

Predicting the consequence of action in digital control state spaces

Emmanuel Daucé

▶ To cite this version:

Emmanuel Daucé. Predicting the consequence of action in digital control state spaces. 2016. hal-01374155

HAL Id: hal-01374155 https://hal.science/hal-01374155

Preprint submitted on 29 Sep 2016

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 4.0 International License

Predicting the consequence of action in digital control state spaces

Emmanuel Daucé

September 29, 2016

Abstract

The objective of this dissertation is to shed light on some fundamental impediments in learning control laws in continuous state spaces. In particular, if one wants to build artificial devices capable to learn motor tasks the same way they learn to classify signals and images, one needs to establish control rules that *do not necessitate* comparisons between quantities of the surrounding space. We propose, in that context, to take inspiration from the "end effector control" principle, as suggested by neuroscience studies, as opposed to the "displacement control" principle used in the classical control theory.

1 Problem statement

Most machine learning techniques rely on a pattern matching principle, with representative vectors constructed from many passes over large datasets of examples. The resulting set of prototypes then works as a dictionary of shapes, providing a reduced description of the input data. The projection of the input vectors on this redescription space is then expected to facilitate further data processing and interpretation. Learning then coincides with the construction of a vocabulary of shapes that serve as a key to interpret the data. There are however few examples in the literature where auto-encoding algorithms end up in learning a control law in a continuous actions space. This difficulty seems to stem from the conceptual separation between, on one side, the dynamical systems approach and, on the other side, the pattern matching principles.

From the dynamical system perspective, a usual simplification considers a controller and its environment as a single dynamical system, with a reciprocal coupling between an internal part devoted to data processing and control and an external part consisting in material objects. The transformations from the internal state space toward the external state space are done through actuators. Reciprocally, the transformations from the external state space toward the internal state space are done through captors . When the question of learning and updating actions is considered, some inconsistencies may however rapidly show off.

2 Quick mathematic overview

2.1 Closed-loop control

Closed-loop control systems are generally described using two state spaces, namely the controller space \mathcal{X}_{in} , whose state at time t is $\boldsymbol{x}_t^{in} \in \mathcal{X}_{in}$ and the physical space \mathcal{X}_{out} , whose state at time t is $\boldsymbol{x}_t^{\text{out}} \in \mathcal{X}_{\text{out}}$. The general evolution of the system is provided by a set of 4 stochastic update equations, i.e.:

$$\boldsymbol{x}_{t+dt}^{\text{out}} \sim P_{\text{out}}(\boldsymbol{x}_t^{\text{out}}, \boldsymbol{u}_t) \tag{1}$$

$$\boldsymbol{I}_t \sim Q_{\rm in}(\boldsymbol{x}_t^{\rm out}) \tag{2}$$

$$\boldsymbol{x}_{t+dt}^{\text{in}} \sim P_{\text{in}}(\boldsymbol{x}_t^{\text{in}}, \boldsymbol{I}_t) \tag{3}$$

$$\boldsymbol{u}_t \sim Q_{\text{out}}(\boldsymbol{x}_t^{\text{in}}) \tag{4}$$

where eq. (1) is the external update, eq. (2) the sensors actualization, eq. (3) the internal update and eq. (4) the motor realization, with I_t the sensory input, u_t the command, P_{out} the (external) generative process, Q_{in} the measure, P_{in} the internal program, and Q_{out} the command transmission chain.

2.2 Active sensing

In most cases, the external process is considered hidden, and only the sensory consequence I_t of the hidden state is observable. Then, learning is mainly inferring the parameters Θ that describe both the (inferred) generative and observation processes, \hat{P} and \hat{Q} , i.e.

$$\max_{\boldsymbol{x},\Theta} L(\boldsymbol{x},\Theta) \tag{5}$$

with $L(\boldsymbol{x}, \Theta)$ the likelihood of the observation (given parameters Θ and a past estimation on \boldsymbol{x}_t):

$$L(\boldsymbol{x},\Theta) = \hat{Q}(\boldsymbol{I}_{t+dt}|\boldsymbol{x},\Theta)\hat{P}(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{u}_t,\Theta)$$
(6)

which is intractable in the general case and must be decomposed in a marginalization "expectation" step:

$$\hat{f}(\boldsymbol{x}) = L(\boldsymbol{x}|\Theta)$$

and a parameter maximization step (Dempster et al., 1977):

$$\max_{\Theta'} \int L(\boldsymbol{x},\Theta') \hat{f}(\boldsymbol{x}) d\boldsymbol{x}$$

The active inference paradigm, as proposed by (Friston & Kiebel, 2009), mainly innovate from this classical scheme by including the command in the optimization, i.e.

$$\max_{\boldsymbol{x},\boldsymbol{u},\Theta} L(\boldsymbol{x},\boldsymbol{u},\Theta) \tag{7}$$

which tells in short that the meaning of action is to participate in the likelihood maximization process by sampling regions of the sensory space that should reduce the uncertainty about the hidden state. The choice of action u_t is there related to the maximization of the likelihood of the future hidden state estimate x_{t+dt} in the next observation step (see eq.6), with action outcome estimated from current \hat{P} and \hat{Q} . A particular modeling effort needs thus to be put on estimating the sensory prediction of action (forward model of action), which is the essence of the predictive coding scheme. The approach, however, doesn't provide a insight on which motor control schemes are effectively suitable to implement such prediction.

Classical motor control schemes consider for instance actions as changes taking place in the physical (Euclidian) space where objects are defined by specific 3D coordinates. Active sensing means thus moving the eyes, neck and trunk in order to estimate the nature and location of the

objects that compose the scene. Many referentials may be defined in relation to the different degrees of freedom of the main articulations taking part in scanning the visual environment, with e.g. the egocentric referential (trunk referential), the head-centered referential (relative to the direction of the head) and the retinocentric referential (relative to the direction of the sight). Since relevant referentials and DOFs are defined, an action mainly reduces to a referential change on an invariant scene. For instance, moving the eyes to the right translates the visual field to the left in the direction opposite to the angular command. The linear transformation associated with action has a predictable consequence on where the objects will project on it after change.

3 Learning to predict

Realizing and implementing those predictive transformations in a neuronal circuit (suitable for learning) happens however to be more difficult than expected. This difficulty is mainly due to the absence of a definite (operational) distance in the control space that implements learning, on contrary to the physical space that implements action.

When learning is considered, the universality of a learning scheme (its capability to learn any transformation) mostly relies on using non-linear transfer functions at one step (at least) of the data processing pathway, with, e.g. threshold functions, argmax, sigmoid, kernels etc. (Haykin, 1999). The many modern learning techniques used in pattern recognition rely on those non-linear steps to gain enough discriminative power and extract common features from many samples. The pattern matching approach thus does not guarantee the proportionality of treatments between close-by and distant parts of the physical space. In other words, the distance that can be defined into the redescription space may not respect the distance and topology of the initial space. The redescription space is generally considered as a "qualitative" space (or "digital" space), containing features, by contrast with the external "quantitative" space.

In any case, as soon as the initial distances between the physical objects are not preserved, predicting the consequence of an action may necessitate a specific training (instead of considering explicit geometric transformations). If every sensory modality of the physical space is distributed on many neurons (whose receptive field may only respond to a particular interval or "pixel"), then simple transformations such as translations, rotations (and even identity) need to be learned "by heart" (Pouget & Sejnowski, 1997) i.e. may not result from local operations on graded firing rate inputs. Even feasible at arbitrary precision in theory, learning "by heart" the sensory consequences of action may thus suffers from a superlinear (squared) complexity when distances or referential transformations are considered. Learning, for instance, to predict the visual field consequence of visual orientation commands would necessitate to learn every combination of objects under the many possible eye and head orientation with trunk-head and head-eye referential transfers, which may expectedly be very costly (Pouget & Sejnowski, 1997; Pouget et al., 2002).

4 Alternate view

If one expects build artificial devices that learn motor tasks the same way they may learn to classify images and signals, one needs to reconsider motor control from the basements, in a way that renders compatible the auto-encoding and the predictive coding principles. This first implies a digital description of motor commands and motor acts, i.e a decomposition of the task space in elementary components, the same way the sensory state space may de decomposed in elementary constituents. In counterpart, the isometric transformations and quantitative comparisons used in classical control theory should be eliminated when possible to avoid combinatorial explosion in learning "by heart" the linear (euclidian) transformations and distances in the controller state space.

In order to do predictive coding without learning transformations and distances, what one needs to do is directly learn a mapping from the motor command space toward the sensory space. In complex bodies made of many joints, a motor space is made of many dimensions, and the spatial disposition of the different segments of the body (posture) are unambiguously defined by their respective coordinates. In particular, the orientation of the trunk, the neck and the sight (relatively to the ground) are unambiguously set.

When visual prediction is considered, a simple mapping of the postural information toward the visual field is possible. If, for simplicity, only the hand and vision are considered, many visual orientations may produce predictable changes in the visual field for the same invariant hand position. Reversely, moving the hand or the fingers while keeping the eyes fixed also induces predictable changes. The visual consequences of possible combination of eyes and hand position may thus be learned in a straightforward way without intermediary transformation.

This has important consequences on the way motor control should be considered. A traditional view considers motor acts as displacements that are proportional to a certain positional error detected by the senses. In that perspective, a *difference* between the desired external state x^* and the actual estimate \hat{x} needs to be calculated to issue a displacement command u. This is the *error-based* control framework and the *displacement-based* command. There is thus no direct mapping from a displacement u toward its visual consequence I. One needs to pass through an intermediary object, the "external state" \hat{x} to perform a prediction.

An alternate view is to consider the *end-effector*-based (or "ballistic") control framework. Direct stimulations of the motor cortex have shown converging movements toward a same final posture whatever the initial position (Graziano et al., 2002). The presence of motor vector fields were also observed in many species, and vector-field-based control was theorized by e.g. (Mussa-Ivaldi & Solla, 2004). The vector field approach to control considers the end-position of the limb as the relevant control level over which the brain operates. An end-position mainly consists in the distal spatial coordinates of certain effector (tentacle, tongue or arm in animals). A motor vector field is a definite set of flexion/extension commands whose combination mechanically conducts the effector toward a particular end-point. To control an end-effector in 2D coordinates for instance, a simple referential can be set with three fields (whose end-points are not colinear). Then, controlling the effector within the triangle made by the three end-points is done by linear combination. A dot in the coordinate system is defined by three weights affected to the different fields. This principle is generalizable to 3D physical spaces (four dots at least) and of course to the more complex (many degrees-based) postural states like, e.g., the fingers of the hand.

An interesting consequence, when the arm is considered, is the correspondence of every motor command with a particular position of the hand (end-effector) in the peripheral space. One can thus postulate the hand-and-arm-controlling motor cortex to be organized like a map, where every particular position in the map codes for a specific "motor primitive" (a linear combination of motor fields) that may conduct the hand toward a particular position in the peripheral space. This (approximate) correspondence between a motor command and a corresponding 3D coordinate in peripheral space renders possible a direct prediction from the command $u_{\rm arm}$ toward its visually expected consequence I (as far as the hand is considered).

Despite fewer biological evidence, the very same reasoning can be applied on eye orientation effectors. Orienting the eye toward the peripheral space can be seen as setting the angular direction u_{eye} of the central part of the retina (fovea) in rotational coordinates. As far as only the hand and the sight are considered, and ignoring all possible luminance and light incidence

changes, a direct mapping from $(\boldsymbol{u}_{arm}, \boldsymbol{u}_{eye})$ toward $\boldsymbol{I}_{visual field}$ can be considered for learning, from a systematic scan of every possible eye and arm command. This is the equivalent of the "babbling" seen in young humans and animals, also addressed in epigenetic robotics (Andry et al., 2001).

A direct mapping and prediction from the action space toward the sensory space is thus made possible in a digital control space system provided specific ballistic control principles are considered. This first and preliminary insight is not a definite answer to the addressed question, but may help to reorient and radically simplify the learning control framework in order to address more challenging questions.

5 Concluding remarks

A first remark, as soon as brain modeling is considered, is the probable combination of several control principles serving different purposes in biological systems. If end-effector control happen to have interesting predictive capabilities, a counterpart is a certain lack of adaptivity in changing conditions. Changes in body orientation, muscle fatigue, or any effector flaw may have a dramatic consequence on the final precision of movement. A parallel and fast-adapting control is expected to take place in the cerebellum, with the capability to deviate the motor orders on the fly at the time it is sent to the muscles. The cerebellum has been shown to be sensible to the motor error, and capable to calculate a displacement correction (Wolpert et al., 1998; Haruno et al., 2001; Dean et al., 2010). The huge number of neurons implied in that calculation may reflect the combinatorial problem encountered as soon as positional differences need to be calculated to set a command.

A follow-up of the hand/eye end-effector learning perspective is the learning the manipulation of objects by hand. The many different visual orientations of the different objects in the surrounding space are directly linkable to the way they can be handled by the hand. This is the essence of object affordances postulated by Gibson (Gibson, 1979). Handling an object by hand is here an essential part of its identity (or category from the machine learning perspective). Scanning an object identity means here trying different manipulation schemes and inferring their visual and proprioceptive predictible consequences. An object in space is here assimilated with a grasping operation, with a hand aperture and orientation compatible with its actual position and orientation. From this motor perspective, the many visual perspective may associate unambiguously with a single invariant position in space.

If any single object of the visual field may, in principle, be learned from this end-effector prospect, the detailed linkage and decomposition in features and components is still an openended question. For instance, as soon as hand manipulation is concerned, a detailed and exhaustive exploration of each and every end-effector configuration should reveal impossible in practice, and a random-driven exploration may rather be considered. Surprise-driven exploration (with unexpected changes preferred to known expectations) may there be formalized and adapted to that specific prospect.

A final perspective may be the modeling the construction of objects, i.e. the assimilation (or generalization) of many perspectives in a single motor framework, as postulated by Piaget in the sensori-motor stage (Piaget, 1973). This generalization capability, from multiple views toward a single concept, may correspond to the effective interpretation of an external "cause" \hat{x} to the actual sensation.

References

- Andry, P., Gaussier, P., Moga, S., Banquet, J.-P. & Nadel, J. (2001). Learning and communication via imitation: An autonomous robot perspective. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on, 31(5), 431–442.
- Buesing, L., Bill, J., Nessler, B. & Maass, W. (2011). Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput. Biol*, $\gamma(11)$, e1002211.
- Dean, P., Porrill, J., Ekerot, C.-F. & Jörntell, H. (2010). The cerebellar microcircuit as an adaptive filter: experimental and computational evidence. *Nature Reviews Neuroscience*, 11(1), 30–43.
- Dempster, A. P., Laird, N. M. & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society. Series B (Methodological) (pp. 1–38).
- Friston, K. & Kiebel, S. (2009). Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society B: Biological Sciences, 364(1521), 1211–1221.
- Gibson, J. J. (1979). The ecological approach to visual perception. Boston: Houghton-Mifflin.
- Graziano, M. S., Taylor, C. S. & Moore, T. (2002). Complex movements evoked by microstimulation of precentral cortex. *Neuron*, 34(5), 841–851.
- Haruno, M., Wolpert, D. & Kawato, M. (2001). Mosaic model for sensorimotor learning and control. Neural computation, 13(10), 2201–2220.
- Haykin, S. (1999). Neural Networks : a comprehensive fundation. Prentice-Hall.
- Mussa-Ivaldi, F. A. & Solla, S. A. (2004). Neural primitives for motion control. Oceanic Engineering, IEEE Journal of, 29(3), 640–650.
- Piaget, J. (1973). Biologie et connaissance. Gallimard idées.
- Pouget, A., Deneve, S. & Duhamel, J.-R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience*, 3(9), 741–747.
- Pouget, A. & Sejnowski, T. (1997). Spatial transformations in the parietal cortex using basis functions. Cognitive Neuroscience, Journal of, 9(2), 222–237.
- Wolpert, D. M., Miall, R. C. & Kawato, M. (1998). Internal models in the cerebellum. Trends in cognitive sciences, 2(9), 338–347.