output of sensel c_i as a random variable over the pre-action outputs of all sensels c_j . It consequently quantifies the surprise remaining in the determination of which (if any) sensel is the actual successor of c_i . Finally, the total criterion C is therefore an average measure of certainty in the discovery of successor sensel pairs.

We can deduce from the usual properties of entropy (Wehrl, 1978) that if all coefficients $M_{i,j}$ are equal to 1, the corresponding value of the criterion is $C(M) = 0$; on the other hand if on every row i there is *at most* one nonzero coefficient $M_{i,j} = 1$ at column j , then $H_i = 0$ (for all $1 \leq i \leq N_c$), which further entails $C(M) = 1$. This shows that, following update rule (3.21), the sequences $C_k = (C_k[t_n])_{n \in \mathbb{N}} = (C(M_{a_k}[t_n]))_{n \in \mathbb{N}}$ all follow nondecreasing trajectories from 0 to 1 during the sensorimotor experience.

3.3.3.2 Sources of variance in the dynamics of discovery

There are a number of factors at play in determining the dynamics with which the agent learns its empirical permutation matrices M_{a_k} . Indeed, since at any timestep t_n of the sensorimotor interaction the update rule of Eq. (3.21) only modifies the matrix M_{a_k} of action $a_k = a[t_n]$ which is being run, it appears that the sequences $(M_{a_k}[t_n])_{n \in \mathbb{N}}$ depend on the draws of actions $(a[t_n])_{n \in \mathbb{N}}$. For example, let us consider one particular sensorimotor experience in which the agent only ever performs one particular action $a_{k_0} \in \mathcal{A}_{\text{init}}$, that is we have

$$\forall n \in \mathbb{N}, a[t_n] = a_{k_0}.$$

Clearly then, following the update rule we have that all matrices $M_{a_{k'}}$ with $k' \neq k_0$ will indefinitely stay at their initial value; consequently, their associated values for criterion C , $C_{k'}[t_n]$, will also stay indefinitely at 0, making no progress whatsoever. Therefore, we should consider this variance induced by the sequence of *action selections* in the examination of the dynamics.

The line of reasoning followed in the proofs of the previous simulation hint at another distinct source of variance in the speed at which the agent learns its sensel permutations, provided by the environment state the agent interacts with. Indeed, Theorem 3 shows that the eventual convergence of the empirical matrices M_{a_k} to their desired values M_{σ_a} hinges on the environment being “not periodic” enough; in fact, for the $\vec{d}_{a,c,c'}$ vectors of the translations considered here (i.e. a_k with $2 \leq k \leq 5$), its aperiodicity requirement only fails for the uniformly black (or gray, or white...) environment previously mentioned. This condition on environment states might seem artificial should one ask for it *everywhere* (for all $(x, y) \in \mathbb{R}^2$) and *everytime* (for all t_n). However it shows how the practical discovery of the M_{σ_a} depends on the characteristics of the environment state the agent samples, in particular how it depends on local ambiguity in state ε at distinct points of space. This in fact very much aligns with our intuition: consider a new example

case where the visual environment is entirely (and uniformly) black *except* for one small white region of space. Although by Theorem 3 we know that the agent can (and will) *eventually* discover its sensel permutations M_{σ_a} , it is clear that any interaction which entirely occurs in the dark region of space will not directly contribute to its learning. This shows that we should expect an influence on the dynamics of learning the M_{a_k} matrices from the *local richness of the environment* state the agent interacts with.

3.3.3.3 Experimental setup

The setup is largely similar to that of the previous experiment; in particular the agent, its ambient space, and its initial motor capabilities $\mathcal{A}_{\text{init}}$ are entirely identical. The main difference lies in the statistical survey approach we take to deal with the aforementioned sources of variance. Indeed, this time the agent performs a large number of explorations:

1. For each fixed set of starting conditions (i.e. environment state and initial configuration in this environment state) it performs $R = 1000$ distinct and independent explorations. Each corresponds to a sequence of draws $(a[t_n])_{n \in \mathbb{N}}$ so that the differences between these R runs are subject to the first source of variance, that of action selection.
2. These explorations are conducted for $N_\epsilon = 575$ distinct starting conditions. In this experiment the environment state is set constant (to the still picture depicted in Fig. 3.11a) to better control for its local variance, therefore each starting condition corresponds to the configuration (position and orientation) the agent starts at for its exploration. Because the agent and its displacements are small relative to the scale of the environment, as long as the explorations occur in a short timeframe they limit its interaction with the neighborhood of its initial configuration. Then, especially when one accounts for the variance of action selection of the R runs, the differences between the explorations for various starting conditions reflect the influence local properties of the environment have on learning the permutation matrices.

The agent still performs its learning *via* motor babbling and computing update rule (3.21) for any of its $R \times N_\epsilon = 575,000$ explorations; its empirical matrices M_{a_k} are reset to their initial values between each exploration. At any timestep t_n , $n \in \mathbb{N}$ of its interaction it records the values of the entropy criterion associated to each of its empirical matrices M_{a_k} , $a_k \in \mathcal{A}_{\text{init}}$, yielding




FIGURE 3.10: Representation of the criterion $C_k = C(M_{a_k})$ for the seven actions a_k . Each jump in this figure corresponds to a reevaluation of the criterion happening at a timestep when the corresponding action has been drawn in the set in $\mathcal{A}_{\text{init}}$. As expected, the criterion starts from 0 to reach 1, indicating that all possible permutations have been found.

the sequences $(C_k[t_n])_{n \in \mathbb{N}}$. Any given exploration is set to stop when all C_k values have reached 1 (corresponding to the agent completing discovery of its sensel permutations) or at timestep t_{3000} , whichever happens first.

3.3.3.4 Results

Fig. 3.10 presents the typical evolution of the C_k values of the entropy criterion during a sensorimotor exploration. Each curve in this graph is plotted against the timestep t_n , $n \in \mathbb{N}$ of the interaction, and corresponds to the seven distinct actions $a_k \in \mathcal{A}_{\text{init}}$ following the same color convention as in the description of the system. Following the previous discussion, we can observe that every curve goes from 0 to 1 as the agent makes progress in discovering its permutations matrices $M_{a_k}[t_n]$; moreover the increases in the values of the C_k occur in sparse and irregular jumps, corresponding to the timestep at which specific actions were selected in this particular exploration. This plot shows that a relatively small number of jumps suffices to reach the final limit state: in this particular exploration, each action a_k needs to be performed around 7 times for the agent to fully determine its corresponding sensel permutation $M_{\sigma_{a_k}}$. Furthermore we see that for any given action $a_k \in \mathcal{A}_{\text{init}}$, the amplitude of jumps in its C_k curve decreases during the sensorimotor experience: after only one execution of a_k , its corresponding value already approximately reaches $C_k[t_n] \approx 0.7$ (i.e. 70% of the total increase).

On the other hand, Fig. 3.11 presents statistical measures over all the explorations. Fig. 3.11a depicts the environment used for this experiment,




FIGURE 3.11: Statistical survey of the dynamics in the internal discovery of sensel permutations.

where each point in the superimposed grid figures one of the $N_\varepsilon = 575$ possible starting conditions. Following our discussion of Fig. 3.10, for each exploration E and each action a_k the number of jumps J_{E,a_k} in the curve C_k is taken as a measure of the difficulty in discovering the permutation $M_{\sigma_{a_k}}$ in the particular exploration E . Then the average of all such values for a given starting condition $J = \frac{1}{L} \sum_E \sum_k J_{E,a_k}$ (with averaging factor $L = R \times n_A = 7,000$) is depicted as the color of the point corresponding to that starting condition. Green points correspond to a low value of J (starting from approximately 4) whereas red ones represent conditions where more executions proved necessary (up to about 17 on average). Pictured circled are the points corresponding to the extremal starting conditions, that is those respectively associated to the lowest (for the green circle) and the highest values of J .

We can readily observe that the various starting conditions are overwhelmingly green: that is, the discovery process carried out by the matrix update rule generally completes rapidly¹⁹. On the contrary, the red dots are tied to some specific regions of the picture. In line with the previous intuition, these correspond to areas with somewhat low local contrast. Observe indeed that they all occur in large patches of sky, shadow, snow; it is especially clear in the reflections thereof, which tend to introduce blurring and therefore reduce local sharpness.

¹⁹In fact, the algorithm terminated on the condition that $t_n = t_{3000}$ only for some explorations at the condition with highest J ; in all other cases, the agent completely determined its sensel permutation matrices.

The statistics of the jumps J_{E,a_k} are further detailed in the histograms of Fig. 3.11. The topmost (Fig. 3.11b) breaks down the distribution of the J_{E,a_k} for the two particular extremal starting conditions (with each condition corresponding to the plot sharing the color of its dot). Consequently each histogram considered separately accounts for the variance induced by the selection of action at each particular starting condition, while the comparison between the distributions reflects that of the environmental influence on the dynamics of learning the M_{a_k} . The two distributions are unimodal, roughly symmetric, and separated: the maximal J_{E,a_k} value at the green starting condition (across *all* runs and all actions), being 4, is lower than even the minimal value at the red starting condition (that is 8). This shows in a strong sense that it is significantly easier to discover the sensel permutations in the first case as compared to the latter. Interestingly the internal variance is visibly different between the conditions, the spread of the corresponding distributions increasing with the associated difficulty values J (which by definition is equal to the mean of their respective distribution). Indeed, all explorations at the easiest condition have the update process of the M_{a_k} terminate at roughly the same pace: each action needs to be performed between 3 and 5 times ($J = 4$ on average) before its associated sensel permutation is correctly mapped. Instead, the same range is far more diffuse in the hard condition. There, actions can take between 8 and 35 executions (with an average of $J = 17$) before their matrices are fully determined. This shows that, at any given starting condition, observing the variance in the time needed to determine the permutations $M_{\sigma_{a_k}}$ provides information about the local richness of the environment to the agent even when it cannot directly experience other conditions.

On the other hand the bottommost histogram (Fig. 3.11c) is that of the distribution of the $R \times n_A = 575,000$ values of J_{E,a_k} , each corresponding to one of all independent explorations the agent does during the experiment. We see that the overall distribution is still unimodal but with a large positive skew; in fact, its mode is almost equal to its minimum. This corroborates our earlier observations that most starting conditions lead to quick convergence of the empirical matrices M_{a_k} . More precisely we can see that randomly selecting over all explorations (or equivalently over starting conditions), there is more than a 66% chance that permutation matrices $M_{\sigma_{a_k}}$ be determined before the corresponding action a_k has been performed 5 times. This imbalance also factors in the skew of the global distribution. Indeed the typical

convergence (that is the one observed in the “easy” conditions, such as pictured in Fig. 3.10) relies on the C_k values substantially increasing on the first few times the a_k are performed, a start-up behavior which is known to cause right-skew.

3.3.4 Simulation 3: internally building the conservative sub-algebra

The previous simulations all examined how the agent could discover the sensory outcomes of its motor actions by effectively performing them, that is by *acting* (J. K. O’Regan and Block, 2012). Accordingly, it only immediately leads to discovering properties of those actions it has actually performed in its sensorimotor interaction with its environment. These actions, however, may represent only a fraction of its motor capabilities; this is all the more true as the embodiment of the agent grows more complex. On the contrary, we now detail a simulation where the agent *cognizes*, using the sensory predictions it empirically learned and their compositional properties to infer global properties of its motor capabilities.

Once it has learned its (partial) permutation matrices M_{a_k} , $a_k \in \mathcal{A}_{\text{init}}$, the agent computes the set of all of their products. As per Eq. (2.11), this set is precisely the *submonoid* (of the monoid of sensory predictions $\mathfrak{S}_S$) *generated by the M_{a_k}* ²⁰. Furthermore it naturally carries a graph structure determined by the product law of matrices, which following our discussion in 3.2.2.2 can be transposed to that of their corresponding motor actions. This entails in particular that properties of the graph of sensory predictions the agent internally builds correspond to combinatorial properties of its action capabilities.

3.3.4.1 Generative Dijkstra-like algorithm for computing the group of sensory predictions

The relation between any starting set of motor actions $\mathcal{A}_{\text{init}}$ and the group $\langle \mathcal{A}_{\text{init}} \rangle$ it generates is encoded by their *Cayley graph* $\Gamma(\langle \mathcal{A}_{\text{init}} \rangle, \mathcal{A}_{\text{init}})$ (Cayley, 1878; Magnus, Karrass, and Solitar, 2004). It is defined as the (directed) graph which vertices are all elements of $\langle \mathcal{A}_{\text{init}} \rangle$ (that is all finite products over $\mathcal{A}_{\text{init}}$ up to inverses), and with edges a_k of $\mathcal{A}_{\text{init}}$ from vertex $a \in$

²⁰In fact, it is even a (sub)group as long as the agent can perform at least one rotation: in the following, we will therefore always treat it as a group.

$\langle \mathcal{A}_{\text{init}} \rangle$ to $a' \in \langle \mathcal{A}_{\text{init}} \rangle$ precisely when $a' = a_k a$. This definition makes the Cayley graph a formal correlate of the visualization of actions as vector fields in 2.3.2.1. Indeed let us think of the directed graph $\Gamma(\langle \mathcal{A}_{\text{init}} \rangle, \mathcal{A}_{\text{init}})$ as its visual representation, made of points (for the vertices) and arrows between these points (for the directed edges). Then it is precisely the collection of the (affine) vector fields associated with the actions a_k of $\mathcal{A}_{\text{init}}$, plotted starting at each of the points in the (much larger) set $\langle \mathcal{A}_{\text{init}} \rangle$. The main difference lies in the first case generally using motor configurations, that is positional data, as the set of its points; on the other hand, in the $\Gamma(\langle \mathcal{A}_{\text{init}} \rangle, \mathcal{A}_{\text{init}})$ graph arrows denote transformations between points which as motor actions are themselves displacements. The underlying analogy is the same as that under which one equates the point in the plane P which coordinates are $x_P = 2, y_P = 0$ and the planar displacement vector $\vec{u}_P = (2 \ 0)$ ²¹. One can for example verify that moving P by $\vec{v} = (0 \ 3)$ yields $x_{P+\vec{v}} = 2$ and $y_{P+\vec{v}} = 3$, which coincides with $\vec{v} + \vec{u}_P = (2 \ 3)$: although both viewpoints originate from entirely different objects, they correspond to the same thought process.

By definition (see 2.3.2.1) $\langle \mathcal{A}_{\text{init}} \rangle$ contains precisely those actions which can be written as words $a_{k_n} \dots a_{k_1}$ over $\mathcal{A}_{\text{init}}$. In terms of graph this immediately entails that $\Gamma(\langle \mathcal{A}_{\text{init}} \rangle, \mathcal{A}_{\text{init}})$ is *weakly connected* (Pemmaraju and Skiena, 2003), that is between any two of its vertices one can find an undirected path. In fact whenever $\langle \mathcal{A}_{\text{init}} \rangle$ is a group (such as in these cases) it is even strongly connected so that the path can be chosen to respect the orientation of its edges²². The graph also makes possible the definition of a metric on the group $\langle \mathcal{A}_{\text{init}} \rangle$, the *word metric* (associated to $\mathcal{A}_{\text{init}}$) $d_{\mathcal{A}_{\text{init}}}$. In algebraic terms, the value of $d_{\mathcal{A}_{\text{init}}}(a) = d_{\mathcal{A}_{\text{init}}}(e, a)$ for any a of $\langle \mathcal{A}_{\text{init}} \rangle$ is defined as the minimal whole number N such that a can be written as a word $a = a_N \dots a_1$ over the generating set $\mathcal{A}_{\text{init}}$. Importantly this graphically corresponds to the minimal path length (measured in number of edges traversed) over all paths from the vertex corresponding to $e \in \langle \mathcal{A}_{\text{init}} \rangle$ to that of a . This property makes apparent the fact that $d_{\mathcal{A}_{\text{init}}}(a, a')$, which accordingly measures the minimal path length from a to a' (corresponding to a minimal decomposition of $a'a^{-1}$), is indeed a metric in the usual sense (Sher and Daverman, 2001).

²¹Which –using the language introduced in the previous sections– informally reflects an *isomorphism as vector spaces* of planar positions and planar translations, both identified as $\mathbb{R}^2$.

²²This is essentially true because if one takes any two vertices a and $a' = a'_n \dots a'_1$ in $\langle \mathcal{A}_{\text{init}} \rangle$, $a^{-1} = a_m^{-1} \dots a_1^{-1}$ (which lies in $\langle \mathcal{A}_{\text{init}} \rangle$ since it is a group) provides a path

This examination provides the agent with a way to *internally* compute the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ which in turn carries combinatorial properties of composite actions $a = a_{k_n} \dots a_{k_1} \notin \mathcal{A}_{\text{init}}$, that is actions it did *not* actively learn during its previous sensorimotor experience. Indeed once the agent has completed the discovery of its sensel permutation matrices $M_{\sigma_{a_k}}$, it can further compute the result of successive predictions $M_{\sigma_{a_{k'}}} \times \dots \times M_{\sigma_{a_k}}$ as the composition of internally computable mappings following the argument made in 2.3.2.3. In terms of the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ this means that the agent can start at any vertex M_0 of the $M_{\sigma_{a_k}}$ matrices (e.g. the $M_{\sigma_{a_1}} = I_{N_c}$ matrix of the identity action) and explore its direct neighborhood in the graph by computing the $M_{\sigma_{a_k}} M_0$, yielding the collection N_{M_0} of all neighbours of M_0 in $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ as the new ends of the explored graph. This process of enumerating the neighbouring vertices can then be iteratively applied to these newly discovered ends to progress in the exploration of $< \mathcal{A}_{\text{init}} >$. This is realized by Algorithm 1 which implements (in pseudo-code) a Breadth First computation (akin to the more usual Dijkstra's algorithm (Dijkstra, 1959)) and construction of the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$. Since we know that $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ is strongly connected, there is in particular a directed path from M_0 to each of the vertices of $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$; therefore, we know from the properties of Breadth First Search (or more generally breadth-first traversals) that Algorithm 1 will eventually and successively reach *all* of the vertices of $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$. Of course, the theoretical Cayley graph is often infinite²³ so that a desired stopping length D is provided: in terms of the previous definitions, it means that the algorithm will return the ball (for the word metric $d_{\mathcal{A}_{\text{init}}}$) of radius D centered on M_0 .

of (oriented) edges in $\mathcal{A}_{\text{init}}$ from a to e while a' provides a path from e to a' . Therefore the concatenation a'' of the words $a' a^{-1} = (a'_n \dots a'_1) (a_m^{-1} \dots a_1^{-1}) = a'_n \dots a_1^{-1}$ gives a path of oriented edges such that $a'' a = a' a^{-1} a = a'$, that is a path from a to a' .

²³For example, the composite translations a_3^N that “go forward by N units” go on and on for arbitrarily large values of N .

Algorithm 1 Algorithm followed to build the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$

Input

- A The set of all matrices $M_{\sigma_{a_k}}$ learned in the previous experience
- D A bound on length of matrix combinations used
- O A reference matrix around which to explore

Output

- G A local view of the combinatorial graph of matrix products around O, using edges in A

Add O to collection U

O.depth $\leftarrow$ 0

Add node O to G

while U is not empty **do** $\triangleright$ True iff the neighborhood of some node K is still Unexplored

 K $\leftarrow$ node in U

for all M_a in A **do** $\triangleright$ Test all learned predictions starting from node K

 P $\leftarrow M_a K$

 P.depth $\leftarrow$ K.depth + 1

if P.depth $\leq$ D **then**

 B $\leftarrow$ False

for all node C in G **do** $\triangleright$ Test previously discovered nodes for equality

if predictions for P and C match **then**

 B $\leftarrow$ True

 Set edge M_a : K $\rightarrow$ C in G

end if

end for $\triangleright$ END for all node C in G

if B is False **then** $\triangleright$ Branch taken iff P := $M_a K$ was not previously discovered

 Add P to U

 Add node P to G

 Set edge M_a : K $\rightarrow$ P in G

end if

end if $\triangleright$ END if P.depth $\leq$ D

end for $\triangleright$ END for all a in A

 Remove K from U

end while $\triangleright$ END while U is not empty

3.3.4.2 Managing loss of information in nonconservative actions

The point underlying this experiment is that, following the development made in 3.2.2, the combinatorial properties of the sensory predictions $M_{\sigma_{a_k}}$ encoded in $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ are somewhat equivalent to that of their corresponding actions a_k in $< \mathcal{A}_{\text{init}} >$. Indeed for conservative actions, we know that

$$M_{\sigma_{a'}} M_{\sigma_a} = M_{\sigma_{a'a}}. \quad (3.25)$$

Instead, this relation is found in a weaker form in a non conservative setting: as seen in 3.2.2.3, we can only hope for $M_{\sigma_{a'}} M_{\sigma_a} \preceq M_{\sigma_{a'a}}$. This introduces a technical difficulty in Algorithm 1 compared to the usual breadth first traversal. Indeed when it processes some vertex M recently discovered to enumerate the collection of its neighbours N_M , it must take care to discard all elements of N_M it has previously discovered instead of adding them to the list of vertices it will next process. In particular, provided M was not the starting vertex M_0 it must have been reached by an edge $M_{\sigma_{a_k}}$ originating from some neighbouring vertex M' , so that $M = M_{\sigma_{a_k}} M'$; then it is necessary for the algorithm to discern that the $M_{\sigma_{a_k}^{-1}}$ edge starting from M' actually ends at M , even though for non conservative actions we may have $M_{\sigma_{a_k}^{-1}} M' = M_{\sigma_{a_k}^{-1}} M_{\sigma_{a_k}} M \neq M$. It therefore searches for *compatibility* (in the sense of semigroups discussed in 3.2.2.3) of prediction matrices rather than their strict equality. This is reflected in Algorithm 1 by branching on the condition "predictions for P and C match", which denotes the boolean test of assertion " $P \sim C$ "²⁴.

This solution however raises another issue. From the same theoretical discussion we are assured that

$$M_{\sigma_{a_1}} M_{\sigma_{a_2}} \sim M_{\sigma_{a'_1}} M_{\sigma_{a'_2}} \text{ whenever } a_1 a_2 = a'_1 a'_2 \quad (3.26)$$

by which Algorithm 1 will correctly identify all matrix products $M_{\sigma_{a_{k_n}}} \dots M_{\sigma_{a_{k_1}}}$ which represent the same originating action $a \in < \mathcal{A}_{\text{init}} >$ as such. However the converse is not necessarily true: owing to the loss of information in the generalized successor maps of non conservative actions a_k , the agent

²⁴In terms of implementations, this test can be carried out with ease on the matrix representation of partial bijections discussed in 3.3.2.1. Indeed two matrices of partial bijections represent contradictory predictions if (and only if) at any one row they both have a nonzero coefficient, the two of which do not occur at the same column. This can be easily and efficiently detected by taking the maximum (or equivalently the logical *or*) of the two matrices, and checking that each row of this result sums to at most one.

may find several matrix products that are only spuriously compatible. In this case, the unconditional use of Algorithm 1 will result in erroneously equating vertices that originate from distinct motor actions. We can however get some quantified conditions on the rate at which this loss accumulates as the agent stacks successive predictions. Indeed denote by $l_{a_k} \in \mathbb{N}$ the number of sensels which *fail* to have a successor under each of the starting actions $a_k \in \mathcal{A}_{\text{init}}$. From the definition of the (partial) permutation matrices $M_{\sigma_{a_k}}$ in 3.3.2.1, we know that l_{a_k} corresponds to the numbers of rows in $M_{\sigma_{a_k}}$ where *all* coefficients are null. But then temporarily interpreting $M_{\sigma_{a_k}}$ as the more usual matrix of a linear mapping shows that l_{a_k} is, in fact, the dimension of its kernel. Therefore it follows from the rank-nullity theorem²⁵ that

Proposition 11. *Let a_k and $a_{k'}$ be any two starting actions of $\mathcal{A}_{\text{init}}$, and let $l_{a_{k'}, a_k}$ denote the number of sensels which lack successor under (composite) predictions $M_{\sigma_{a_{k'}}} M_{\sigma_{a_k}}$. Then we have*

$$\min(l_{a_k}, l_{a_{k'}}) \leq l_{a_{k'}, a_k} \leq l_{a_k} + l_{a_{k'}}. \quad (3.27)$$

Moreover observe that for this particular agent, all actions $a_k \in \mathcal{A}_{\text{init}}$ result in *rigid* displacements of its total receptive field $F_{\mathcal{C}}$ in the 2-d environment²⁶. Because the composition of rigid transformations is itself rigid, we are certain that every composite action $a \in \langle \mathcal{A}_{\text{init}} \rangle$ also possesses this property. But, in the plane, such displacements are entirely determined by their effect on two distinct points. Therefore, it suffices that any two products $M_{\sigma_{a_{k_p}}} \dots M_{\sigma_{a_{k_1}}}$ and $M_{\sigma_{a_{k'_q}}} \dots M_{\sigma_{a_{k'_1}}}$ of partial bijection matrices coincide on at least two nonzero coefficients for them to partialize the same action. Indeed any M_{σ_a} matrix has nonzero coefficient at coordinates (i, j) if and only if its displacement of receptive fields sends F_{c_i} on F_{c_j} , by which two nonzero coefficients would provide the two points determining the displacement. Finally, by definition of the starting actions (see 3.3.1) we know that $l_{a_k} = 0$ for the rotations of the agent (which are actually conservative) and $l_{a_k} = W_c$ or $H_c = 10$ for its elementary translations. From this it follows that in the case where the agent successively computes its prediction associated to $a_{k_1}, a_{k_2}, \dots, a_{k_n}$, by successive application of the right hand side in Prop. 3.27 we have

$$l_{a_{k_n}, \dots, a_{k_1}} \leq 10n^{27}.$$

²⁵see (Axler, 2014) for a reference on linear algebra (from which the result originates), although only a minuscule part is needed here.

²⁶More precisely these displacement even preserve handedness, that is they do not include reflections over axes.

Now let us consider two products of permutation matrices M and M' , respectively corresponding to the products over the sequences $a_{k_p}, \dots, a_{k_1}$ and $a_{k'_p}, \dots, a_{k'_1}$, and assume they result in compatible (in the sense of semigroups as in 3.2.2.3) matrices. Then by the pigeonhole²⁸ principle whenever $\text{rank } M + \text{rank } M' \geq N_c + 2$ the two product matrices must coincide on at least two nonzero coefficients, which in turn entails that they in fact partialize the same product $a_{k_p} \dots a_{k_1} = a_{k'_p} \dots a_{k'_1}$.

This yields some conditions for the agent such that the converse of Eq. (3.26) holds as per

Theorem 4. *Let p, q be two natural numbers such that $p + q < \min(W_c, H_c) = 10$, let $a_{k_1}, \dots, a_{k_p}$ and $a_{k'_q}, \dots, a_{k'_1}$ be any starting actions. Then*

$$\left(M_{\sigma_{a_{k_p}}} \dots M_{\sigma_{a_{k_1}}} \sim M_{\sigma_{a_{k'_q}}} \dots M_{\sigma_{a_{k'_1}}} \right) \Rightarrow \left(a_{k_p} \dots a_{k_1} = a_{k'_q} \dots a_{k'_1} \right).$$

Proof. Let us respectively denote M and M' the matrices corresponding to the products of the $M_{\sigma_{a_k}}$ over the sequences $a_{k_p}, \dots, a_{k_1}$ and $a_{k'_q}, \dots, a_{k'_1}$. By the rank-nullity theorem we have $\text{rank } M = N_c - l_{a_{k_p}, \dots, a_{k_1}}$ and $\text{rank } M' = N_c - l_{a_{k'_q}, \dots, a_{k'_1}}$ so that

$$\text{rank } M + \text{rank } M' = 2N_c - \left(l_{a_{k_p}, \dots, a_{k_1}} + l_{a_{k'_q}, \dots, a_{k'_1}} \right)$$

but by the previous remarks

$$\left(l_{a_{k_p}, \dots, a_{k_1}} + l_{a_{k'_q}, \dots, a_{k'_1}} \right) \leq 10p + 10q = 10(p + q).$$

Since p and q are whole numbers, so is their sum so that $p + q < 10 \Rightarrow p + q \leq 9$, therefore $10(p + q) < 90 = N_c - 10$.

But then

$$2N_c - \left(l_{a_{k_p}, \dots, a_{k_1}} + l_{a_{k'_q}, \dots, a_{k'_1}} \right) > 2N_c - (N_c - 10) = N_c + 10 > N_c + 2.$$

Therefore the restriction on p and q suffices to ensure that the previous argument holds. $\square$

²⁷The upper bound can even be chosen lower at $10n_t$ where n_t is the number of translations amongst the a_{k_i} . Although this information is available to the agent which –at this point– has discovered its M_{σ_a} matrices, it bears not on the crux of the argument.

²⁸alternatively (Dirichlet's) drawer principle

which in particular entails for its Cayley graph that

Corollary 5. *For a given radius parameter D , Algorithm 1 starting at the identity introduces no spurious compatibilities between the product of permutation matrices under the condition that*

$$D \leq \left\lfloor \frac{\min(W_c, H_c) - 1}{2} \right\rfloor.$$

Therefore, for such radii Algorithm 1 correctly returns the expected result, that is the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ ²⁹ in the usual sense of groups.

This result ensures that although we cannot hope for an unbounded isomorphism between the a_k and the $M_{\sigma_{a_k}}$ in nonconservative settings, the process carried out by Algorithm 1 nevertheless yields a valid correspondence up to a specific radius entirely dependent on the sensory apparatus of the agent.

3.3.4.3 Force directed methods for plotting the Cayley graphs


Algorithm 1 only results in an abstract graph, which itself acts for the agent as an understanding of the compositional law in $< \mathcal{A}_{\text{init}} >$. However, we still need to plot it for external visualisation. To this end we use a force-directed algorithm³⁰, which attempts to determine an embedding of $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ in $\mathbb{R}^2$ or $\mathbb{R}^3$ via a generic physics-based process. It operates on the vertices of the graph as if they were atomic particles linked by springs, with the rules that

- Any two vertices exert a *repulsive* force on each other, modelled on Coulomb's law for charged particles,
- Any two vertices connected by an *edge* are *attracted* to one another, with the interaction being that of a square law of the distance separating the vertices.

The drawing algorithm initializes the spatial embedding by randomly placing all vertices in $\mathbb{R}^d$; then it keeps updating their positions according to the aforementioned interactions, stopping once it reaches an equilibrium configuration.

²⁹Of course, to be precise it correctly returns the subgraph of $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ delimited by the corresponding ball (of radius D) of.

³⁰In this particular case, we use the force setting of MATLAB's graph plot which draws on the method described in (Fruchterman and Reingold, 1991).

FIGURE 3.12: Directed graph of empirical permutation matrices M_{a_k}

This type of algorithm has been found to meet a number of desirable “aesthetic criteria” (Kobourov, 2012). In particular, it graphically reflects the internal symmetries of the graph as well as tends to produce uniform edge lengths. Moreover it operates with no *a priori* knowledge of the semantics of the graph; instead, it only probes the connectivity at each vertex. Therefore it is especially suited to our approach outlined in 2.3.2.3 since its results illustrate what a “naive” observer might gather from the structure of $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$.

3.3.4.4 Results

Figure 3.12 shows the Cayley graphs we obtain by applying Algorithm 1 to the prediction matrices $M_{\sigma_{a_k}}$ from the first part of the experiment. Both graphs have been built starting from $M_{\sigma_{a_1}} = I_{N_c}$; their radius parameter $D = 3$ used as input is purposefully kept low enough to make the graphs small, increasing their readability. Moreover, the color convention used in the plot is chosen consistently with the illustration of the setup in Fig. 3.8.

The first graph Fig. 3.12a is obtained when one restricts its attention to the sole combinations of *translations* $a_1, \dots, a_5$. If one denotes $\mathcal{A}_{\text{init}}^t \subset \mathcal{A}_{\text{init}}$ their collection, then it can alternatively be characterized as the Cayley graph $\Gamma(< \mathcal{A}_{\text{init}}^t >, \mathcal{A}_{\text{init}}^t)$ of the group $< \mathcal{A}_{\text{init}}^t >$ they generate. The central node is that of $M_{\sigma_{a_1}}$ from which the exploration of the graph branched out; the choice of radius is then apparent in that every vertex is separated from this central

node by at most $D = 3$ edges. The plot clearly shows rotational symmetry of order 4 corresponding to every four directions of translation a_2, a_3, a_4, a_5 . Furthermore, it mirrors many algebraic properties of the $M_{\sigma_{a_k}}$ prediction matrices the agent empirically learns. Indeed it appears that light blue arrows always correspond to self-loops in the graph: this directly mirrors them being the arrow associated to $M_{\sigma_{a_1}}$, that is the matrix identity I_{N_c} . Moreover, we can also distinguish some particular pairs of edges such as the purple and orange ones: indeed it appears that (up to the borders of the graph) whenever a purple edge sends vertex M to M' , then a corresponding orange one must send M' to M . Therefore successively following the purple arrow and then the red one results in a self loop $M (\rightarrow M') \rightarrow M$: this thus identifies the corresponding predictions matrices $M_{\sigma_{a_k}}$ and $M_{\sigma_{a_{k'}}$ as reciprocal inverses. Importantly this relation is somewhat homogeneous: whenever such pairs compose to a self-loop at any one vertex they must do so at *every* one of them, coinciding with the algebraic notion of inverse (see 2.3.2.1). Finally the last two edges (green and yellow) similarly correspond to another pair of inverse predictions.

This discussion was made as one of properties of the *sensory predictions* $M_{\sigma_{a_k}}$ since, by construction, Algorithm 1 builds the Cayley graph of the group generated by these matrices. However it is interesting to take it as a graph of their originating *motor actions*, along the result of Theorem 4. Then, one can think of the properties they visually assess as algebraic assertions about the motor capabilities of the agent. And indeed the previous observations about self-loops and pairs of inverse edges directly reflect the motor relationships of their corresponding actions. For example the (orange, purple) edges, which are graphical inverses, respectively denote the forward and backward translations; so do the yellow and green inverse edges for the rightward and leftward ones. Yet another observable property can be found in the four central squares. Indeed, one can graphically check that taking the yellow edge and following with the purple one always leads to the same endpoint that purple first, yellow second. This reflects the intuitive idea that, excluding rotations, moving “rightward then backward” shares its result with “backward then rightward”³¹. These assertions might seem painfully obvious to an external observer; however, crucially they did not appear in any way in the *a priori* knowledge the agent was given. Instead, they emerge as the results of the incrementally diverse internal computations it can perform as its

sensorimotor experience unfolds.

The same examination can be made of the second graph Fig. 3.12b, in which all seven starting actions $a_k \in \mathcal{A}_{\text{init}}$ are considered. Note that the graph $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ obtained *via* Algorithm 1 has been further trimmed at its boundaries: this is made necessary because the addition of rotations make the plot “inherently” 3-dimensional, by which it very quickly loses clarity as the number of vertices grow. Therefore it does not fully represent the ball of radius $D = 3$ in $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$, in which case it should in particular contain the full first graph as a subgraph. Still, we can see that it shows the previous four central squares in an “equatorial” plane of translations³². However this plane is now supplemented by a third dimension depicting the change of orientation caused by the rotations a_6 and a_7 , corresponding to the vivid red and blue vertical arrows. In fact its global shape is roughly the union of four disjoint planes of translations, each corresponding to one fixed orientation. These planes are in turn linked by the rotation actions according to their effect on the orientation of the agent. One can indeed observe that each of the subgraphs of translation planes is similar to one another up to a rotation of its edges, depicted by their colors (or directions of arrows) shifting between any two of those planes.

These additional paths introduce several graphical properties. First, we can observe that new squares of paths have appeared. For example following the red edge and then the green one always leads at the same result as purple first, red second. As before, this is made clearer by thinking in terms of motor actions: then the previous trajectories respectively correspond to “turn left then leftward” and “backward then turn left”. Importantly, these pairs of equivalent paths are not as trivial as in the previous case of translations. Indeed, the previous example shows that the graph encodes the information that $a_2 a_7 = a_7 a_4$; in particular it also shows that $a_2 a_7 \neq a_7 a_2$, which reflects the fact that translations and rotations do *not* commute. The layout of the 4 planes of translations also makes apparent that rotations form cycles (of order 4), changing orientation at constant position in the planar subgraphs.

³¹Algebraically this corresponds to the commutativity property of translations, that is $a_{k'} a_k = a_k a_{k'}$ for $a_k, a_{k'} \in \mathcal{A}_{\text{init}}^t$. This is not a trivial property in the context of groups (as opposed to say vector spaces), as can be seen by the matrix product being generally non-commutative.

³²The orientation of the plane –contrary to it being a plane in itself– is only an effect of the external viewpoint chosen, e.g. for added convenience. In particular, it bears no consequence on the knowledge brought by the abstract $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$ graph.




FIGURE 3.13: Graphs of predictions of translations obtained *via* Algorithm 1 for higher values of radius parameter D . The “central” node M_{a_1} used to start the algorithm is circled in black. For these graphs (and these graphs only) the camera considered is 20-by-20 instead of 10-by-10 otherwise.

Such a cycle is highlighted in Figure 3.12b with the 4 circled nodes, corresponding to the 4 rotational orientations at the central position of $M_{\sigma_{a_1}}$ in the translation graph. In fact, the visual property that edges of rotations form cycles in the graph also mirrors a corresponding algebraic one of a_6 and a_7 : it reflects the external observations that successively following four $\frac{\pi}{2}$ (or $-\frac{\pi}{2}$) rotations takes one back to its initial configuration. This specific property characterizes rotations among starting actions, thus providing the agent a way to internally distinguish them from translations from that point on.

Folding of the resulting graph for higher radii Figure 3.13 shows the Cayley graphs of translations that can be obtained for higher radii D . Note that it corresponds to the experience of an agent identical to that presented in 3.3.1 except that the pixel layout of its camera is 20×20 , in order to better visualise the effects of scale on the resulting structure. Moreover, only translations have been used in building this graph (as in 3.12a before) since the multiplication of arrows otherwise makes reading any such large graph tedious. The first graph, Fig. 3.13a, is obtained for a depth parameter D equal to 9. We see that it is largely similar to that previously shown in Fig. 3.12a: all observations about inverse pairs and commutative squares are homogeneously found at every interior node of the graph. The force-directed algorithm has automatically chosen to embed it as a square³³ grid, clearly adhering to our

intuition about the structure of the translations available to the agent.

The situation is much different in Fig. 3.13b, which depicts the corresponding graph for depth parameter $D = 10$. The same grid of displacements –including intact relations between edges– can be discerned, particularly around the origin point M_{a_1} circled in black. However its shape is now clearly 3-dimensional, with the center node corresponding to a saddle point of the embedding. Furthermore two particular vertices, circled in red, have appeared in the graph. Let us focus on the vertex appearing at the “top”³⁴ of the graph, which we will denote V . One path going from the central black node to V , taking the “right” direction, is that made of a succession of orange arrows. Precisely, it shows that composing M_{a_5} a total of 10 times leads to vertex V . On the other hand, another similar path reaching V goes “left”: it corresponds to $M_{a_3}^{10}$. Recall that from the definition of the setup illustrated in Fig. 3.8, actions a_3^{10} and a_5^{10} respectively denote going forward and backward by 10 units. Therefore vertex V (and the similar red vertex at the bottom, which corresponds to the rightward/leftward actions) illustrates how compositions of predictions M_{a_k} can result in the grouping of products that correspond to distinct, and in fact diametrically opposed, actions.

Figure 3.14 illustrates the product predictions at play. Note that the results are once again presented for a 10-by-10 camera to increase legibility; moreover, as opposed to those previously seen in 3.3.2, none of the two matrices is that of an $a_k \in \mathcal{A}_{\text{init}}$ the agent has actually learned in its active sensorimotor experience. Instead, they are the matrices of *composite* predictions it cognizes in building the graph of Algorithm 1. As before, a_3^5 and a_5^5 respectively correspond to going forward and backward by 5 displacement units. And as a matter of fact the matrix compositions are visibly distinct, with the external visualisation of the bottom row (realized as in Fig. 3.9) reporting the expected effect. However what happens is that the two matrices are somewhat *too* different, so much so that they do not map any common sensel. Indeed, observe that $M_{a_3}^5$ precisely sends the bottom half of the sensel layout on the top half via a rigid translation; of course, $M_{a_5}^5$ performs the reciprocal mapping of the top half onto the bottom half. But consequently, these predictions

³³A small degree of curvature (which, although much less extreme, resembles that of Fig. 3.13b) can be seen, for example in the way distances seem to “shrink” on the borders of the graph. It is entirely a side effect of the method attempting to find a 3-d embedding, by which distances between extremal points can be further increased by having them take opposite directions in the third dimension.

³⁴All such spatial indications are of course only defined in, and relevant to, our *external* examination of the graph.




FIGURE 3.14: Visualisation of compatible predictions resulting from distinct actions.

are entirely *noncontradictory*. In particular there exists a total permutation of sensels, namely that which exchanges the top and bottom halves, which extends both partial predictions. This permutation is illustrated in Fig. 3.15. Notice that following the cutting and gluing procedure followed one gets a valid permutation of the blocks, which in turn corresponds to a relation between input sensation and output sensory prediction.

Consequently, the two predictions of Fig. 3.14 are compatible (in the sense of semigroups) even though they correspond to distinct motor actions. Importantly this prediction is entirely virtual. We externally know that it cannot actually exist, for example from our argument that the embodiment of the agent enforces rigid displacements of its receptive fields. The distinction occurs because what Algorithm 1 builds is actually a graph of *equivalence classes* (for the compatibility relation) of compositions of partial permutations $M_{\sigma_{a_k}}$, and not one of actions themselves. It is also essentially the issue Corollary 5 guards us against by formulating conditions under which the two notions can safely be equated³⁵. However, this examination shows that the “error” observed in the graph is a natural consequence of the same idea that created




FIGURE 3.15: Illustration of a sensel permutation that would extend both “spuriously” compatible predictions. It does not correspond to an *actual* valid action of the agent due to its embodiment restricting movements of its field of view to rigid ones.

it in the first place provided the agent is naive. Indeed the way in which we could determine when and how it occurs here made explicit use of our external knowledge of the system; in particular, it relied on knowing that the sensory apparatus of the agent is but one rigid body (so that the movements of its field of view are themselves rigid) and that it is essentially sampling a 2-dimensional surface. Without it, we would be unable to determine for sure that the permutation of Fig. 3.15 is in fact impossible. Consequently, we cannot reasonably expect the agent to be *a priori* able to determine this. What is generally and internally true, then, is that Algorithm 1 outputs the expected result for “small enough” distances from the origin, after which it yields such erroneous compatibilities. It should nevertheless be highlighted that vertices in the graph do not correspond to a fixed configuration in ambient space as much as a relative one (with respect to the “current” one). As a consequence, the agent may still make use of the farther regions of the graph. Although it may not validly *cognize* about arbitrarily large compositions of action predictions it may instead settle on a smaller composition, that is one small enough to fall in the ball of Corollary 5. It may then perform the corresponding movement –which it is assured will correspond to its intended effect– and reevaluate its sensory predictions at the new sensory output it has thus acquired. Therefore it may still use the graph structure of $\langle \mathcal{A}_{\text{init}} \rangle$ to navigate its full range of actions by performing such short hops.

The locality requirement of this technique also reflects the agent having no memory, no *internal representation* of its environment or their interaction. It can only access the sensory output it is currently experiencing, which in turn informs it about the immediate state of its interaction with the world. For a clear example of this, consider the agent choosing to perform the a_3^5

“move forward by 5 units” action (which prediction we examined earlier in Fig. 3.14). Once it has reached its target configuration, it can then proceed and predict the sensory outcome of now performing a_5^5 , that is the inverse action “move backward by 5 units”. If it remembered what it just saw before its first jump, it would obviously reach a perfect and total prediction (correctly) concluding that this second hop will take it back to its starting configuration; on the contrary, it is restricted to having the partial prediction seen in Fig. 3.14 because it does not keep such information in an internal representation. A corresponding viewpoint, instantiating a form of *vehicle externalism* (S. Hurley, 2012; Rowlands, Lau, and Deutsch, 2020) and notably outlined in (J. K. O'Regan, 1992), is that the external world functionally realizes a record storing the necessary “memory”. As for the agent, what it keeps instead is the ability to access and leverage the properties learned from its action graph at *any* future stage of its sensorimotor experience to further probe this record.

3.3.5 Simulation 4: reparameterizing basis actions

All previous simulations have been run with the set $\mathcal{A}_{\text{init}}$ of initial motor actions performed by the agent taken equal to A , that of the a_k actions we externally defined in 3.3.1. This is in itself a very strong assumption about the *a priori* effective motor capabilities of the agent, which the previous results would ideally not depend upon. We now turn to reproducing the final results of the experiment with this assumption no longer holding: in the following, $n_{\mathcal{A}}$ available starting actions $a \in \mathcal{A}_{\text{init}}$ are drawn in the collection of *combinations* of the $a_k \in A$. The agent therefore first learns the permutation maps M_{σ_a} of its particular starting actions $a \in \mathcal{A}_{\text{init}}$ (as in 3.3.2) before computing their corresponding Cayley graph (as in 3.3.4).

3.3.5.1 Dependence of the graph with respect to the starting action set

The discussion in 3.3.4.1 and its notation itself highlight that the Cayley graph $\Gamma(<\mathcal{A}_{\text{init}}>, \mathcal{A}_{\text{init}})$ we seek to obtain intrinsically depends on the initial starting action set $\mathcal{A}_{\text{init}}$. This influence is twofold. First, the vertices of the graph are precisely the elements of $<\mathcal{A}_{\text{init}}>$. Therefore it will appear in the graph whether the initial motor capabilities of the agent allow it to generate its full

³⁵Note especially that, for the 20-by-20 camera (that is $W_c = H_c = 20$), the loss of correspondence precisely happens as D goes from 9 ($< \lfloor \frac{20-1}{2} \rfloor$) to 10 ($\geq \lfloor \frac{20-1}{2} \rfloor$).

range of displacements. In fact, this corresponds to the graphs of displacement we showed in 3.3.4: the graphs built only on the translation matrices (which we denoted $\mathcal{A}_{\text{init}}^t$) “lack” the nodes corresponding to distinct orientations of the agent. Nevertheless, we described in 2.3.2.1 how distinct sets of motor actions could generate the same monoid; we know that whichever starting actions $a \in \mathcal{A}_{\text{init}}$ are drawn, their generated monoid $\langle \mathcal{A}_{\text{init}} \rangle$ cannot end up larger than that generated by the seven $a_k \in A$ we studied in the previous experiment³⁶.

Under the assumptions that the graphs from two starting motor sets cover the same vertices, another dependency appears concerning their arrows. The $\Gamma(\langle \mathcal{A}_{\text{init}} \rangle, \mathcal{A}_{\text{init}})$ graph is such that at each vertex $n_{\mathcal{A}}$ go outward (and as many inwards), one corresponding to each starting action (as seen in the previous graphs of Fig. 3.12 and Fig. 3.13). Consequently we know that altering the starting motor actions $\mathcal{A}_{\text{init}}$ will alter the degrees of the graph, which in turn affects the $d_{\mathcal{A}_{\text{init}}}$ word metric³⁷. For a draw of starting actions $\mathcal{A}_{\text{init}}$, three main situations may occur:

- Some of the “external” a_k may be missing from $\mathcal{A}_{\text{init}}$, that is the agent cannot perform them during its active sensorimotor exploration;
- Some a_k may have been drawn multiple times, so that several actions of $\mathcal{A}_{\text{init}}$ are in fact duplicates of one another;
- Some starting actions $a \in \mathcal{A}_{\text{init}}$ may be composite actions themselves, occurring only as composition products of the a_k of A .


Figure 3.16 presents the output of Algorithm 1 for two distinct runs of the simulations, each with its own $\mathcal{A}_{\text{init}}$ set. The first graph, Fig. 3.16a, corresponds to

$$\mathcal{A}_{\text{init}1} = \{a_1, a_3, a_4, a_5, a_6, a_7, a_3\}.$$


We can observe a particular pair of arrows, displayed yellow and black, which run parallel; contrary to the previous examples, these arrows also share their orientation. This reflects the fact that a_3 appears two times in

³⁶This is a consequence of $\langle \mathcal{A}_{\text{init}} \rangle$ being the *smallest* monoid containing (all) the actions a of $\mathcal{A}_{\text{init}}$: we know that $\langle A \rangle$ is such a monoid, therefore it must in particular be larger than $\langle \mathcal{A}_{\text{init}} \rangle$.

³⁷The point about degree is sufficient to show that the graphs for two distinct $\mathcal{A}_{\text{init}}$ are not isomorphic as graphs as well as not even homeomorphic. However they are always *quasi-isometric* (Sher and Daverman, 2001), by which the ratios of distances between vertices in the two instances are controlled by a single global upper bound.



(A) Graph obtained with a duplicate starting action (depicted in black) and a missing action (a_2 , green in other figures). A path made of 3 edges (red, orange, blue) equivalent to the missing starting action is highlighted, providing an inverse to the (also highlighted) yellow edge.



(B) Graph obtained by adding a combination (depicted in black) to the set of starting actions. Here, the action set was limited to translations to keep a clear visual.

FIGURE 3.16: Effects on the Cayley graph of various random draws of starting actions.

$\mathcal{A}_{\text{init}1}$, which the agent could not determine at the start of the experiment³⁸. The arrows have been grouped because although the two occurrences of a_3 may initially seem distinct to the agent, they cause the same effects in the sensorimotor interaction. Their corresponding prediction matrices M_{σ_a} it learned during its active exploration have therefore converged to the same result, by which they *a fortiori* take the same paths in $\Gamma(< \mathcal{A}_{\text{init}} >, \mathcal{A}_{\text{init}})$. Besides, these arrows are missing an inverse one, corresponding to the (absent) green edge of the previous graphs. This is caused by a_2 missing from $\mathcal{A}_{\text{init}1}$. This especially means that the agent could not initially perform leftward translation *via* $\mathcal{A}_{\text{init}}$. Nevertheless, we can see that the structure brought about by the graph provides a motor alternative. Indeed, the “red $\rightarrow$ orange $\rightarrow$ blue” path (highlighted in its leftmost instance) acts as an inverse to the yellow (and black) edges. While the latter allows the agent to go from the red vertex to the blue one, the former allows it to close the loop back to its initial configuration. Illustrated as motor displacements, this corresponds to “rotate clockwise, move backward and then rotate counter-clockwise” being

³⁸As seen in 3.3.1, its ability to compare actions for equality applies to occurrences in the multiset $\mathcal{A}_{\text{init}}$.

an alternative to “move leftward” and indeed we can see that it is a correct inference.

Figure 3.16b, which corresponds to

$$\mathcal{A}_{\text{init}2} = \{a_1, a_2, a_3, a_4, a_5, a_5a_3\},$$

instead falls under the last of the 3 situations. It presents the familiar structure of translation graphs, although this time a new edge (colored black) appears as a diagonal in its squares. It is the edge associated with the additional action $a = a_5a_3$, “forward then rightward” as a motor displacement. Because of our assumptions about *a priori* knowledge, the agent had no initial clue about this action being the product of two other available actions (a_3 and a_5 , respectively displayed in yellow and orange). It however ends up appearing in the resulting graph, because the sensory consequences of both alternatives coincide. As before, that provides a way for the agent to determine which actions (if any) of its available repertoire $\mathcal{A}_{\text{init}}$ are somewhat redundant: indeed, the graph structure makes clear that it can prune all such black edges without losing any of the range of reachable states.

3.4 Conclusion

A fundamental property of the embodied account of perception is that it is a localized process. Indeed, at any one point of its sensorimotor experience the agent is interacting with some restricted part of the environment, which lies in some corresponding restricted part of space. Accordingly, changing with which part of its environment it interacts requires the agent to leverage its motor capabilities and perform displacements.

In this chapter, we have shown how the shifts of an agent’s receptive fields as it moves during its interaction lead to certain typical sensory sequences. A theoretical discussion, first made for the ideal conservative actions, demonstrated that these shifts correspond to permutations of the outputs of the agent’s sensels; we also further showed that considering the loss of information exactly generalizes the result to partial permutations. What this establishes, in essence, is that the phenomenon of shifts of receptive fields can be picked up on by agents which are naive in the very restrictive sense detailed in 2.3.2.3 *via* their sensory images as sensel permutations.

Moreover, we also discussed how the agent can discover the algebraic structure of its motor actions by developing the capability to predict their

outcomes. Indeed we showed that, whatever the particular “state” in terms of which it may predict action effects, the composition law of predictions must necessarily match that of their originating actions. As a result, such predictions may serve for an agent as a proxy to the algebraic structure of its motor actions. This sort of result has previously been experimentally observed (see e.g. (Alban Laflaquière and Hafner, 2019)); nonetheless, the theoretical development we made lends streamlined insight into this property, which we find to be entirely independent of the means by which the agent achieves action effect prediction.

Taken together, the two previous points show that the agent can *effectively* (that is, internally) discover (partial) sensory predictions associated with its displacements and subsequently infer the combinatorial structure of these displacements. We illustrated these two main results in corresponding simulations we detailed in 3.3.2 and 3.3.4. First, they demonstrated that a naive agent can indeed internally discover the sensel permutations associated with its actions. Notwithstanding their possibly being inscrutable to us depending on their specific implementation, they nevertheless exhibit the desired algebraic structure mirroring that of motor actions. Accordingly, the experiment showed how they allow the agent to infer a number of previously hidden combinatorial properties of its actions; for instance, it can establish which pairs of actions commute as well as which are made of mutual inverses.

Several extensions naturally follow from this examination. First, one could pursue determining how the internal discovery of its sensory predictions provides the agent with new knowledge of its interaction. Indeed we only illustrated the main result (see Corollary 2) through the straightforward application of 3.3.5, whereby the agent leveraged its learned sensory predictions to better choose its effective action repertoire. However, numerous works show that the internal knowledge of action effects predictions entails further capabilities (Shin, Proctor, and Capaldi, 2010) such as object perception (Le Hir, Sigaud, and Alban Laflaquière, 2018; Maye and Andreas K. Engel, 2011), action selection and motor planning (Maye and Andreas K. Engel, 2012), motor control (Schröder-Schetelig, Manoonpong, and Wörgötter, 2010) or motor sequence compression (Ortiz and Alban Laflaquière, 2018). Accordingly it should be made explicit how one can go about formulating such results in the proposed formalism, and especially to what extent they can be made to follow from Corollary 2.

Second, one could try to improve on the first part of the experiment, that

is the internal determination by the agent of its sensory predictions. Indeed, the proposed implementation crucially relies on ideal conditions which can only be guaranteed in carefully controlled environments such as computer-simulated ones. One such condition, imposed by the reliance on categorical judgements of *equality*, is the absence of any noise whatsoever. In particular, the experimental setup required that the agent precisely repeat the same displacements over and over; that any given environment state deterministically result in a fixed sensory output, without any sensory noise; and that the environment state and agent's motor configuration only be allowed to change consecutively. A classical solution proposal would involve shifting to approximate judgements, by which the agent could hope to filter out said variance. We explore how it can begin to form such approximate judgements in the next part (see 4). Moreover, the particular mechanism the simulated agent leveraged to discover its sensory predictions is that of conservativity. It is however far from the only such mechanism, and not without its limitations. We showed that conservativity mainly characterizes (rigid) displacements of the agent, which form a very specific subset of its motor capabilities. Furthermore, the experiment demonstrated in 3.3.4 that nonideal conservative prediction only stays valid for (very) small movements, predictably yielding incorrect results at larger scales. As a result, one may be interested in circumventing this limitation (e.g. by introducing some internal memory to establish constancy) or exploring other types of sensorimotor predictions.

Chapter 4

Towards a naive continuity of sensorimotor experiences

Contents

4.1	Speaking from experience: the need for continuity	131
4.2	Finding smoothness in statistical regularities	133
4.2.1	Characterizing continuity: the topological approach	133
4.2.1.1	Let's get real: the initial model case	133
4.2.1.2	Abstracting topology out of the real example	135
4.2.1.3	Application to the sensorimotor experience	138
4.2.1.4	Discrete spaces: so close yet so far...	141
4.2.2	A naive process for building quantified continuity	143
4.2.2.1	Ascribing experiential meaning to metric values	143
4.2.2.2	Building the sensory metric	144
4.2.3	Some properties of the sensory metric	147
4.2.3.1	Fitting spatial and sensory dynamics in the exploration	147
4.2.3.2	Using the metric in regularity assessment	150
4.3	Experiment: building the continuity map under perturbation	151
4.3.1	Description of the experimental setup	151
4.3.2	Computing the sensory distance	154
4.3.3	Spatially representing the sensory metric	156
4.3.3.1	The use of dimensionality reduction	156
4.3.3.2	ISOMAP	158

4.3.4	Results	161
4.3.4.1	The sensory metric	161
4.3.4.2	Crawling in my film: Effects of amplitude in the agent's movements on the discovered regularities	165
4.3.4.3	Introducing noise in the sensorimotor in- teraction	169
4.4	Conclusion	174

4.1 Speaking from experience: the need for continuity

One crucial property of many (psychological) perceptual processes is that they generally seem continuous (Dainton, 2018); in point of fact, this intuition is strong enough that it is the converse situations where it visibly is *not* that earn explicit mentions, such as that of Categorical Perception (Fugate, 2013; Toscano et al., 2010). However such continuity does not trivially follow from our knowledge of how perceptual processes are materially –e.g. neurally– mediated (Herwig, 2015; Stroud, 1956; VanRullen and Koch, 2003). In the instance of visual perception, for example, it is known that the eye only acquires very partial snapshots of visual information due to the sparse layout of (discrete) photoreceptors on its retina as well as the typical trajectories of ocular saccades.

Nevertheless, the continuity of perception sensorimotor agents subjectively experience is undeniably useful, allowing for formulation and exploitation of several powerful ideas. One such idea, for instance, is that of inter- and extrapolation. Consider an agent which has repeatedly rung a gong, learning the laws of their interaction so that it can recognize the sound it emits when struck. Then we should hope for this agent to eventually learn that the intensity of the sound created evolves *continuously* following a hit, *slowly* decreasing until it fades back into silence. Furthermore, it could also learn that (and how) the sound changes *ever so slightly* as it initiates the interaction *a bit* differently, going as far as learning that “gongs of *similar* shapes, size and material emit *similar* sounds”. Importantly, all of these hypothetical bits of knowledge appeal to a shared idea that the agent have the capabilities to distinguish “similar” things, be it external objects (the gong), sensory attributes (intensity of the sound), or perceptual items (the sound in itself). These capabilities may in turn provide grounds for the emergence of its felt continuity of perception. As for its functional roles, the previous informal example shows that it can for instance be associated with increased scalability and robustness. Indeed, the aforementioned agent could determine that two occurrences of (it hearing) the sound of any one gong are indeed results of a “same” interaction, even though the two sounds heard are not perfectly equal¹; instead, it could make this judgement by way of the sounds being similar enough. Besides it could try and guess the result of a *slightly* different interaction, be it one with a *slightly* smaller gong or by hitting it *slightly*

more slowly, at the very least predicting that the resulting experience will be similar to that it has previously encountered.

Such properties are usually leveraged in robotic settings *via* the mathematical continuity of mappings $\mathbb{R}^n \rightarrow \mathbb{R}^m$ since the data available to the agent is then represented numerically. It has proved a powerful tool for bootstrapping (Censi, 2012) and in SMCT contexts (Alban Laflaquière, 2017, 2020; V. Marcel, Garcia, et al., 2015), notably granting robustness and scalability to the results. On the contrary, the approach we outlined in 2.3 (and particularly in 2.3.2.3) does not *a priori* include continuity of perception; in particular, it does not causes the agent to have any such preexisting concept for its sensorimotor experience. Instead, all of its initially available assessments are entirely categorical: either some sensory value s_c is *equal* to some other s'_c or it is not. Although it is a severe limitation, in that it directly locks one out of leveraging the previous desirable properties, this is entirely by design. Indeed, we hypothesize that whatever subjective sense of perceptual continuity should actually emerge as a contingency out of the sensorimotor flow of the agent. This would in particular, as opposed to the previous continuity of mappings, have the felt continuity be entirely that of perceptual processes regardless of the numeric (or neural) codes mediating them (J K O'Regan and Noë, 2001)².

In the following part, we therefore examine how internal notions of continuity in sensorimotor experiences can appear in the formalism of 2.3. We start by investigating the insight granted by the formal notion of continuity developed in the mathematical branch of topology, showing how it effectively hinges on a dual notion of similarity. Following the failure of pure topology to fully account for the properties of the subjective continuity, we particularize to the particular subcase of metric topology. In particular, we find it allows to formulate a way for the agent to internally form a spatial structure of its sensory readings that make its typical sensorimotor experience formally continuous. We then adapt the resulting procedure to an experimental setup similar to that of 3.3.1, assess which result topology it gives to color spectra. Importantly we find that the metric topology derived from the exigence of *formal* continuity coincides with that of our intuition, which itself follows from (our) *subjective* continuity.

¹which is all but certain to happen in realistic settings, especially for more chaotic interactions.

²(from the citation:) “there is no *a priori* reason why similar neural processes should generate similar percepts.”

4.2 Finding smoothness in statistical regularities

4.2.1 Characterizing continuity: the topological approach


The introduction highlighted that we are in fact investigating two distinct, although hopefully related, notions of continuity. The first is that which we intuitively know to characterize many of our perceptual experiences, which we previously called *subjective* continuity. The other then is a mathematical property we conversely referred to as *formal* continuity. To make the development more fluid, in the following “continuity” (or “continuous”) is also used to denote the fuzzy notion of subjective continuity, while the unquoted and unqualified versions are restricted to the formal one.

4.2.1.1 Let’s get real: the initial model case

The modern examination of continuity and related problems is the subject of *topology* (Morris, 2020): it is a field of mathematics which is precisely devoted to the study of what it means for something to be continuous, as well as the corresponding notion of “being continuously equivalent”³. Because of this explicit objective it has previously been used as the model framework for geometric ideas in several SMCT works (Laflaquière et al., 2015; V. Marcel, Argentieri, and Gas, 2017; Philipona, J K O’Regan, and Nadal, 2003), in particular those that attempted internally establishing properties of external space. It can fruitfully be found to emerge as a formal generalization of the classical case found in real numbers, illustrated in Fig. 4.1. In this case, continuity is a property of mappings $f : \mathbb{R} \rightarrow \mathbb{R}$ that obey the particular property depicted in Fig. 4.1a. This property essentially asks of f that it send no two “similar” points $x_1, x_2 \in \mathbb{R}$ to images that would in turn be “too dissimilar”. More precisely it asks that for any output $y_0 = f(x_0)$ of f , taking any one neighborhood interval N_{y_0} —however small—of y one may be able to find a corresponding neighborhood N_{x_0} of x_0 such that all x in this last interval can only end as removed from x_0 as N_{y_0} allows. Using the characterization of such intervals using the familiar modulus operator $|\cdot|$ over $\mathbb{R}$ it yields the classical definition of continuity

$$\forall \epsilon > 0, \exists \eta > 0 \text{ such that } \forall x \in \mathbb{R}, |x - x_0| \leq \eta \Rightarrow |f(x) - f(x_0)| \leq \epsilon. \quad (4.1)$$

³(see 2.2.2.1)

FIGURE 4.1: Illustration of metric continuity: the case of maps $\mathbb{R} \rightarrow \mathbb{R}$.

Consider the function graph displayed in Fig. 4.1a, where arbitrary x_0 and y_0 have been selected. The horizontal orange strip denotes an (also arbitrary) neighborhood interval N_{y_0} of the image $y_0 = f(x_0)$, which corresponds as a subset of real numbers to the interval it intersects the y axis by. It is once this interval is given that a corresponding N_{x_0} neighborhood of the input x_0 can be determined. The assertion of Eq. (4.1) asks that each point of the graph of f falling in the (blue) strip corresponding to N_{x_0} should *also* be in the (orange) strip of N_{y_0} , by which one can assess that this property actually holds in the configuration depicted. Importantly one should also visualise that it being true does not depend on the initial N_{y_0} orange strip, and that instead whatever its choice (in particular however small it is) one could always have

found a corresponding interval N_{x_0} around x_0 . It is precisely this last observation that characterizes the *formal* continuity (at x_0) of f .

Fig. 4.1b then depicts the opposite case. It displays a function that is intuitively discontinuous at $x_0 \in \mathbb{R}$, corresponding to the jump in its curve at this input value. As before, we select an arbitrary neighborhood interval of y_0 by the orange strip N_{y_0} . However, observe now the part of the other neighborhood N_{x_0} colored in red: none of the graph of f in this region can simultaneously lie in N_{y_0} . One should once again imagine how changing the width η of the vertical N_{x_0} strip, provided it stays positive, does not alter this outcome. Instead, however small one may choose η will always result in the corresponding N_{x_0} containing a part of the graph of f in the red subregion. Accordingly, the depicted function is not formally continuous at x_0 .

4.2.1.2 Abstracting topology out of the real example

These examples help cement the intuition that some meaning of what we perceive as continuity, at least that pertaining to $\mathbb{R} \rightarrow \mathbb{R}$ maps, is indeed formally captured by the above property. It turns out that the underlying idea can be made substantially more general. A candidate first step resides in observing that what the two $|\cdot - \cdot| \leq \delta$ inequalities in Eq. (4.1) denote is that the corresponding points be “at most δ apart”. Accordingly, the definition can easily be retrofitted to conform to any metric $d_{\mathbb{R}}$ between points; it can even further accommodate distinct metrics d_x and d_y for the two inequalities⁴. Continuity can thus readily be thought of as a property of maps of *metric spaces*: whenever one can think of distances between *things*, so does it make sense to think of continuity of *transformations* thereof. The intuition developed in Fig. 4.1 can however be pushed even further. Indeed, the core of the argument can be found in the neighborhoods we denoted N_{y_0} and N_{x_0} . Toying with the definitions, one may observe that the notion of continuity obtained can alternatively be characterized as f “inversely preserving” such subsets. Although we won’t detail the specific technicalities⁵, the resulting property can be written

$$f : \mathbb{R} \rightarrow \mathbb{R} \text{ is continuous if and only if } \forall O \subset \mathbb{R} \text{ open, } f^{-1}(O) \subset \mathbb{R} \text{ is open.} \quad (4.2)$$

⁴with the resulting assertion transparently becoming $\forall \epsilon > 0, \exists \eta > 0 \mid \forall x \in \mathbb{R}, d_x(x, x_0) \leq \eta \Rightarrow d_y(f(x), f(x_0)) \leq \epsilon$.

⁵see any topology textbook, such as (again) Morris, 2020

From there one can follow the same reasoning as that made for metric spaces: as long as one can find the “openness” properties for certain subsets of a set X , then it immediately entails a corresponding notion of continuity for maps $X \rightarrow X$ (and just as easily for maps $X \rightarrow Y$). The sets X where such a notion of “open sets” $\mathcal{O}_X$ is defined are thus made *topological* spaces, and the collection $\mathcal{O}_X = \{O_X \subset X \text{ open}\}$ of all their open sets is their topology. Correspondingly, a map between topological spaces X and Y is continuous if and only if it is, in some way encoded by the above property, compliant with their respective topologies $\mathcal{O}_Y$ and $\mathcal{O}_X$. We have mentioned before in 2.2.2.1 how these different contexts related to one another: the general metric case is (strictly) more general than the model $\mathbb{R}$ one, where the general topological case is in turn (strictly) more general than the metric one. In particular, extensive study of the topological account of continuity has since repeatedly reinforced the intuition that it describes all deformations that neither break nor tear spaces which they are applied to; conversely it has also shown that topological spaces are after all a natural framework for such a study, despite their abstract definition.

Making continuity a guideline Although the jump from continuity for $\mathbb{R} \rightarrow \mathbb{R}$ maps to that of topological spaces may seem unnecessarily abstract, they in fact do offer some practical insight. Indeed, observe that it technically cannot be asserted whether a map $f : X \rightarrow Y$ of topological spaces is continuous in itself. Instead, one also needs to know of the topologies $\mathcal{O}_X$ and $\mathcal{O}_Y$ involved in the statement. Consider for instance the identity mapping on $X = [0; 2\pi[$

$$\begin{aligned} f : [0; 2\pi[&\rightarrow \mathbb{R} \\ x &\mapsto x \end{aligned} \tag{4.3}$$

One can see it as the restriction of the real identity map $\text{id}_{\mathbb{R}} : x \in \mathbb{R} \mapsto x \in \mathbb{R}$ to the closed interval $[0; 2\pi[\subset \mathbb{R}$, making it trivially continuous. However $[0; 2\pi[$ is also classically used as denoting the circle, in which case $x \in X$ denotes the principal measure of an angle in radians. The distinction between the two cases is found at the level of their topologies since that of the circle essentially encodes that the point of x gets vanishingly close to that of 0 as x tends to 2π ⁶. Importantly, the choice of the circular topology would then make f discontinuous. Indeed, consider its behavior around the point

⁶The informal arguments generally appeal to metric intuition since abstract topology is far less amenable to assertions in natural languages. This idea can however be formulated

of the circle labelled by $x = 0$ (which one can also think as $x = 2\pi$): for $x_\epsilon^+ = x + \epsilon$ with $\epsilon > 0$ sufficiently small, we have $f(x_\epsilon^+) \approx 0$. On the other hand, for $x_\epsilon^- = 2\pi - \epsilon$ with the same restrictions on ϵ we have $f(x_\epsilon^-) \approx 2\pi$. But as ϵ vanishes to 0, x_ϵ^- and x_ϵ^+ all converge to the same point in the circle. Consequently f is discontinuous at x for the circular topology, thus showing how the *formal* property of continuity of maps depends not only on the sets on which they operate but also on the topology we assume for them. As a result of this fact one should precise which topologies they consider when they make statements about continuity, for example turning them into the sort “ $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a continuous map”⁷.

Because of the dependency of continuity on the underlying topology, we can then attend to the converse problem. Let us consider an arbitrary function $f : X \rightarrow Y$, and further assume that a particular topology $\mathcal{O}_X$ is given on X . Then for any possible topology $\mathcal{O}_Y$ on Y one can wonder whether f is continuous relative to $\mathcal{O}_X$ and $\mathcal{O}_Y$, the answer to that question varying with $\mathcal{O}_Y$. This formally defines a collection

$$\mathcal{O}_Y^f = \{\mathcal{O}_Y \subset \mathfrak{P}(Y) \mid f \text{ is a continuous map w.r.t. } \mathcal{O}_X \text{ and } \mathcal{O}_Y.\}$$

of all topologies on Y that make f continuous. Observe that this runs opposite to our usual intuition: what this idea formalizes is us *assuming* that a given transformation of “things” is continuous and *inferring* what structure on “things” follows from it. Fortunately a classical result gives a comprehensive answer to this inquiry in the context of topological spaces as per

Proposition 12. *Let $(X, \mathcal{O}_X)$ be a topological space, Y a set and $F = \{f_i\}$ a collection of maps $X \rightarrow Y$.*

There exists a topology $\mathcal{O}_Y^F$ on Y such that

- *Every map $f_i : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y^F)$ of F is continuous, and*
- *Any topology $\mathcal{O}_Y$ that makes all f_i continuous is coarser⁸ than $\mathcal{O}_Y^F$.*

entirely in terms of neighborhoods and open sets: it corresponds to every neighborhood of 0 (however small) containing every $2\pi - \epsilon$ for ϵ sufficiently small.

⁷In virtually all mathematical practice properties are not actually written as such. Instead one generally refers to a topological space $(X, \mathcal{O}_X)$ merely as X ; the topology, which is either “natural” or explicitly discussed once, is correspondingly kept implicit. This is however entirely a matter of reading convenience, and does not change the fact that topological (and continuity) assertions examine properties of *topologies*. In particular, “[0; 2π[” as an interval and as a circle are two distinct topological entities in the sense discussed in 2.2.2.1.




FIGURE 4.2: Illustration of the final topology in the case of the toy model.

$\mathcal{O}_Y^F$ is called the final topology of the f_i . Additionally, it can be constructively given as

$$\mathcal{O}_Y^F = \left\{ Y' \subset Y \text{ such that for all } f_i \in F, f_i^{-1}(Y') \in \mathcal{O}_X \right\}$$

What the second property of Proposition 12 establishes is that $\mathcal{O}_Y^F$ is in some way the most general topology that makes all the f_i continuous. Indeed from the characterization of general continuity we outlined in Eq. (4.2) as (inverse) maps of open subsets, one can see that eliminating open sets from the topology on Y only makes the condition easier to fulfill. Therefore, changing topologies to coarser ones always has continuous functions remain so; conversely, this makes the finer such topology the limit case. Thus, the result shows a precise correspondence between certain continuities –that of the f_i maps– and the spatial properties, as captured by the topologies, of that Y which varies continuously.

4.2.1.3 Application to the sensorimotor experience

Figure 4.2 illustrates how the previous discussion applies to our model robot. Consider the interaction depicted in Fig. 4.2a, in which the robot is sampling the color environment as it moves right. Two distinct environment states ϵ_1 and ϵ_2 are pictured. Observe that the layout of colors in the first is somewhat typical of what we may expect: blues only fade into purple and cyan tones,

⁸Given two topologies $\mathcal{O}_Y = \{O \subset Y\}$ and $\mathcal{O}'_Y = \{O' \subset Y\}$ on a same set Y , $\mathcal{O}_Y$ is said to be coarser than $\mathcal{O}'_Y$ if and only if $\mathcal{O}_Y \subset \mathcal{O}'_Y$. This equivalently states that every set $O \subset Y$ that is open for $\mathcal{O}_Y$ must also be open for $\mathcal{O}'_Y$.

orange to red and yellow ones, and so on. Consequently, one would expect the agent to find that the succession of colors its sensors report is “continuous” provided it had such a notion. Observe also that these successions follow that of the usual representation of the hue spectrum, displayed top in Fig. 4.2b. If one assumes that colors are themselves topologically related thus⁹, then the sensory sequence of the agent in ε_1 is *formally* continuous. Consider now an agent which assumes this usual topology of hues for its sensory values in $\mathcal{S}_c$, and which experience now occurs in environment state ε_2 . This environment state especially contains a particular atypical succession of colors, highlighted by an arrow on Fig. 4.2a, between green and purple shades. These two particular shades are pictured in the hue spectrum of Fig. 4.2b as corresponding dotted lines: one can clearly see that they indeed correspond to two distant values of hue in the topological sense. Consequently, should the agent rigidly adhere to the usual hue topology during its experience, it would then conclude that it observed a discontinuity in its sensory sequence.

The converse view we entertain instead is that this experience be continuous so that the assumed topology of hues must be made to fit this particular sequence. As a consequence of the sequence it observes in ε_2 , the agent must thus conclude that these green and purple tones are close (as colors) to one another since they occur at close points of its internal experience. A visualisation of the corresponding topological process is proposed in Fig. 4.2b with the three layouts of the hue spectrum. Indeed the folding of the line into the bottommost “ $\times$ ”-like configuration has the green and purple tones become neighbours so that there is now a continuous path from one another which does not need to cross intermediate colors. Note that this entails further new relations between other colors: for example, there is now a continuous path between orange and indigo which does not cross red or blue colors. Importantly, this folding introduces the *minimal* amount of modifications to the previous topology needed to make the experience in ε_2 continuous. It could have instead collapsed every color thus, by which they would all be mutual neighbours; the resulting topology would trivially make all color sequences continuous, *a fortiori* that in ε_2 . This corresponds to the resulting topology in Fig. 4.2b being the finest such topology, by which Proposition 12 determines it is the final topology of the experience of the agent.

⁹Note that the two red boundaries are however implicitly thought of as glued, corresponding to a “continuous” path of color orange $\rightarrow$ red $\rightarrow$ purple: it could actually be pictured as a loop to account for this property.

This argument can be made more general in the formalism of 2.3. Because of the viewpoint about what the *a priori* internal knowledge of the agent developed in 2.3.2.3, the experience it will assume is continuous is that given by the readings $\mathbf{s}_{\mathcal{E}}$ its sensors report. More precisely, this continuity is that of its experience unfolding with the “time” $\mathcal{T}$ during which it occurs. Accordingly this corresponds to a preexisting topology $\mathcal{O}_{\mathcal{T}}$ of time, wherewith the agent may distinguish whether a sequence of timesteps denotes a continuous interval of time or not. Because of the previous assumption that $\mathcal{T}$ be a linearly ordered space such a natural topology is its *order topology*, which is determined by the basic intervals $I_{t_0, t_1} = \{t \in \mathcal{T} \text{ such that } t_0 \leq_{\mathcal{T}} t \leq_{\mathcal{T}} t_1\}$. Therefore what we specifically ask for here is that the agent *internally* have this knowledge.

The definitions proposed back in 2.3 described snapshots of the sensorimotor interaction. Nevertheless they can easily be adapted in order to better address how its state changes with time: for example, all possible “temporal” sequences of environment states can be written as functions

$$\gamma_{\epsilon} : t \in \mathcal{T} \mapsto \epsilon_t = \gamma_{\epsilon}(t) \in \mathcal{E}\mathcal{X}$$

where ϵ_t denotes the particular environment state at timestep t of the sensorimotor interaction. The collection of all such maps $\mathcal{T} \rightarrow \mathcal{E}\mathcal{X}$ may however include certain “impossible” sequences of environmental states, for instances because of what we would externally formulate as constraints from the external laws of physics. As a result, we denote a particular subset

$$\mathcal{E}\mathcal{X}_{\mathcal{T}} \subset \mathcal{F}(\mathcal{T}, \mathcal{E}\mathcal{X})$$

of all such trajectories of environment spaces which are effectively realizable. We similarly consider the set $\mathcal{B}_{\mathcal{T}} \subset \mathcal{F}(\mathcal{T}, \mathcal{B})$ of all trajectories $\gamma_{\mathbf{b}}$ of absolute configuration $\mathbf{b}$ the agent can effectively achieve during the sensorimotor interaction. But then it follows that all sequences of sensory outputs it may experience are realized under the form

$$\begin{aligned} \gamma_{\mathbf{s}} = \gamma_{\mathbf{b}, \epsilon} : \mathcal{T} &\rightarrow \mathcal{S}_{\mathcal{E}} \\ t &\mapsto \psi_{\mathcal{E}}(\gamma_{\mathbf{b}}(t), \gamma_{\epsilon}(t)) \end{aligned} \tag{4.4}$$

for $\gamma_{\mathbf{b}} \in \mathcal{B}_{\mathcal{T}}$ and $\gamma_{\epsilon} \in \mathcal{E}\mathcal{X}_{\mathcal{T}}$. We now accordingly denote $\mathcal{S}_{\mathcal{T}}$ the set of all these sensory trajectories that can occur during the sensorimotor interaction.

By construction, these $\gamma_s \in \mathcal{S}_{\mathcal{T}}$ trajectories precisely represent which sensory experiences the agent may encounter while interacting with its environment. Following the illustration in Fig. 4.2, it is especially them it should eventually find to be “continuous”. But from the way we formally defined them and the previous discussion about formal topology, we can then conclude that them being continuous corresponds to there being a particular topology on $\mathcal{S}_{\mathcal{C}}$. More precisely, provided the “temporal” topology $\mathcal{O}_{\mathcal{T}}$, all sensory experiences $\gamma_s \in \mathcal{S}_{\mathcal{T}}$ can be continuous if and only if $\mathcal{S}_{\mathcal{C}}$ has a topology *coarser* than that final for $\mathcal{S}_{\mathcal{T}}$, that is $\mathcal{O}_{\mathcal{S}_{\mathcal{C}}}^{\mathcal{S}_{\mathcal{T}}}$. Moreover, if the topology constraints imposed on sensory values $s \in \mathcal{S}_{\mathcal{C}}$ only extend to that making the sensory experiences continuous, then the corresponding topological structure is precisely that of $\mathcal{O}_{\mathcal{S}_{\mathcal{C}}}^{\mathcal{S}_{\mathcal{T}}}$. This shows that searching to establish the continuity of the sensory experiences the agent may encounter during its explorations is equivalent to finding a suitable topology for its sensory values, that is one that makes the aforementioned experiences continuous. Better still, Proposition 12 provides it with a constructive way¹⁰ to determine this sensory topology. Indeed, if it has “topological” knowledge of time so that it can distinguish the open sets $O_{\mathcal{T}}$ of $\mathcal{T}$ then it can internally build

$$\mathcal{O}_{\mathcal{S}_{\mathcal{C}}}^{\mathcal{S}_{\mathcal{T}}} = \left\{ O_{\mathcal{S}} \subset \mathcal{S}_{\mathcal{C}} \text{ such that } \forall \gamma_s \in \mathcal{S}_{\mathcal{T}}, O_{\mathcal{T}} = \gamma_s^{-1}(O_{\mathcal{S}}) \text{ is open} \right\}. \quad (4.5)$$

4.2.1.4 Discrete spaces: so close yet so far...

Robotic settings often assume that computation occurs as a discrete time system; in fact, numerical simulations go as far as having their events occur in discrete time. This topologically corresponds to *all* subsets $O_{\mathcal{T}}$ of time $\mathcal{T}$ being part of its topology $\mathcal{O}_{\mathcal{T}}$, that is $\mathcal{O}_{\mathcal{T}} = \mathfrak{P}(\mathcal{T})$. Indeed, consider the prototypic discrete space $\mathbb{N}$ with its usual topology: drawing a sufficiently small ball around it (e.g. one of radius $r \leq \frac{1}{3}$), one can see that any point $n \in \mathbb{N}$ is in and of itself a neighborhood of itself. But then since all subsets of $\mathcal{T}$ are open, the restricting condition that appears in the brackets of Eq. (4.5) is trivially verified regardless of $O_{\mathcal{S}}$. Therefore in discrete time one necessarily has

$$\mathcal{O}_{\mathcal{S}_{\mathcal{C}}}^{\mathcal{S}_{\mathcal{T}}} = \{O_{\mathcal{S}} \subset \mathcal{S}_{\mathcal{C}}\} = \mathfrak{P}(\mathcal{S}_{\mathcal{C}}),$$

so that the resulting sensory topology makes $\mathcal{S}_{\mathcal{C}}$ discrete as well.

As a matter of fact, this result is in line with the intuition formulated as we examined Fig. 4.2. Indeed, if the agent “felt” its experience occurring

¹⁰(in as much as general topology can be)




FIGURE 4.3: Illustration of the limited capabilities of topology applied to discrete spaces.

as jumps in times such that no instant follows continuously from the previous one, so need it not introduce new continuities in its sensations to make their succession continuous. It also mirrors a limitation of the formal framework of topology relative to our considerations, which we illustrate in Figure 4.3. The leftmost column depicts two distinct curves of maps $\mathbb{R} \rightarrow \mathbb{R}$, which are continuous relative to its usual topology. Importantly any one of these curves can be obtained from the other *via* a continuous (and even homeomorphic, that is topologically equivalent) reparameterization of their argument $x \in \mathbb{R}$. In particular they are therefore topologically equivalent themselves, by which *any* topological assertion is true of one if and only if it is true of the other. And indeed though one may say that the bottom curve is visibly, “quantitatively”, steeper in its rightmost portion it is also clear that they stand “qualitatively” alike as for their continuity. The rightmost column, then, depicts some discrete samples $(x, f(x))$ taken along these curves. These samples are chosen at fixed intervals of the argument x , corresponding to a fixed sampling frequency. The resulting collections of samples are finite and discrete. What the previous result tells us is that the two are themselves topologically equivalent: thus, one may not hope to tell them from one another using solely topological means. But then again it is clearly apparent

that the two visualisations of the rightmost column *seem* different. Indeed, let us forget the “continuous” curves they originate from and focus only on the discrete samples. Certainly, one would say that the topmost graph displays a continuous process. On the other hand, they would at least entertain the thought that there be a discontinuity –as denoted by the red dotted line– in the bottom one. This is therefore a perceivable distinction that cannot be accounted for by the tools of topology. To reiterate, should one want the agent to make this kind of judgement about discrete samples by way of its *subjective* sense of continuity, then this sense cannot be entirely grounded in topology; in particular, it cannot be reduced to that of *formal* continuity.

4.2.2 A naive process for building quantified continuity

The previous example shows that subjective continuity involves a notion of “closeness” between the items it orders, such as the samples of Fig. 4.3b. General topology is, in some sense, *too* general for such a concept to be formulated: for instance, we saw that topological equivalences do not preserve it. To have it become available, one could instead turn to other related contexts such as that of uniform geometry (James, 1987) or coarse geometry (Roe, 2006). In the following we will assume that the experiment occurs in discrete time and further particularize to the case of *metric geometry*, which relation to topology we described earlier in 4.2.1.2. Moreover the intuition drawn from Fig. 4.3 suggests that, in the discrete case, the subjective continuity underlying a collection of samples can only be assessed relative to a quantitative closeness of the timesteps at which they were sampled. Accordingly we will additionally assume that $\mathcal{T}$ is endowed with a metric structure itself, so that there exist durations $d_{\mathcal{T}}(t, t')$ between any two timesteps t, t' . In fact, since $\mathcal{O}_{\mathcal{T}}$ follows from the linear order $\leq_{\mathcal{T}}$ on timesteps we are assured such a metric $d_{\mathcal{T}}$ of durations can be found (Lutzer, 1969). Consequently the real content of this assumption is that the agent knows of this metric, that is provided two timesteps $t, t' \in \mathcal{T}$ it can compute the duration between them $d_{\mathcal{T}}$ as well as compare any two durations for equality.

4.2.2.1 Ascribing experiential meaning to metric values

The topological argument we pursued, e.g. in Fig. 4.2, was actually metric. Following it informally and tuning a *sensory metric* instead of a sensory topology, the agent ought to set to null any distance between sensory values that

“immediately” follow one another. In fact, this is precisely the surgery procedure depicted in Fig. 4.2b for the green and purple hues. However, this behavior is not actually advisable for two reasons. First, we ought to expect the agent to encounter “atypical” patterns in the color sequences it experiences, for instance because they correspond to the borders of various objects in its visual scene. It should consequently find that these patterns are color discontinuities, and crucially have them remain such. Second, in a discrete setting it will successively experience colors that follow one another *immediately* but not *continuously*. Consider again the samples seen in Fig. 4.3b: although certain pairs of samples occur consecutively, the quantified length of the temporal intervals between their samplings is central to how we perceive them. In particular the triangle inequality tells us that if we set all these pairwise distances to zero, then the distance between any two samples should necessarily be null as well.

We therefore need to find another guideline to design the sensory metric. The first point in the rebuttal above points at the metric instead quantifying the “typicality” of a sensorimotor experience. Thus, the previous topological intuition becomes that sensory values which are reliably experienced in close succession are taken to be close to one another; further still, the closeness between two sensory values s, s' should accordingly increase (i.e. the distance between them should decrease) as their successive experiences grow more frequent. As a result, the design proposed is inherently statistical. Moreover, since it only relies on assessing the frequencies at which certain sensory events occur, this procedure may be amenable to be internally computable as per 2.3.2.3. Finally, should the agent determine a sensory metric possessing such properties then by the converse argument so would it necessarily double as an internal inferential measure of typicality in its experience.

4.2.2.2 Building the sensory metric

In accordance with the previous discussion, we now deal with statistical properties of the sensorimotor flow. In the following we will assume that the law of the sensorimotor experiences the agent can observe is *time homogeneous* in some sense: let us consider any sensory experience γ_s , and let us denote for any time isometry $\tau_{\mathcal{T}}$

$$\tau_{\mathcal{T}}\gamma_s : t \in \mathcal{T} \mapsto \gamma_s(\tau_{\mathcal{T}}t) \in \mathcal{S}_{\mathcal{C}}$$

the sensory trajectory mimicking γ_s up to a shift of $\tau_{\mathcal{T}}$ in time. This property – which we previously examined in 2.3.1.1 for ambient space $\mathcal{X}$ – is summarized in $\forall t \in \mathcal{T}, (\tau_{\mathcal{T}}\gamma_s)(t) = \gamma_s(\tau_{\mathcal{T}}t)$, which holds by definition. Note that the condition that $\tau_{\mathcal{T}}$ be an isometry specifies that it does not alter durations, since for any two timesteps t and t' we have $d_{\mathcal{T}}(\tau_{\mathcal{T}}t', \tau_{\mathcal{T}}t) = d_{\mathcal{T}}(t', t)$. Consider for instance the familiar case where such timesteps are denoted by integers $p, q \in \mathbb{Z}$: the corresponding isometries are the timeshifts by k units of time $\Delta k : n \mapsto n + k$. But accordingly, the duration between any two timesteps p and q does not change if both of them are shifted by such a Δk :

$$\begin{aligned} d_{\mathbb{Z}}((\Delta k)p, (\Delta k)q) &= |(q + k) - (p + k)| \\ &= |q - p + k - k| = |q - p| = d_{\mathbb{Z}}(p, q). \end{aligned}$$

Then the time homogeneous assumption specifies that for any measurable statistical event $X \subset \mathcal{S}_{\mathcal{T}}$ and every time isometry $\tau_{\mathcal{T}}$ we have

$$\mathbb{P}(\gamma_s \in X) = \mathbb{P}(\tau_{\mathcal{T}}\gamma_s \in X). \quad (4.6)$$

This property denotes that no statistical measurement the agent can empirically obtain from its sensorimotor experience may depend on the absolute value of the timestep indexing the interaction. In particular it should be a natural consequence of the particular choice of timestep being an entirely external convention, implementing a sort of independence of choice of reference. To illustrate this point, let us consider an agent which is experiencing a given sensory trajectory γ_s corresponding to a choice of reference timestep $t_0 \in \mathcal{T}$. Then, any observation it may make of its sensory experience at timestep t would equivalently be found at any timestep $t' = \tau_{\mathcal{T}}t$ in the sensory trajectory $\tau_{\mathcal{T}}\gamma_s$. Contrapositively, following the principle of indifference (E. T. Jaynes and Bretthorst, 2003; E. T. Jaynes, 1973) if Eq. (4.6) did not hold it would provide a statistical internally observable signature of which timestep is the particular reference t_0 . Therefore, the assumption of Eq. (4.6) corresponds to there being no such preference in the choice of reference.

Now to effectively define a sensory metric, let us again consider $\mathcal{S}_{\mathcal{T}}$ the set of all sensory trajectories the agent can experience. For any two sensory values $\mathbf{s}$ and $\mathbf{s}'$ let us denote

$$P_{\mathbf{s}, \mathbf{s}'} = \mathbb{P}(\gamma_s(t+1) = \mathbf{s}' | \gamma_s(t) = \mathbf{s}) \quad (4.7)$$

the likelihood over all possible experiences that $\mathbf{s}'$ immediately follows $\mathbf{s}$ in

the sensorimotor flow of the agent. Note that the t variable in Eq. (4.7) need not be specified. Technically, the probabilistic events

$$E_s(t) = \{\gamma_s(t) = \mathbf{s}\} \subset \mathcal{S}_{\mathcal{T}}$$

generally varies with t since the sensorimotor experiences γ_s that verify the corresponding conditions of $E_s(t)$ and $E_s(t')$ may differ when $t \neq t'$. However any two such events are related only by a time shift, so that by the time homogeneity condition of Eq. (4.6) their associated probabilities must be equal. As a result, $P_{s,s'} = \frac{\mathbb{P}(E_{s'}(t+1) \cap E_s(t))}{\mathbb{P}(E_s(t))}$ can be defined as in Eq. (4.7) by the constant value of the right hand side as t varies.

To proceed with the path we set out on, we must obtain a suitable notion of closeness from the succession likelihoods $P_{s,s'}$. Accordingly, we define

$$\delta_f(\mathbf{s}, \mathbf{s}') = f(P_{s,s'}) \quad (4.8)$$

for any two sensory values $\mathbf{s}, \mathbf{s}'$ and nonincreasing $f : [0; 1] \rightarrow \mathbb{R}_+$. f being nonincreasing ensures that $\delta_f(\mathbf{s}, \mathbf{s}')$ grows smaller as $P_{s,s'}$ increases, fulfilling the aim discussed in 4.2.2.1 that it grade the typicality of sensory transitions. This thus defines a map $\delta_f : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_+$. However δ_f need not be a metric since it may fail to obey the corresponding axioms. To solve this issue, let us consider the graph $G_{\mathcal{S}}^f$ which vertices are all sensory values $\mathbf{s}$ and which has all edges $\mathbf{s} \rightarrow \mathbf{s}'$ with corresponding weights $\delta_f(\mathbf{s}, \mathbf{s}')$. Then it defines a function on $\mathcal{S} \times \mathcal{S}$ by taking the length of the shortest weighted (directed) path between any two vertices $\mathbf{s}, \mathbf{s}'$. Alternatively, this function can be formally defined as

$$d_f(\mathbf{s}, \mathbf{s}') = \inf \left\{ \sum_{k=0}^{n-1} \delta_f(\mathbf{s}_k, \mathbf{s}_{k+1}), n \in \mathbb{N}, \mathbf{s}_0 = \mathbf{s} \text{ and } \mathbf{s}_n = \mathbf{s}' \right\}. \quad (4.9)$$

We can already assert several metric axioms of d_f from its characterization as the length of shortest paths in $G_{\mathcal{S}}^f$. Indeed, it must necessarily verify the triangle inequality, that is

$$\forall \mathbf{s}, \mathbf{s}', \mathbf{s}'' \quad d_f(\mathbf{s}, \mathbf{s}'') \leq d_f(\mathbf{s}, \mathbf{s}') + d_f(\mathbf{s}', \mathbf{s}'').$$

Moreover, because for $\mathbf{s}' = \mathbf{s}$ one can take $n = 0$ and $\mathbf{s}_0 = \mathbf{s}$ in Eq. (4.9), it follows that $d_f(\mathbf{s}, \mathbf{s}) = 0$ for any sensory value $\mathbf{s}$. Therefore d need only be

symmetric for it to be a metric. In fact, because it records the shortest distances of paths of *directed* edges it is not necessarily such. However, we will also assume that any sensory transition from value $\mathbf{s}$ to value $\mathbf{s}'$ is equally as probable as that from $\mathbf{s}'$ to $\mathbf{s}$ taken over all possible experiences; in particular it can follow from Eq. (4.6) whenever we restrict our attention to reversible experiences. This condition precisely specifies that $P_{\mathbf{s},\mathbf{s}'} = P_{\mathbf{s}',\mathbf{s}}$, so that $G_{\mathcal{S}}^f$ reduces to an undirected graph¹¹. This then especially entails that d_f is symmetric, by which it must be a metric on $\mathcal{S}$.

4.2.3 Some properties of the sensory metric

Now that we described how an internal metric of sensory values could be defined to account for the subjective sense of continuity in the sensorimotor experiences, we turn to examining some of its properties the agent may leverage during its interaction.

4.2.3.1 Fitting spatial and sensory dynamics in the exploration

The intuition developed from 4.2.1.3 onwards is that subjective continuity –and the corresponding closeness of sensations– the agent *feels* in its experience derives from certain successions of experiences being more likely than others. In the example of Fig. 4.2, for instance, this corresponds to the transition of experiences “Red $\rightarrow$ Orange” being more representative of a “typical” interaction than “Red $\rightarrow$ Cyan”. But what the atypical environment state ε_2 of the figure showed is that this preference depends in particular on the environment states and the successive configurations with which the agent samples it. More precisely it hinges on the idea that typical environment states display certain typical patterns themselves¹². This implies that the environment state is somewhat locally coherent, that is the local restrictions $\varepsilon|_{\mathcal{X}'} = \gamma_\varepsilon(t)|_{\mathcal{X}'}$ may mutually condition one another. One such condition is spatial, mandating that e.g. immediately next to a red region $\mathcal{X}'$ of ambient space $\mathcal{X}$ it is more likely to be another region $\mathcal{X}''$ that is orange than cyan itself. As a result, we would generally expect the two events

$$\left\{ \gamma_\varepsilon(t)|_{\mathcal{X}'} = \varepsilon_0 \right\} \text{ and } \left\{ \gamma_\varepsilon(t)|_{\mathcal{X}''} = \varepsilon_1 \right\}$$

¹¹ Alternatively we could deem it sufficient since a nonsymmetric d_f already makes $\mathcal{S}$ into a quasimetric space (Wilson, 1931), especially since they can be studied using Finslerian geometry (Dahl, 2006) much as by Riemannian geometry do ordinary metric spaces.

¹² which, externally, we would intuitively enunciate by “environment states are (mostly) continuous” (both in time and space).

to largely depend on one another when $\mathcal{X}'$ and $\mathcal{X}''$ denote close (and small) regions of space. Furthermore, our proposal relies on an identical assumption for local *temporal* coherence. Indeed, we expect that environment state at any one localization $\mathcal{X}' \subset \mathcal{X}$ does not immediately change too randomly so that $\{\gamma_\varepsilon(t + \Delta t)_{|\mathcal{X}'} = \varepsilon_1\}$ is effectively conditioned by $\{\gamma_\varepsilon(t)_{|\mathcal{X}'} = \varepsilon_0\}$ when Δt remains sufficiently small¹³. This represents a nontrivial assumption about there being some structure in the environment states the agent interacts with, as well as the manner in which it conducts its interaction, by which such a notion of certain transitions being “more typical than” others can emerge. It is this assumption which allows the agent to draw meaningful information from comparing sensory samples it experienced at two distinct timesteps of its interaction and thus with two distinct environment states $\gamma_\varepsilon(t + 1)$ and $\gamma_\varepsilon(t)$. Thus, as we demonstrate later in 4.3 and as opposed to the case in 3.3, it allows us to have them both “act” *simultaneously*. Importantly local coherences fail when the agent interacts with a white noise environment state instead since by definition its values are all (spatially and temporally) independent. However we surmise that this assumption should be effectively verified in all but these tailored counterexamples, enough so that the resulting metric can capture some interesting information of sensations. We especially expect that it hold for most “natural”¹⁴ interactions in which an embodied agent moves continuously and interacts with a “natural” stimulus.

This coherence property however should only be local. Indeed, it is clear that for instance the color at any one point $x \in \mathcal{X}$ and timestep $t \in \mathcal{T}$ does not weigh much –if at all really– on which color is apparent five kilometers, three and a half days from there. But this entails that, for the procedure of 4.2.2.2 to succeed, it is not only the environment states that should be locally coherent but also the successive sensory samples in the sensorimotor flow of the agent. For instance, consider a particular agent as in Fig. 4.2a that moves fast enough that it crosses five kilometers between any two samples its camera provides it. As a consequence its successive sensory states report on very distant (local) states of the environment, which are not necessarily coherent; but this further entails that there is no more reason why certain

¹³Because we consider agents which interact with their environments in discrete time here, what this actually denotes is that the *duration between any two sensory samplings* of the agent must be small compared to that with which the environment changes.

¹⁴(that is, relative to our own usual sensorimotor experiences)

particular sensory transitions should emerge as more coherent in its sensorimotor flow. To reiterate, although the environment may exhibit the desired local coherence property, the agent cannot pick up on it from the way it conducts its exploration.

This observation can be accounted for using the definitions of 4.2.2.2. Let us slightly generalize the previous sensory transition probabilities by introducing, for any collection of motor trajectories $\mathcal{B}'_{\mathcal{T}} \subset \mathcal{B}_{\mathcal{T}}$,

$$P_{\mathbf{s}, \mathbf{s}'}^{\mathcal{B}'_{\mathcal{T}}} = \{ \gamma_{\mathbf{s}}(t+1) = \mathbf{s}' \mid \gamma_{\mathbf{s}}(t) = \mathbf{s} \text{ and } \gamma_{\mathbf{s}} = \gamma_{\mathbf{b}, \varepsilon} \text{ with } \gamma_{\mathbf{b}} \in \mathcal{B}'_{\mathcal{T}} \}.$$

This serves to denote the fact that the ways $\gamma_{\mathbf{b}} \in \mathcal{B}'_{\mathcal{T}}$ in which the agent explores its environment condition the sensory transitions it will find in its sensorimotor flow. In particular, it subsumes the previous $P_{\mathbf{s}, \mathbf{s}'}$ since by definition we have $P_{\mathbf{s}, \mathbf{s}'} = P_{\mathbf{s}, \mathbf{s}'}^{\mathcal{B}_{\mathcal{T}}}$. Notwithstanding, we can now restrict our attention to certain particular explorations. Following with the informal argument, let us consider $\mathcal{B}'_{\mathcal{T}}$ a collection of motor explorations $\gamma_{\mathbf{b}}$ such that receptive fields $F_c(\gamma_{\mathbf{b}}(t))$ and $F_c(\gamma_{\mathbf{b}}(t+1))$, which determine successive sensory outputs $\gamma_{\mathbf{s}}(t)$ and $\gamma_{\mathbf{s}}(t+1)$, fall far apart from one another. Consequently, in the limit case the corresponding local environment states

$$\gamma_{\varepsilon|F_c(\gamma_{\mathbf{b}}(t+1))}(t+1) \text{ and } \gamma_{\varepsilon|F_c(\gamma_{\mathbf{b}}(t))}(t)$$

are independent. But then it follows that

$$\gamma_{\mathbf{s}}(t+1) = f_c(\gamma_{\varepsilon|F_c(\gamma_{\mathbf{b}}(t+1))}(t+1))$$

and

$$\gamma_{\mathbf{s}}(t) = f_c(\gamma_{\varepsilon|F_c(\gamma_{\mathbf{b}}(t))}(t))$$

are independent themselves. As a result we have

$$\begin{aligned} P_{\mathbf{s}, \mathbf{s}'}^{\mathcal{B}'_{\mathcal{T}}} &= \{ \gamma_{\mathbf{s}}(t+1) = \mathbf{s}' \mid \gamma_{\mathbf{s}}(t) = \mathbf{s}, \gamma_{\mathbf{b}} \in \mathcal{B}'_{\mathcal{T}} \} \\ &= \{ \gamma_{\mathbf{s}}(t+1) = \mathbf{s}' \mid \gamma_{\mathbf{b}} \in \mathcal{B}'_{\mathcal{T}} \}, \end{aligned} \tag{4.10}$$

by which the transition probability $P_{\mathbf{s}, \mathbf{s}'}^{\mathcal{B}'_{\mathcal{T}}}$ does not actually depend on previous sensory output $\mathbf{s}$; instead, it simply replicates the *unconditional* probability that the agent experience the particular sensory value $\mathbf{s}'$ (provided its motor exploration $\gamma_{\mathbf{b}}$ obeys the condition of $\mathcal{B}'_{\mathcal{T}}$). To the agent, this means that the knowledge of which sensation $\mathbf{s}$ it experiences at timestep t does

not give it *any* information on which sensation $\mathbf{s}'$ it is poised to experience at $t + 1$. Importantly this shows that suitable choices of motor explorations are required for learning a valid sensory metric, as well as giving an internal observation to assess whether this condition fails *via* Eq. (4.10).

4.2.3.2 Using the metric in regularity assessment

We conclude by quickly detailing how the definitions enunciate the idea, previously outlined in 4.1 and 4.2.2.1, that the internal sensory metric allows the agent to grade the typicality of its sensory experience. To begin with, consider Eq. (4.8) by which $\delta_f(\mathbf{s}, \mathbf{s}')$ is defined in terms of the sensory transition probabilities $P_{\mathbf{s}, \mathbf{s}'}$. It is trivially equivalent to

$$\forall \mathbf{s}, \mathbf{s}' \in \mathcal{S}, \mathbb{P}(\gamma_{\mathbf{s}}(t+1) = \mathbf{s}' \mid \gamma_{\mathbf{s}}(t) = \mathbf{s}) \in f^{-1}(\delta_f(\mathbf{s}, \mathbf{s}')) \subset \mathbb{R}_+.$$

Additionally the “ $\in f^{-1}(\delta_f(\mathbf{s}, \mathbf{s}'))$ ” becomes equality “ $= f^{-1}(\delta_f(\mathbf{s}, \mathbf{s}'))$ ” as soon as f is injective¹⁵, which we will assume in the following. But because f is nonincreasing, so must f^{-1} be; this entails that the probability of any sensory transition from $\mathbf{s}$ to $\mathbf{s}'$ is a decreasing function of the sensory distance between them.

Of course this is entirely expected, since that was the point in requiring f to be nonincreasing in the first place. However, we also know from defining the resulting metric d_f from shortest paths in Eq. (4.9) that

$$\forall \mathbf{s}, \mathbf{s}' \in \mathcal{S}, d_f(\mathbf{s}, \mathbf{s}') \leq \delta_f(\mathbf{s}, \mathbf{s}').$$

It then immediately follows that

$$\forall \mathbf{s}, \mathbf{s}' \in \mathcal{S}, P_{\mathbf{s}, \mathbf{s}'} = \mathbb{P}(\gamma_{\mathbf{s}}(t+1) = \mathbf{s}' \mid \gamma_{\mathbf{s}}(t) = \mathbf{s}) \leq f^{-1}(d_f(\mathbf{s}, \mathbf{s}')). \quad (4.11)$$

Importantly, this shows that the internal knowledge of d_f informs the agent about which transitions it should expect. Indeed at any timestep $t \in \mathcal{T}$ of its interaction where it is experiencing a particular $\mathbf{s} \in \mathcal{S}$, for any sensory value $\mathbf{s}'$ it can justifiably expect to next experience it with likelihood (at most) $f^{-1}(d_f(\mathbf{s}, \mathbf{s}'))$. Conversely, it can examine *ex post* the transitions from $\gamma_{\mathbf{s}}(t)$ to $\gamma_{\mathbf{s}}(t+1)$ it has (if only just) experienced and rate their typicality as this likelihood. Going back to the example developed in Figure 4.2 we would

¹⁵(in particular this is guaranteed whenever it is *decreasing*, and the two properties are in fact equivalent here)

expect this likelihood to be low for the “light green $\rightarrow$ purple” transition, whereas it should be high for the “purple $\rightarrow$ dark blue” one.

Finally one should note that Eq. (4.11) is merely an inequality, as opposed to the corresponding equality for $f^{-1}(\delta_f)$. To the agent, this means that there may be some particular transitions from $\mathbf{s}$ to $\mathbf{s}'$ which are unlikely even while the agent has found $d_f(\mathbf{s}, \mathbf{s}')$ to be small in Eq. (4.9). Accordingly, the above reasoning will have the agent possibly *overestimate* the typicality of certain sensory transitions; in turn, it will possibly accept certain of its successive experiences as plausible where they should in fact be deemed rare. What Eq. (4.11) guarantees instead is that whenever the agent judges some sensory transition to be implausible by way of its known metric d_f , then its likelihood must be *at least* as low as that the agent estimated. In a sense, this means that the criterion suggested by Eq. (4.11) for grading the typicality of sensory experiences admits no *false negatives* as it rejects outlier events.

4.3 Experiment: building the continuity map under perturbation

The purpose of the previous theoretical discussion was to detail how metric topology could try to account for the phenomenon of subjective continuity we might expect an agent to feel, even in discrete settings. We now turn to implementing the resulting procedure in a simulated experiment, which will allow us to assess whether the formal construction actually captures the properties we highlighted in the introduction 4.1. The experiment we detail aims in particular to discover continuity in the sensorimotor experience of colors.

4.3.1 Description of the experimental setup

The system involved in the experiment is largely similar to that previously studied in 3.3, featuring a mobile agent observing a 2-d visual scene. This scene displays as continuous stream of (colored) images, supplied by playing back a video file. To obtain an “expected” result, this video file should correspond to a “natural” stimulus since it will form the basis of the resulting sensory topology. The experiment occurs in discrete time, denoted by the index timesteps t_n , $n \in \mathbb{N}$. Importantly these timesteps are regularly spaced (that is, $t_{n+1} - t_n$ is independent of n), which is practically assured

by each timestep corresponding to a frame being displayed in the visual scene. Between any two timesteps t_n, t_{n+1} of its interaction with its environment, the agent may move from $\mathbf{b}[t_n]$ to $\mathbf{b}[t_{n+1}]$ in order to survey different areas of the visual scene. It then locally samples the current image $\varepsilon[t_{n+1}] = \gamma_\varepsilon(t_{n+1})$ using its W_c -by- H_c camera, resulting in a sensory array $\mathbf{s}[t_{n+1}] = (\mathbf{s}_{c_k}[t_{n+1}])_{1 \leq k \leq N_c = W_c \times H_c}$.

The value of a pixel at any single point in the scene is given by an RGB tuple $\varepsilon(x, y) = (R, G, B) \in \llbracket 0; 255 \rrbracket^3$. The sensitivity function (see 2.3.2.2) of every sensel of the camera is given by

$$\begin{aligned} f_c(R, G, B) &= \left\lfloor \frac{1}{32} \begin{pmatrix} R \\ G \\ B \end{pmatrix} \right\rfloor^{16} \cdot \begin{pmatrix} 1 \\ 8 \\ 64 \end{pmatrix} \\ &= h_8(R) \times 8^0 + h_8(G) \times 8^1 + h_8(B) \times 8^2 \end{aligned} \quad (4.12)$$

where $h_8 : x \in \llbracket 0; 255 \rrbracket \mapsto \lfloor \frac{x}{32} \rfloor \in \llbracket 0; 7 \rrbracket$, pictured in Figure 4.4, compresses the initial 256 values of any one component R, G or B into 8 possible output values. Consequently the total output value $f_c(R, G, B)$ of sensel c is a whole number at most equal to $8^3 - 1 = 511$, and which takes all $n_S = 512$ values of $\llbracket 0; 511 \rrbracket$ over the range of its input. In the following, we use $\mathbf{s}_i$ to denote the particular sensory value of $\mathcal{S}$ with label $0 \leq i \leq 511$. Figure 4.5 further illustrates the correspondence between RGB values and their associated labels $f_c(R, G, B)$. It makes apparent that the internal sensory coding performed by f_c does not seem continuous to us, that is the closeness of labels $f_c(\cdot)$ (as numbers) does not necessarily correspond to the closeness we perceive of their corresponding colors.

In this experiment we consider a number of distinct sensorimotor experiences of this setup. Furthermore the dimensions of the camera are kept constant between explorations, with $W_c = H_c = 201$. Any one of these experiments runs thus: at timestep t_n , the agent has absolute configuration $\mathbf{b}[t_n] = (x[t_n], y[t_n]) \in \mathcal{B}$. It samples the current environment state $\varepsilon[t_n]$, resulting in a sensory output

$$\mathbf{s}[t_n] = (\mathbf{s}_{c_k})_{1 \leq k \leq N_c} = \left(f_c(\varepsilon|_{F_{c_k}(\mathbf{b}[t_n])} + \mathcal{W}_c[t_n]) \right)_{1 \leq k \leq N_c}.$$

¹⁶Although customary in computer science this is a slight abuse of notation: it denotes that we take the floor function on a component-wise basis.




FIGURE 4.4: Compression of the input color values by the camera sensor



FIGURE 4.5: Illustration of the internal code of color sensations as numeric labels

$$\mathcal{W}_c[t_n] = \begin{pmatrix} w_1 \\ w_2 \\ w_3 \end{pmatrix} \text{ is a white noise vector, added to the color } \varepsilon_{|F_{c_k}}(\mathbf{b}[t_n]) = \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

effectively displayed in the scene before conversion by the sensitivity function f_c of the pixel. Computationally, all components w_i are independently drawn from the $\mathcal{N}(0, \sigma^2)$ Gaussian distribution with mean 0 and standard deviation σ . The value of σ is kept constant for any one sensorimotor interaction; furthermore, all draws of $\mathcal{W}_c[t_n]$ are also independent as sensel c and timestep t_n vary.

Then the agent moves from $\mathbf{b}[t_n]$ to $\mathbf{b}[t_{n+1}]$ as its experience progresses. The amplitude

$$\Delta \mathbf{b} = \|\mathbf{b}[t_{n+1}] - \mathbf{b}[t_n]\|_2 = \sqrt{(x[t_{n+1}] - x[t_n])^2 + (y[t_{n+1}] - y[t_n])^2}$$

of the movement separating any two consecutive absolute configurations is kept constant in each sensorimotor exploration. It obtains a new sensory reading $\mathbf{s}[t_{n+1}]$ corresponding to its new configuration $\mathbf{b}[t_{n+1}]$ and the updated state of the scene $\varepsilon[t_{n+1}]$. Then it is able to update the relative frequencies $\hat{P}_{\mathbf{s}_i, \mathbf{s}_j}$ by comparing $\mathbf{s}[t_n]$ and $\mathbf{s}[t_{n+1}]$ component by component. Importantly we thus have all sensels contribute to the same frequency count. Although this corresponds once more to an assumption that outputs $\mathbf{s}_c$ and $\mathbf{s}_{c'}$

of distinct sensels are comparable (as opposed to what we outlined in 2.3.2.3), this time it only yields increased computational convenience and does not bear on the fundamental result. Indeed, fusing all counts into one accelerates the estimation of $P_{\mathbf{s}_i, \mathbf{s}_j}$: now each pair of successive timesteps t_n, t_{n+1} produces $N_c = W_c \times H_c = 40401$ events with which the agent can update its observed frequencies, instead of one only. Moreover, keeping separate frequencies $\hat{P}_{c, \mathbf{s}_i, \mathbf{s}_j}$ for distinct sensels $c \in \mathcal{C}$ would involve updating $N_c \times n_{\mathcal{S}} \times n_{\mathcal{S}} = 10,590,879,744$ distinct records, compared to the already sizeable $n_{\mathcal{S}} \times n_{\mathcal{S}} = 262,144$ entries of the $\hat{P}_{\mathbf{s}_i, \mathbf{s}_j}$ fused case. This simplification hinges on the fact that the distinct sensels c_i ought, statistically speaking, to produce the same sensory experiences since they all share the same sensitivity function f_c and are identically affected by displacements of the agent. Therefore, we would expect their separate empirical probabilities $\hat{P}_{c, \mathbf{s}_i, \mathbf{s}_j}$ to all *eventually* converge towards shared values independently of the particular sensel c ; but then this shared value would also be the limit of their mean, that is $\hat{P}_{\mathbf{s}_i, \mathbf{s}_j}$.

This interaction proceeds for $n_t = 4000$ timesteps (corresponding to 4000 frames in the scene), at the end of which we assess the sensory metric resulting from the empirical observations of the agent. This process is further run several times to independently assess the influence of two condition variables. First, we investigate the effect of varying the displacement amplitude $\Delta \mathbf{b}$ by which the agent moves between each timestep, with the base case $\Delta \mathbf{b} = 0$ corresponding to an immobile agent. Second, we have the external white noise $\mathcal{W}$ vary by setting distinct values of standard deviation σ in its original distribution; there the base case is that where no noise is added, that is $\sigma = 0$. Between all sensorimotor explorations, we respectively test for the values

$$\Delta \mathbf{b} \in \{5, 10, 25, 50, 100\}$$

$$\sigma \in \{3, 5, 10, 20, 50\}.$$

4.3.2 Computing the sensory distance

Labelling the sensory outputs of the agent by the single number appearing as output of the sensitivity function f_c makes it so that the $P_{\cdot, \cdot}$ sensory transition probabilities are naturally indexed by pairs of natural numbers. As a result the collection of all such probabilities can be experimentally described as a 512-by-512 matrix M_P , where $(M_P)_{i,j} = P_{\mathbf{s}_i, \mathbf{s}_j}$. Because $P_{\mathbf{s}_i, \mathbf{s}_j}$ is

defined as $\mathbb{P}(\gamma_s[t_{n+1}] = \mathbf{s}_j \mid \gamma_s[t_n] = \mathbf{s}_i)$, we know that at any row i we have $\sum_j (M_P)_{i,j} = \sum_j P_{\mathbf{s}_i, \mathbf{s}_j} = 1$. Therefore, M_P is a Markov matrix.

During the course of its experience, the agent is led to count the number of occurrences of events $\{\mathbf{s}_{c_k}[t_{n+1}] = \mathbf{s}_j \mid \mathbf{s}_{c_k}[t_n] = \mathbf{s}_i\}$ for all possible values $\mathbf{s}_i, \mathbf{s}_j \in \llbracket 0; 511 \rrbracket$. To record these, we therefore define a 512-by-512 matrix M_o it updates as timesteps increase. It is initialized by setting all of its coefficients $o_{i,j} = (M_o)_{i,j}$ to 0, corresponding to the agent not having yet observed any sensory transitions. Then, going from timestep $t_n \in \mathcal{T}$ to t_{n+1} it is updated along

$$o_{i,j}[t_{n+1}] = o_{i,j}[t_n] + \eta_{i,j}[t_n] \quad (4.13)$$

where $\eta_{i,j}[t_n]$ counts the number of sensels $c_k \in \mathcal{C}$ such that $\mathbf{s}_{c_k}[t_n] = \mathbf{s}_i$ and $\mathbf{s}_{c_k}[t_{n+1}] = \mathbf{s}_j$; that is, $\eta_{i,j}[t_n]$ counts the number of transitions from sensory output $\mathbf{s}_i$ to $\mathbf{s}_j$ the agent has observed from t_n to t_{n+1} . Following this definition, the relative transition frequencies $\hat{p}_{i,j} = \hat{P}_{\mathbf{s}_i, \mathbf{s}_j}$ can be obtained from the matrix of occurrences M_o by

$$\hat{p}_{i,j}[t_n] = \frac{o_{i,j}[t_n]}{\sum_j o_{i,j}[t_n]}. \quad (4.14)$$

This consequently defines the matrix $M_{\hat{p}}[t_n]$ of empirical probabilities, which converges towards M_P (as t_n goes to infinity) by the strong law of large numbers (Dekking, 2005).

In the experiment, the function f used to compute the prototypic closeness values $\delta_f(\mathbf{s}_i, \mathbf{s}_j)$ of Eq. (4.8) is set to $-\log$ once and for all; as a result, we dispose of the explicit subscript and instead simply write $\delta(\mathbf{s}_i, \mathbf{s}_j)$. From the previous definitions, it can be applied component by component on $M_{\hat{p}}$ as per

$$\hat{\delta}_{i,j}[t_n] = -\log(\hat{p}_{i,j}[t_n]) \quad (4.15)$$

which further yields the matrix $M_{\hat{\delta}}$. Following the discussion made in 4.2.2.2, it is taken as the adjacency matrix of a graph $G_S^{-\log}$ of sensory values. Accordingly, we solve the All-Pairs Shortest Paths (APSP) problem corresponding to this graph, by which we obtain the quasimetric $\hat{d} = \hat{d}_{-\log}$ of 4.2.2.2 on sensory values. This is done by directly applying Floyd-Warshall's algorithm (Floyd, 1962) on the $M_{\hat{\delta}}[t_n]$ matrix. It outputs $M_{\hat{d}}[t_n]$, the matrix where $\hat{d}_{i,j}[t_n] = (M_{\hat{d}})_{i,j}[t_n]$ is the weighted length of the shortest path from $\mathbf{s}_i$ to $\mathbf{s}_j$ using the edges specified by $M_{\hat{\delta}}[t_n]$. As a consequence, this provides the agent with the result metric we set out to discover *via* $d(\mathbf{s}_i, \mathbf{s}_j) = \hat{d}_{i,j}$.

4.3.3 Spatially representing the sensory metric

4.3.3.1 The use of dimensionality reduction

At this point, the agent has computed a collection of distances $\hat{d}_{i,j} = \hat{d}(\mathbf{s}_i, \mathbf{s}_j)$ separating its sensory values. Although this makes $\mathcal{S}$ a (quasi)metric space, this does not immediately provide a “spatial” layout as we may have wanted it. What such a layout entails, instead, is the search for a subspace $\hat{\mathcal{S}} = \{\hat{\mathbf{s}}_i\}$ of the familiar space $\mathbb{R}^n$ for some n , such that the spatial distance between any two representations $\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j \in \mathbb{R}^n$ matches with the corresponding $\hat{d}(\mathbf{s}_i, \mathbf{s}_j)$. Ideally this representation would be an isometry, showing that the $\hat{\mathcal{S}}$ representation brings no additional use to the agent itself. However, it allows us to assess which metric regularities it has learned from its sensorimotor experience. Following the *manifold hypothesis* (Fefferman, Mitter, and Narayanan, 2013), we assume that the “natural” metric structure learned by the agent is compatible with its sensory space $\mathcal{S}$ being a (Riemannian) manifold. As a result, it can be isometrically embedded into a Euclidean space $\mathbb{R}^n$ (Nash, 1956)¹⁷. The problem of determining a spatial representation $\hat{\mathcal{S}}$ is therefore one of manifold learning (Izenman, 2012), for which a number of algorithms are available (van der Maaten, Postma, and Herik, 2007). However, our use case is somewhat unusual compared to the practice described in these papers. Indeed, a common assumption in manifold learning is that the data we are trying to analyze is already presented as a (sub)manifold. That is, the procedures operate on a set M of N -dimensional samples $\{x_i\}_{1 \leq i \leq n_x}$ where each x_i is itself a vector of $\mathbb{R}^N$; moreover, the distance in M is taken to be the vector distance between the x_i . From there, it searches for a projection $\text{proj}_d : \mathbb{R}^N \rightarrow \mathbb{R}^d$ onto a lower-dimensional space, that is with $d < N$. The projected values $\bar{x}_i = \text{proj}_d(x_i) \in \mathbb{R}^d$ are thought of as their low dimensional representation. Usually, one has to find a suitable d to minimize the representational dimension (and thus approach the intrinsic dimension of the data) while retaining the properties of interest in the dataset. On the contrary, in our case the agent only knows of the *distances* between the x_i data points, that is the property the representation seeks to preserve. We are therefore limited to using algorithms that allow for this sort of limited input, such as ISOMAP (Tenenbaum, 2000) or CDA (J. A. Lee et al., 2000).

¹⁷In the general imbedding theorem very high values of ambient dimension n can be found necessary, depending on the complexity of the metric structure of sensory values. Nonetheless, the manifold hypothesis further states that the *intrinsic* dimension of $\mathcal{S}$ is much lower than that by which it is initially presented, that is $N_c = 201^2$.




FIGURE 4.6: Comparison of the ambient and intrinsic metrics on a Riemannian manifold.

Intrinsic and ambient metrics on a submanifold The criterion for estimating whether a representation manifold $\hat{\mathcal{S}}$ accurately depicts the initial metric structure of $\mathcal{S}$ (given by $\hat{d}$) is, informally speaking, that

$$\forall \mathbf{s}_i, \mathbf{s}_j \in \mathcal{S}, d_{\hat{\mathcal{S}}}(\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j) \approx \hat{d}(\mathbf{s}_i, \mathbf{s}_j)$$

where ideally we would even get an equality. It is therefore important to make explicit which distance $d_{\hat{\mathcal{S}}}$ one takes between representational points $\hat{x}_i$. Indeed, at least two candidates naturally appear from our assuming $\hat{\mathcal{S}}$ is a submanifold in $\mathbb{R}^n$ (Gromov, 1999). This is illustrated in Figure 4.6, in the case where the intrinsic structure of $\mathcal{S}$ is that of a circle embedded in $\mathbb{R}^2$. Let us consider Fig. 4.6a. It pictures the entirety of sensory representation space $\hat{\mathcal{S}}$, among which two arbitrary (representations of) sensory values $\mathbf{s}_i$ and $\mathbf{s}_j$ have been singled out. One way to measure the distance separating these values is by specifying their *ambient* distance as points in the plane $\mathbb{R}^2$: this corresponds to taking the length $d_{\mathbb{R}^2}(\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j)$ of the orange line segment. If one restricts themselves to evaluating the distance separating them along paths *in* $\hat{\mathcal{S}}$, however, they would obtain the *geodesic* distance $d_{\gamma_{\hat{\mathcal{S}}}}(\hat{\mathbf{s}}_i, \hat{\mathbf{s}}_j)$ displayed in blue. This geodesic metric is invariant of the geometry of $\mathcal{S}$, that is independent of which embedding $\hat{\mathcal{S}}$ one may choose.

Fig. 4.6a shows that the two metrics need not coincide. For instance, let us consider $\hat{\mathbf{s}}$ and $\hat{\mathbf{s}}'$ any two antipodal points of $\hat{\mathcal{S}}$. Then we would necessarily have $d_{\mathbb{R}^2}(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = 2R$ where R is the radius of $\hat{\mathcal{S}}$ in $\mathbb{R}^2$; on the other hand, we

would have $d_{\gamma_S}(\hat{s}, \hat{s}t\hat{s}') = \pi R$ since the corresponding path covers exactly one half of $\hat{S}$. In fact, if for any two $\hat{s}_i, \hat{s}_j$ we have $d_{\mathbb{R}^2}(\hat{s}_i, \hat{s}_j) = d_{\gamma_S}(\hat{s}_i, \hat{s}_j)$, then it means that the two (orange and blue) line segments are in fact equal. In particular, that means that the two metrics globally agree if and only if $\hat{S}$ contains every line segment joining any two of its points, that is if and only if $\hat{S}$ is convex. Thus we need to specify which metric on spatial representations $\hat{s}_i \in \hat{S}$ we want to mirror that the agent has internally discovered¹⁸. Recall that in 4.2.2 we introduced the (abstract) sensory metric d_f on S as the distances of shortest paths in some graph G_S^f . These paths correspond to the geodesics in G_S^f seen as a metric space, which are further generalized (for non discrete collection of vertices) by the geodesic distances of manifolds. Thus, we will henceforth consider the intrinsic metric d_{γ_S} as *the* metric we want representation $\hat{S}$ to match with the internal $\hat{d}$ metric.

4.3.3.2 ISOMAP

We now quickly describe how ISOMAP determines a suitable spatial representation $\hat{S}$ of the metric structure the agent learned on its sensory values. Its sensorimotor experience had it learn the pairwise distances $\hat{d}_{i,j} = \hat{d}(\mathbf{s}_i, \mathbf{s}_j)$ between any two sensory values $\mathbf{s}_i, \mathbf{s}_j \in S$. Then,

1. For each sensory value $\mathbf{s}_i \in S$, determine a neighbourhood $N_{\mathbf{s}_i}$. Typically, this is done by selecting either its k nearest neighbours $\mathbf{s}_{j_1}, \dots, \mathbf{s}_{j_k}$ (with k a parameter whole number) or every value that is at most ϵ apart from $\mathbf{s}_i$ for the initial metric $\hat{d}$ (with ϵ a parameter positive number). ISOMAP then assumes that each geodesic distance between values $\mathbf{s}_j, \mathbf{s}_k$ in any one such neighbourhood $N_{\mathbf{s}_i}$ is given by $d_{\gamma_S}(\mathbf{s}_j, \mathbf{s}_k) = \hat{d}(\mathbf{s}_j, \mathbf{s}_k)$. Every other distance, that is those between values which do not share a neighbourhood, is not set at this stage.
2. The distinct neighborhoods are glued together by solving the APSP problem, e.g. using Dijkstra's algorithm. Provided there is a path

$$\mathbf{s}_i = \mathbf{s}_{i_1}, \mathbf{s}_{i_2}, \dots, \mathbf{s}_{i_n} = \mathbf{s}_j$$

¹⁸This is, in essence, the difference between MDS (Borg and Groenen, 1997), which by design operates on the ambient metric, and ISOMAP.

between any two sensory values $\mathbf{s}_i, \mathbf{s}_j$ such that for all k , $\mathbf{s}_{i_k}$ and $\mathbf{s}_{i_{k+1}}$ share a neighbourhood¹⁹, this defines a total (quasi)metric $\mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_+$.

3. Finally, perform classical multidimensional scaling on the matrix of the previous metric $d_{\gamma_{\mathcal{S}}}$. This involves finding the d eigenvectors $e_1, \dots, e_d$ associated to its largest eigenvalues (where d is a parameter whole number): the projections of the $\mathbf{s}_i \in \mathcal{S}$ on $(e_i)_{1 \leq i \leq d}$ give their d -dimensional representations $\hat{\mathbf{s}}_i$.

The intuition behind this process is illustrated in Figure 4.6b, where the distinct metrics on $\hat{\mathcal{S}}$ are illustrated for two close values $\mathbf{s}_i, \mathbf{s}_j$. We can see that the difference between the two distances $d_{\mathbb{R}^2}$ and $d_{\gamma_{\mathcal{S}}}$ vanishes as we restrict our attention to gradually smaller neighborhoods in $\hat{\mathcal{S}}$. This is in fact guaranteed by the manifold being “smooth”, from which both the orange chord and the blue geodesic arc converge towards the tangent direction at $\mathbf{s}_i$ as $\mathbf{s}_j$ tends towards $\mathbf{s}_i$. Therefore the distances calculated in Step 2 occur as the lengths of paths of chords, which increasingly approximate the limit value of the geodesic arc lengths as the length of the chords vanish.

Two parameters appear in any one run of the above procedure: the neighborhood selection parameter k (or ϵ , depending on the selection process) and the output dimension parameter d . In particular, the first controls the quality with which small geodesic lengths are approximated by chord lengths in Step 1; in turn, this directly impacts whether the lengths of paths of selected chords in Step 2 closely follow the geometry of $\mathcal{S}$. Choosing it too small results in missing actual neighbours at any sensory values, introducing “holes” in the resulting structure; on the other hand, if one retains too many neighbours they risk considering chords which “short-circuit” the actual geodesic (such as that seen in Fig. 4.6a). This critical effect of the initial choice of parameter makes ISOMAP topologically unstable (Balasubramanian, 2002). However, methods for algorithmically selecting suitable values have been proposed in further works such as (Choo et al., 2010).

Because ISOMAP ends by applying classical MDS it requires in particular that its input be symmetrical, that is

$$\forall i, j, \hat{d}_{i,j} = \hat{d}_{j,i}.$$

¹⁹(which is especially guaranteed whenever the neighbourhoods are determined using k Nearest Neighbours rules)

This is a very standard requirement in visualisation algorithms, since the resulting plot generally lies in a *metric* (as opposed to quasimetric) space $\mathbb{R}^d$ where apparent distances are themselves symmetric. Nevertheless, this may not be strictly true of the $\hat{d}$ quasimetric the agent has discovered in its sensorimotor experience. We therefore symmetrize it beforehand by defining

$$\hat{d}_{i,j}^{\text{sym}} = \frac{1}{2} (\hat{d}_{i,j} + \hat{d}_{j,i}). \quad (4.16)$$

The resulting (proper) metric $\hat{d}^{\text{sym}}$ is then fed to the ISOMAP algorithm previously described. The reasoning supporting this preprocessing runs thus: if the (theoretical) transition probabilities $P_{\mathbf{s}_i, \mathbf{s}_j}$ are themselves symmetrical, then $\hat{d}$ estimates a function which is symmetrical in its inputs itself. But then $\hat{d}_{i,j}$ and $\hat{d}_{j,i}$ are two estimates of the same value, by which their mean $\hat{d}_{i,j}^{\text{sym}} = \hat{d}_{j,i}^{\text{sym}}$ also estimates this common value $d(\mathbf{s}_i, \mathbf{s}_j)$. Nonetheless, we should also verify that the learned quasimetric $\hat{d}$ is at least *almost* symmetric. To this end, consider any square $n_S \times n_S$ matrix M . It is known that M can be written as the sum of a symmetric part and an antisymmetric one

$$M = M^{\text{sym}} + M^{\text{antisym}}$$

where we further know that $M^{\text{sym}} = \frac{1}{2}(M + M^\top)$ and $M^{\text{antisym}} = \frac{1}{2}(M - M^\top)$. Moreover, these two parts are orthogonal for the matrix inner product $\langle A|B \rangle = \text{tr}(A^\top B)$. This further entails *via* the Pythagorean theorem

$$\|M\|_2^2 = \|M^{\text{sym}}\|_2^2 + \|M^{\text{antisym}}\|_2^2$$

where $\|\cdot\|_2$ is the norm associated with said inner product. Therefore, for any square matrix M we define

$$C_{\text{antisym}}(M) = \frac{\|M^{\text{antisym}}\|_2^2}{\|M\|_2^2}. \quad (4.17)$$

This yields a function which maps all such matrices to the range $[0; 1] \subset \mathbb{R}$, with output 0 (resp. 1) corresponding to symmetric (resp. antisymmetric) matrices. Moreover, if we denote C_{sym} the symmetric equivalent, then we have

$$\forall M, C_{\text{sym}}(M) + C_{\text{antisym}}(M) = 1.$$

Thus, $C_{\text{antisym}}(M)$ measures some proportion by which M is antisymmetric.

We therefore evaluate $C_{\text{antisym}}(M_{\hat{d}})$ before symmetrizing the empirical quasimetric $\hat{d}$ to ascertain whether this preprocessing step introduces significant changes in the $\hat{d}_{i,j}$: if our assumption that it is almost symmetrical actually holds, then the resulting values of C_{antisym} should be close to 0.

4.3.4 Results

4.3.4.1 The sensory metric

Figure 4.7 pictures the result metric the agent discovers in the base case, that is in an interaction where it stayed immobile ($\Delta b = 0$) and where no additional noise was added to the visual scene ($\sigma = 0$). Its first three columns respectively display $M_{\hat{p}}, M_{\hat{\delta}}$ and $M_{\hat{d}}$ the matrices of relative (conditional) frequencies, closeness values $\hat{\delta}(\mathbf{s}_i, \mathbf{s}_j)$ and eventual sensory quasimetric. The correspondence (previously seen in Fig. 4.5) between numerical label i of $\mathbf{s}_i$ and (one of) the color it represents is figured next to their bottom and left sides to better interpret their shape. The color scale denoting the value of the matrix coefficients is kept consistent between the three examples: dark colors correspond to the lowest values (which go as low as 0), while bright ones denote the larger numbers. Because all scales internally use the same (full) range of color, note that a certain shade may not denote the same value in two distinct matrices. For instance the exact middle tone, a dark reddish pink, denotes a value of 0.5 in the first matrix $M_{\hat{p}}$ since all relative frequencies span the entire $[0; 1]$ range; on the other hand, the same color corresponds to a value $\hat{d}_{\text{med}} \approx 17$ in the final matrix $M_{\hat{d}}$. Then the final, fourth column shows the embedding in $\mathbb{R}^3$ ISOMAP determines from the internal (symmetrized) metric $\hat{d}^{\text{sym}}$. We have further evaluated the antisymmetry coefficient $C_{\text{antisym}}(M_{\hat{d}})$ for the quasimetric matrices the agent discovers in all of the explorations described in 4.3.1. Their distribution histogram is pictured in Figure 4.8. It shows that after 4000 timesteps, the antisymmetric proportion of the quasimetric matrix $M_{\hat{d}}$ lies in the range $[0.01; 0.024]$: in particular, it is always smaller than 2.4% over all explorations considered here. Therefore we proceed with the symmetrization into $\hat{d}^{\text{sym}}$ on the grounds that it does not significantly alter the relative (asymmetric) distances $\hat{d}_{i,j}$, instead merely regularizing some discrepancy that arises from empirical estimation.




FIGURE 4.7: Sensory metric matrices resulting from the sensorimotor interaction of the agent in the base case.




FIGURE 4.8: Evaluation of the symmetric component of the quasimetrics empirically discovered

We can see in Fig. 4.7 that $M_{\hat{p}}$ is almost everywhere black. This denotes the property that some particular sensory transitions from $\mathbf{s}_i$ to $\mathbf{s}_j$ are vastly more frequent than most others, and that the conditionals

$$\{\mathbf{s}[t_{n+1}] = \mathbf{s}_j \mid \mathbf{s}[t_n] = \mathbf{s}_i\}$$

are generally extremely unlikely. In particular the diagonal coefficients $(M_{\hat{p}})_{i,i}$ form the only region which is relatively bright, corresponding to sensory outputs $\mathbf{s}_i$ reliably following themselves in the sensorimotor flow. The relative differences between these values appear smoothed in M_{δ} (as well as reversed, since $f = -\log$ is nonincreasing), which has patterns of darker regions appear in the matrix. These patterns visibly form squares, with each row and column being divided in eight equally sized squares. These are purely consequences of the coding of sensory values introduced by the sensitivity f_c in 4.3.1, as the juxtaposed index colorbar shows: these squares correspond to one of the larger “periods” of Fig. 4.5, such as that encompassing $\mathbf{s}_0, \dots, \mathbf{s}_{63}$. The same observation extends to the apparent “dotting” motifs appearing in each of these large squares, which further mirrors the smaller “periods” (of length 8, e.g. that of $\mathbf{s}_0, \dots, \mathbf{s}_7$) of the same figure. Importantly, this shows that the distances discovered between any two numeric labels i, j

is predictably contingent on the specific internal code f_c of sensations²⁰. In fact $M_{\hat{\delta}}$ shows that $\hat{\delta}_{i,j}$ need not be small whenever i and j are close *as internal code*: this is illustrated by the visual discontinuities where dark regions sit directly next to bright ones. In particular, the edges of the aforementioned squares are all such regions where continuity of the labelling codes $i = f_c(\cdot)$ does not coincide with that of sensations the agent is determining. Finally, we can see that making $\hat{\delta}$ into the quasimetric $\hat{d}$ by computing the shortest paths smoothes out the matrix coefficients. In fact, apart from some specific very bright regions corresponding to some sensory values the agent hardly encountered in this specific exploration, most values of $\hat{d}$ fall in the lower values of the range displayed on the color scale. This is because although the agent may have not directly experienced transitions e.g. from pure red (labelled by $i = 7$) to pure green (resp. $j = 56$), it has however experienced those from red to pure yellow (resp. $k = 63$) and from yellow to green. It can therefore accordingly infer that however large $\hat{\delta}_{7,56}$ is, the corresponding metric distance $\hat{d}_{7,56}$ should be much lower since pure red and pure green can be joined in as few as 2 relatively plausible transitions. Still, while the quantitative contrast (as illustrated by the colors used) between values $\hat{d}_{i,j}$ is thus lower than in $M_{\hat{\delta}}$ our previous observations still hold. In particular, the same square patterns show that the continuity of the metric $\hat{d}$ is not that of the labels $i \in [0; 511]$ as numbers.

Let us now consider the spatial embedding displayed in the last column of Figure 4.7. It appears that the manifold is intrinsically 3-dimensional, resembling a ball in $\mathbb{R}^3$. The view in Fig. 4.7 has been chosen to have the border of the apparent circle correspond to the hue parameter; indeed, we can see that distinct colors (e.g. blue, yellow, and red tones) are grouped in their respective angular portions of the manifold. Moreover, the way these portions are arranged is compliant with our intuition about mixing colors: purples lie between blues and reds, oranges between yellows and reds, and cyans between blues and greens. This is entirely a spatial consequence of the pairwise distances internally discovered, e.g. the property that distances between blues and purples being smaller than that between blues and reds, which the agent has learned only by assessing that the experience of blue was more reliably followed by that of purple than by that of red. Although

²⁰This mirrors the development about numberings of sensels previously made in 3.3.2.1, except here we would relabel sensory values instead of sensory elements. A permutation of the labels would accordingly only result in a conjugate of the matrix, which importantly does not bear on the latter computations.

the 2-dimensional projection here makes it somewhat harder to discern, the foreground points are that of pale colors whereas darker shades are symmetrically relegated to the background. Additionally, the interior points of the manifold correspond to desaturated colors, with various shades of grey placed near the spatial center of the ball. This shows that, although the entire process by which the agent inferred this metric structure of sensory values from its naive experience had no assumption about the structure of colors, it nonetheless has a very familiar system of coordinates emerge. Let us indeed consider the parameterization of the manifold *via* spherical coordinates (Weisstein, n.d.) (r, θ, ϕ) relative to a zenith direction passing through the black and white portions (that is, orthogonally through the 2-d view of Fig. 4.7). Then the radial distance r , the azimuthal angle θ , and the polar one ϕ respectively represent saturation, hue, and luminance associated with the sensory values. Therefore this shows how the formal continuity which determined the sensory metric $\hat{d}$ corresponds to our intuitive understanding of colors in terms of these 3 characteristics.

4.3.4.2 Crawling in my film: Effects of amplitude in the agent's movements on the discovered regularities

We now examine how the movements the agent can perform in its environment alter this behavior. More specifically, we consider the internal metrics $\hat{d}^{\Delta \mathbf{b}}$ obtained from the particular conditions wherein any two successive configurations $\mathbf{b}[t_n]$ and $\mathbf{b}[t_{n+1}]$ of the agent are exactly $\Delta \mathbf{b}$ apart. Accordingly higher values of $\Delta \mathbf{b}$ correspond to larger changes of configurations between each time step, which (in a continuous setting) reflects faster movements of the agent.

We saw in 4.2.3.1 that, provided the metric structure discovered by the agent in $\hat{d}$ is mediated by some *local* coherence by which certain particular sensory transitions $\mathbf{s}_i \rightarrow \mathbf{s}_j$ are more likely, then we should expect certain behaviors as its experience loses this local coherence property. In particular, total loss of coherence has the events $\{\mathbf{s}[t_{n+1}] = \mathbf{s}_j\}$ and $\{\mathbf{s}[t_n] = \mathbf{s}_i\}$ be independent. As a result, we would expect in the limit case that all $\hat{p}_{i,j}$ only depend on i , by which this is also the case of the $\hat{d}_{i,j}$. Consider then for each sensory value index $i \in \{0, \dots, 511\}$ the i -th row

$$L_i = \left(\hat{d}_{i,j} \right)_{0 \leq j \leq 511} = \left(\hat{d}_{i,0} \quad \dots \quad \hat{d}_{i,511} \right) \in \mathbb{R}^{512}$$

of the metric matrix $M_{\hat{d}}$. The previous hypothesis of totally incoherent experiences entails that all L_i are equal, independent of subscript i . For any metric matrix $M_{\hat{d}^{\Delta\mathbf{b}}}$ we therefore quantify the extent to which they are dissimilar by computing their variance (as a collection of vectors) $\text{Var}(L_i)$. This is pictured in Figure 4.9. The x axis labels the values of $\Delta\mathbf{b}$ corresponding to each trial, whereas the y axis denotes this measure of row dissimilarity. Although it never reaches the null value of the ideal incoherent case, we can indeed see that the plot displays a decreasing graph corresponding to the rows L_i growing more similar as amplitude $\Delta\mathbf{b}$ increases. Moreover two stages, distinguished by the slopes involved, appear in the pictured graph. The first, corresponding to the initial increases in $\Delta\mathbf{b}$, is steeply decreasing: indeed, the discovered metrics $\hat{d}^{\Delta\mathbf{b}}$ go from $\text{Var}(L_i) \approx 114$ to $\text{Var}(L_i) \approx 94$ as $\Delta\mathbf{b}$ grows from 5 to 25, for an average slope of $\frac{-20}{20} = -1$. On the contrary, for higher values of $\Delta\mathbf{b}$ the decrease is much more limited. In particular, the value $\text{Var}(L_i)$ of row dissimilarity for $\Delta\mathbf{b} = 100$ is around 84, by which the corresponding average slope is $\frac{-10}{75} \approx -0.13$. This shows several properties of the local coherence present in the sensorimotor experience. First, since $\text{Var}(L_i)$ never reaches 0 its effects can be found even as the agent performs very large displacements. Second, from the initial steep decrease we can conclude that the coherence is particularly strong in *very* local experiences, that is experiences where the agent moves very slowly²¹.

The previous discussion shows that movement of the agent naturally introduces increased uncertainty in its observed sensory transitions, with the increase being especially notable starting from the entirely immobile case. As a consequence, we should expect the agent to pick up on distinct “typical” sensory transitions depending on the movement amplitude $\Delta\mathbf{b}$ wherewith it learned their probability laws. Earlier in 4.2.2.1, we introduced the desired sensory metric $d(\mathbf{s}, \mathbf{s}')$ as one which graded such a notion of typicality of sensory experiences; we also further examined in 4.2.3.1 how this becomes formally apparent by relating the distance between $\mathbf{s}_i$ and $\mathbf{s}_j$ and the probability (that is, to the agent, the expectation) that $\mathbf{s}_j$ should follow $\mathbf{s}_i$. But then it follows that if the movement amplitude $\Delta\mathbf{b}$ altered which sensory transitions the agent should find typical of its experience, then so should it affect the capability of the resulting metric $\hat{d}^{\Delta\mathbf{b}}$ to accurately assess the typicality of

²¹Of course, the same intuition tells us that the rate at which displacements of the agent have its experience lose coherence also depends on the environment state. In particular all of the discussion hinges on the environment states having a similar local coherence (both spatial and temporal) themselves, as opposed to them displaying e.g. white noise.




FIGURE 4.9: Quantifying the robustness of the conditioning in sensory transitions $\mathbf{s}[t_{n+1}] = \mathbf{s}_j \mid \mathbf{s}[t_n] = \mathbf{s}_i$ as movement amplitude $\Delta \mathbf{b}$ of the agent varies.

latter experiences depending on their movement conditions.

This is further studied in Figure 4.10. It displays ten 2-d (blue) point cloud plots, each corresponding to two distinct movement amplitudes $\Delta \mathbf{b}_{\text{learning}}$ (constant for each row) and $\Delta \mathbf{b}_{\text{testing}}$ (resp. for each column). Any one graph indexed by $(\Delta \mathbf{b}_{\text{learning}}, \Delta \mathbf{b}_{\text{testing}})$ is that of a distinct exploration. It represents the experience of an agent that has learned its sensory metric in the amplitude condition $\Delta \mathbf{b}_{\text{learning}}$ but which is further interacting with its environment along condition $\Delta \mathbf{b}_{\text{testing}}$. What each plot depicts, on the y axis, is the probability in the *latter* experience (that is, with $\Delta \mathbf{b} = \Delta \mathbf{b}_{\text{testing}}$) at which occur transitions of a certain metric value from the *former* (i.e. with $\Delta \mathbf{b} = \Delta \mathbf{b}_{\text{learning}}$), denoted on the x axis. From the discussion of 4.2.3.1, we know that in the ideal case this point cloud should follow the graph $y = e^{-x}$ ²² to reflect the agent's capability to accurately grade the typicality of sensory transitions. Accordingly, each of the 10 blue point clouds is fitted by a function of the form $y_{a,b} : x \mapsto ae^{-bx}$ (with a and b as parameters) using the least squares method. For each separate plot, the graph of the fit determined thus is plotted in red against the point cloud; additionally, the values of parameters a, b are displayed on top of the corresponding graph.

²²The theoretical fit being an exponential is entirely dependent on us choosing $f = -\log$ in the experimental setup; in the more general (but bijective) case, we would have fit against $y = f^{-1}(x)$.




FIGURE 4.10: Comparison of the effectiveness of $\hat{d}^{\Delta b}$ at grading typicality of sensory experiences for distinct values of Δb .

For now let us focus on the topmost row of Fig. 4.10, which corresponds to a fixed movement amplitude $\Delta b_{\text{learning}} = 5$ in the learned stage. We can see that the shape of the graph varies between columns: it appears flatter as $\Delta b_{\text{testing}}$ increases. The difference between the two extremes amplitude values (5 and 100) is particularly apparent. In the first one, the point cloud clearly follows a decreasing shape: this denotes the agent observing that transitions which are “larger” –as per its learned internal metric $\hat{d}$, on the x axis– occur comparatively less frequently than those which are smaller. In particular, the exponential fit is indeed close to the theoretical $y = e^{-x}$ (i.e. $a = b = 1$) in the condition $\Delta b_{\text{testing}} = \Delta b_{\text{learning}} = 5$. On the contrary the rightmost graph is almost horizontal, as reflected by the very small value of corresponding parameter b . This shows that the distance values $\hat{d}_{i,j}^5$ the agent has learned while moving slowly do not provide it reliable information about which sensory transitions are typical when its displacements are much larger (i.e. at $\Delta b = 100$). We can further observe that this flattening occurs gradually as $\Delta b_{\text{testing}}$ goes from 5 to 100; moreover, the successive values of fit parameter b show that most of it occurs in the first increases from 5 to 10 and from 10 to 25. This precisely corroborates our earlier assessment in Figure 4.9: in particular, it shows that the structure of which sensory transitions are typical is influenced by a strong local coherence effect which quickly vanishes as Δb passes the 25 threshold. To the agent this means that its capability to discern which transitions are typical, which it derives from its sensory metric $\hat{d}^5$, cannot efficiently account for the sensorimotor experiences it further experiences if it does not keep its displacements accordingly small.

The bottommost row, for which the movement amplitude while learning is fixed at $\Delta \mathbf{b}_{\text{learning}} = 100$, displays the converse situation. There again, the condition $\Delta \mathbf{b}_{\text{learning}} = \Delta \mathbf{b}_{\text{testing}} = 100$ shows the expected result, with an exponential fit close to the theoretical $y = e^{-x}$. Similarly, the ability of the agent to grade the typicality of transitions in experiences where it moves by various values of $\Delta \mathbf{b}_{\text{testing}}$ degrades as they grow farther from $\Delta \mathbf{b}_{\text{learning}} = 100$. However the worst case performance, achieved for minimal $\Delta \mathbf{b}_{\text{testing}} = 5$, is noticeably better than the reverse $\Delta \mathbf{b}_{\text{learning}} = 5, \Delta \mathbf{b}_{\text{testing}} = 100$. In fact there is only a small difference in shape of the point cloud as $\Delta \mathbf{b}_{\text{testing}}$ shrinks from 25 to 10 and, further, the corresponding difference from 10 to 5 is almost negligible as reported by the fit parameters a and b . This shows that the statistical regularities in sensory transitions the agent has learned while performing large displacements generalize better to the (almost) immobile case than the converse do. Therefore, although the agent is better able to distinguish which of its sensory transitions are “typical” whenever the conditions of its sensorimotor experience match that of its previous learning, this shows an implicit advantage of learning these metric regularities during *active* exploration.

4.3.4.3 Introducing noise in the sensorimotor interaction

We now investigate how the internal metric $\hat{d}$ discovered by the agent is affected by the addition of random sensory noise in its learning exploration. To this end, we compare several experiences of the agent where an artificial Gaussian noise, centered and parameterized by its standard deviation $\sigma > 0$, is applied to each color it samples in its environment state. Figure 4.11 displays the results for each case, ranging from $\sigma = 3$ to $\sigma = 50$. The top row presents the resulting distance matrices $M_{\hat{d}\sigma}$, matched below by their respective 3-d ISOMAP embeddings. Each column corresponds to a single value of the standard deviation σ in the noise distribution, specified at the bottom.




FIGURE 4.11: Influence of the amount of noise on the discovered sensory metric. Each column corresponds to a fixed amount of external sensory noise, parameterized by the standard deviation σ of its source normal distribution $\mathcal{N}(0, \sigma^2)$. The topmost row pictures the resulting distances matrices M_d^σ , whereas the bottom one displays the corresponding 3-d ISOMAP embeddings.

We can see that all distance matrices $M_{\hat{d}_\sigma}$ exhibit the shape described earlier in 4.7, with subdivisions into periodic squares and local discontinuities corresponding to the internal coding of sensations of f_c . Similarly all but the last embedding, corresponding to $\sigma = 50$, resemble what we previously examined. They are all viewed from a similar perspective, determined by blue shades being in the foreground and black ones appearing in the topmost region. Interestingly we can see that the orientation –or *handedness*– of the manifold is not constant. Indeed going from cyans to greens follows a clockwise path in the embedding of the first column, whereas it plots out counterclockwise trajectories in the next three. This is in fact an expected result: since each of these manifolds and their respective mirror images (which reverses handedness) are mutually isometric, orientation does not unambiguously follow from the metric data $\hat{d}$ the agent has learned. Therefore this observation helps reinforce the claim that these embeddings faithfully depict that “spatial” knowledge the agent has effectively acquired by computing $\hat{d}$, without introducing external structure in the process.

On the other hand, the embedding corresponding to the maximal noise case $\sigma = 50$ seems clearly distinct from all others. Indeed, the correspondence between the colors of representational points and their spatial groupings is far less obvious than that described in 4.7. For instance, observe that very diverse colors (among them reds, yellows, some greens and blacks) are all clumped in the leftmost region of the embedding; furthermore, purple and blue shades we experience as very similar have been put at distant places in the representation. Moreover the resulting space $\hat{\mathcal{S}}$ is clearly no longer a ball: although it may still locally be 3-dimensional, it is mostly devoid of points near its center. It all results in the disappearance of the previous identification we found between the formal spatial structure imposed by the metric, which we could leverage using spherical coordinates (r, θ, ϕ) , and the subjective qualities of colors as we perceive them. This corresponds to the distance matrices $M_{\hat{d}_\sigma}$ being increasingly “smoothed” as σ grows, which is made especially apparent in $M_{\hat{d}_{50}}$. For example, we can see that the range of distance values appearing in $M_{\hat{d}_\sigma}$ shrinks as σ increases: the maximal value goes from about $\hat{d}_{\max} \approx 33$ for $\sigma = 3$ to ≈ 27 for $\sigma \in \{5, 10, 20\}$, before falling to about 17 as $\sigma = 50$. Therefore, increasing the amount of noise has the result of making distinct sensory transitions from $\mathbf{s}_i$ to $\mathbf{s}_j$ less discernible by simply comparing their metric values. Indeed the noise added to the visual scene decreases the influence of the local coherence property leveraged by





FIGURE 4.12: Evolution of the distributions of learned sensory distances $\hat{d}^\sigma(\hat{s}_i, \hat{s}_j)$ as σ varies.

the procedure, especially since we specified its samples to be (both spatially and temporally) independent. Therefore the conditioning of $\mathbf{s}[t_{n+1}]$ relative to $\mathbf{s}[t_n]$ becomes weaker as the spread of its sampling distribution, parameterized by σ , grows. Moreover, the experienced sensory output $\mathbf{s}[t_n]$ becomes increasingly determined by the white noise as opposed to the actual environment state $\varepsilon[t_n]$ as the signal-to-noise ratio shrinks. As a result, in the limit case of a flat distribution (corresponding to arbitrarily large values of σ), the experienced sensations behave as independent draws from a uniform random distribution. In turn, this corresponds to a limit state of $M_{\hat{d}^\sigma}$ where all coefficients are equal, since we would have $\hat{p}_{i,j} = \frac{1}{n_S} = \frac{1}{512}$ irrespective of i and j .

Figure 4.12 further develops and quantifies this reasoning. It presents two series—one per row—of 5 histograms, with each column corresponding to one of the experiences with a fixed value of σ . These histograms depict the distribution densities of the distance values $\hat{d}_{i,j}$ in the corresponding matrices $M_{\hat{d}^\sigma}$, with the bottom one being that of the cumulative sum. The horizontal ranges have been chosen equal between all plots to better visualize the way the shape evolves as σ varies; additionally, the vertical shape is also fixed for the cumulative histograms since by definition they precisely span the $[0; 1]$ range. Consider the density distributions pictured on the top row. Each of them is unimodal and all but the last, for $\sigma = 50$, is at most slightly right-skewed. We can see however that the spread of the distributions and the location of their mode gradually (and monotonically) varies with σ . More

FIGURE 4.13: Evolution of the variance in distances $\hat{d}_{i,j}^\sigma$ as σ varies.

precisely, the distribution of the $\hat{d}_{i,j}^\sigma$ gets narrower as σ increases. This is illustrated by the leftmost histogram being spread over 18 bins (excluding outliers), whereas the rightmost one only necessitates 7²³. This is congruent with the distribution of the $\hat{d}^\sigma$ gradually shifting to the previously discussed Dirac distribution of the totally uniform experience. This phenomenon is also clearly visible on the cumulative distributions of the bottom row, where the extremal case of $\sigma = 50$ displays a very steep increase. The limit theoretical case here is that of the Heaviside step function, to which the cumulative distributions grow closer as σ increases. In particular, the bottom right plot clearly shows that the range effectively used by the $\hat{d}_{i,j}^{50}$ has shrunk compared to that of the $\hat{d}_{i,j}^3$: to the agent, this means that its capability to grade the typicality of sensory experiences covers a much more narrow range of cases.

The extent to which this spread evolves with the amount of noise is detailed in Figure 4.13, which plots the variance of all entries $\hat{d}_{i,j}^\sigma$ in $M_{\hat{d}^\sigma}$ against σ . But the variance of the $\hat{d}_{i,j}^\sigma$ is precisely the variance of the corresponding distribution in Fig. 4.12, which quantifies its spread. And indeed Fig. 4.13 shows that the variance of $\hat{d}_{i,j}^\sigma$ is a decreasing function of σ . In fact, it is approximately decreased sixfold (from 24.3 to 3.9) as σ goes from 3 to 50. Furthermore, we see that the rate at which the variance decreases is not uniform. Instead the first few increases in σ , from 3 to 10, have minimal effect on $\text{Var}(\hat{d}_{i,j})$: it only goes from 24.3 to 23.4 in this region of the graph, for an

²³Since the horizontal range of the plot as well as the bin width are kept constant, the number of bins directly mirrors the actual range spanned by the distribution.

average slope of -0.13 . On the other hand the decrease is much sharper as σ goes from 10 to 50, corresponding to an average slope of -0.5 .

This helps cement the observation, previously made when we examined Fig. 4.11, about how the procedure used for internally determining the sensory metric $\hat{d}$ is affected by the introduction of external noise. To begin with, it shows that it is robust to the addition of some “sufficiently small” amount of noise as parameterized by σ , which does not significantly alter its resulting output (as reported by the first embeddings of Fig. 4.11 or the variance of Fig. 4.13 for low σ). On the other hand it shows how, once it crosses a certain threshold, this external noise can drown out the desired regular sensory transitions in the sensorimotor experience of the agent. This entails that the metric it learns does not effectively support the capability to grade the typicality of sensory transitions (as seen in Fig. 4.12 for large values of σ); but this is in turn reflected on the spatial understanding of closeness of sensations, which are not faithfully laid out in the last embedding of Fig. 4.11. In point of fact, this limitation is not surprising at all considering how this sensory metric was entirely designed as an internal image of the likelihood of pairwise sensory transitions. What was less obvious, then, is that it *can* produce valid results in noisy settings as long as the noise is small enough. Recalling our discussion in 2.2.3.3, this is a substantial improvement in the context of our processing sensorimotor signals as uninterpreted symbols (described in 2.3.2.3). In particular, this procedure can both be carried out in a noisy setting and further endow the agent with the capability to process noise in its subsequent experiences.

4.4 Conclusion

The SMCT context has agents face significant orders of magnitude from the onset of their cognitive development. Indeed, it has the robot cognize at the subsymbolic level of its bare sensorimotor flow, which is a very rich if unstructured stream of information. Consequently, the wealth of interactions an agent may have in realistic conditions can only result in an immense number of possible states of its sensorimotor flow. Forasmuch as its development may also be open-ended, this requires that the agent be able to transfer knowledge along related interactions.

This chapter was devoted to studying a “topological” notion of similarity between interactions that could mediate such a capability. It relies on the

intuition that two “similar” interactions (i.e. interactions that involve “similar” states of the environment and actuations of the agent) must result in two “similar” sensorimotor experiences, a quality that the agent should be able to ascertain. As such it corresponds to a certain notion of subjective continuity for the agent, which precisely asserts that its sensorimotor interaction is continuous.

We first detailed how the mathematical notion of continuity yields some insight into this relationship between similarity and continuity. In particular, we detailed how topology (i.e. precisely the study of general mathematical continuity) establishes a bidirectional correspondence between the two points of view *via* the introduction of final topologies. Consequently, we described how one could go about framing the continuity of an agent’s internal sensorimotor experience not as an experimental fact dependent on which specific interaction it took part in, but as an *a priori* enforced constraint from which to deduce similarity relations on its sensory outputs.

However, we also showed that general topology could not thereby provide a comprehensive and satisfactory account of our subjective notion of sensory similarity. The analogy between formal and felt continuity especially breaks down in the discrete case, a case where we necessarily leverage additional structure. We therefore transposed the previous idea to the context of metric topology, resulting in the definition of a sensory metric. Importantly, this sensory metric grades the similarity of sensations as the likelihood of their being successively experienced so that typical sensorimotor experiences result in continuous sensory trajectories.

We illustrated the construction in a simulated experiment where the agent was made to discover the sensory metric of its color sensations. We found that the proposed process resulted in the discovery of a 3-dimensional topology of colors coherent with our intuition, hence showing that formal (metric) continuity provides satisfactory grounds for the desired subjective similarity. Importantly we demonstrated that the similarity the agent learned is entirely that of its experienced sensations, as opposed to that of their internal codes. This demonstrates that metric (and topological) relations between sensations can be found as sensorimotor contingencies and that we do not need to assume that the agent *a priori* know of them. Besides, it provides the agent with capabilities to tackle approximate judgements instead of the previously purely categorical ones, thus possibly giving a lead to bridge the numericity gap outlined in 2.2.3.3. We also showed that the resulting color topology is

robust to the introduction of (some) sensory noise, potentially enabling the agent to face more realistic setups at last.

This development opens several immediate perspectives. On a theoretical level, we only showed that metric spaces provided a sufficient toolset to implement the subjective sense of continuity discussed, whereas topological spaces did not. Thus it remains to be determined however if the full expressive capabilities of metrics are in fact necessary, or if one could settle on an intermediary context such as that of uniform spaces. Additionally, we did not ensure that the definition of the internal sensory metric make it an internal object in the sense of 2.3.2.3. Indeed, it is not necessarily true that the conversion from transition likelihood $P_{s,s'}$ to closeness value $\delta_f(s, s')$ can be internally performed²⁴. A tentative workaround may be found by noticing that the $P_{s,s'}$ appear as if they were state transition probabilities of a Markov process. Following this interpretation then, the sensory metric is essentially asked to be a metric on the Markov state space such that more likely transitions drag their corresponding states closer. As a result, one may investigate a number of related metric constructions on Markov processes (Boyd et al., 2020) such as commute time, hitting time, or diffusion distance (Chebotarev and Deza, 2020; Coifman and Lafon, 2006; Rozinas, 2010); crucially these may be entirely computed in terms of experienced statistical observations of sensory events, which would make them internally available. On the other hand, practical improvements could be sought in the empirical estimation of the sensory transition likelihoods. Indeed, our simulation relied on all sensels pooling their contribution to a shared estimated law of sensory transitions. However, computational convenience alone justified this approach which may yield incorrect results for agents with heterogeneous sensors. In the general case, two natural options arise: either one can estimate the (joint) transition law for sensations of the total sensory apparatus $\mathcal{C}$, or they can instead separately estimate the $|\mathcal{C}|$ separate transition laws corresponding to each sensel $c \in \mathcal{C}$. Both cases raise issues, however, with the first providing relatively few samples of transitions in the huge total state space $\mathcal{S}_{\mathcal{C}}$ whereas the latter corresponds to the numerous marginal probabilities of transitions in the smaller $\mathcal{S}_c$. Else, one could instead generalize the pooling our agent performed provided they internally determined which sensors actually estimate common laws (e.g. by monitoring their pairwise correlations as in (Censi and Scaramuzza, 2013; Grossmann, Gaspar, and Orabona, 2010)

²⁴(unless, of course, f is internally available itself)

or the latter parts of (Valentin Marcel, 2020)). This would potentially allow agents endowed with complex sensory apparatuses to determine their sensory topologies within a reasonable time.

Conclusion

This thesis aimed to give a formal account of the *sensorimotor contingencies theory* of perception, with a particular emphasis on its application in robotic settings. The classical approach to understanding perception in such contexts has historically been both passive and symbolic. The corresponding paradigms have robotic agents seek perception in internal representations of the world and their interaction. These agents obtain said representations by applying feature extraction techniques to their sensory “inputs”; importantly these symbolic representations are both that of our conscious understanding and that on which to form adapted sensorimotor behaviours. The SMCT approach breaks from this viewpoint in several aspects. First, it is adamantly antirepresentationalist and instead surmises that sensorimotor cognitive functions operate on a subsymbolic level, namely that of the bare sensorimotor flow. Second, it has perception depend on crucial constraints on the actual sensorimotor experiences of a sensorimotor agent. Notably, this involves constraints that result from its interaction being embodied, situated, and enactive. Finally, it frames perceptive capabilities as something the agent can develop instead of some signal processing techniques it must *a priori* know.

In addition to the insight it gives into the study of perception and its development, these components lend several desirable properties to the SMCT paradigm in the context of robotics. Following its principles, perceiving only requires robots to have the capabilities to recognize the contingencies that appear in their sensorimotor flow as their experience unfolds. Hence this developmental process is independent of the specific sensorimotor capabilities of the agent and its environment. Consequently, the sensorimotor contingencies theory considers very generic laws agents can leverage to achieve autonomous sensorimotor behaviour adapted to the particular conditions they are effectively experiencing. Several previous works then investigated how agents could discover key characteristics of their interaction (e.g. their body

schema or the dimensionality of the ambient space) by observing certain particular events in their sensorimotor flow (respectively self contact and compensable movements). In doing so, they highlighted that care needed to be taken in the formal modelling of these interactions since it makes these characteristics into variables of interest. As a result, various frameworks have been set up to study sensorimotor contingencies.

The first contribution of this work lies in Chapter 2 where we proposed such a formalisation. It followed from an examination of how the choices of formal language we make reflect on the content it can address. In particular, we showed that choosing which mathematical structures the formalism recruits accounts for considerations about *a priori* knowledge of the interaction (whether by the observer or the agent itself). This argument enabled us to formalise the claims about genericity of the SMCT approach by comparing the *a priori* assumptions between the distinct frameworks. We also discussed how limitations can result from properties of the formalisation where they are, in a sense, hardcoded. For instance, we detailed how previous parameterizations of the sensorimotor interaction logically excluded any locomotion.

Our proposal followed three main objectives. First, it intended to provide a unified and (more) definitive basis for the previous iterations we built it upon; it especially had to be at least as generic as to subsume them. To this end, we proved that it led to statements (and thus results) equivalent to that found in V. Marcel's works, which previously strived to maintain extreme genericity. Second, it attempted to transparently capture and reproduce the key concepts exposed in (J K O'Regan and Noë, 2001). Our framework reflected this consideration in its definitions, which emphasize changes in the sensorimotor *flow* (e.g. motor actions and sensory transitions) as well as on separating the basis of subjective experience from its "hardware" implementation or "neural correlates" (e.g. sensors and their internal semantics). Finally, it had to address the limitations we outlined, which presupposes that we avoided their formal causes. The discussion of these causes made this point straightforward: our introduction of motor actions and adjustment of the formalisation of the kinesthetic sense allowed us to add mobile and dynamically-controlled agents to the scope of our models. Moreover, we introduced a formal account of the longstanding notion of knowledge of the agent, based on its compositional structure. It allowed us to formalise the constructivist aspect of SMCT by which agents gradually accrue perceptive capabilities. While it has been a subject of interest in related works, what

knowledge the agent (*a priori*) possesses, how, and why it does so was generally kept informal. As a matter of fact, to our knowledge it is the first time that an explicit characterization of this property is provided.

Chapter 3 followed our first exploitation of the expressive capabilities afforded by our framework. In it, we studied the shifts of receptive fields of an agent's sensors as it moves in its environment. We established in a theoretical exposition that such shifts appeared at the sensory level as permutations of the sensory array. These permutations ideally mirror the sequences of conservative actions that cause them; we further proved that generalizing to inverse semigroups of partial permutations transparently accounts for the non-ideal case. Importantly these permutations are invariant for a given agent, therefore their occurring at the sensory level it can naively access means they must form predictable patterns. Besides, we showed in great generality that predictions of the effects of actions must necessarily carry the algebraic composition structure of their corresponding actions. Hence we demonstrated that a naive agent can recognize the sensorimotor contingency of sensory permutations associated with shifts of receptive fields to probe the combinatorial structure of its motor actions. In particular, it provides agents with an alternative characterization of their compensable displacements. As a result, following the argument made by H. Poincaré (and developed since by D. Philipona, A. Terekhov and A. Laflaquière) it lends knowledge about the geometry of their ambient space to agents that learn the associated sensorimotor predictions.

Chapter 4 had us take another direction and define an internal notion of similarity between sensorimotor experiences. We discussed how such ideas are generally formalised using spatial (i.e. geometrical, or more generally topological) means. We found that topology could not entirely account for the approximate judgements we aimed to produce. Nonetheless, we pursued its insight on the relation between similarity of experiences and the agent's experienced continuity of its sensorimotor flow in the more favourable metric setting. This particular insight essentially resides in an inversion of the intuitive viewpoint: there, continuity goes from an empirical assessment made of one's experience to an axiom dictating the structure of its content. Accordingly, it resulted in a sensory metric whereby sensations are deemed close whenever they occur at closely points in the sensorimotor experience of the agent. Our experiment showed that this construction resulted in an agent's internal assessment of similarities between colors remarkably coherent with

ours. Moreover, it enabled the agent to determine that its typical sensorimotor experiences were continuous. We therefore demonstrated how a naive agent might infer both similarity between sensations and continuity of one's experience from the statistical observation of its sensorimotor transitions. What is more, the approximate judgements afforded by a gradual similarity (as opposed to the categorical one of binary equality comparison) pave the way for agents to face realistic and noisy settings. In fact, the same experiment established that the metric construction was robust to the introduction of some measure of noise. As a result this similarity might improve the agent's cognitive functions with scalability and robustness properties, thus addressing a most urgent issue of our model.

Using the techniques described in this work, a naive agent can recognize two new sorts of sensorimotor contingencies. These contingencies appear as regularities in its (strongly) uninterpreted sensorimotor flow as per the specifications of the theory. Besides, they correspond to some important detail of the sensorimotor interaction so that their mastery by the agent is precisely it perceiving said detail. As important and perhaps more novel, we argue, is the way these statements manifest in our formalisation. Indeed the main results of Chapters 3 and 4 separate their answers depending on which of 3 angles we choose. Our characterization about *a priori* knowledge of the agent establishes *how* it can recognize the sensorimotor contingencies we considered, that is generally by monitoring the successive values output by its sensels. Our external description of the system elucidates *why* these specific sensel permutations (respectively typical sensation transitions) occur: it is the internal image of the shifts of receptive fields (resp. of the local spatiotemporal coherence) that occur in the embodied interaction. On the contrary, the answer as to the *what* is more partial. Indeed we never claimed to determine precisely what image of the contingencies and form thereof the agent would capture; we later observed that both the sensel permutation and sensory transition matrices needed not be understandable to us. What mattered was not these internal subsymbolic images *per se* but the structures they formed (e.g. the algebraic composition structure for sensory predictions or the metric spatial structure for sensory similarity) and the relations between internal structures and that of the interaction. Therefore, we followed a path between the total intelligibility of symbolic cognition on the one hand and the "black box" paradigms such as deep learning on the other one. Accordingly, we can hope to avoid the classical pitfalls of symbolic approaches

while maintaining (internal and external) explainability.

In addition to the various developments we outlined in each of the parts, several lines that would use elaboration emerge over the entire work. First, it should be further detailed to what extent our proposed framework can effectively and efficiently account for other SMCT works. Indeed we only established that fact in a sufficient capacity relative to some select contributions; it remains to conduct the same examination for many others. Given the variety of backgrounds and practices they represent, this would likely prove a significant endeavour. We can jot a second axis on a more practical note. Indeed, although some strides have been taken to consider more complex and realistic situations (see for one the whole of Chapter 4), attentive readers will surely have noted that those described in this work are still remarkably simplified. A natural extension would then be to try and adapt both the formal framework of Chapter 2 and the developments of Chapters 3 and 4 to more complex setups such as physical robots. In particular, both an extension of the formalism to dynamically-controlled agents (which we briefly sketched) and a practical implementation of the determination of the sensory similarity experiment for complex agents would be of great interest in pursuing this goal.

Bibliography

- Asada, M. et al. (May 2009). "Cognitive Developmental Robotics: A Survey". In: *IEEE Transactions on Autonomous Mental Development* 1.1, pp. 12–34. ISSN: 1943-0604, 1943-0612. DOI: [10.1109/TAMD.2009.2021702](https://doi.org/10.1109/TAMD.2009.2021702). URL: <http://ieeexplore.ieee.org/document/4895715/> (cit. on p. 6).
- Axler, Sheldon (2014). *Linear Algebra Done Right*. New York: Springer. ISBN: 978-3-319-11079-0 (cit. on p. 113).
- Bach-y-Rita, P. (1972). *Brain Mechanisms in Sensory Substitution*. Acad. Press (cit. on pp. 7, 43).
- Bach-y-Rita, P. et al. (Mar. 1969). "Vision Substitution by Tactile Image Projection". In: *Nature* 221.5184, pp. 963–964 (cit. on p. 7).
- Bach-y-Rita, Paul (June 1, 2005). "Emerging Concepts of Brain Function". In: *Journal of Integrative Neuroscience* 04.02, pp. 183–205. ISSN: 0219-6352. DOI: [10.1142/S0219635205000768](https://doi.org/10.1142/S0219635205000768). URL: <https://www.worldscientific.com/doi/abs/10.1142/S0219635205000768> (cit. on p. 7).
- Bach-y-Rita, Paul et al. (2005). "Late Human Brain Plasticity: Vestibular Substitution with a Tongue BrainPort Human-Machine Interface". In: *Intellectica. Revue de l'Association pour la Recherche Cognitive* 40.1, pp. 115–122. ISSN: 0769-4113. DOI: [10.3406/intel.2005.1362](https://doi.org/10.3406/intel.2005.1362). URL: https://www.persee.fr/doc/intel_0769-4113_2005_num_40_1_1362 (cit. on p. 8).
- Balasubramanian, M. (Jan. 4, 2002). "The Isomap Algorithm and Topological Stability". In: *Science* 295.5552, 7a–7. ISSN: 00368075, 10959203. DOI: [10.1126/science.295.5552.7a](https://doi.org/10.1126/science.295.5552.7a). URL: <https://www.sciencemag.org/lookup/doi/10.1126/science.295.5552.7a> (cit. on p. 159).
- Belinfante, J. G., B. Kolman, and H. A. Smith (Jan. 1966). "An Introduction to Lie Groups and Lie Algebras, with Applications on JSTOR". In: *SIAM Review* 8.1, pp. 11–46. JSTOR: 2028170 (cit. on p. 42).
- Berger, Nicolas (2009). "Arithmétique ordinale et cardinale et hiérarchies sur les ensembles". Lausanne: UNIL-HEC. 46 pp. URL: <http://www.hec.unil.ch/logique/recherche/berger.pdf> (cit. on p. 28).

- Bhanpuri, Nasir H., Allison M. Okamura, and Amy J. Bastian (2013). "Predictive Modeling by the Cerebellum Improves Proprioception". In: *Journal of Neuroscience* 33.36, pp. 14301–14306. ISSN: 0270-6474. DOI: [10.1523/JNEUROSCI.0784-13.2013](https://doi.org/10.1523/JNEUROSCI.0784-13.2013). URL: <https://www.jneurosci.org/content/33/36/14301> (cit. on p. 55).
- Billing, Erik, Robert Lowe, and Yulia Sandamirskaya (Oct. 2015). "Simultaneous Planning and Action: Neural-Dynamic Sequencing of Elementary Behaviors in Robot Navigation". In: *Adaptive Behavior* 23.5, pp. 243–264. ISSN: 1059-7123, 1741-2633. DOI: [10.1177/1059712315601188](https://doi.org/10.1177/1059712315601188). URL: <http://journals.sagepub.com/doi/10.1177/1059712315601188> (cit. on p. 39).
- Bohg, Jeannette et al. (Dec. 2017). "Interactive Perception: Leveraging Action in Perception and Perception in Action". In: *IEEE Transactions on Robotics* 33.6, pp. 1273–1291. ISSN: 1552-3098, 1941-0468. DOI: [10.1109/RO.2017.2721939](https://doi.org/10.1109/RO.2017.2721939). arXiv: [1604.03670](https://arxiv.org/abs/1604.03670). URL: <http://arxiv.org/abs/1604.03670> (visited on 11/03/2020) (cit. on p. 9).
- Borg, Ingwer and Patrick J. F. Groenen (1997). *Modern Multidimensional Scaling: Theory and Applications*. Springer Series in Statistics. New York: Springer. 471 pp. ISBN: 978-0-387-94845-4 (cit. on p. 158).
- Botvinick, Matthew M. (May 1, 2008). "Hierarchical Models of Behavior and Prefrontal Function". In: *Trends in Cognitive Sciences* 12.5, pp. 201–208. ISSN: 1364-6613. DOI: [10.1016/j.tics.2008.02.009](https://doi.org/10.1016/j.tics.2008.02.009). URL: <https://www.sciencedirect.com/science/article/pii/S1364661308000880> (cit. on p. 39).
- Bouman, Amanda et al. (Nov. 30, 2020). *Autonomous Spot: Long-Range Autonomous Exploration of Extreme Environments with Legged Locomotion*. arXiv: [2010.09259](https://arxiv.org/abs/2010.09259) [cs]. URL: <http://arxiv.org/abs/2010.09259> (cit. on p. 3).
- Boyd, Zachary M. et al. (June 25, 2020). *A Metric on Directed Graphs and Markov Chains Based on Hitting Probabilities*. arXiv: [2006.14482](https://arxiv.org/abs/2006.14482) [cs, math, stat]. URL: <http://arxiv.org/abs/2006.14482> (visited on 11/05/2020) (cit. on p. 176).
- Braitenberg, Valentino (1986). *Vehicles: Experiments in Synthetic Psychology*. MIT Press. 174 pp. ISBN: 978-0-262-52112-3. Google Books: [7KkUAT_q_sQC](https://books.google.com/books?id=7KkUAT_q_sQC) (cit. on p. 43).
- Bridgeman, Bruce (July 2007). "Efference Copy and Its Limitations". In: *Computers in Biology and Medicine* 37.7, pp. 924–929. ISSN: 00104825. DOI: [10.1016/j.cbm.2007.05.009](https://doi.org/10.1016/j.cbm.2007.05.009) (cit. on p. 176).

- 1016/j.combiomed.2006.07.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010482506001211> (cit. on p. 55).
- Brooks, Rodney A. (June 1, 1990). "Elephants Don't Play Chess". In: *Robotics and Autonomous Systems*. Designing Autonomous Agents 6.1, pp. 3–15. ISSN: 0921-8890. DOI: 10.1016/S0921-8890(05)80025-9. URL: <https://www.sciencedirect.com/science/article/pii/S0921889005800259> (cit. on pp. 4, 5).
- (Aug. 1991a). "Intelligence Without Reason". In: *Proceedings of the Twelfth International Conference on Artificial Intelligence: Darling Harbour, Sydney, Australia, 24 - 30 August 1991*. IJCAI. Ed. by John Mylopoulos and Ray Reiter. Sydney, Australia: Morgan Kaufmann, pp. 569–595. ISBN: 1-55860-160-0 (cit. on p. 4).
- (Jan. 1991b). "Intelligence without Representation". In: *Artificial Intelligence* 47.1-3, pp. 139–159. ISSN: 00043702. DOI: 10.1016/0004-3702(91)90053-M. URL: <https://linkinghub.elsevier.com/retrieve/pii/000437029190053M> (cit. on p. 4).
- (1999). *Cambrian Intelligence: The Early History of the New AI*. A Bradford Book. Cambridge, Mass.: MIT Press. 199 pp. ISBN: 978-0-262-02468-6 (cit. on p. 3).
- (June 1, 2008). *I, Rodney Brooks, Am a Robot*. IEEE Spectrum. URL: <https://spectrum.ieee.org/i-rodney-brooks-am-a-robot> (cit. on p. 3).
- Buhrmann, Thomas, Ezequiel Di Paolo, and Xabier Barandiaran (2013). "A Dynamical Systems Account of Sensorimotor Contingencies". In: *Frontiers in Psychology* 4, p. 285. ISSN: 1664-1078. DOI: 10.3389/fpsyg.2013.00285. URL: <https://www.frontiersin.org/article/10.3389/fpsyg.2013.00285> (cit. on p. 12).
- Burr, Christopher (Mar. 8, 2017). "Embodied Decisions and the Predictive Brain". In: DOI: 10.15502/9783958573086 (cit. on pp. 2, 74).
- Cayley, Arthur (1878). "Desiderata and Suggestions: No. 2. The Theory of Groups: Graphical Representation". In: *American Journal of Mathematics* 1.2, p. 174. ISSN: 00029327. DOI: 10.2307/2369306. JSTOR: 2369306 (cit. on p. 108).
- Censi, Andrea (June 27, 2012). "Bootstrapping Vehicles: A Formal Approach to Unsupervised Sensorimotor Learning Based on Invariance". PhD thesis. Pasadena, CA: California Institute of Technology. 385 pp. URL: https://censi.science/pub/research/preprints/bootstrapping_vehicles-1.3-a4.pdf (cit. on pp. 12, 19, 53, 132).

- Censi, Andrea and Richard M. Murray (May 2012). "Learning Diffeomorphism Models of Robotic Sensorimotor Cascades". In: *2012 IEEE International Conference on Robotics and Automation*. 2012 IEEE International Conference on Robotics and Automation (ICRA). St Paul, MN, USA: IEEE, pp. 3657–3664. ISBN: 978-1-4673-1405-3. DOI: [10.1109/ICRA.2012.6225318](https://doi.org/10.1109/ICRA.2012.6225318). URL: <http://ieeexplore.ieee.org/document/6225318/> (cit. on p. 6).
- Censi, Andrea and Davide Scaramuzza (2013). "Calibration by Correlation Using Metric Embedding from Nonmetric Similarities". In: *IEEE transactions on pattern analysis and machine intelligence* 35.10, pp. 2357–2370 (cit. on p. 176).
- Chalmers, David J. (1995). "Facing Up to the Problem of Consciousness". In: *Journal of Consciousness Studies* 2.3, p. 27 (cit. on p. 10).
- Chebotarev, Pavel and Elena Deza (Mar. 2020). "Hitting Time Quasi-Metric and Its Forest Representation". In: *Optimization Letters* 14.2, pp. 291–307. ISSN: 1862-4472, 1862-4480. DOI: [10.1007/s11590-018-1314-2](https://doi.org/10.1007/s11590-018-1314-2). arXiv: [1801.00413](https://arxiv.org/abs/1801.00413). URL: <http://arxiv.org/abs/1801.00413> (visited on 11/05/2020) (cit. on p. 176).
- Choo, Jaegul et al. (Apr. 29, 2010). "P-ISOMAP: An Efficient Parametric Update for ISOMAP for Visual Analytics". In: *Proceedings of the 2010 SIAM International Conference on Data Mining*. Proceedings of the 2010 SIAM International Conference on Data Mining. Society for Industrial and Applied Mathematics, pp. 502–513. ISBN: 978-0-89871-703-7. DOI: [10.1137/1.9781611972801.44](https://doi.org/10.1137/1.9781611972801.44). URL: <https://epubs.siam.org/doi/10.1137/1.9781611972801.44> (cit. on p. 159).
- Cisek, Paul (Sept. 29, 2007). "Cortical Mechanisms of Action Selection: The Affordance Competition Hypothesis". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 362.1485, pp. 1585–1599. ISSN: 0962-8436. DOI: [10.1098/rstb.2007.2054](https://doi.org/10.1098/rstb.2007.2054). pmid: 17428779. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2440773/> (cit. on p. 4).
- Cisek, Paul and Alexandre Pastor-Bernier (Nov. 5, 2014). "On the Challenges and Mechanisms of Embodied Decisions". In: *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 369.1655, p. 20130479. ISSN: 1471-2970. DOI: [10.1098/rstb.2013.0479](https://doi.org/10.1098/rstb.2013.0479). pmid: 25267821 (cit. on p. 4).
- Coifman, Ronald R. and Stéphane Lafon (July 1, 2006). "Diffusion Maps". In: *Applied and Computational Harmonic Analysis*. Special Issue: Diffusion Maps and Wavelets 21.1, pp. 5–30. ISSN: 1063-5203. DOI: [10.1016/j.acha](https://doi.org/10.1016/j.acha).

- 2006.04.006. URL: <https://www.sciencedirect.com/science/article/pii/S1063520306000546> (visited on 02/07/2021) (cit. on p. 176).
- Dahl, Matias (July 12, 2006). "An Brief Introduction to Finsler Geometry". In: p. 39 (cit. on p. 147).
- Dainton, Barry (2018). "Temporal Consciousness". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2018. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/win2018/entries/consciousness-temporal/> (cit. on p. 131).
- Dekking, Michel, ed. (2005). *A Modern Introduction to Probability and Statistics: Understanding Why and How*. Springer Texts in Statistics. London: Springer. 486 pp. ISBN: 978-1-85233-896-1 (cit. on p. 155).
- Dijkstra, E. W. (1959). "A Note on Two Problems in Connexion with Graphs". In: *Numerische Mathematik* 1.1, pp. 269–271. ISSN: 0945-3245. DOI: [10.1007/BF01386390](https://doi.org/10.1007/BF01386390) (cit. on p. 110).
- Doncieux, Stephane et al. (2015). "Evolutionary Robotics: What, Why, and Where To". In: *Frontiers in Robotics and AI* 2, p. 4. ISSN: 2296-9144. DOI: [10.3389/frobt.2015.00004](https://doi.org/10.3389/frobt.2015.00004). URL: <https://www.frontiersin.org/article/10.3389/frobt.2015.00004> (cit. on p. 7).
- Dreyfus, Hubert L. (Jan. 1, 1965). *Alchemy and Artificial Intelligence*. RAND Corporation. URL: <https://www.rand.org/pubs/papers/P3244.html> (cit. on p. 4).
- (1986). *What Computers Can't Do: The Limits of Artificial Intelligence*. Rev. ed. 7th print. Perennial Library. New York: Harper & Row. 354 pp. ISBN: 978-0-06-090613-9 (cit. on pp. 4, 5).
- Dummit, David Steven and Richard M. Foote (2009). *Abstract Algebra*. 3. ed., [Nachdr.] New York: Wiley. 932 pp. ISBN: 978-0-471-45234-8 (cit. on p. 94).
- Exel, Ruy (1998). "Partial Actions of Groups and Actions of Inverse Semigroups". In: *Proceedings of the American Mathematical Society* 126.12, pp. 3481–3494. ISSN: 0002-9939, 1088-6826. DOI: [10.1090/S0002-9939-98-04575-4](https://doi.org/10.1090/S0002-9939-98-04575-4). URL: <http://www.ams.org/proc/1998-126-12/S0002-9939-98-04575-4/> (cit. on p. 85).
- Farber, Michael, Serge Tabachnikov, and Sergey Yuzvinsky (Jan. 1, 2003). "Topological Robotics: Motion Planning in Projective Spaces". In: *International Mathematics Research Notices* 2003.34, p. 18. DOI: [10.1155/S1073792803210035](https://doi.org/10.1155/S1073792803210035) (cit. on p. 16).

- Fefferman, Charles, Sanjoy Mitter, and Hariharan Narayanan (Dec. 19, 2013). *Testing the Manifold Hypothesis*. arXiv: 1310.0425 [math, stat]. URL: <http://arxiv.org/abs/1310.0425> (cit. on p. 156).
- Festinger, L. and J. D. Holtzman (Nov. 1978). "Retinal Image Smear as a Source of Information about Magnitude of Eye Movement". In: *Journal of Experimental Psychology. Human Perception and Performance* 4.4, pp. 573–585. ISSN: 0096-1523. DOI: 10.1037//0096-1523.4.4.573. pmid: 722249 (cit. on p. 64).
- Flanagan, J. Randall et al. (Oct. 15, 1999). "Composition and Decomposition of Internal Models in Motor Learning under Altered Kinematic and Dynamic Environments". In: *The Journal of Neuroscience* 19.20, RC34–RC34. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.19-20-j0005.1999. URL: <https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.19-20-j0005.1999> (cit. on p. 39).
- Floyd, Robert W. (June 1962). "Algorithm 97: Shortest Path". In: *Communications of the ACM* 5.6, p. 345. ISSN: 0001-0782, 1557-7317. DOI: 10.1145/367766.368168. URL: <https://dl.acm.org/doi/10.1145/367766.368168> (cit. on p. 155).
- Fruchterman, Thomas M. J. and Edward M. Reingold (Nov. 1991). "Graph Drawing by Force-Directed Placement". In: *Software: Practice and Experience* 21.11, pp. 1129–1164. ISSN: 00380644, 1097024X. DOI: 10.1002/spe.4380211102. URL: <https://onlinelibrary.wiley.com/doi/10.1002/spe.4380211102> (cit. on p. 115).
- Fugate, Jennifer M. B. (Jan. 2013). "Categorical Perception for Emotional Faces". In: *Emotion review : journal of the International Society for Research on Emotion* 5.1, pp. 84–89. ISSN: 1754-0739. DOI: 10.1177/1754073912451350. pmid: 25525458. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4267261/> (cit. on p. 131).
- Gapenne, Olivier (June 12, 2014). "The Co-Constitution of the Self and the World: Action and Proprioceptive Coupling". In: *Frontiers in Psychology* 5. ISSN: 1664-1078. DOI: 10.3389/fpsyg.2014.00594. pmid: 24971073. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4054590/> (cit. on p. 55).
- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Boston: Houghton (cit. on p. 74).
- (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton (cit. on p. 63).

- González-Grandón, Ximena, Andrea Falcón-Cortés, and Gabriel Ramos-Fernández (2021). "Proprioception in Action: A Matter of Ecological and Social Interaction". In: *Frontiers in Psychology* 11, p. 3372. ISSN: 1664-1078. DOI: [10.3389/fpsyg.2020.569403](https://doi.org/10.3389/fpsyg.2020.569403). URL: <https://www.frontiersin.org/article/10.3389/fpsyg.2020.569403> (cit. on p. 55).
- Grabiner, Judith V. (Oct. 1, 1986). "Computers and the Nature of Man: A Historian's Perspective on Controversies About Artificial Intelligence". In: *Bulletin of the American Mathematical Society* 15.2, pp. 113–127. ISSN: 0273-0979. DOI: [10.1090/S0273-0979-1986-15461-3](https://doi.org/10.1090/S0273-0979-1986-15461-3). URL: <http://www.ams.org/journal-getitem?pii=S0273-0979-1986-15461-3> (cit. on p. 4).
- Gries, David and Fred B. Schneider (Oct. 22, 1993). *A Logical Approach to Discrete Math*. Springer Science & Business Media. 536 pp. ISBN: 978-0-387-94115-8. Google Books: [ZWTDQ6H6gsUC](https://books.google.com/books?id=ZWTDQ6H6gsUC) (cit. on p. 58).
- Gromov, Mikhael (1999). *Metric Structures for Riemannian and Non-Riemannian Spaces*. Progress in Mathematics v. 152. Boston: Birkhäuser. 585 pp. ISBN: 978-0-8176-3898-6 (cit. on p. 157).
- Grossmann, E., J. Gaspar, and Francesco Orabona (2010). "Discrete Camera Calibration from Pixel Streams". In: *Comput. Vis. Image Underst.* DOI: [10.1016/j.cviu.2009.03.009](https://doi.org/10.1016/j.cviu.2009.03.009) (cit. on p. 176).
- Harnad, Stevan (June 1990). "The Symbol Grounding Problem". Version 1. In: *Physica D: Nonlinear Phenomena* 42.1-3, pp. 335–346. ISSN: 01672789. DOI: [10.1016/0167-2789\(90\)90087-6](https://doi.org/10.1016/0167-2789(90)90087-6). arXiv: [cs/9906002](https://arxiv.org/abs/cs/9906002). URL: <http://arxiv.org/abs/cs/9906002> (cit. on p. 5).
- Hartline, H. K. (Jan. 31, 1938). "The Response of Single Optic Nerve Fibers of the Vertebrate Eye to Illumination of the Retina". In: *American Journal of Physiology-Legacy Content* 121.2, pp. 400–415. ISSN: 0002-9513. DOI: [10.1152/ajplegacy.1938.121.2.400](https://doi.org/10.1152/ajplegacy.1938.121.2.400). URL: <https://journals.physiology.org/doi/abs/10.1152/ajplegacy.1938.121.2.400> (cit. on p. 45).
- Held, Richard and Alan Hein (Nov. 1963). "Movement-Produced Stimulation in the Development of Visually Guided Behavior". In: *Journal of comparative and physiological psychology* 56, pp. 872–6. DOI: [10.1037/h0040546](https://doi.org/10.1037/h0040546) (cit. on p. 7).
- Helmholtz, Hermann von (Sept. 4, 2013). *Treatise on Physiological Optics*. Courier Corporation. 1257 pp. ISBN: 978-0-486-17470-9. Google Books: [cSjEAgAAQBAJ](https://books.google.com/books?id=cSjEAgAAQBAJ) (cit. on p. 74).
- Herwig, Arvid (Dec. 1, 2015). "Transsaccadic Integration and Perceptual Continuity". In: *Journal of Vision* 15.16, pp. 7–7. ISSN: 1534-7362. DOI: [10.5796/journalofvision.15.16.7](https://doi.org/10.5796/journalofvision.15.16.7)

- 1167/15.16.7. URL: <https://jov.arvojournals.org/article.aspx?articleid=2475388> (cit. on p. 131).
- Hoffmann, Matej (2014). "Minimally Cognitive Robotics: Body Schema, Forward Models, and Sensorimotor Contingencies in a Quadruped Machine". In: *Contemporary Sensorimotor Theory*. Ed. by John Mark Bishop and Andrew Owen Martin. Cham: Springer International Publishing, pp. 209–233. ISBN: 978-3-319-05107-9 (cit. on p. 11).
- Hrbacek, Karel and Thomas J. Jech (1999). *Introduction to Set Theory*. 3rd ed., rev. and expanded. Monographs and Textbooks in Pure and Applied Mathematics 220. New York: M. Dekker. 291 pp. ISBN: 978-0-8247-7915-3 (cit. on p. 34).
- Hurley, S. L. (2002). *Consciousness in Action*. Consciousness in Action. Cambridge, MA, US: Harvard University Press, pp. xii, 506. xii, 506. ISBN: 978-0-674-16420-8 (cit. on p. 2).
- Hurley, Susan (Jan. 1, 2012). "Varieties of Externalism". In: *The Extended Mind*, pp. 101–153. ISSN: 9780262014038. DOI: [10.7551/mitpress/9780262014038.003.0006](https://doi.org/10.7551/mitpress/9780262014038.003.0006) (cit. on p. 123).
- Imamizu, Hiroshi (May 2010). "Prediction of Sensorimotor Feedback from the Efference Copy of Motor Commands: A Review of Behavioral and Functional Neuroimaging Studies". In: *Japanese Psychological Research* 52, pp. 107–120. DOI: [10.1111/j.1468-5884.2010.00428.x](https://doi.org/10.1111/j.1468-5884.2010.00428.x) (cit. on p. 55).
- Izenman, Alan Julian (2012). "Introduction to Manifold Learning". In: *WIREs Computational Statistics* 4.5, pp. 439–446. ISSN: 1939-0068. DOI: [10.1002/wics.1222](https://doi.org/10.1002/wics.1222). URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/wics.1222> (cit. on p. 156).
- James, I. M. (1987). *Topological and Uniform Spaces*. Red. by J. H. Ewing, F. W. Gehring, and P. R. Halmos. Undergraduate Texts in Mathematics. New York, NY: Springer New York. ISBN: 978-1-4612-9128-2. DOI: [10.1007/978-1-4612-4716-6](https://doi.org/10.1007/978-1-4612-4716-6). URL: <http://link.springer.com/10.1007/978-1-4612-4716-6> (cit. on p. 143).
- Jaynes, E. T and G. Larry Bretthorst (2003). *Probability Theory the Logic of Science*. Cambridge, UK; New York, NY: Cambridge University Press. ISBN: 978-1-107-12745-6 (cit. on p. 145).
- Jaynes, E. T. (Dec. 1, 1973). "The Well-Posed Problem". In: *Foundations of Physics* 3.4, pp. 477–492. ISSN: 1572-9516. DOI: [10.1007/BF00709116](https://doi.org/10.1007/BF00709116). URL: <https://doi.org/10.1007/BF00709116> (cit. on p. 145).

- Jeffries, Ailsa, Nathaniel Killian, and John Pezaris (Jan. 1, 2013). "Mapping the Primate Lateral Geniculate Nucleus: A Review of Experiments and Methods". In: *Journal of Physiology-Paris* 108. DOI: [10.1016/j.jphysparis.2013.10.001](https://doi.org/10.1016/j.jphysparis.2013.10.001) (cit. on p. 45).
- Jonschkowski, Rico and Oliver Brock (Oct. 2015). "Learning State Representations with Robotic Priors". In: *Autonomous Robots* 39.3, pp. 407–428. ISSN: 0929-5593, 1573-7527. DOI: [10.1007/s10514-015-9459-7](https://doi.org/10.1007/s10514-015-9459-7). URL: <http://link.springer.com/10.1007/s10514-015-9459-7> (cit. on p. 6).
- Kellendonk, J. and Mark V. Lawson (Feb. 1, 2004). "Partial Actions of Groups". In: *International Journal of Algebra and Computation* 14.01, pp. 87–114. ISSN: 0218-1967. DOI: [10.1142/S0218196704001657](https://doi.org/10.1142/S0218196704001657). URL: <https://www.worldscientific.com/doi/abs/10.1142/S0218196704001657> (cit. on p. 85).
- Kleene, S. C. (Dec. 1, 1936). "General Recursive Functions of Natural Numbers". In: *Mathematische Annalen* 112.1, pp. 727–742. ISSN: 1432-1807. DOI: [10.1007/BF01565439](https://doi.org/10.1007/BF01565439). URL: <https://doi.org/10.1007/BF01565439> (cit. on p. 57).
- Klein, Felix (Mar. 1, 1893). "Vergleichende Betrachtungen über neuere geometrische Forschungen". In: *Mathematische Annalen* 43.1, pp. 63–100. ISSN: 1432-1807. DOI: [10.1007/BF01446615](https://doi.org/10.1007/BF01446615). URL: <https://doi.org/10.1007/BF01446615> (cit. on p. 29).
- Knuth, Donald Ervin (1981). *The Art of Computer Programming*. 2nd ed. Vol. 2. 4 vols. Reading, Mass: Addison-Wesley. 3 pp. ISBN: 978-0-201-89683-1 (cit. on p. 90).
- Kobourov, Stephen G. (Jan. 14, 2012). *Spring Embedders and Force Directed Graph Drawing Algorithms*. arXiv: [1201.3011](https://arxiv.org/abs/1201.3011) [cs]. URL: <http://arxiv.org/abs/1201.3011> (cit. on p. 116).
- König, Sabine U. et al. (Dec. 13, 2016). "Learning New Sensorimotor Contingencies: Effects of Long-Term Use of Sensory Augmentation on the Brain and Conscious Perception". In: *PLOS ONE* 11.12, e0166647. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0166647](https://doi.org/10.1371/journal.pone.0166647). URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0166647> (cit. on p. 12).
- Kuipers, Benjamin (May 2000). "The Spatial Semantic Hierarchy". In: *Artificial Intelligence* 119.1-2, pp. 191–233. ISSN: 00043702. DOI: [10.1016/S0004-3702\(00\)00017-5](https://doi.org/10.1016/S0004-3702(00)00017-5). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0004370200000175> (visited on 03/17/2021) (cit. on p. 6).

- Kuipers, Benjamin (2008). "An Intellectual History of the Spatial Semantic Hierarchy". In: *Robotics and Cognitive Approaches to Spatial Mapping*. Ed. by Margaret E. Jefferies and Wai-Kiang Yeap. Vol. 38. Springer Tracts in Advanced Robotics. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 243–264. ISBN: 978-3-540-75386-5. DOI: [10.1007/978-3-540-75388-9_15](https://doi.org/10.1007/978-3-540-75388-9_15). URL: http://link.springer.com/10.1007/978-3-540-75388-9_15 (visited on 03/17/2021) (cit. on p. 6).
- Kuipers, Benjamin and Yung-Tai Byun (Nov. 1991). "A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations". In: *Robotics and Autonomous Systems* 8.1-2, pp. 47–63. ISSN: 09218890. DOI: [10.1016/0921-8890\(91\)90014-C](https://doi.org/10.1016/0921-8890(91)90014-C). URL: <https://linkinghub.elsevier.com/retrieve/pii/092188909190014C> (visited on 03/17/2021) (cit. on p. 6).
- Kuipers, Benjamin J. et al. (June 2006). "Bootstrap Learning of Foundational Representations". In: *Connection Science* 18.2, pp. 145–158. ISSN: 0954-0091, 1360-0494. DOI: [10.1080/09540090600768484](https://doi.org/10.1080/09540090600768484). URL: <http://www.tandfonline.com/doi/abs/10.1080/09540090600768484> (cit. on pp. 6, 53).
- Laflaquière, A et al. (Sept. 2015). "Learning Agents Spatial Configuration from Sensorimotor Invariants". In: *Robotics and Autonomous Systems* 71, pp. 49–59 (cit. on pp. 21, 133).
- Laflaquière, Alban (July 19, 2013). "Approche sensorimotrice de la perception de l'espace pour la robotique autonome". PhD thesis. Université Pierre et Marie Curie - Paris VI. URL: <https://tel.archives-ouvertes.fr/tel-00865091> (cit. on pp. 12, 21).
- (Dec. 2017). "Grounding the Experience of a Visual Field through Sensorimotor Contingencies". In: *Neurocomputing* 268, pp. 142–152. ISSN: 09252312. DOI: [10.1016/j.neucom.2016.11.085](https://doi.org/10.1016/j.neucom.2016.11.085). arXiv: 1810.01871. URL: <http://arxiv.org/abs/1810.01871> (cit. on pp. 21, 25, 26, 67, 132).
- (Oct. 29, 2020). *Emergence of Spatial Coordinates via Exploration*. arXiv: 2010.15469 [cs]. URL: <http://arxiv.org/abs/2010.15469> (cit. on p. 132).
- Laflaquière, Alban and Verena V. Hafner (June 3, 2019). *Self-Supervised Body Image Acquisition Using a Deep Neural Network for Sensorimotor Prediction*. arXiv: 1906.00825 [cs, stat]. URL: <http://arxiv.org/abs/1906.00825> (cit. on pp. 21, 127).
- Laflaquière, Alban, J. Kevin O'Regan, et al. (Sept. 2015). "Learning Agent's Spatial Configuration from Sensorimotor Invariants". In: *Robotics and Autonomous Systems* 71, pp. 49–59. ISSN: 09218890. DOI: [10.1016/j.robot.](https://doi.org/10.1016/j.robot.)

- 2015.01.003. arXiv: 1810.01872. URL: <http://arxiv.org/abs/1810.01872> (visited on 11/02/2020) (cit. on p. 12).
- Laflaquière, Alban and Michael Garcia Ortiz (2019). “Unsupervised Emergence of Egocentric Spatial Structure from Sensorimotor Prediction”. In: *NeurIPS*, pp. 7158–7168 (cit. on pp. 17, 21).
- Lawson, Mark V (Apr. 2018). “3. Inverse Semigroups”. In: . *INVERSE SEMI-GROUPS*, p. 25 (cit. on p. 73).
- Le Hir, Nicolas, Olivier Sigaud, and Alban Laflaquière (2018). “Identification of Invariant Sensorimotor Structures as a Prerequisite for the Discovery of Objects”. In: *Frontiers in Robotics and AI* 5, p. 70. ISSN: 2296-9144. DOI: 10.3389/frobt.2018.00070. URL: <https://www.frontiersin.org/articles/10.3389/frobt.2018.00070/full> (cit. on pp. 12, 25, 26, 67, 127).
- Lee, Chiwon et al. (Jan. 2017). “Soft Robot Review”. In: *International Journal of Control, Automation and Systems* 15. DOI: 10.1007/s12555-016-0462-3 (cit. on p. 4).
- Lee, John Aldo et al. (2000). “A robust nonlinear projection method”. In: European Symposium on Artificial Neural Networks (ESANN’2000). URL: <https://dial.uclouvain.be/pr/boreal/object/boreal:111353> (visited on 03/25/2021) (cit. on p. 156).
- Lesort, Timothée et al. (Dec. 2018). “State Representation Learning for Control: An Overview”. In: *Neural Networks* 108, pp. 379–392. ISSN: 08936080. DOI: 10.1016/j.neunet.2018.07.006. arXiv: 1802.04181. URL: <http://arxiv.org/abs/1802.04181> (cit. on p. 75).
- Levine, Joseph (Oct. 1983). “Materialism and Qualia: The Explanatory Gap”. In: *Pacific Philosophical Quarterly* 64.4, pp. 354–361. ISSN: 02790750. DOI: 10.1111/j.1468-0114.1983.tb00207.x. URL: <http://doi.wiley.com/10.1111/j.1468-0114.1983.tb00207.x> (visited on 11/03/2020) (cit. on p. 10).
- Li, Siqi, Hao Zhu, and Xing Tian (Nov. 1, 2020). “Corollary Discharge Versus Efference Copy: Distinct Neural Signals in Speech Preparation Differentially Modulate Auditory Responses”. In: *Cerebral Cortex* 30.11, pp. 5806–5820. ISSN: 1047-3211. DOI: 10.1093/cercor/bhaa154. URL: <https://doi.org/10.1093/cercor/bhaa154> (cit. on p. 55).
- Lighthill, Sir James (July 1973). “Artificial Intelligence: A General Survey”. In: *Artificial Intelligence: A Paper Symposium*. British Science Research

- Council. URL: <https://aitopics.org/doc/classics:D8235CF9/> (cit. on p. 4).
- Lungarella, Max, Fumiya Iida, et al. (2007). "AI in the 21st Century – With Historical Reflections". In: *50 Years of Artificial Intelligence*. Ed. by Max Lungarella et al. Vol. 4850. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 1–8. ISBN: 978-3-540-77295-8. DOI: [10.1007/978-3-540-77296-5_1](https://doi.org/10.1007/978-3-540-77296-5_1). URL: http://link.springer.com/10.1007/978-3-540-77296-5_1 (cit. on p. 4).
- Lungarella, Max, Giorgio Metta, et al. (Dec. 1, 2003). "Developmental Robotics: A Survey". In: *Connection Science* 15.4, pp. 151–190. ISSN: 0954-0091. DOI: [10.1080/09540090310001655110](https://doi.org/10.1080/09540090310001655110). URL: <https://doi.org/10.1080/09540090310001655110> (cit. on p. 6).
- Lutzer, David J. (Feb. 1, 1969). "A Metrization Theorem for Linearly Orderable Spaces". In: *Proceedings of the American Mathematical Society* 22.2, pp. 557–557. ISSN: 0002-9939. DOI: [10.1090/S0002-9939-1969-0248761-1](https://doi.org/10.1090/S0002-9939-1969-0248761-1). URL: <http://www.ams.org/jourcgi/jour-getitem?pii=S0002-9939-1969-0248761-1> (cit. on p. 143).
- Mac Lane, Saunders (1978). *Categories for the Working Mathematician*. 2nd ed. Graduate Texts in Mathematics. New York: Springer-Verlag. ISBN: 978-0-387-98403-2. DOI: [10.1007/978-1-4757-4721-8](https://doi.org/10.1007/978-1-4757-4721-8). URL: <https://www.springer.com/gp/book/9780387984032> (cit. on p. 20).
- Magnus, Wilhelm, Abraham Karrass, and Donald Solitar (2004). *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Mineola, N.Y: Dover Publications. 444 pp. ISBN: 978-0-486-43830-6 (cit. on p. 108).
- Marcel, V., S. Argentieri, and B. Gas (June 2017). "Building a Sensorimotor Representation of a Naive Agent's Tactile Space". In: *IEEE Transactions on Cognitive and Developmental Systems* 9.2, pp. 141–152. ISSN: 2379-8920. DOI: [10.1109/TCDS.2016.2617922](https://doi.org/10.1109/TCDS.2016.2617922) (cit. on pp. 26, 133).
- (2019). "Where Do I Move My Sensors? Emergence of a Topological Representation of Sensors Poses from the Sensorimotor Flow". In: *IEEE Transactions on Cognitive and Developmental Systems*, pp. 1–1. ISSN: 2379-8939. DOI: [10.1109/TCDS.2019.2959915](https://doi.org/10.1109/TCDS.2019.2959915) (cit. on pp. 15, 17, 21, 54, 75).
- Marcel, V., B. Garcia, et al. (2015). "Building the Representation of an Agent Body from Its Sensorimotor Invariants". In: *Workshop on Sensorimotor Contingencies for Robotics, IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'15* (cit. on pp. 26, 132).

- Marcel, Valentin (Feb. 4, 2020). "A Formal Approach to the Emergence of Spatial Representations from Sensorimotor Inputs in Robotics". Sorbonne Université (cit. on pp. 12, 22, 23, 54, 177).
- Marconi, L., R. Naldi, and L. Gentili (2011). "Modelling and Control of a Flying Robot Interacting with the Environment". In: *Autom.* 47, pp. 2571–2583 (cit. on p. 3).
- Maturana, Humberto R. and Francisco J. Varela (1980). *Autopoiesis and Cognition: The Realization of the Living*. Boston Studies in the Philosophy of Science v. 42. Dordrecht, Holland ; Boston: D. Reidel Pub. Co. 141 pp. ISBN: 978-90-277-1015-4 (cit. on p. 5).
- Maye, Alexander and Andreas Engel (May 1, 2016). "The Sensorimotor Account of Sensory Consciousness Implications for Machine Consciousness". In: *Journal of Consciousness Studies* 23.5-6, pp. 177–202 (cit. on p. 51).
- Maye, Alexander and Andreas K Engel (Dec. 1, 2013). "Extending Sensorimotor Contingency Theory: Prediction, Planning, and Action Generation". In: *Adaptive Behavior* 21.6, pp. 423–436. ISSN: 1059-7123. DOI: 10.1177/1059712313497975. URL: <http://40hz.net/maye-2013-adaptive-behavior.pdf> (cit. on pp. 12, 26).
- (May 2011). "A Discrete Computational Model of Sensorimotor Contingencies for Object Perception and Control of Behavior". In: *2011 IEEE International Conference on Robotics and Automation*. 2011 IEEE International Conference on Robotics and Automation (ICRA). Shanghai, China: IEEE, pp. 3810–3815. ISBN: 978-1-61284-386-5. DOI: 10.1109/ICRA.2011.5979919. URL: <http://ieeexplore.ieee.org/document/5979919/> (cit. on pp. 12, 26, 127).
- (2012). "Using Sensorimotor Contingencies for Prediction and Action Planning". In: *From Animals to Animats 12*. Ed. by Tom Ziemke, Christian Balke, and John Hallam. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 106–116. ISBN: 978-3-642-33093-3. DOI: 10.1007/978-3-642-33093-3_11 (cit. on p. 127).
- McCorduck, Pamela (2004). *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 25th anniversary update. Natick, Mass: A.K. Peters. 565 pp. ISBN: 978-1-56881-205-2 (cit. on p. 3).
- McCrea, W. H. and F. J. W. Whipple (1940). "XXII.—Random Paths in Two and Three Dimensions". In: *Proceedings of the Royal Society of Edinburgh* 60.3, pp. 281–298. ISSN: 0370-1646. DOI: 10.1017/S0370164600020265. URL: <https://www.cambridge.org/core/journals/proceedings-of->

- [the-royal-society-of-edinburgh/article/abs/xxiirandom-paths-in-two-and-three-dimensions/F36FD46233C0BD785CCC813687503BE8](https://www.royal-society-of-edinburgh/article/abs/xxiirandom-paths-in-two-and-three-dimensions/F36FD46233C0BD785CCC813687503BE8) (cit. on p. 99).
- Melnik, Andrew et al. (2018). "The World as an External Memory: The Price of Saccades in a Sensorimotor Task". In: *Frontiers in Behavioral Neuroscience* 12, p. 253. ISSN: 1662-5153. DOI: [10.3389/fnbeh.2018.00253](https://doi.org/10.3389/fnbeh.2018.00253). URL: <https://www.frontiersin.org/article/10.3389/fnbeh.2018.00253> (visited on 09/28/2021) (cit. on pp. 10, 12).
- Montone, Guglielmo, J. Kevin O'Regan, and Alexander V. Terekhov (2015). "Unsupervised Model-Free Camera Calibration Algorithm for Robotic Applications". In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3058–3063 (cit. on pp. 25, 64, 67).
- Morgan, Charles (Jan. 2006). "Axiomatic Set Theory - Chapter 4. Cardinal Arithmetic". URL: https://www.ucl.ac.uk/~ucahcjm/ast/ast_notes_4.pdf (cit. on p. 28).
- Morgese, Giorgia, Giovanni Lombardo, and Vilfredo De Pascalis (July 17, 2017). "Localizationism, Antilocalizationism, and the Emergence of the Unitary Construct of Consciousness in Luigi Luciani (1840-1919)". In: *History of Psychology* 20. DOI: [10.1037/hop0000073](https://doi.org/10.1037/hop0000073) (cit. on p. 44).
- Morris, Sidney A (2020). "Topology Without Tears". In: p. 727 (cit. on pp. 133, 135).
- Morse, Anthony F. et al. (Dec. 1, 2011). "The Role of Robotic Modelling in Cognitive Science". In: *New Ideas in Psychology*. Special Issue: Cognitive Robotics and Reevaluation of Piaget Concept of Egocentrism 29.3, pp. 312–324. ISSN: 0732-118X. DOI: [10.1016/j.newideapsych.2011.02.001](https://doi.org/10.1016/j.newideapsych.2011.02.001). URL: <https://www.sciencedirect.com/science/article/pii/S0732118X11000109> (cit. on p. 4).
- Mutambara, A.G.O. and Jonathan Litt (1998). *A Framework for a Supervisory Expert System for Robotic Manipulators with Joint-Position Limits and Joint-Rate Limits*. NASA, Cleveland, OH Research Center. URL: <https://ntrs.nasa.gov/citations/19990018026> (cit. on p. 3).
- Nagel, Thomas (Oct. 1974). "What Is It Like to Be a Bat?" In: *The Philosophical Review* 83.4, p. 435. ISSN: 00318108. DOI: [10.2307/2183914](https://doi.org/10.2307/2183914). JSTOR: [2183914](https://www.jstor.org/stable/2183914) (cit. on p. 10).
- NASA (n.d.). *Curiosity Instrument Summary*. NASA's Mars Exploration Program. URL: <https://mars.nasa.gov/msl/spacecraft/instruments/summary> (cit. on p. 43).

- Nash, John (1956). "The Imbedding Problem for Riemannian Manifolds". In: *Annals of Mathematics* 63.1, pp. 20–63. ISSN: 0003-486X. DOI: [10.2307/1969989](#). JSTOR: [1969989](#) (cit. on p. [156](#)).
- Newell, Allen and H.A. Simon (1963). "GPS, a Program That Simulates Human Thought". In: *Readings in Cognitive Science*. Elsevier, pp. 453–460. ISBN: 978-1-4832-1446-7. DOI: [10.1016/B978-1-4832-1446-7.50040-6](#). URL: <https://linkinghub.elsevier.com/retrieve/pii/B9781483214467500406> (cit. on p. [3](#)).
- Nguyen, Tuan, Sarath Sreedharan, and Subbarao Kambhampati (2017). "Robust Planning with Incomplete Domain Models". In: *Artificial Intelligence* 245, pp. 134–161. ISSN: 0004-3702. DOI: [10.1016/j.artint.2016.12.003](#) (cit. on p. [4](#)).
- Nicolay, Samuel and Laurent Simons (Aug. 28, 2014). "Building Cantor's Bijection". In: (cit. on p. [28](#)).
- Nicolescu, Monica N. and Maja J. Matarić (2002). "A Hierarchical Architecture for Behavior-Based Robots". In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems Part 1 - AAMAS '02*. The First International Joint Conference. Bologna, Italy: ACM Press, p. 227. ISBN: 978-1-58113-480-3. DOI: [10.1145/544741.544798](#). URL: <http://portal.acm.org/citation.cfm?doid=544741.544798> (cit. on p. [39](#)).
- Niziolek, Caroline A., Srikantan S. Nagarajan, and John F. Houde (Oct. 9, 2013). "What Does Motor Efference Copy Represent? Evidence from Speech Production". In: *The Journal of Neuroscience* 33.41, pp. 16110–16116. ISSN: 0270-6474. DOI: [10.1523/JNEUROSCI.2137-13.2013](#). pmid: [24107944](#). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3792453/> (cit. on p. [55](#)).
- Noë, A. (2004). *Action in Perception*. The MIT Press (cit. on pp. [6](#), [10](#)).
- O'Regan, J K and A Noë (Oct. 2001). "A Sensorimotor Account of Vision and Visual Consciousness." In: *The Behavioral and brain sciences* 24.5, 939–973, discussion 973–1031. ISSN: 0140-525X. pmid: [12239892](#) (cit. on pp. [9](#), [10](#), [12](#), [19](#), [26](#), [63](#), [79](#), [132](#), [180](#)).
- O'Regan, J. K. (1992). "Solving the "Real" Mysteries of Visual Perception: The World as an Outside Memory". In: *Canadian Journal of Psychology* 46.3, pp. 461–488 (cit. on pp. [10](#), [123](#)).

- O'Regan, J. (Jan. 1, 2012). "Why Red Doesn't Sound Like A Bell: Understanding the Feel of Consciousness". In: *Why Red Doesn't Sound Like A Bell: Understanding the feel of Consciousness*, pp. 1–224. ISSN: 9780199775224. DOI: [10.1093/acprof:oso/9780199775224.001.0001](https://doi.org/10.1093/acprof:oso/9780199775224.001.0001) (cit. on p. 9).
- O'Regan, J. Kevin and Ned Block (Mar. 2012). "Discussion of J. Kevin O'Regan's "Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness"". In: *Review of Philosophy and Psychology* 3.1, pp. 89–108. ISSN: 1878-5158, 1878-5166. DOI: [10.1007/s13164-012-0090-7](https://doi.org/10.1007/s13164-012-0090-7). URL: <http://link.springer.com/10.1007/s13164-012-0090-7> (cit. on pp. 26, 108).
- Oddi, Angelo et al. (2020). "Integrating Open-Ended Learning in the Sense-Plan-Act Robot Control Paradigm". In: *Santiago de Compostela*, p. 8 (cit. on p. 3).
- Ortiz, Michael Garcia and Alban Laflaquière (May 16, 2018). "Learning Representations of Spatial Displacement through Sensorimotor Prediction". In: arXiv: [1805.06250](https://arxiv.org/abs/1805.06250) [cs]. URL: <http://arxiv.org/abs/1805.06250> (visited on 11/02/2020) (cit. on p. 127).
- Pemmaraju, Sriram and Steven Skiena (Dec. 8, 2003). *Computational Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*®. 1st edition. Cambridge, U.K. ; New York: Cambridge University Press. 494 pp. ISBN: 978-0-521-80686-2 (cit. on p. 109).
- Petrich, Mario (Jan. 1, 1984). *Inverse Semigroups*. New York: Wiley. 674 pp. ISBN: 978-0-471-87545-1 (cit. on p. 86).
- Pfeifer, R. and J. Bongard (2006). "How the Body Shapes the Way We Think - a New View on Intelligence". In: DOI: [10.7551/mitpress/3585.001.0001](https://doi.org/10.7551/mitpress/3585.001.0001) (cit. on p. 5).
- Philipona, D, J K O'Regan, and J.-P. Nadal (2003). "Is There Something out There?: Inferring Space from Sensorimotor Dependencies". In: *Neural Comput.* 15.9, pp. 2029–2049. ISSN: 0899-7667. DOI: [10.1162/089976603322297278](https://doi.org/10.1162/089976603322297278) (cit. on pp. 11, 26, 133).
- Philipona, David et al. (2004). "Perception of the Structure of the Physical World Using Unknown Sensors and Effectors". In: *Advances in Neural Information Processing Systems* 15 (cit. on p. 53).
- Philipona, David L. and J. Kevin O'Regan (May 2006). "Color Naming, Unique Hues, and Hue Cancellation Predicted from Singularities in Reflection Properties". In: *Visual Neuroscience* 23 3-4.3-4, pp. 331–339. ISSN: 0952-5238, 1469-8714. DOI: [10.1017/S0952523806233182](https://doi.org/10.1017/S0952523806233182). URL: <https://www.>

- [cambridge.org/core/product/identifier/S0952523806233182/type/journal_article](https://www.cambridge.org/core/product/identifier/S0952523806233182/type/journal_article) (visited on 11/03/2020) (cit. on p. 11).
- Piaget, J. (1937). *La Construction Du Réel Chez l'enfant*. [The Construction of the Real in the Child.] La Construction Du Réel Chez l'enfant. Oxford, England: Delachaux & Niestle, p. 398. 398 pp. (cit. on p. 6).
- Pierce, David and Benjamin J. Kuipers (May 1997). "Map Learning with Uninterpreted Sensors and Effectors". In: *Artificial Intelligence* 92.1-2, pp. 169–227. ISSN: 00043702. DOI: [10.1016/S0004-3702\(96\)00051-3](https://doi.org/10.1016/S0004-3702(96)00051-3). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0004370296000513> (visited on 03/17/2021) (cit. on p. 53).
- Poincaré, H (1895). "L'espace Et La Géométrie". In: *Revue de Métaphysique Et de Morale* 3.6, pp. 631–646 (cit. on p. 30).
- (1902). *La Science et l'hypothèse*. Flammarion (cit. on pp. 11, 30).
- Preston, G. B. (Oct. 1954). "Inverse Semi-Groups". In: *Journal of the London Mathematical Society* s1-29.4, pp. 396–403. ISSN: 00246107. DOI: [10.1112/jlms/s1-29.4.396](https://doi.org/10.1112/jlms/s1-29.4.396). URL: <http://doi.wiley.com/10.1112/jlms/s1-29.4.396> (cit. on p. 73).
- Prinz, Wolfgang (June 1997). "Perception and Action Planning". In: *European Journal of Cognitive Psychology* 9.2, pp. 129–154. ISSN: 0954-1446, 1464-0635. DOI: [10.1080/713752551](https://doi.org/10.1080/713752551). URL: <http://www.tandfonline.com/doi/abs/10.1080/713752551> (cit. on p. 78).
- Pynn, Laura and Joseph DeSouza (Nov. 2012). "The Function of Efference Copy Signals: Implications for Symptoms of Schizophrenia." In: *Vision Research* 76. DOI: [10.1016/j.visres.2012.10.019](https://doi.org/10.1016/j.visres.2012.10.019) (cit. on p. 55).
- Reck, Erich and Georg Schiemer (2020). "Structuralism in the Philosophy of Mathematics". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2020. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/spr2020/entries/structuralism-mathematics/> (cit. on p. 19).
- Remolina, Emilio and Benjamin Kuipers (Jan. 1, 2004). "Towards a General Theory of Topological Maps". In: *Artificial Intelligence* 152.1, pp. 47–104. ISSN: 0004-3702. DOI: [10.1016/S0004-3702\(03\)00114-0](https://doi.org/10.1016/S0004-3702(03)00114-0). URL: <https://www.sciencedirect.com/science/article/pii/S0004370203001140> (cit. on p. 6).

- Rescorla, Michael (2020). "The Computational Theory of Mind". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2020. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/fall2020/entries/computational-mind/> (cit. on p. 3).
- Roe, John (June 2006). "What Is... a Coarse Space?" In: *Notices of the AMS*, p. 2. URL: <https://www.ams.org/notices/200606/whatis-roe.pdf> (cit. on p. 143).
- Roschin, V. Y. and A.A. Frolov (2011). "A Neural Network Model for the Acquisition of a Spatial Body Scheme Through Sensorimotor Interaction". In: *Neural Computation* 23, pp. 1821–1834 (cit. on p. 11).
- Rowlands, Mark, Joe Lau, and Max Deutsch (2020). "Externalism About the Mind". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2020. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/win2020/entries/content-externalism/> (cit. on p. 123).
- Rozinas, M. R. (Apr. 24, 2010). *Metric on State Space of Markov Chain*. arXiv: 1004.4264 [math]. URL: <http://arxiv.org/abs/1004.4264> (visited on 02/07/2021) (cit. on p. 176).
- Russell, Stuart J. and Peter Norvig (1995). *Artificial Intelligence: A Modern Approach*. Prentice Hall Series in Artificial Intelligence. Englewood Cliffs, N.J: Prentice Hall. 932 pp. ISBN: 978-0-13-103805-9. URL: <http://aima.cs.berkeley.edu/> (cit. on p. 2).
- Sampaio, E. (1994). "Les substitutions sensorielles adaptées aux déficits visuels importants." In: *Le déficit visuel. Des fondements neurophysiologiques à la pratique de la réadaptation*. Ed. by Avinoam B Safran and André Assimacopoulos. Paris: Masson. ISBN: 978-2-225-84768-4 (cit. on p. 8).
- Schmidt, Nico M. et al. (2013). "Bootstrapping Perception Using Information Theory: Case Studies in a Quadruped Robot Running on Different Grounds". In: *Advances in Complex Systems (ACS)* 16 (02n03), pp. 1–25. ISSN: 1793-6802. URL: https://econpapers.repec.org/article/wsiacsxxx/v_3a16_3ay_3a2013_3ai_3a02n03_3an_3as0219525912500786.htm (cit. on p. 55).
- Schröder-Schetelig, Johannes, Poramate Manoonpong, and Florentin Wörgötter (Nov. 2010). "Using Efference Copy and a Forward Internal Model for Adaptive Biped Walking". In: *Autonomous Robots* 29.3, pp. 357–366. ISSN: 1573-7527. DOI: 10.1007/s10514-010-9199-7. URL: <https://doi.org/10.1007/s10514-010-9199-7> (cit. on p. 127).

- Seth, Anil K. (Apr. 3, 2014). "A Predictive Processing Theory of Sensorimotor Contingencies: Explaining the Puzzle of Perceptual Presence and Its Absence in Synesthesia". In: *Cognitive Neuroscience* 5.2, pp. 97–118. ISSN: 1758-8928. DOI: [10.1080/17588928.2013.877880](https://doi.org/10.1080/17588928.2013.877880). pmid: [24446823](https://pubmed.ncbi.nlm.nih.gov/24446823/). URL: <https://doi.org/10.1080/17588928.2013.877880> (visited on 05/30/2021) (cit. on p. 12).
- Sexton, Miriam E. (Jan. 1, 1983). "The Development of the Understanding of Causality in Infancy". In: *Infant Behavior and Development* 6.2, pp. 201–210. ISSN: 0163-6383. DOI: [10.1016/S0163-6383\(83\)80027-7](https://www.sciencedirect.com/science/article/pii/S0163638383800277). URL: <https://www.sciencedirect.com/science/article/pii/S0163638383800277> (cit. on p. 6).
- Sharpe, R. W. (Nov. 21, 2000). *Differential Geometry: Cartan's Generalization of Klein's Erlangen Program*. Springer Science & Business Media. 452 pp. ISBN: 978-0-387-94732-7. Google Books: [Ytqs4xU5QKAC](#) (cit. on p. 29).
- Sher, R. B. and R. J. Daverman (Dec. 20, 2001). *Handbook of Geometric Topology*. Elsevier. 1145 pp. ISBN: 978-0-08-053285-1. Google Books: [80YxdADnhZoC](#) (cit. on pp. 109, 124).
- Shin, Yun Kyoung, Robert W. Proctor, and E. J. Capaldi (Nov. 2010). "A Review of Contemporary Ideomotor Theory". In: *Psychological Bulletin* 136.6, pp. 943–974. ISSN: 1939-1455. DOI: [10.1037/a0020541](https://doi.org/10.1037/a0020541). pmid: [20822210](https://pubmed.ncbi.nlm.nih.gov/20822210/) (cit. on p. 127).
- Siegel, Mel (July 5, 2003). "The Sense-Think-Act Paradigm Revisited". In: 5 pp. ISBN: 978-0-7803-8109-4. DOI: [10.1109/ROSE.2003.1218700](https://doi.org/10.1109/ROSE.2003.1218700) (cit. on p. 3).
- Smith, Linda B. and Esther Thelen (July 12, 1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Ed. by Stephen E. Palmer. Cognitive Psychology. Cambridge, MA, USA: A Bradford Book. 402 pp. ISBN: 978-0-262-20095-0 (cit. on p. 6).
- Sperry, R. W. (Dec. 1950). "Neural Basis of the Spontaneous Optokinetic Response Produced by Visual Inversion". In: *J Comp Physiol Psychol* 43.6, pp. 482–489 (cit. on p. 55).
- Spitzer, Frank (Mar. 14, 2013). *Principles of Random Walk*. Springer Science & Business Media. 419 pp. ISBN: 978-1-4757-4229-9. Google Books: [hcveBwAAQBAJ](#) (cit. on p. 99).
- Srivastava, Ankit (Feb. 5, 2019). *Sense-Plan-Act in Robotic Applications*. DOI: [10.13140/RG.2.2.21308.36481](https://doi.org/10.13140/RG.2.2.21308.36481) (cit. on p. 3).

- Stewart, John, Olivier Gapenne, and Ezequiel Di Paolo (Jan. 1, 2010). *Enaction: A New Paradigm for Cognitive Science*. ISBN: 978-0-262-01460-1. DOI: [10 . 7551/mitpress/9780262014601.001.0001](https://doi.org/10.7551/mitpress/9780262014601.001.0001) (cit. on p. 6).
- Stock, A. and C. Stock (Apr. 2004). "A Short History of Ideo-Motor Action". In: *Psychol Res* 68.2-3, pp. 176–188 (cit. on p. 78).
- Stroud, John M. (1956). "The Fine Structure of Psychological Time." In: (cit. on p. 131).
- Taubes, Clifford Henry (Jan. 1987). "Gauge Theory on Asymptotically Periodic {4}-Manifolds". In: *Journal of Differential Geometry* 25.3, pp. 363–430. ISSN: 0022-040X. DOI: [10 . 4310 / jd g / 1214440981](https://doi.org/10.4310/jdg/1214440981). URL: [https : / / projecteuclid . org / journals / journal - of - differential - geometry / volume - 25 / issue - 3 / Gauge - theory - on - asymptotically - periodic - 4 - manifolds / 10 . 4310 / jd g / 1214440981 . full](https://projecteuclid.org/journals/journal-of-differential-geometry/volume-25/issue-3/Gauge-theory-on-asymptotically-periodic-4-manifolds/10.4310/jdg/1214440981.full) (cit. on p. 19).
- Tenenbaum, J. B. (Dec. 22, 2000). "A Global Geometric Framework for Non-linear Dimensionality Reduction". In: *Science* 290.5500, pp. 2319–2323. ISSN: 00368075, 10959203. DOI: [10 . 1126 / science . 290 . 5500 . 2319](https://doi.org/10.1126/science.290.5500.2319). URL: [https : / / www . sciencemag . org / lookup / doi / 10 . 1126 / science . 290 . 5500 . 2319](https://www.sciencemag.org/lookup/doi/10.1126/science.290.5500.2319) (cit. on p. 156).
- "Introduction" (1975). In: *Sheaf Theory*. Ed. by B. R. Tennison. London Mathematical Society Lecture Note Series. Cambridge: Cambridge University Press, pp. v–vi. ISBN: 978-0-521-20784-3. DOI: [10 . 1017 / CB09780511661761 . 001](https://doi.org/10.1017/CB09780511661761.001). URL: [https : / / www . cambridge . org / core / books / sheaf - theory / introduction / 2626DD0FB99F3A8CEE2200F14ACFB5C](https://www.cambridge.org/core/books/sheaf-theory/introduction/2626DD0FB99F3A8CEE2200F14ACFB5C) (cit. on p. 31).
- Terekhov, A. V. and J. K. O'Regan (Aug. 2013). *Space as an Invention of Biological Organisms*. arXiv: [1308.2124v1](https://arxiv.org/abs/1308.2124v1) [cs.AI] (cit. on pp. 11, 56).
- Terekhov, Alexander V. and J. Kevin O'Regan (2016). "Space as an Invention of Active Agents". In: *Frontiers in Robotics and AI* 3, p. 4. ISSN: 2296-9144. DOI: [10 . 3389 / frobt . 2016 . 00004](https://doi.org/10.3389/frobt.2016.00004). URL: [https : / / www . frontiersin . org / articles / 10 . 3389 / frobt . 2016 . 00004 / full](https://www.frontiersin.org/articles/10.3389/frobt.2016.00004/full) (cit. on pp. 19, 27).
- Toscano, Joseph C. et al. (Oct. 2010). "Continuous Perception and Graded Categorization: Electrophysiological Evidence for a Linear Relationship between the Acoustic Signal and Perceptual Encoding of Speech". In: *Psychological science* 21.10, pp. 1532–1540. ISSN: 0956-7976. DOI: [10 . 1177 / 0956797610384142](https://doi.org/10.1177/0956797610384142). pmid: 20935168. URL: [https : / / www . ncbi . nlm . nih . gov / pmc / articles / PMC3523688 /](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3523688/) (cit. on p. 131).

- Turing, A. M. (1937). "Computability and λ -Definability". In: *The Journal of Symbolic Logic* 2.4, pp. 153–163. ISSN: 0022-4812. DOI: [10.2307/2268280](https://doi.org/10.2307/2268280). JSTOR: [2268280](https://www.jstor.org/stable/2268280) (cit. on p. 57).
- Urbanik, David (n.d.). "A Brief Introduction to Schemes and Sheaves". In: (), p. 16 (cit. on p. 31).
- Van der Maaten, Laurens, Eric Postma, and H. Herik (Jan. 1, 2007). "Dimensionality Reduction: A Comparative Review". In: *Journal of Machine Learning Research - JMLR* 10 (cit. on p. 156).
- VanRullen, Rufin and Christof Koch (May 1, 2003). "Is Perception Discrete or Continuous?" In: *Trends in Cognitive Sciences* 7.5, pp. 207–213. ISSN: 1364-6613, 1879-307X. DOI: [10.1016/S1364-6613\(03\)00095-0](https://doi.org/10.1016/S1364-6613(03)00095-0). pmid: [12757822](https://pubmed.ncbi.nlm.nih.gov/12757822/). URL: [https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613\(03\)00095-0](https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(03)00095-0) (cit. on p. 131).
- Varela, Francisco J., Eleanor Rosch, and Evan Thompson (1993). *The Embodied Mind* (cit. on p. 5).
- Von Holst, Erich and Horst Mittelstaedt (Jan. 1950). "Das Reafferenzprinzip". In: *Naturwissenschaften* 37.20, pp. 464–476. ISSN: 1432-1904. DOI: [10.1007/BF00622503](https://doi.org/10.1007/BF00622503). URL: <https://doi.org/10.1007/BF00622503> (cit. on p. 55).
- Wehrl, Alfred (Apr. 1, 1978). "General Properties of Entropy". In: *Reviews of Modern Physics* 50.2, pp. 221–260. ISSN: 0034-6861. DOI: [10.1103/RevModPhys.50.221](https://doi.org/10.1103/RevModPhys.50.221). URL: <https://link.aps.org/doi/10.1103/RevModPhys.50.221> (cit. on p. 103).
- Weisstein, Eric W. (n.d.). *Spherical Coordinates*. URL: <https://mathworld.wolfram.com/SphericalCoordinates.html> (cit. on p. 165).
- Wilson, W. A. (1931). "On Quasi-Metric Spaces". In: *American Journal of Mathematics* 53.3, pp. 675–684. ISSN: 0002-9327. DOI: [10.2307/2371174](https://doi.org/10.2307/2371174). JSTOR: [2371174](https://www.jstor.org/stable/2371174) (cit. on p. 147).
- Wolpert, D.M. and M. Kawato (Oct. 1998). "Multiple Paired Forward and Inverse Models for Motor Control". In: *Neural Networks* 11.7-8, pp. 1317–1329. ISSN: 08936080. DOI: [10.1016/S0893-6080\(98\)00066-5](https://doi.org/10.1016/S0893-6080(98)00066-5). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0893608098000665> (cit. on p. 39).
- Ziemke, Tom (2001). "The Construction of 'Reality' in the Robot: Constructivist Perspectives on Situated Artificial Intelligence and Adaptive Robotics". In: p. 55. DOI: [10.1023/A:1011394317088](https://doi.org/10.1023/A:1011394317088) (cit. on p. 6).