



HAL
open science

Action of the Euclidean versus Projective group on an agent's internal space in curiosity driven exploration: a formal analysis

Grégoire Sergeant-Perthuis, David Rudrauf, Dimitri Ognibene, Yvain Tisserand

► To cite this version:

Grégoire Sergeant-Perthuis, David Rudrauf, Dimitri Ognibene, Yvain Tisserand. Action of the Euclidean versus Projective group on an agent's internal space in curiosity driven exploration: a formal analysis. 2023. hal-04054736v1

HAL Id: hal-04054736

<https://hal.science/hal-04054736v1>

Preprint submitted on 1 Apr 2023 (v1), last revised 13 Feb 2024 (v2)

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.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Action of the Euclidean versus Projective group on an agent’s internal space in curiosity driven exploration: a formal analysis

Grégoire Sergeant-Perthuis¹, David Rudrauf², Dimitri Ognibene³, Yvain Tisserand⁴

Abstract—In human spatial awareness, information appears to be represented according to 3-D projective geometry. It structures information integration and action planning within an internal representation space. The way different first person perspectives of an agent relate to each other, through transformations of a world model, defines a specific perception scheme for the agent. In mathematics, this collection of transformations is called a ‘group’ and it characterizes a geometric space by acting on it. We propose that imbuing world models with a ‘geometric’ structure, given by a group, is one way to capture different perception schemes of agents. We explore how changing the geometric structure of a world model impacts the behavior of an agent.

In particular, we focus on how such geometrical operations transform the formal expression of epistemic value in active inference as driving an agent’s curiosity about its environment, and impact exploration behaviors accordingly. We used group action as a special class of policies for perspective-dependent control. We compared the Euclidean versus projective groups. We formally demonstrate that the groups induce distinct behaviors. The projective group induces nonlinear contraction and dilatation that transform entropy and epistemic value as a function of the choice of frame, which fosters exploration behaviors. This contribution opens research avenues in which a geometry structures *a priori* an agent’s internal representation space for information integration and action planning.

I. INTRODUCTION

In artificial agent learning and control, intrinsic and extrinsic rewards can be combined to optimize the balance between exploration and exploitation.

Intrinsic rewards in Reinforcement Learning (RL) [Hester and Stone, 2017], [Merckling et al., 2022], [Oudeyer et al., 2007] or terms of epistemic value in active inference [Friston et al., 2015] have been brought forth as mechanisms mimicking curiosity and driving exploration, e.g. by integrating prediction error or uncertainty to drive actions favoring their reduction. However, efficient exploration is a computationally hard task. Recent neural planning models have increased planning flexibility and generality [Sekar et al., 2020]. Yet, it is well-known that models’ structures heavily impact planning performance, computational cost and tractability [Geffner and Bonet, 2013]. A good representation may improve learning and search efficiency.

These issues are particularly salient for computation-heavy, highly recursive machine learning algorithms and

applications, e.g. reinforcement learning (RL) in artificial agents [Bonet and Geffner, 2019]. Although generic neural world models can support exploration-related processes, incorporating prior knowledge that shapes internal representations to more effectively support exploration across a broad range of environments, such as 3-D environments, may enable autonomous agents to explore more complex and realistic settings on a larger scale. The exploration planning problem can thus be approached by considering how the structure of representation impacts exploration behaviors.

Here, we do not consider mechanisms of representation learning, e.g. in which world dynamics and action effects need to be learned and represented, as typically done in RL. We specifically consider control and execution when object locations, world states, and maps may not be known but dynamics, rewards, and action effects already are. We focus on action selection for environment exploration and mapping.

We adopt the active inference framework, i.e. an implementation of the Bayesian Brain Hypothesis aimed at generating adaptive behaviors in agents [Friston et al., 2006], [Timsit and Sergeant-Perthuis, 2021]. It relies on an internal representation of the environment that an agent is driven to explore and exploit. The agent continually updates its beliefs about plausible competing internal hypotheses on the environment state. Under common sensory limitations, active inference relates to Partially Observable Markov Decision Process (POMDP) [Da Costa et al., 2020], [Ognibene et al., 2019]. The epistemic value of states is a quantity that arises in active inference [Friston et al., 2015]. Its maximization drives the agent’s curiosity and actions.

Defining states’ embedding capturing both the environment’s local coherence and global topology is key for exploration optimization [Merckling et al., 2022]. For exploration or search in 3-D space, it is warranted to consider how geometrical principles could be embedded into efficient control mechanisms, to regularize the internal representation of information and mediate exploration under a drive of uncertainty reduction or information maximization. Geometrical considerations have previously been integrated into a variety of optimization and machine learning approaches, such as RL, active inference, and Bayesian inference (See Related Work below), but not in the specific perspective we introduce herein.

We build upon the hypothesis that 3-D internal representations of space in agents performing active inference may correspond to specific geometries, with properties that can be mathematically analyzed and

¹OURAGAN team, Inria Paris & IMJ-PRG, Paris, France. gregoire.sergeant-perthuis@inria.fr

² CIAMS, Université Paris-Saclay, Orsay & Université d’Orléans, Orléans, France.

³ Dipartimento di Psicologia, Università Milano-Bicocca, Milan, Italy.

⁴ CISA, University of Geneva, Geneva, Switzerland.

demonstrated. More specifically, we consider how different first person perspectives may relate to each other, through transformations of a world model, as a specific perception (and imagination) scheme for agents. This entails considering the action, on the spatial distribution of information experienced and encoded by agents internally, of geometrical groups of transformations (in the mathematical sense of the concept in Group theory; see Section III-A). The question is whether such group action could contribute to information gain estimation and maximization, as an internal planning or perspective-dependent control mechanism. Certain geometrical groups might imply internal representations, policies, value functions and principles of action that are particularly relevant for search and exploration. More specifically, we wish to compare how different groups impact the quantification of epistemic value. We then wish to characterize how the optimization of action from those different groups may yield different exploration behaviors. We contrasted two separate toy models of an agent performing a simple search task using active inference based solely on epistemic drives. One model used Euclidean geometry and the other projective geometry for the agent’s internal space. We compared the two models in terms of resulting exploration behaviors.

We chose to compare Euclidean versus projective geometry based on previous work, leveraging psychological research on the phenomenology of spatial consciousness and its role in the control of behaviors [Rudrauf et al., 2017], [Rudrauf et al., 2020], [Rudrauf et al., 2022], [Rudrauf et al., 2023]. This research suggested that 3-D projective geometry plays a central role in human cognition and decision-making by shaping information representation and subsequent drives. It also offers a mechanism of changes of points of view and perspectives on a world model, including for perspective taking in social cognition, which is critical for the development of strategies of action planning in humans (see [Rudrauf et al., 2022], [Rudrauf et al., 2023]). We used Euclidean geometry as a standard baseline geometry for comparison [Ognibene and Demiris, 2013]. Our geometrical rationale implies a different understanding of how agents’ actions in their environment (here in the behavioral sense of the term) are implemented compared to usual active inference. Agents’ actions can naturally be seen as dual to internal changes of perspective, i.e. group actions, in their representation space. We thus used group actions as a predictive model of actual behavioral actions.

The approach allowed us to formally study and demonstrate how the geometry governing the internal representation space may directly impact the computation of epistemic value and ensuing exploratory behaviors. Projective geometry *versus* Euclidean geometry demonstrated remarkable properties of information integration for motion planning under epistemic drive.

II. RELATED WORKS

A. Representation of space and exploration, in the context of machine learning and control

The integration of geometrical mapping in machine learning has been proposed to reduce the high-dimensionality of input spaces and provide efficient solutions for action selection and navigation. Seminal neurally inspired models used projections on 2-D manifolds for representation learning of complex spatial information and self-motion effects [Arleo et al., 2004]. The impact of changes of perspective in exploration has long been of interest. [Ognibene and Demiris, 2013]. [Ferreira et al., 2013] proposed an internal 3-D egocentric, spherical representation of space, to modulate information sampling and uncertainty as a function of distance, and control a robot attention through Bayesian inference. This was a seminal example of how geometrical rationale could suggest solutions to integrate perception and action planning.

Exploration methods must often maintain high-resolution representations of space to maximize information gain following action. This may hinder exploration efficiency, in particular in large-scale environments. 3-D topological representations of ambient space have been proposed as part of an abstract planning scheme, showing promising improvements of exploration efficiency [Yang et al., 2021].

Active vision principles, combined with curiosity-based algorithms and RL, were applied to the learning of saliency maps in the context of autonomous robots’ navigation [Craye et al., 2016]. The approach yielded promising optimization solutions to both adaptive learning of task-independent, spatial representations, and efficient exploration policies, which could serve as prior to support long-term, task-oriented, utility-driven RL mechanisms [Craye et al., 2016] (see also [Ognibene and Baldassare, 2014], [Sperati and Baldassarre, 2017]).

Complex control tasks with continuous state and action spaces have been solved using deep reinforcement learning (DRL) with joint learning of representations and predictions. Such approach may entail non-stationarity, risks of instability and slow convergence, in particular in control tasks with active vision. Separating representation learning and policies’ computations may mitigate the issues, but may also lead to inefficient information representations. [Merckling et al., 2022] have sought to build compact and meaningful representations based on task-agnostic and reward-free agent-environment interactions. They used (recursive) state representation learning (SRL) while jointly learning a state transition estimator with near-future prediction objective, to contextually remove distracting information and reduce the exploration problem complexity. Positive outcomes were maximized through inverse predictive modeling, and prediction error was used to favor actions reducing uncertainty, which improved subsequent performance in RL tasks. The authors emphasized that dealing with partial observability through memory and active vision may require

new solutions to both representation learning of hidden information and exploration strategies.

Uncertainty-based methods using intrinsic reward and exploration bonuses to plan trajectories have been criticized for inducing non-stationary decaying surprise, and for being hard to structure and optimize [Guo et al., 2021]. Maximum State-Visitation Entropy (MSVE) was introduced to maximize state exploration uniformity, but optimization has been often challenging for large state spaces. [Guo et al., 2021] have introduced Geometric Entropy Maximization, which leverages geometry-aware entropy based on Adjacency Regularization (AR) and a similarity function, in order to optimize the MSVE problem at scale.

Geometrical constraints considered across these related works were not integrated into a global model, and were somewhat *ad hoc*. They pertained to a lower level of processing than the one we are concerned with here. However, they emphasize the current needs and challenges for integrating geometry in learning, control and navigation.

Methods and algorithms combining computer vision, machine learning and optimization, e.g. for robotic planning, have integrated group theoretic concepts to obtain, for instance, invariance to rotation and translation in image processing [Lee and Moore, 2004], [Qin et al., 2019], [Meng et al., 2017]. Likewise, the leveraging of geometrical concepts, in Deep learning, e.g. for learning manifolds and graphs, has been growing in recent years, demonstrating very promising results for representation learning [Gerken et al., 2021], [Cao et al., 2022]. The approach introduces combinatorial structures to leverage prior knowledge of geometry on the data of interest, e.g. applying ‘convolutional Neural Networks’ to non-Euclidean space. However, the Euclidean group E_3 , or more specifically SE_3 (see [Lee and Moore, 2004]), which includes translations and rotations, but excludes reflections, or simply SO_3 , the 3-dimensional rotation group [Gerken et al., 2021], are the groups being typically considered.

Here, in addition to the Euclidean group, we also consider PGL_3 , the projective general linear group in 3-D (which has more degrees of freedom than E_3). The projective group is central to computer vision, for instance to generate 2-D images from 3-D information, but is used in such context in a restricted manner. Our approach, which is based on cognitive science, considers spatial cognition and its relations to action at a higher level of integration, and do not reduce to the visual modality.

B. Projective Consciousness Model (PCM) and active inference

We have shown that geometrically constrained active inference can be used as a framework to understand and model central aspects of human spatial consciousness, through our Projective Consciousness Model (PCM) [Rudrauf et al., 2017], [Rudrauf et al., 2022]. Consciousness accesses and represents multimodal information through a Global Workspace [Dehaene et al., 2017] within which subjective

perspectives on an internal world model can be taken. The process contributes to appraise possible actions based on their expected utility and epistemic value [Rudrauf et al., 2022]. In [Rudrauf et al., 2017], [Williford et al., 2018], [Rudrauf et al., 2022], [Rudrauf et al., 2023], [Williford et al., 2022], it is hypothesized that such internal representation space is geometrically structured as a $3-D$ projective space, denoted $P_3(\mathbb{R})$. Changes of perspective then correspond to the choice of a projective transformation ψ , i.e. an action from PGL_3 . A projective transformation can also be modeled as a linear isomorphism $M_\psi \in GL_4(\mathbb{R})$ up to a multiplicative constant. The model yielded an explanation for the Moon illusion [Rudrauf et al., 2020] with, for the first time, falsifiable predictions on how strong the effect should be depending on context, as well as for the generation of adaptive and maladaptive behaviors, consistent with developmental and clinical psychology (see [Rudrauf et al., 2022]). Though essential in integrative spatial cognition, notably for understanding multi-agent social interactions, perspective taking is rarely integral to existing models of consciousness or formally implemented [Koch et al., 2016], [Kleiner and Tull, 2021], [Mashour et al., 2020], [Dehaene et al., 2017], [Merker et al., 2022]. The PCM assumes that projective mechanisms of perspective changes are integral to the global workspace of consciousness, both in non-social and social contexts. The advantages of mechanisms of perspective taking for cybernetics remains to be fully formulated (see [Rudrauf et al., 2022]).

III. METHODOLOGY

The experiment we consider is that of an agent, denoted as a , which is looking for an object O in the ‘real world’, the 3-D Euclidean space $E_3 := \mathbb{R}^3$. The set of moves of the agent is denoted M . The position of O is denoted $o \in E_3$. The agent ‘represents’ the position of the object O inside its ‘internal world model’. We consider ‘internal world models’, spaces denoted W , that are such that there is a group acting on them; we call such spaces, group structured world models. This group accounts for the change of coordinates that each movement of the agent induces when the positions of the object are expressed in the agents reference frame. We consider two spaces in particular:

- 1) *Euclidean case*: W is the 3-D vector space, $W = \mathbb{R}^3$
- 2) *Projective case*: W is the 3-D projective space, denoted as $P_3(\mathbb{R})$

We will denote $\mathcal{B}(W)$ the Borel σ -algebra of the respective topological spaces.

In Section III-B, we explain how the ‘real world’ and the ‘internal world model’ are related to one another in both the Euclidean and Projective case. Figure 1 illustrates the setup of the toy model and main transformations considered. The agent’s internal beliefs about the position of the object are encoded by a probability measure on W that the agent updates through observations. The agent explores its environment through the computation of an epistemic value, the maximization of which captures curiosity-based exploration.




Fig. 1. Toy model setup and main transformations

Upper-tier. Agent a simulates move m in Euclidean space E . \mathcal{R}^0 and \mathcal{R} are its frames in E before and after the move, oriented toward object O . Vertical arrows indicate transformations ϕ from the external to the internal space. *Lower-tier.* Rendering of the effect of the internal group action $\psi(m)$ corresponding to move m in the Euclidean *versus* projective case. (Made with Unity).

In Section III-C, we explain how epistemic value is defined for group structured internal representations. In Section III-D we give the details of the exploration algorithm.

A. Group structured world model

Let us first recall what a group is.

Definition 1 (Group, §2 Chapter 1 [Lang, 2012]). A group is a set G with an operation $\cdot : G \times G \rightarrow G$ that is associative, such that there is an element $e \in G$ for which $e.g = g$ for any $g \in G$, and any $g \in G$ has an inverse denoted g^{-1} defined as satisfying, $g.g^{-1} = g^{-1}.g = e$.

We call a group structured world model, a world model provided with a group action; we now make this statement formal.

Definition 2 (Group structured world model). W is a group structured world model for the group G when there is a map $h : G \times W \rightarrow W$ denoted as $h(g, x) = g.x$ for $g \in G$ and $x \in X$, such that,

- 1) $(g.g_1).x = g.(g_1.x)$ for all $g, g_1 \in G, x \in W$
- 2) $e.x = x$, for all $x \in W$

In the *Euclidean case* the group structured world model, W , is the 3-D vector space \mathbb{R}^3 ; it is structured by the group of

invertible matrices $GL_3(\mathbb{R})$. In the *Projective case*, the group structured world model, W , is the projective space $P_3(\mathbb{R})$; it is structured by the group of projective linear transformations $PGL(\mathbb{R}^3)$.

B. Relating the ‘real world’ to the ‘internal world model’

We assume that the ‘real world’ is the 3-D Euclidean space, E_3 . We assume that the ‘real world’ comes with with an Euclidean frame \mathcal{R}_E , i.e. a point C and three independent vectors e_0, e_1, e_2 . This frame is used to set up the experiment: the configurations of the object and agent across time are encoded in this frame; it is fixed once and for all before starting the experiment. Therefore we now identify E_3 with \mathbb{R}^3 , C with $(0, 0, 0)$ and e_0, e_1, e_2 with the respective basis vectors, $(1, 0, 0), (0, 1, 0), (0, 0, 1)$.

The agent, denoted as a , is modeled as a solid in the ‘real world’; it has its own Euclidean frame (the solid reference frame), $\mathcal{R} := (\mathcal{P}, u_0, u_1, u_2)$, with \mathcal{P} the center of a and u_0, u_1, u_2 three unitary vectors that form a basis.

In the *Euclidean case*, the map that relates E_3 and its group structured world model, W , is the affine map, $\phi_{\mathcal{R}}$, that changes the coordinate in \mathcal{R}_E to coordinates in \mathcal{R} .

In the *Projective case*, this map is a projective transformation. The choice of such a projective transformation is dictated by Proposition A.1 [Rudrauf et al., 2022]. Let us now recall some facts about that transformation.

Let for any $(x, y, z) \in \mathbb{R}^3$,

$$\rho(x, y, z) = \left(\frac{x}{\gamma z + 1}, \frac{y}{\gamma z + 1}, \frac{z}{\gamma z + 1} \right) \quad (1)$$

with $\gamma \in \mathbb{R}_+$ a strictly positive parameter. The (projective) transformation $\phi_{\mathcal{R}}^p$, from E_3 to W , which relates the ‘real world’ to the ‘internal world model’ in the projective case, is posed to be $\phi_{\mathcal{R}}^p := \rho \circ \phi_{\mathcal{R}}$.

Proposition 1. *When the agent a makes the move $m \in M$, its solid reference frame changes from \mathcal{R} to \mathcal{R}^m . In the Euclidean case this move induces an invertible linear map, $\psi_m \in GL_3(\mathbb{R})$ from the ‘internal world model’ to itself. In the Projective case it induces a projective transformation, $\psi_m \in PGL(\mathbb{R}^3)$.*

We will denote 1_U or $x \rightarrow 1[x \in U]$ the characteristic function of subset U , i.e. the functions that is equal to 1 for $x \in U$ and 0 elsewhere.

Remark 1. In both cases there is a dense open subset, U , of W which is in continuous bijection with \mathbb{R}^3 . From the Lebesgue measure dx on \mathbb{R}^3 , we define the following measure on W , $d\lambda := 1_U dx$. In what follows we do not raise this technical point anymore and simply refer to $d\lambda$ as the Lebesgue measure on W .

C. Beliefs, policies and epistemic value

1) *Beliefs:* The agent a keeps internal beliefs about the position of the object represented in its ‘internal world model’; these beliefs are encoded by a probability measure $Q_X \in \mathbb{P}(W)$, where $\mathbb{P}(W)$ denotes the set of probability measures on W . Probability measures will be denoted with

upper case letters and their densities with lower case letters. These beliefs are updated according to noisy sensory observations of the position of O . ‘Markov Kernels’ can be used to formalize noisy sensors. Let us recall their definition.

A ‘Markov Kernel’ Π from Ω_1 to Ω_2 is a map $\Pi : \Omega_1 \times \Omega_2 \rightarrow [0, 1]$ such that for any $\omega_1 \in \Omega_1$, $\sum_{\omega_2 \in \Omega_2} \Pi(\omega_2 | \omega_1) = 1$, i.e. a map that sends any $\omega_1 \in \Omega_1$ to a probability measure $\Pi_{|\omega_1} \in \mathbb{P}(\Omega_2)$.

The uncertainty on the sensors of a is captured by a Markov kernel $P_{Y|X}$ from W to W . It is a parameter of the experiment: it is fixed before the agent starts looking for O . The couple $(P_{Y|X}, Q_X)$ defines the following probability density, $P_{X,Y} \in \mathbb{P}(W \times W)$: for any $x, y \in W$,

$$P_{X,Y}(dx, dy) := p_{Y|X}(y|x)q_X(x)dx dy \quad (2)$$

where dx is the Lebesgue measure on W . An observation of the position of the object $y^o \in W$ triggers an update of the belief Q_X to the belief with density

$$Q_{X|y^o} = \frac{p_{X,Y}(x|y^o)q_X(x)dx}{\int_{x \in W} p_{X,Y}(x|y^o)q_X(x)dx} \quad (3)$$

2) *Policies*: Recall that the agent has a set of moves it can make M ; moves $m \in M$ are associated with the group action $\psi_m : W \rightarrow W$ (Proposition 1). The agent plans the consequence of its moves on its internal world model one step ahead: each change of frame induces the following Markov Kernel, for any $m \in M$, $A \in \mathcal{B}(W)$, and $x_0 \in W$,

$$p_{X_1|X_0, m}(A|x_0, m) = 1[\psi_m(x_0) \in A] \quad (4)$$

Each move m spreads a prior Q_X on X_0 into the following prior on X_1 : $\forall A \in \mathcal{B}(W)$,

$$\psi_{m,*}Q_X(A) := \int 1[\psi_m(x_0) \in A]q_X(x_0)dx \quad (5)$$

$$= Q_X(\psi_m^{-1}A) \quad (6)$$

We chose to denote this probability measure as $\psi_{m,*}Q_X$, because it is the standard mathematical notation for the ‘pushforward measure’ by ψ_m . The generative model the agent uses to plan its future actions is summarized in Figure 2.




Fig. 2. $m \in M$ is a move of the agent a , $1[\psi_m(x_0) \in A]$ defines the kernel induced by move m , $P_{Y|X}$ is the noisy sensor. The diagram constituted of solid arrows defines the generative model the agent uses to plan its actions. o is the position of the object in the ‘real world’, $y^o \in W$ is the representation of o in the ‘internal world model’ of a with respect to the solid reference frame \mathcal{R} , y_m^o is the same for the reference frame \mathcal{R}^m after move m .

3) *Epistemic value*: The following definition is a restatement of the *epistemic value* introduced in [Friston et al., 2015] in the case of the kernel $P_{Y|X} : W \rightarrow W$.

Definition 3 (Epistemic Value). For any probability measure $Q_X \in \mathbb{P}(W)$, the epistemic value of this measure is:

$$C(Q_X) := \mathbb{E}_{P_Y} [H(P_{X|Y}|Q_X)] \quad (7)$$

$$= \int p_Y(y)dy \int p_{X|Y}(x|y) \ln \frac{p_{X|Y}(x|y)}{q_X(x)} dx \quad (8)$$

H is the relative entropy, also called Kullback-Leibler divergence.

Reexpressing Equation 7, it becomes apparent that epistemic value is simply a mutual information:

$$C(Q_X) = \int p_{X,Y} \ln \frac{p_{X,Y}(x,y)}{p_Y(y)q_X(x)} dx dy \quad (9)$$

We propose to define the epistemic value of move m as the epistemic value of the induced prior on X_1 ,

$$C(m) := C(\psi_{m,*}Q_X) \quad (10)$$

D. Exploration algorithm

Let us now put the previous elements together to describe the exploration behavior programmed in our agent. The agent a is initialized in a configuration of the ‘real world’, with solid reference frame \mathcal{R}^0 ; the object O is positioned at $o \in E_3$. a starts with an initial belief $Q_X^0 \in \mathbb{P}(W)$ on the position of O . It plans one step ahead the consequence of move m ; move m induces a group action $\psi_m : W \rightarrow W$ that pushes forward the belief Q_X^0 to $\psi_{m,*}Q_X^0$. The agent then evaluates the epistemic value of $(P_{Y|X}, \psi_{m,*}Q_X^0)$ for each move m and chooses the move that maximizes this value, m^* . a executes the move m^* which transforms its solid reference frame \mathcal{R}_0 to \mathcal{R} . It can then observe (with its ‘noisy sensors’) the position of O which is $y^o := \phi_{\mathcal{R}}(o)$ in its internal world model, which triggers the update of prior $\psi_{m^o,*}Q_X^0$ to the distribution conditioned on the observation: $(\psi_{m^o,*}Q_X^0)_{|y_m^o}$. The process is iterated with this new prior. The exploration algorithm is summarized in Algorithm 1.

Algorithm 1: Curiosity based Exploration for agent a

Data: Initialization: Q_X^0 initial belief, \mathcal{R}^0 initial solid reference frame of a

- 1 $Q_X \leftarrow Q_X^0$;
 - 2 **while** True **do**
 - 3 $m^* \leftarrow \operatorname{argmax}_{m \in M} C(\psi_{m,*}Q_X)$;
 - 4 $\mathcal{R} \leftarrow$ solid reference frame of a after move m^* ;
 - 5 $Q_X \leftarrow \psi_{m^o,*}Q_X$;
 - 6 $y^o \leftarrow \phi_{\mathcal{R}}(o)$;
 - 7 $Q_X \leftarrow Q_{X|y^o}$;
 - 8 **end**
-

IV. THEORETICAL PREDICTIONS

We wish to understand how the group by which the internal world model is structured influences the exploration behavior of the agent. The *Euclidean case* serves as the reference model; in this case the world model shares the same structure as the real world: it is the ‘classical’ way of modeling this exploration problem. The *Projective case* corresponds to the hypothesis underlying the PCM. The following Theorem states that this experiment allows us to discriminate when the behavior of the agent is dictated by ‘objective’ perspectives (Euclidean change of frame) *versus* by ‘subjective’ perspectives (projective change of frame) on its environment.

We consider the following noisy sensor, for any $x, y \in \mathbb{R}^3$,

$$P_{Y|X}(y|x) = \frac{3}{4\pi\epsilon^3} 1[\|x - y\| \leq \epsilon] \quad (11)$$

where $\|\cdot\|$ designates the Euclidean norm on \mathbb{R}^3 , i.e. $\|x\|^2 = x_0^2 + x_1^2 + x_2^2$; $\epsilon > 0$ is a strictly positive real number.

Theorem (Discrimination of behavior with respect to internal representations). *Let us assume that staying still is always a possible move for the agent.*

Euclidean case: when the agent has an objective representation of its environment, given by an affine map, the agent stays still.

Projective case: Assume now that the set of moves M is finite; assume furthermore that after any possible move, the agent faces O , in other words, we assume that the agent knows in which direction to look in order to find the object but is still uncertain on where the object is exactly. If it has a ‘subjective’ perspectives, i.e. its representation is given through a projective transformation, it will choose the moves that allows it to approach O (for any ϵ small enough).

Proof: The details of the proof are given in Appendix A. Let us here give an idea of the proof. The agent circumscribes a region of space in which it believes it is likely to find the object. This region corresponds to the error the agent tolerates on the measurement it makes of the position of O ; we can also see it as the precision up to which the agent measures the position of O . In the *Euclidean case*, the region in which the agent circumscribes the object appears to always be of the same size, irrespective of the agent’s configuration with respect to the object. Therefore not moving ends up being an optimal option and the agent will not approach the object without additional extrinsic reward. In the *Projective case*, the agent can ‘zoom’ on this region in order to gain more precision in measuring o ; the configurations of the agent in which this region is magnified are more informative regarding the position of O and therefore preferred by the agent. The only way for the agent to actually zoom onto this area is to approach the location it believes O is likely to be, therefore the agent will end up approaching O .

Remark 2. This particular choice of Markov kernel (Equation 11) allows for an explicit expression of epistemic value

which simplifies the proof of the result; however we expect the result to hold for a larger class of kernels.

V. CONCLUSION AND FUTURE WORK

In this paper we proposed a generative model for environment exploration based on first person perspective in which actions are encoded as changes of perspective. The family of all possible perspective taking on the environment structures the representation of sensory evidence inside the world model of the agent. In other words each family corresponds to a specific perception scheme for the agent. We encoded each of such family as a group acting on the internal world model of the agent, i.e. in the geometric properties of this internal world model. We showed that different geometries induce different behaviors, focusing on two case: when the internal world model of the agent followed Euclidean geometry *versus* projective geometry. The study of world models with projective geometries was motivated by ongoing work in computational psychology aimed at reproducing features of consciousness. Although preliminary, this result contributes to understanding how integrative geometrical processing and principles can play a central role in cybernetics. In our approach, the geometry of the world model links perception and representation with action and behavior.

In the future, we will further examine how the geometry of a latent space intertwines with information processing. One motivation is theoretical, as we would like to assess how geometry changes learning behavior. In this contribution, we have discarded representation learning *per se*, as it was beyond its scope. In future work, we wish to use deep learning to learn group structured representations. However, it is important to note that such approach differs from geometric deep learning [Bronstein et al., 2021] as we do not seek to learn equivariant representations: a group structure will only be considered for the internal world model but none will be presupposed on the observation side. Likewise, we will examine how geometry may play a role in overt and covert attention. Projective geometry induces effects of magnification and focalization that appear immediately relevant for spatial attention, and more generally for contextual salience. Another motivation for this research is more practical, as we would like to use such principles to design virtual and robotic artificial agents mimicking human cognition and behaviors following [Rudrauf et al., 2022], [Rudrauf et al., 2023].

REFERENCES

- [Arleo et al., 2004] Arleo, A., Smeraldi, F., and Gerstner, W. (2004). Cognitive navigation based on nonuniform gabor space sampling, unsupervised growing networks, and reinforcement learning. *IEEE Transactions on Neural Networks*, 15(3):639–652.
- [Bonet and Geffner, 2019] Bonet, B. and Geffner, H. (2019). Learning first-order symbolic representations for planning from the structure of the state space. *arXiv preprint arXiv:1909.05546*.
- [Bronstein et al., 2021] Bronstein, M. M., Bruna, J., Cohen, T., and Velicković, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges.
- [Cao et al., 2022] Cao, W., Zheng, C., Yan, Z., and Xie, W. (2022). Geometric deep learning: progress, applications and challenges. *Science China Information Sciences*, 65(2):126101.

- [Craye et al., 2016] Craye, C., Filliat, D., and Goudou, J.-F. (2016). RLiAc: An exploration policy for online saliency learning on an autonomous mobile robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4877–4884.
- [Da Costa et al., 2020] Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99:102447.
- [Dehaene et al., 2017] Dehaene, S., Lau, H., and Kouider, S. (2017). What is consciousness, and could machines have it? *Science*, 358(6362):486–492.
- [Ferreira et al., 2013] Ferreira, J., Lobo, J., Bessiere, P., Castelo-Branco, M., and Dias, J. (2013). A bayesian framework for active artificial perception. *IEEE transactions on cybernetics*, 43(2):699–711.
- [Friston et al., 2006] Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1):70–87. Theoretical and Computational Neuroscience: Understanding Brain Functions.
- [Friston et al., 2015] Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4):187–214.
- [Geffner and Bonet, 2013] Geffner, H. and Bonet, B. (2013). A concise introduction to models and methods for automated planning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(1):1–141.
- [Gerken et al., 2021] Gerken, J. E., Aronsson, J., Carlsson, O., Linander, H., Ohlsson, F., Petersson, C., and Persson, D. (2021). Geometric deep learning and equivariant neural networks. *arXiv preprint arXiv:2105.13926*.
- [Guo et al., 2021] Guo, Z. D., Azar, M. G., Saade, A., Thakoor, S., Piot, B., Pires, B. A., Valko, M., Mesnard, T., Lattimore, T., and Munos, R. (2021). Geometric entropic exploration. *arXiv preprint arXiv:2101.02055*.
- [Hester and Stone, 2017] Hester, T. and Stone, P. (2017). Intrinsically motivated model learning for developing curious robots. *Artificial Intelligence*, 247:170–186.
- [Kleiner and Tull, 2021] Kleiner, J. and Tull, S. (2021). The mathematical structure of integrated information theory. *Frontiers in Applied Mathematics and Statistics*, 6.
- [Koch et al., 2016] Koch, C., Massimini, M., Boly, M., and Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*.
- [Lang, 2012] Lang, S. (2012). *Algebra*, volume 211. Springer Science & Business Media.
- [Lee and Moore, 2004] Lee, P. Y. and Moore, J. B. (2004). Geometric optimization for 3d pose estimation of quadratic surfaces. In *Conference Record of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers, 2004.*, volume 1, pages 131–135. IEEE.
- [Mashour et al., 2020] Mashour, G. A., Roelfsema, P., Changeux, J.-P., and Dehaene, S. (2020). Conscious processing and the global neuronal workspace hypothesis. *Neuron*, 105(5):776–798.
- [Meng et al., 2017] Meng, Z., Sun, H., Qin, H., Chen, Z., Zhou, C., and Ang, M. H. (2017). Intelligent robotic system for autonomous exploration and active slam in unknown environments. In *2017 IEEE/SICE International Symposium on System Integration (SII)*, pages 651–656. IEEE.
- [Merckling et al., 2022] Merckling, A., Perrin-Gilbert, N., Coninx, A., and Doncieux, S. (2022). Exploratory state representation learning. *Frontiers in Robotics and AI*, 9.
- [Merker et al., 2022] Merker, B., Williford, K., and Rudrauf, D. (2022). The integrated information theory of consciousness: a case of mistaken identity. *Behavioral and Brain Sciences*, 45.
- [Ognibene and Baldassare, 2014] Ognibene, D. and Baldassare, G. (2014). Ecological active vision: four bioinspired principles to integrate bottom-up and adaptive top-down attention tested with a simple camera-arm robot. *IEEE transactions on autonomous mental development*, 7(1):3–25.
- [Ognibene and Demiris, 2013] Ognibene, D. and Demiris, Y. (2013). Towards active event recognition. In *Twenty-Third International Joint Conference on Artificial Intelligence*, pages 2495–2501.
- [Ognibene et al., 2019] Ognibene, D., Mirante, L., and Marchegiani, L. (2019). Proactive intention recognition for joint human-robot search and rescue missions through monte-carlo planning in pomdp environments. In Salichs, M. A., Ge, S. S., Barakova, E. I., Cabibihan, J.-J., Wagner, A. R., Castro-González, Á., and He, H., editors, *Social Robotics*, pages 332–343. Cham. Springer International Publishing.
- [Oudeyer et al., 2007] Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286.
- [Qin et al., 2019] Qin, H., Meng, Z., Meng, W., Chen, X., Sun, H., Lin, F., and Ang, M. H. (2019). Autonomous exploration and mapping system using heterogeneous uavs and uavs in gps-denied environments. *IEEE Transactions on Vehicular Technology*, 68(2):1339–1350.
- [Rudrauf et al., 2017] Rudrauf, D., Bennequin, D., Granic, I., Landini, G., Friston, K., and Williford, K. (2017). A mathematical model of embodied consciousness. *Journal of theoretical biology*, 428:106–131.
- [Rudrauf et al., 2020] Rudrauf, D., Bennequin, D., and Williford, K. (2020). The moon illusion explained by the projective consciousness model. *Journal of Theoretical Biology*, 507:110455.
- [Rudrauf et al., 2023] Rudrauf, D., Sergeant-Perthuis, G., Tisserand, Y., Monnor, T., De Gevigney, V., and Belli, O. (2023). Combining the Projective Consciousness Model and Virtual Humans for immersive psychological research: a proof-of-concept simulating a ToM assessment. *ACM Transactions on Interactive Intelligent Systems*.
- [Rudrauf et al., 2022] Rudrauf, D., Sergeant-Perthuis, G., Belli, O., Tisserand, Y., and Serugendo, G. D. M. (2022). Modeling the subjective perspective of consciousness and its role in the control of behaviours. *Journal of Theoretical Biology*, 534:110957.
- [Sekar et al., 2020] Sekar, R., Rybkin, O., Daniilidis, K., Abbeel, P., Hafner, D., and Pathak, D. (2020). Planning to explore via self-supervised world models. In *International Conference on Machine Learning*, pages 8583–8592. PMLR.
- [Sperati and Baldassarre, 2017] Sperati, V. and Baldassarre, G. (2017). Bio-inspired model learning visual goals and attention skills through contingencies and intrinsic motivations. *IEEE Transactions on Cognitive and Developmental Systems*, 10(2):326–344.
- [Timsit and Sergeant-Perthuis, 2021] Timsit, Y. and Sergeant-Perthuis, G. (2021). Towards the idea of molecular brains. *International Journal of Molecular Sciences*, 22(21).
- [Williford et al., 2018] Williford, K., Bennequin, D., Friston, K., and Rudrauf, D. (2018). The projective consciousness model and phenomenal selfhood. *Frontiers in Psychology*, 9:2571.
- [Williford et al., 2022] Williford, K., Bennequin, D., and Rudrauf, D. (2022). Pre-reflective self-consciousness & projective geometry. *Review of Philosophy and Psychology*, 13(2):365–396.
- [Yang et al., 2021] Yang, F., Lee, D.-H., Keller, J., and Scherer, S. (2021). Graph-based topological exploration planning in large-scale 3d environments. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12730–12736.

APPENDIX

A. Proof of Proposition and Theorem

1) *Proof of Proposition 1: Euclidean case:* For any reference frame of the agent \mathcal{R} , $\phi_{\mathcal{R}}$ is an affine map and $\phi_{\mathcal{R}}^{-1}(0)$ is the center of the agent, \mathcal{P} , in E_3 . In particular, $\phi_{\mathcal{R}^m} \circ \phi_{\mathcal{R}}^{-1}$ is an affine map that sends 0 to 0; it is therefore a linear map. By construction $\psi_m = \phi_{\mathcal{R}^m} \circ \phi_{\mathcal{R}}^{-1}$, therefore ψ_m is a linear map.

Any 3-D affine transformation is encoded by a matrix $M = (m_{i,j}; i, j = 1..3)$ and a vector $(m_{j,4}; j = 1..3)$; let $(m_{i,j}^{\mathcal{R}}; i, j = 1..3)$ be the matrix associated to $\phi_{\mathcal{R}}$ and $(m_{4,j}^{\mathcal{R}}; j = 1..3)$ its vector.

Projective case: $\phi_{\mathcal{R}}^p = \rho \circ \phi_{\mathcal{R}}$ is the projective map with expression in homogeneous coordinates given by the matrix,

$$\begin{pmatrix} m_{1,1} & m_{1,2} & m_{1,3} & m_{1,4} \\ m_{2,1} & m_{2,2} & m_{2,3} & m_{2,4} \\ m_{3,1} & m_{3,2} & m_{3,3} & m_{3,4} \\ 0 & 0 & \gamma & 1 \end{pmatrix}$$

By construction, the transition map in the projective case, ψ_m^p , is $\phi_{\mathcal{R}^m}^p \circ \phi_{\mathcal{R}}^p^{-1}$; it is the composition of two projective transformations, therefore it is a projective transformation.

2) *Proof of Theorem:* We will denote B_y^ϵ the Euclidean ball of radius 1 around $y \in \mathbb{R}^3$, i.e. $B_y^\epsilon = \{x \in \mathbb{R}^3 \mid \|x - y\| \leq \epsilon\}$.

Lemma 1. For any $Q \in \mathbb{P}(W)$, both in Euclidean and Projective cases, for any affine map or projective transformation $\psi : W \rightarrow W$,

$$C(\psi_*Q) = - \int dy Q(\psi^{-1}(B_y^\epsilon)) \ln Q(\psi^{-1}(B_y^\epsilon)) \quad (12)$$

Proof:

$$\begin{aligned} C(\psi_*Q) &= \frac{3}{4\pi\epsilon^3} \times \\ &\int \psi_*Q(dx_1) \int dy 1[x_1 \in B_y^\epsilon] \ln \frac{1[x_1 \in B_y^\epsilon]}{\int \psi_*Q(dx_1) 1[x_1 \in B_y^\epsilon]} \\ &= -\frac{3}{4\pi\epsilon^3} \int dy \ln Q(\psi^{-1}(B_y^\epsilon)) \int \psi_*Q(dx_1) 1[x_1 \in B_y^\epsilon] \\ &= -\frac{3}{4\pi\epsilon^3} \int dy Q(\psi^{-1}(B_y^\epsilon)) \ln Q(\psi^{-1}(B_y^\epsilon)) \end{aligned} \quad (13)$$

□

Proof of Theorem:

Euclidean case: for any set of moves M , and for any $m \in M$, ψ_m is a rotation; therefore for any $y \in W$, $\psi_m^{-1}(B_y^\epsilon) = B_{\psi_m^{-1}(y)}^\epsilon$. Then, for any prior $Q \in \mathbb{P}(W)$,

$$\begin{aligned} C(\psi_{*,m}Q) &= -\frac{3}{4\pi\epsilon^3} \int dy Q(B_{\psi_m^{-1}(y)}^\epsilon) \ln Q(B_{\psi_m^{-1}(y)}^\epsilon) \\ &= -\frac{3}{4\pi\epsilon^3} \int dy Q(B_y^\epsilon) \ln Q(B_y^\epsilon) \end{aligned} \quad (14)$$

In this case, the epistemic value is independent from the change of Euclidean frame, and not moving is a perfectly valid choice for the agent to maximize it, at each time step of the exploration algorithm (Algorithm 1).

Remark 3. The fact that staying still is a valid strategy arises as the agent assumes (or believes) that it has access to the whole configuration space of O . If it knew it had limited access to it, through for example limited sight, we expect the agent would look around until the object O would be in sight, and then stop moving.

Projective case: Consider two projective transformations $\psi, \psi_1 : W \rightarrow W$, if for any $y \in W$,

$$\psi^{-1}(B_y^\epsilon) \subseteq \psi_1^{-1}(B_y^\epsilon) \quad (15)$$

then,

$$-Q(\psi^{-1}(B_y^\epsilon)) \ln Q(\psi^{-1}(B_y^\epsilon)) \quad (16)$$

$$\geq -Q(\psi_1^{-1}(B_y^\epsilon)) \ln Q(\psi_1^{-1}(B_y^\epsilon)) \quad (17)$$

This suggests that the moves that maximize epistemic value are those where ψ_m^{-1} shrinks the zone around $y^o = \rho(\phi_{\mathcal{R}}(o))$, which is the representation of O in the internal

world of the agent. In particular, it means magnifying the zone around $\rho(\phi_{\mathcal{R}^m}(o))$ in the agent's new frame, \mathcal{R}^m , after move m . The only way to do so is to select moves that bring the agent closer to O . Let us denote $y_m^o := \rho(\phi_{\mathcal{R}^m}(o))$. Let us now make the previous argument more formal. We assume that the set of moves M is finite. Let $Q_0 = q_0 d\lambda$ be any initial prior on $W = P_3(\mathbb{R})$, at stating time $t = 0$. After one step, move m_1 is chosen and the agent updates its prior as,

$$q_1(x) \cong 1[x \in B_{y_{m_1}^o}^\epsilon] q_0(x) \quad (18)$$

where \cong means proportional to. The prior we now consider is Q_1 denoted simply as Q . One shows that there is $\alpha > 0$, such that for all $m \in M$, and $\epsilon > 0$ small enough,

$$\begin{aligned} C(\psi_{*,m}Q) &= -\frac{3}{4\pi\epsilon^3} \times \\ &\int dy 1[y \in B_{y_m^o}^{\alpha\epsilon}] Q(\psi_m^{-1}(B_y^\epsilon)) \ln Q(\psi_m^{-1}(B_y^\epsilon)) \end{aligned} \quad (19)$$

$$(20)$$

Let \approx stand for 'approximately equal to' (equal at first order in development in powers of ϵ). Then from the previous statement the summand can be approximated by its value in y_m^o :

$$C(\psi_{*,m}Q) \approx -\alpha^3 Q(\psi_m^{-1}(B_{y_m^o}^\epsilon)) \ln Q(\psi_m^{-1}(B_{y_m^o}^\epsilon)) \quad (21)$$

Furthermore, $Q(\psi_m^{-1}(B_{y_m^o}^\epsilon)) \approx \frac{4\pi\epsilon^3}{3} \frac{q_1(y^o)}{|\det \nabla \psi_m|(y^o)}$, where $|\det \nabla \psi_m|(y^o)$ is the absolute value of the Jacobian determinant of ψ_m at y^o . The epistemic value is maximized when $|\det \nabla \psi_m|(y^o)$ is maximized. By definition, $\psi_m = \rho \circ \phi_{\mathcal{R}^m} \circ \phi_{\mathcal{R}^{-1}} \circ \rho^{-1}$, therefore, by the chain rule of differentiation

$$\begin{aligned} |\det \nabla \psi_m|(y^o) &= |\det \nabla \rho|(\phi_{\mathcal{R}^m}(o)) \cdot |\det \phi_{\mathcal{R}^m}|(o) \cdot |\det \nabla[\phi_{\mathcal{R}^{-1}} \circ \rho^{-1}]|(y^o) \end{aligned} \quad (22)$$

Let us make explicit each terms in the previous equation. $\phi_{\mathcal{R}^m}$ is a rigid movement therefore, $|\det \phi_{\mathcal{R}^m}|(o) = 1$. $|\det \nabla[\phi_{\mathcal{R}^{-1}} \circ \rho^{-1}]|(y^o)$ does not depend on m so we can label it as a constant C . $\phi_{\mathcal{R}^m}(o)$ is the coordinate of O in the Euclidean frame \mathcal{R}^m ; let us denote (x^m, y^m, z^m) these coordinates, i.e. $(x^m, y^m, z^m) := \phi_{\mathcal{R}^m}(o)$. Then,

$$|\det \nabla \rho|(x^m, y^m, z^m) = \frac{1}{(\gamma z^m + 1)^4} \quad (23)$$

Therefore, $|\det \nabla \psi_m|(y^o) = C \frac{1}{(\gamma z^m + 1)^4}$.

As we assumed that for any move $m \in M$, the object O is always in front of the agent, then $z^m \geq 0$; in this case, z^m is also the distance of the agent to the object. Epistemic value is maximized when z^m is minimized and therefore the agent selects moves that reduce its distance to the object. Denote one of such move m^* ; the argument then loops back with the new reference frame \mathcal{R}^{m^*} and updated belief $q \leftarrow \psi_{m^*} q|_{y_{m^o}^o}$. □