INTRINSIC-ENERGY JOINT EMBEDDING PREDICTIVE ARCHITECTURES INDUCE QUASIMETRIC SPACES

Anonymous authors
Paper under double-blind review

ABSTRACT

Joint-Embedding Predictive Architectures (JEPAs) aim to learn representations by predicting target embeddings from context embeddings, inducing a scalar compatibility energy in a latent space. In contrast, Quasimetric Reinforcement Learning (QRL) studies goal-conditioned control through *directed* distance values (cost-to-go) that support reaching goals under asymmetric dynamics. In this short article, we connect these viewpoints by restricting attention to a principled class of JEPA energy functions : *intrinsic (least-action) energies*, defined as infima of accumulated local effort over admissible trajectories between two states. Under mild closure and additivity assumptions, any intrinsic energy is a quasimetric. In goal-reaching control, optimal cost-to-go functions admit exactly this intrinsic form ; inversely, JEPAs trained to model intrinsic energies lie in the quasimetric value class targeted by QRL. Moreover, we observe why symmetric finite energies are structurally mismatched with one-way reachability, motivating asymmetric (quasimetric) energies when directionality matters.

1 INTRODUCTION

Joint-Embedding Predictive Architectures (**JEPAs**) (LeCun, 2022; Assran et al., 2023) recently emerged as a compelling self-supervised paradigm for representation learning in which models predict *target embeddings* from *context embeddings* rather than reconstructing raw observations. From a control viewpoint, the resulting training score can be interpreted as a scalar *compatibility energy* between inputs : low energy meaning *compatible*, while high energy means *incompatible*.

Independently, **Quasimetric Reinforcement Learning** (**QRL**) (Wang et al., 2023) represents goal-conditioned control through the geometry of a learned cost-to-go. A central fact is that reaching costs induce a *directed* notion of distance : going from s to g can be easy, while returning from g to s may be difficult or impossible, and multi-step composition should be consistent. Hence, QRL formalizes this by targeting function classes with *quasimetric* structure (Wang & Isola, 2022), and proposes a learning framework and objectives designed for that regime.

This short article addresses a narrow structural question :

Under what conditions does a JEPA-induced energy behave like a cost-to-go ?

A key insight is that sequential compositionality is not a modelling preference, but usually required when reasoning (Johnson-Laird, 2010; Plaat et al., 2025). If a state u (e.g., a position, a proposition, a concept) is accessible (e.g., reachable, demonstrable, explainable) on the way from x to y, then an energy meant to support multi-step reasoning should satisfy a consistency inequality of the form :

$$E(x, g) \;\leq\; E(x, u) + E(u, g),$$

which is precisely a triangle inequality.

In physics and optimal control, the canonical way to obtain such a global, compositional energy is to define it intrinsically via a *least-action* principle: the energy between two states is the infimum of accumulated local effort over admissible trajectories connecting them (De Sapio et al., 2008; Ho, 2018; Terekhovich, 2018). This construction yields triangle inequality by definition (via concatenation), and it naturally produces asymmetry whenever admissibility or local effort is directed.

Our main message is simple :

If a JEPA is defined according to an intrinsic (least-action) energy,
then it induces a quasimetric space.

In goal-conditioned tasks, costs-to-go can be interpreted as intrinsic energies, hence such *Intrinsic-Energy JEPAs* fall into the same quasimetric value-functions class that QRL is designed to model.

2 Minimal Background

JEPA (Latent Representations and Induced Energies). Joint-Embedding Predictive Architectures learn representations by predicting *target embeddings* from *context embeddings* (LeCun, 2022). Given a context encoder f_ϕ and a target encoder $f_{\bar{\phi}}$, JEPA forms $z_x = f_\phi(x), z_y = f_{\bar{\phi}}(y)$, and a predictor p_θ produces a predicted target embedding $\hat{z}_y = p_\theta(z_x;\, c)$, where c denotes a conditioning.

Training uses a comparator function $D(\cdot,\cdot)$ in the embedding space :

$$\mathcal{L}_{\text{JEPA}}(\phi, \bar{\phi}, \theta) \;=\; \mathbb{E}_{(x,y)\sim\mathcal{D}}[\, D(\hat{z}_y, \mathtt{sg}(z_y))] \,,$$

Even when the comparator form is fixed, the model learns an *induced energy landscape* over inputs through the learned representation and predictor: pairs (x, y) that are *compatible* are precisely those for which the latent prediction error is small. This motivates interpreting JEPA training as learning a scalar energy (compatibility) between x and y, in line with the energy-based perspective on JEPA.

QRL (Goal-Reaching values are Quasimetrics). Quasimetric Reinforcement Learning studies goal-conditioned control through directed distances that compose over time (Wang et al., 2023). In reaching-cost problems, the optimal value function is $V^\star(s, g)$, usually negative (up to sign conventions), and satisfies a triangle inequality in (s, g), yielding a quasimetric structure :

$$d^\star(s, g) \;\triangleq\; -V^\star(s, g), \qquad d^\star(s, g) \le d^\star(s, w) + d^\star(w, g),$$

with $d^\star(s, g) = +\infty$ for unreachable pairs. QRL leverages this structure by learning d_θ within quasimetric function classes (Wang & Isola, 2022), enforcing local constraints from observed transitions and using the triangle inequality to propagate these constraints to long horizons; the resulting objectives come with recovery guarantees under suitable data coverage conditions. For our purposes, the key takeaway is that QRL treats the goal-conditioned value as a *directed geometry*.

Value-Guided Action Planning with JEPA World Models. Value-guided JEPA planning uses JEPA-like world models for control by shaping representation spaces so that an embedding-space cost aligns with a goal-reaching value, enabling planning via minimizing a representation-space objective (Destrade et al., 2025). Their contribution is primarily algorithmic and empirical: learning representations that make planning effective. Our focus is orthogonal and structural: we do not propose planners or additional experiments. Instead, we isolate a hypothesis-class condition on the *energy* itself under which a JEPA-induced energy necessarily satisfies quasimetric inequalities, directly connecting JEPA energies to the quasimetric value viewpoint formalized in QRL.

Intrinsic Energies and the Least-Action Principle. In physics and control, it is standard to define a global cost between configurations as the infimum of an *action functional*, typically an integral of a local effort along admissible trajectories, following a least-action viewpoint (Siburg, 2004). This perspective is used in modern ML to endow learned representations with physically meaningful structure (Greydanus et al., 2019; Cranmer et al., 2020), and more recently to motivate learning from variational constructions (Guo & Schölkopf, 2025). In optimal control, path-integral methods also build directly on trajectory functionals of the form $\int_0^T \ell(x(t), u(t))\, dt$, optimizing them to produce actions, which makes the "least accumulated effort over trajectories" interpretation operational in control pipelines (Asmar et al., 2023; Zhai et al., 2025). In this article, we adopt the same high-level principle but in a representation learning setting: we interpret the JEPA score as an *energy* and restrict attention to energies that admit an intrinsic (least-action) representation.

3 Intrinsic Energy Functions and Quasimetrics

Definition 1 (Quasimetric)**.** *A function $d : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a* quasimetric *if, for all $x, y, z \in \mathcal{X}$: (i : reflexivity) $d(x,x) = 0$, (ii : non-negativity) $d(x,y) \geq 0$, (iii : Identity of indiscernibles) if $d(x,y) = 0$ then $x = y$, (iv : triangular inequality) $d(x,z) \leq d(x,y) + d(y,z)$.*

Definition 2 (Intrinsic (*Least-Action*) Energy)**.** *Let $\mathcal{X}$ be a path-connected state space. For $x, y \in \mathcal{X}$, let $\Gamma(x \to y)$ denote a set of admissible C^1 trajectories $\gamma : [0,T] \to \mathcal{X}$, with $\gamma(0) = x$ and $\gamma(T) = y$, $\forall T > 0$. Let $L : \mathrm{T}\mathcal{X} \to \mathbb{R}^+$ be a local effort density, veryfing $L(x,v) \geq c \cdot \|v\|$ for some norm $\|\cdot\|$ and a constant $c > 0$. Defining the action of a trajectory as :*

$$\mathsf{Act}(\gamma) \;=\; \int_0^T L(\gamma(t), \dot{\gamma}(t))\, dt.$$

the intrinsic energy is $E(x,y) \;=\; \inf_{\gamma \in \Gamma(x \to y)} \mathsf{Act}(\gamma)$, with $E(x,y) = +\infty$ if $\Gamma(x \to y) = \emptyset$.

Remark 1 (Physics and Control Grounding)**.** *The least-action form is standard: it defines global energies from local effort and yields Euler–Lagrange dynamics under suitable regularity. In optimal control, cost-to-go functions arise as infima of accumulated running costs over feasible trajectories.*

Theorem 1 (Intrinsic Energy is a Quasimetric)**.** *E is a quasimetric on $\mathcal{X}$.*

Proof. Non-negativity : This property holds since $L \geq 0$.

Reflexivity : Considering the constant trajectory $\gamma(t) \equiv x$ it comes that $\dot{\gamma}(t) = 0$, $L(x,0)$ is independent of t, and for any T, $\mathsf{Act}(\gamma) = L(x,0) \cdot T$. Thus $E(x,x) = 0$ considering the infimum.

Identity of indiscernibles : Let's consider $x, y \in \mathcal{X}$ such that $E(x,y) = 0$. Since for any trajectory γ :

$$\mathsf{Act}(\gamma) \;=\; \int_0^T L(\gamma(t), \dot{\gamma}(t))\, dt \geq \int_0^T c \cdot \|\dot{\gamma}(t)\|\, dt \geq c \cdot \left\| \int_0^T \dot{\gamma}(t)\, dt \right\| \geq c \cdot \left\| \gamma(T) - \gamma(0) \right\| .$$

With $\|\gamma(T) - \gamma(0)\| = \|x - y\|$, it comes that $E(x,y) = 0 \geq c \cdot \|x - y\| \geq 0$, consequently $x = y$.

Triangle inequality : Let's consider $x, y, z \in \mathcal{X}$. We choose $\gamma_{xy} \in \Gamma(x \to y)$, $\gamma_{yz} \in \Gamma(y \to z)$, and $\varepsilon > 0$ such that $\mathsf{Act}(\gamma_{xy}) \leq E(x,y) + \varepsilon$ and $\mathsf{Act}(\gamma_{yz}) \leq E(y,z) + \varepsilon$. Then considering the concatenation $\gamma_{xz} = \gamma_{xy} \star \gamma_{yz} \in \Gamma(x \to z)$ and use additivity of the integral, we have :

$$E(x,z) \leq \mathsf{Act}(\gamma_{xz}) = \mathsf{Act}(\gamma_{xy}) + \mathsf{Act}(\gamma_{yz}) \leq E(x,y) + E(y,z) + 2\varepsilon \;.$$

Since it is true for any $\varepsilon > 0$, with $\varepsilon \to 0$ the triangular inequality holds. ☐

Proposition 1 (Asymmetry is Generic)**.** *If either (i) admissibility is directed, i.e., for $x, y \in \mathcal{X}$, $f : t \to \gamma(t) \in \Gamma(x \to y) \;\not\Rightarrow\; \bar{f} : t \to \gamma(T - t) \in \Gamma(y \to x)$, or (ii) local effort is anisotropic (e.g., $L(x,v) \neq L(x,-v)$), then in general $E(x,y) \neq E(y,x)$.*

Remark 2 (A Minimal Picture)**.** *Theorem 1 formalizes the statement "valid world energies must compose across time": the energy of $x \to z$ cannot exceed the energy of $x \to y$ plus the energy of $y \to z$. This is the triangle inequality that QRL builds into its value geometry (Wang et al., 2023).*

4 Intrinsic-Energy JEPA and the QRL Hypothesis Class

Definition 3 (Intrinsic-Energy JEPA (IE-JEPA))**.** *Let $f_\phi : \mathcal{X} \to \mathbb{R}^d$ be an encoder, and let $\mathcal{C}$ be a prediction rule producing a scalar score from embeddings (e.g., $\|p_\theta(f_\phi(x)) - f_{\bar{\phi}}(y)\|^2$ as in I-JEPA (Assran et al., 2023)). We say a JEPA induces an energy $E_{\phi,\theta} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$ if its evaluation score can be interpreted as $E_{\phi,\theta}(x,y)$. We call it an* Intrinsic-Energy JEPA *if $E_{\phi,\theta}$, in the sense of Definition 2, i.e., a least-action energy consistent with concatenation and local effort accumulation.*

Corollary 1 (IE-JEPA Energies are Quasimetrics)**.** *If a JEPA-induced energy $E_{\phi,\theta}$ is intrinsic, then it is a quasimetric. Figure 1 illustrates this parallel.*

Remark 3 (Does JEPA "Learn the energy" ?)**.** *In common JEPA instantiations (e.g., I-JEPA), the comparator is often fixed (e.g., ℓ_2 in embedding space), while the encoders and predictors are learned. In that sense, the* analytic form *of the comparator can be given, yet the* induced energy landscape *over inputs is learned through the learned representation and predictor. Conversely, QRL explicitly parameterizes a quasimetric cost and learns it with objectives tailored to a quasimetric structure.*

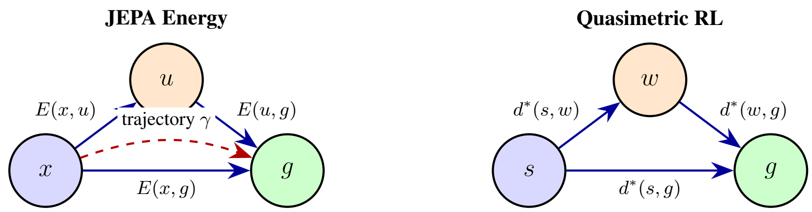

Figure 1: Intrinsic-energy JEPAs induce quasimetric structure.

Goal-reaching control as intrinsic energy. Consider a reaching-cost problem on $\mathcal{X}$: for each feasible trajectory γ from x to g, define accumulated cost $\mathsf{Act}(\gamma) = \int_0^T c(\gamma(t), \dot{\gamma}(t))dt$ with $c \geq 0$ (continuous-time) or $\sum_t c(s_t, s_{t+1})$ (discrete-time). The optimal cost-to-go is

$$V^\star(x, g) = \inf_{\gamma \in \Gamma(x \to g)} \mathsf{Act}(\gamma),$$

which is exactly an intrinsic energy (Definition 2). This is the standard variational/optimal control object underlying dynamic programming Bertsekas (2012).

Corollary 2 (IE-JEPA $\subseteq$ QRL hypothesis class (goal-reaching setting))**.** *In goal-reaching problems where the optimal cost-to-go is an intrinsic energy, any IE-JEPA energy that approximates this intrinsic energy is a quasimetric cost-to-go. This places IE-JEPA energies within the same class of quasimetric value functions formalized and targeted by QRL Wang et al. (2023).*

Why symmetry fails for directed reachability ?

To avoid implying that asymmetry is an aesthetic choice, we include the following basic obstruction.

Proposition 2 (Symmetric Finite Energies cannot Represent Directed Reachability)**.** *Let $R \subseteq \mathcal{X} \times \mathcal{X}$ be a directed reachability relation, where $(x, y) \in R$ means "y is reachable from x". Suppose $E : \mathcal{X} \times \mathcal{X} \to \bar{\mathbb{R}}^+$ satisfies: (i) $E(x, y) < +\infty$ iff $(x, y) \in R$ (finite energy on reachable pairs), and (ii) E is symmetric: $E(x, y) = E(y, x)$. Then R must be symmetric: $(x, y) \in R \Rightarrow (y, x) \in R$. Hence, no symmetric finite energy can encode one-way reachability.*

Proof. If $(x, y) \in R$, by (i) $E(x, y) < +\infty$. By (ii) $E(y, x) = E(x, y) < +\infty$, and $(y, x) \in R$. $\square$

Relation to value-guided JEPA planning. Destrade et al. (2025) shape JEPA representation spaces so that (quasi-)distances approximate negative goal-conditioned values, improving planning performance. Our focus is narrower: we introduce neither planners nor new experiments, but instead identify a structural condition (intrinsic energy) under which a JEPA score is necessarily a quasimetric cost-to-go, directly linking JEPA world-model scoring to the QRL geometry perspective.

5 Discussion and Scope

What this note does *not* claim. We do not claim that all JEPAs are quasimetrics, nor that the JEPA learning framework matches the QRL one unconditionally. Our equivalence is *conditional* : it holds for the family intrinsic energies (least-action functionals). This is a meaningful restriction since intrinsic energies are a canonical compositional energies in physics and optimal control.

Why the least-action form is not ad hoc. The least-action definition is the minimal way to obtain : (i) temporal compositionality, (ii) compatibility with irreversible systems, and (iii) identification with cost-to-go. Under this view, quasimetric regularization is the underlying consistency property.

Beyond RL. The intrinsic-energy view applies whenever there are admissible transformations and accumulated local efforts. This includes (a) deformation models in diffeomorphic image registration where distances are defined as least-action energies, and (b) reasoning structures where implications form directed relations with ordered-embedding approaches explicitly modelling directionality.

REFERENCES

Dylan M Asmar, Ransalu Senanayake, Shawn Manuel, and Mykel J Kochenderfer. Model predictive optimized path integral strategies. *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. *IEEE Conference on Computer Vision and Pattern Recognition*, 2023.

Dimitri Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 4. Athena scientific, 2012.

Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *ICLR Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.

Vincent De Sapio, Oussama Khatib, and Scott Delp. Least action principles and their application to constrained and task-level problems in robotics and biomechanics. *Multibody System Dynamics*, 2008.

Matthieu Destrade, Oumayma Bounou, Quentin Le Lidec, Jean Ponce, and Yann LeCun. Value-guided action planning with jepa world models. *arXiv preprint arXiv:2601.00844*, 2025.

Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 2019.

Siyuan Guo and Bernhard Schölkopf. Physics of learning: A lagrangian perspective to different learning paradigms. *Preprint (ArXiv)*, 2025.

Vu B Ho. On the principle of least action. *International Journal of Physics*, 2018.

Philip N Johnson-Laird. Mental models and human reasoning. *National Academy of Sciences*, 2010.

Yann LeCun. A path towards autonomous machine intelligence. *Preprint (Position Paper)*, 2022.

Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki Van Stein, and Thomas Bäck. Multi-step reasoning with large language models, a survey. *ACM Computing Surveys*, 2025.

Karl Friedrich Siburg. *The principle of least action in geometry and dynamics*. Springer Science & Business Media, 2004.

Vladislav Terekhovich. Metaphysics of the principle of least action. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2018.

Tongzhou Wang and Phillip Isola. Improved representation of asymmetrical distances with interval quasimetric embeddings. *NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, 2022.

Tongzhou Wang, Antonio Torralba, Phillip Isola, and Amy Zhang. Optimal goal-reaching reinforcement learning via quasimetric learning. *International Conference on Machine Learning*, 2023.

Yifan Zhai, Rudolf Reiter, and Davide Scaramuzza. Pa-mppi: Perception-aware model predictive path integral control for quadrotor navigation in unknown environments. *Preprint (ArXiv)*, 2025.