Influence of the geometry of the feature space on curiosity based exploration

Grégoire Sergeant-Perthuis[1], **Nils Ruet**[2], **David Rudrauf**[2], **Dimitri Ognibene**[3], **Yvain Tisserand**[4]

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. This collection of transformations makes 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 acting on the space, 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 (mutual information), a quantity known in active inference for driving an agent's curiosity about its environment, and the impact on 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.

1 Introduction

In previous work it has been shown that geometrically constrained active inference can be used as a framework to understand and model central aspects of human spatial consciousness (see the Projective Consciousness Model (PCM) (1; 2)). Consciousness accesses and represents multimodal information through a Global Workspace (3) 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 (2). In (1; 4; 2; 5; 6), it was hypothesized that such internal representation space is geometrically structured as a 3-D projective space, denoted $P_3(\mathbb{R})$. A change of perspective then corresponds to the choice of a projective transformation ψ, i.e. the action of an element of the group of projective transformations PGL_4. A projective transformation is a linear isomorphism $M_\psi \in GL_4(\mathbb{R})$ up to a multiplicative constant. The model yielded an explanation for the Moon illusion (7) 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 (2)). 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 (8; 9; 10; 3; 11). 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

*[1] LCQB Sorbonne Université & OURAGAN team, Inria Paris gregoire.sergeant-perthuis@sorbonne-universite.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.

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

cybernetics remain to be fully formulated (see (2)). When applying the PCM to agency and action planning, previous work (2; 5) lacked a formal framework. In recent work (12) such ideas have been extended and replaced within stochastic optimal control with state spaces, *world models*, structured as homogeneous spaces. This formal framework allows to further theoretical and algorithmic inquiries, as we propose to do in this article.

2 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 (see Figure 2 Appendix D). 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 perspective that each movement of the agent induces; a perspective on the world is analogous to a choice of a (non linear) coordinate system, when the agent moves this coordinate system changes and the changes of coordinates are the group actions. 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 $P_3(\mathbb{R})$

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 (mutual information), the maximization of which captures curiosity-based exploration. In Section 2.2, we explain how epistemic value is defined for group structured internal representations. In Section 2.3 we give the details of the exploration algorithm.

2.1 Group structured world model and relating the 'real world' to the 'internal world model'

In general, groups can be seen as encoding transformations or changes of frame (see Definition 3 Appendix A). We call a group structured world model, a world model provided with a group action; we now make this statement formal.

Definition 1. *W is a group structured world model for the group G when there is a map $h : G \times W \to W$ 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 $x \in W$*

The group structured world model is usually called a $G-$space, but we will keep this denomination for now to recall the context we consider, that of an agent exploring its environment with noisy sensors. In the *Euclidean case* the group structured world model, W, is the 3-D vector space $\mathbb{R}^3$: the Euclidean space with additional information of a center and three axis on front, on the right and above; it is structured by the group of invertible affine transformations $GL_3(\mathbb{R}) \ltimes \mathbb{R}^3$; we will only consider its subgroup $E(3)$, the Euclidean group of invertible affine transformations from the space into itself preserving the Euclidean metric. 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 $PGL(\mathbb{R}^3)$.

We assume that the 'real world' is the 3-D Euclidean space, E_3; it is the space that is used to set up the experiment: the conformation of the agent and of the place of the object. We assume that the 'real world' comes with an Euclidean frame $\mathcal{R}_E$, i.e. a point $\mathcal{C}$ 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 $\mathbb{R}^3$, $\mathcal{C}$ 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 (2). The (projective) transformation $\phi^p_{\mathcal{R}}$, from E_3 to W, which relates the 'real world' to the 'internal world model' in the projective case, makes use of a projective map ρ, detailed in Appendix A, and then $\phi^p_{\mathcal{R}} := \rho \circ \phi_{\mathcal{R}}$.

2.2 Beliefs, policies and epistemic value

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. These beliefs are updated according to noisy sensory observations of the position of O. 'Markov Kernels' can be used to formalize noisy sensors (Definition 3 Appendix A). Let us recall their definition. 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 $x, y \in W$, $P_{X,Y}(dx, dy) := p_{Y|X}(y|x)q_X(x)dxdy$, 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

$$Q_{X|y^o} = \frac{p_{Y|X}(y^o|x)q_X(x)dx}{\int_{x\in W} p_{Y|X}(y^o|x)q_X(x)dx} \tag{1}$$

Policies: The agent has a set of moves it can make M; a move $m \in M$ is associated to the action of the group element $\psi_m : W \to W$ (Appendix 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 $m \in M$, $A \subseteq W$ a (Borel) subset of W, and $x_0 \in W$, $p_{X_1|X_0,m}(A|x_0, m) = 1[\psi_m(x_0) \in A]$. Each move m spreads a prior Q_X on X_0 into the following prior on X_1: $\forall A \in \mathcal{B}(W)$,

$$\psi_{m,*}Q_X(A) := \int 1[\psi_m(x_0) \in A]q_X(x_0)dx = Q_X(\psi_m^{-1}A) \tag{2}$$

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 3.

Epistemic value: Following (13),

Definition 2 (Epistemic Value)**.** *For any probability measure $Q_X \in \mathbb{P}(W)$, the epistemic value of this measure is:*

$$\begin{aligned} C(Q_X) :=& \mathbb{E}_{P_Y}\left[H(P_{X|Y}|Q_X)\right] && (3)\\ =& \int p_Y(y)dy \int p_{X|Y}(x|y) \ln \frac{p_{X|Y}(x|y)}{q_X(x)}dx && (4)\end{aligned}$$

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

Reexpressing Equation 3, it becomes apparent that epistemic value is simply a mutual information:

$$C(Q_X) = \int p_{X,Y} \ln \frac{p_{X,Y}(x,y)}{p_Y(y)q_X(x)}dxdy \tag{5}$$

We propose to define the epistemic value of move m as the epistemic value of the induced prior on X_1, $C(m) := C(\psi_{m,*}Q_X)$.

2.3 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 $\psi_m : W \to W$ 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, $\overline{m}$. a executes the move $\overline{m}$ 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 $\psi_{\overline{m},*}Q_X^0$ to the distribution conditioned on the observation: $\left(\psi_{\overline{m},*}Q_X^0\right)_{|y^o_{\overline{m}}}$. The process is iterated with this new prior. The exploration algorithm is summarized in Appendix Algorithm 1.

3 Theoretical predictions

We prove that 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. We consider the following noisy sensor, for any $x, y \in \mathbb{R}^3$, $P_{Y|X}(y|x) = \frac{3}{4\pi\epsilon^3}\, 1[\|x - y\| \leq \epsilon]$ where $\|.\|$ 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 1. *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 B.1. 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. □

Implementation of this experiment are provided in Appendix C.

4 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. One motivation is theoretical, as we would like to assess how geometry changes learning behavior and information processing. 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 (14; 15) 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. 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 (2; 5).

References

[1] D. Rudrauf, D. Bennequin, I. Granic, G. Landini, K. Friston, and K. Williford, "A mathematical model of embodied consciousness," *Journal of Theoretical Biology*, vol. 428, pp. 106–131, 2017.

[2] D. Rudrauf, G. Sergeant-Perthuis, O. Belli, Y. Tisserand, and G. D. M. Serugendo, "Modeling the subjective perspective of consciousness and its role in the control of behaviours," *Journal of Theoretical Biology*, vol. 534, p. 110957, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0022519321003763

[3] S. Dehaene, H. Lau, and S. Kouider, "What is consciousness, and could machines have it?" *Science*, vol. 358, no. 6362, pp. 486–492, 2017.

[4] K. Williford, D. Bennequin, K. Friston, and D. Rudrauf, "The projective consciousness model and phenomenal selfhood," *Frontiers in Psychology*, vol. 9, p. 2571, 2018.

[5] D. Rudrauf, G. Sergeant-Perthuis, Y. Tisserand, T. Monnor, V. De Gevigney, and O. Belli, "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*, 2023.

[6] K. Williford, D. Bennequin, and D. Rudrauf, "Pre-reflective self-consciousness & projective geometry," *Review of Philosophy and Psychology*, vol. 13, no. 2, pp. 365–396, 2022.

[7] D. Rudrauf, D. Bennequin, and K. Williford, "The moon illusion explained by the projective consciousness model," *Journal of Theoretical Biology*, vol. 507, p. 110455, 2020.

[8] C. Koch, M. Massimini, M. Boly, and G. Tononi, "Neural correlates of consciousness: progress and problems," *Nature Reviews Neuroscience*, 2016.

[9] J. Kleiner and S. Tull, "The mathematical structure of integrated information theory," *Frontiers in Applied Mathematics and Statistics*, vol. 6, 2021. [Online]. Available: https://www.frontiersin.org/article/10.3389/fams.2020.602973

[10] G. A. Mashour, P. Roelfsema, J.-P. Changeux, and S. Dehaene, "Conscious processing and the global neuronal workspace hypothesis," *Neuron*, vol. 105, no. 5, pp. 776–798, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0896627320300520

[11] B. Merker, K. Williford, and D. Rudrauf, "The integrated information theory of consciousness: a case of mistaken identity," *Behavioral and Brain Sciences*, vol. 45, 2022.

[12] D. Rudrauf, G. Sergeant-Perthuis, Y. Tisserand, G. Poloudenny, K. Williford, and M.-A. Amorim, "The Projective Consciousness Model: projective geometry at the core of consciousness and the integration of perception, imagination, motivation, emotion, social cognition and action ," 2023, to appear Brain sciences MDPI.

[13] K. Friston, F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, and G. Pezzulo, "Active inference and epistemic value," *Cognitive Neuroscience*, vol. 6, no. 4, pp. 187–214, 2015.

[14] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges," 2021. [Online]. Available: https://arxiv.org/abs/2104.13478

[15] G. Sergeant-Perthuis, J. Maier, J. Bruna, and E. Oyallon, “On non-linear operators for geometric deep learning,” in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 10 984–10 995. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/474815daf1d4096ff78b7e4fdd2086a5-Paper-Conference.pdf

[16] S. Lang, *Algebra*. Springer Science & Business Media, 2012, vol. 211.

[17] M. L. Eaton, “Multivariate statistics: A vector space approach,” *Lecture Notes-Monograph Series*, vol. 53, pp. i–512, 2007. [Online]. Available: http://www.jstor.org/stable/20461449

[18] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience, 2006.

A Definitions

A.1 Group and representing the real world in the internal world

Let us first recall what a group is.

Definition 3 (Group, §2 Chapter 1 (16))**.** *A group is a set G with an operation $. : G \times G \to G$ 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, $g.g^{-1} = g^{-1}.g = e$.*

For a given solid frame $\mathcal{R}$ of the agent, the map $\phi^p_{\mathcal{R}} : E_3 \to W$ maps the real world to the internal world of the agent. Such a map depends on a (arbitrary) choice of a projective transformation ρ that allows to embed the moves of the agent, i.e. the Euclidean group E_3, inside the projective linear group. In Proposition A.1 of(2), moves of the agents are related to projective transformations using a set of axioms based on phenomenal experience; we propose that the subject is centered in its internal world, and around it proportions are faithful to reality. These axioms imply a unique family of projective transformations, denoted as ρ_γ in the main text, which we will now define.

Definition 4 (Projective transformation)**.** *Let for any $(x, y, z) \in \mathbb{R}^3$,*

$$\rho(x, y, z) = \left(\frac{x}{\gamma z + 1}, \frac{y}{\gamma z + 1}, \frac{z}{\gamma z + 1} \right) \tag{6}$$

The map that relates real world to internal world is then $\phi^p_{\mathcal{R}} := \rho \circ \phi_{\mathcal{R}}$.

A.2 Markov kernel

Markov kernels play a central role in agency because actions can introduce control errors, and similarly, sensory inputs can be noisy. Sensors may introduce noise, for instance, when one focuses on a specific location in the hope of locating an object. In these situations, sensor-related uncertainty is inevitable, leading to a certain "radius" of uncertainty (typically represented by the standard deviation) around the expected object position. We model this uncertainty using stochastic maps also called Markov kernels, which we will define now.

Definition 5 (Markov Kernel)**.** *A ‘Markov Kernel’ Π from Ω_1 to Ω is a map $\Pi : \Omega \times \Omega_1 \to [0, 1]$ such that for any $\omega_1 \in \Omega_1$, $\sum_{\omega \in \Omega} P(\omega|\omega_1) = 1$, i.e. a map that sends any $\omega_1 \in \Omega_1$ to a probability measure $\Pi_{|\omega_1} \in \mathbb{P}(\Omega)$.*

B Proposition and Theorem and proofs

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 invertible affine transformations $\psi_m \in E(3)$ from the ‘internal world model’ to itself. In the* Projective case *it induces a projective transformation, $\psi_m \in PGL(\mathbb{R}^3)$.*

Proof. Euclidean case:

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 $(m^{\mathcal{R}}_{i,j}; i, j = 1..3)$ be the matrix associated to $\phi_{\mathcal{R}}$ and $(m^{\mathcal{R}}_{4,j}; j = 1..3)$ its vector.

Projective case: $\phi^p_{\mathcal{R}} = \rho \circ \phi_{\mathcal{R}}$ 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^p_m, is $\phi^p_{\mathcal{R}^m} \circ {\phi^p_{\mathcal{R}}}^{-1}$; it is the composition of two projective transformations, therefore it is a projective transformation.

□

B.1 Proof of Main Theorem

We will denote B^ϵ_y the Euclidean ball of radius 1 around $y \in \mathbb{R}^3$, i.e. $B^\epsilon_y = \{x \in \mathbb{R}^3 | \quad \|x - y\| \leq \epsilon\}$.

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

$$C(\psi_* Q) = -\int dy Q(\psi^{-1}(B^\epsilon_y)) \ln Q(\psi^{-1}(B^\epsilon_y)) \tag{7}$$

Proof.

$$\begin{aligned} C(\psi_* Q) = &\frac{3}{4\pi\epsilon^3} \times \\ &\int \psi_* Q(dx_1) \int dy\, 1[x_1 \in B^\epsilon_y] \ln \frac{1[x_1 \in B^\epsilon_y]}{\int \psi_* Q(dx_1)\, 1[x_1 \in B^\epsilon_y]} \\ \\ &= -\frac{3}{4\pi\epsilon^3} \int dy \ln Q(\psi^{-1}(B^\epsilon_y) \int \psi_* Q(dx_1)\, 1[x_1 \in B^\epsilon_y] \\ &= -\frac{3}{4\pi\epsilon^3} \int dy Q(\psi^{-1}(B^\epsilon_y)) \ln Q(\psi^{-1}(B^\epsilon_y)) \end{aligned} \tag{8}$$

□

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 $y \in W$, $\psi^{-1}_m(B^\epsilon_y) = B^\epsilon_{\psi^{-1}_m(y)}$. Then, for any prior $Q \in \mathbb{P}(W)$,

$$\begin{aligned} C(\psi_{m,*} Q) &= -\frac{3}{4\pi\epsilon^3} \int dy Q(B^\epsilon_{\psi^{-1}_m(y)}) \ln Q(B^\epsilon_{\psi^{-1}_m(y)}) \\ &= -\frac{3}{4\pi\epsilon^3} \int dy Q(B^\epsilon_y) \ln Q(B^\epsilon_y) \end{aligned} \tag{9}$$

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 1. *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 $\psi, \psi_1 : W \to W$, if for any $y \in W$,

$$\psi^{-1}(B_y^\epsilon) \subseteq \psi_1^{-1}(B_y^\epsilon) \tag{10}$$

then,

$$-Q(\psi^{-1}(B_y^\epsilon)) \ln Q(\psi^{-1}(B_y^\epsilon)) \tag{11}$$

$$\geq -Q(\psi_1^{-1}(B_y^\epsilon)) \ln Q(\psi_1^{-1}(B_y^\epsilon)) \tag{12}$$

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,

$$q_1(x) \cong 1[x \in B_{y_{m_1}^o}^\epsilon] q_0(x) \tag{13}$$

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,

$$C(\psi_{m,*}Q) = -\frac{3}{4\pi\epsilon^3} \times$$

$$\int dy\, 1[y \in B_{y_m^o}^{\alpha\epsilon}] Q(\psi_m^{-1}(B_y^\epsilon)) \ln Q(\psi_m^{-1}(B_y^\epsilon)) \tag{14}$$

$$\tag{15}$$

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:

$$C(\psi_{m,*}Q) \approx -\alpha^3 Q(\psi_m^{-1}(B_{y_m^o}^\epsilon)) \ln Q(\psi_m^{-1}(B_{y_m^o}^\epsilon)) \tag{16}$$

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

$$|\det \nabla\psi_m|(y^o)$$

$$= |\det \nabla\rho|(\phi_{\mathcal{R}^m}(o)).|\det \phi_{\mathcal{R}^m}|(o).|\det \nabla[\phi_{\mathcal{R}}^{-1} \circ \rho^{-1}]|(y^o) \tag{17}$$

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,

$$|\det \nabla\rho|(x^m, y^m, z^m) = \frac{1}{(\gamma z^m + 1)^4} \tag{18}$$

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 $\overline{m}$; the argument then loops back with the new reference frame $\mathcal{R}^{\overline{m}}$ and updated belief $q \leftarrow \psi_{\overline{m},*}q_{|y^o_m}$. □

C Implementation

Algorithm 1 is implemented in the following manner. Beliefs and the Markov kernel corresponding to sensors were considered to be multivariate normal distributions, that is $P_{Y|X} \sim \mathcal{N}(\mu_{Y|X},\, \Sigma_{Y|X})$ and $Q_X \sim \mathcal{N}(\mu_X,\, \Sigma_X)$. Belief update through the action of a group was approximated using a Gaussian distribution; a projective transformation changes a Gaussian distribution into a non Gaussian one which is difficult to describe. Therefore we replace this non-Gaussian distribution with a Gaussian distribution with same mean and variance. We assumed $\mu_{Y|X} = x$ (which implies $\mu_y = \mu_x$) and $\Sigma_{Y|X} = \epsilon^2\mathbb{I}$ where $\mathbb{I}$ is the identity matrix and $\epsilon > 0$ a positive real number. As a result, for a given observation y^o, $Q_{X|y^o}$ and $C(\psi_{m,*}Q_X)$ can be computed efficiently. The joint distribution P on X, Y is a Gaussian distribution:

$$P(x, y) = p(y|x)p(x)$$

$$P(x, y) \sim \mathcal{N}(\mu_{X,Y},\, \Sigma_{X,Y})$$

with $\mu_{X,Y} = (\mu_X, \mu_X)$ and

$$\Sigma_{X,Y} = \begin{pmatrix} \Sigma_{XX} & \Sigma_{XX} \\ \Sigma_{XX} & \epsilon^2\mathbb{I} + \Sigma_{XX} \end{pmatrix} \tag{19}$$

The variance of Y is $\Sigma_{YY} = \epsilon^2\mathbb{I} + \Sigma_{XX}$.

The joint distribution being Gaussian entails that the distribution of X conditioned on $y = y^o$ is also Gaussian, thus $Q_{X|y^o} \sim \mathcal{N}(\mu_{X|y^o},\, \Sigma_{X|y^o})$. Applying Proposition 3.13. (17) to our setting, the mean and covariance of the conditioned distribution are given by:

$$\mu_{X|y^o} = \mu_X + \Sigma_{XX}\Sigma_{YY}^{-1}(y^o - \mu_X) \tag{20}$$

$$\Sigma_{X|y^o} = \Sigma_{XX} - \Sigma_{XX}\Sigma_{YY}^{-1}\Sigma_{XX} \tag{21}$$

Epistemic value is computed using the Kullback-Leibler divergence. With full knowledge of the joint distribution, in the Gaussian case, following the expression of entropy for gaussian vectors (Chapter 12 Equation 12.39, (18)) it is computed as:

$$C(Q_X) = I(X;Y) = \frac{1}{2}\ln\frac{(\det\Sigma_{XX})(\det\Sigma_{YY})}{\det\Sigma_{X,Y}} \tag{22}$$

The set of moves that can be selected by the agent is restricted to translations as the agent must always face the object. The set of possible translations is composed of eight translations with the same norm, with evenly distributed angles (one of them being oriented toward the object irrespective of the position of the agent), and also contained an idle state, i.e. no translation. Here the angles correspond to the angles of the translation and not a rotation angle of the solid frame of the agent, as the agent always faces the object.

We approximated the belief after the action m of a given group using a Gaussian distribution, $\psi_{m,*}Q_X \sim \mathcal{N}(\mu_m,\, \Sigma_m)$. The mean and covariance matrix are approximated using numerical integration:

$$\mu_m = \int xp(\psi_m^{-1}(x))\frac{1}{|\det J_{\psi_m}(\psi_m^{-1}(x))|}dx \tag{23}$$

$$\Sigma_m = \int (x-\mu_m)(x-\mu_m)^T p(\psi_m^{-1}(x))\frac{1}{|\det J_{\psi_m}(\psi_m^{-1}(x))|}dx \tag{24}$$

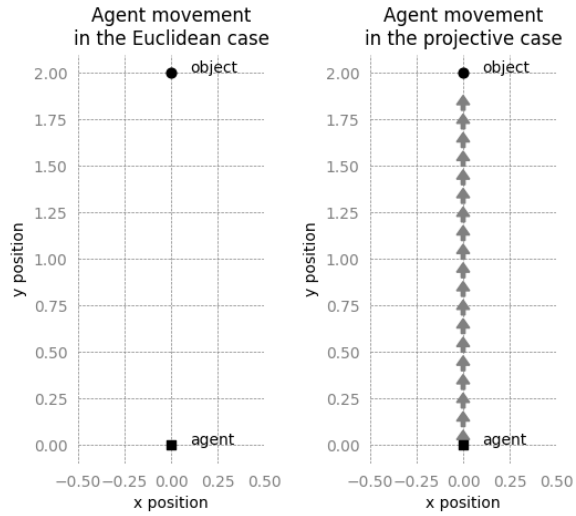

Figure 1: **Movement of the agent in the Euclidean vs projective case**

The experimental results of this implementation show the agent exhibits different behaviors depending on the group used to structure its internal world model (Figure 1). The agent started from an initial position with the object always located at a fixed position, and the algorithm was applied for 20 iterations, for both the Euclidean and Projective internal spaces. The agent started at $(0, 0)$ and the object was located at $(0, 2)$ in the world frame E_2 spanning the agent's displacement floor. If all translations were associated with epistemic values that only varied within a small range ($\pm 1e - 4$) as compared to the epistemic value of the idle state, reflecting numerical imprecision, the idle state was selected (the agent did not move).

D Figures, Algorithm

Algorithm 1: Curiosity based Exploration for agent a

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

$Q_X \leftarrow Q_X^0$;
while *True* **do**
  $\overline{m} \leftarrow \mathrm{argmax}_{m \in M}\, C(\psi_{m,*} Q_X)$;
  $\mathcal{R} \leftarrow$ solid reference frame of a after move $\overline{m}$;
  $Q_X \leftarrow \psi_{\overline{m},*} Q_X$;
  $y^o \leftarrow \phi_{\mathcal{R}}(o)$;
  $Q_X \leftarrow Q_{X|y^o}$;
end

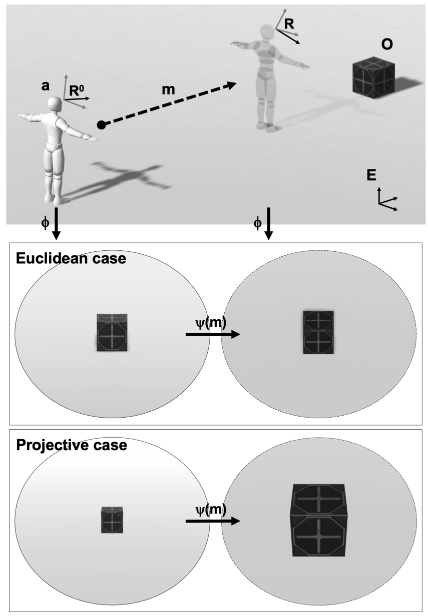

Figure 2: **Toy model setup and main transformations**

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).

$$
\begin{tikzcd}
& X_0 = W \arrow[r, "1[\psi_m(x_0) \in A]"] \arrow[d, "P_{Y|X}"'] & X_1 \arrow[d] \\
o \in E_3 \arrow[r, dashed, "\phi_{\mathcal{R}}"'] \arrow[rr, dashed, bend left=60, "\phi_{\mathcal{R}^m}"] & Y = W \ni y^o & Y \ni y^o_m
\end{tikzcd}
$$

Figure 3: $m \in M$ 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^o_m is the same for the reference frame $\mathcal{R}^m$ after move m.