

Graph Distance as Surprise: Free Energy Minimization in Knowledge Graph Reasoning

Gaganpreet Jhajj^{1,*}, Fuhua Lin¹

¹School of Computing and Information Systems, Athabasca University, Canada

Abstract

In this work, we propose that reasoning in knowledge graph (KG) networks can be guided by surprise minimization. Entities that are close in graph distance will have lower surprise than those farther apart. This connects the Free Energy Principle (FEP) [1] from neuroscience to KG systems, where the KG serves as the agent's generative model. We formalize surprise using the shortest-path distance in directed graphs and provide a framework for KG-based agents. Graph distance appears in graph neural networks as message passing depth and in model-based reinforcement learning as world model trajectories. This work-in-progress study explores whether distance-based surprise can extend recent work showing that syntax minimizes surprise and free energy via tree structures [2].

Keywords

Knowledge Graphs, Graph Neural Networks, Active Inference, Semantic Grounding, Agents

1. Introduction

The Free Energy Principle (FEP) suggests that biological systems minimize surprise by maintaining accurate world models [1, 3, 4]. Recently, Murphy et al. [2] demonstrated that syntactic operations minimize surprise through shallow tree structures. They quantified surprise via tree depth (geometric complexity) and Kolmogorov complexity (algorithmic complexity), approximated through Lempel-Ziv compression [5, 6].

In FEP, agents minimize variational free energy $F = -\log P(o, s) - H[Q(s)]$, where o represents observations, s hidden states, P the generative model, and Q the agent's beliefs [1]. The first term, $-\log P(o, s)$, quantifies surprise: entities with high probability under the generative model (high $P(o, s)$) yield low surprise (low $-\log P(o, s)$). For syntactic trees, Murphy et al. [2] used tree depth to proxy this probability; we extend this principle to general graphs using shortest-path distance.

In active inference, minimizing free energy drives both perception (updating beliefs $Q(s)$) and action (selecting policies that reduce uncertainty) [3]. We apply this principle to KG reasoning: entities at shorter graph distances have a higher probability under the agent's graph-based generative model. The central question we address is: given a KG serving as an agent's generative model, which entity groundings are plausible for a query in context? We propose one principled approach: plausibility inversely correlates with graph distance.

Knowledge graphs (KGs) are increasingly integrated with modern AI agents, with the ability to improve reasoning, memory, and planning [7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. Unlike syntactic tree structures, KGs are directed graphs that can contain cycles and multiple paths between nodes (entities). In this preliminary work, we propose that surprise in KG reasoning corresponds to graph distance, where the KG serves as the agent's generative model. Entities that require shorter paths from context are unsurprising, whereas distant or disconnected entities are more surprising. This is unlike surprise-driven exploration in RL [17, 18], where agents maximize surprise to explore, FEP agents minimize surprise by maintaining accurate generative models. Our work connects the FEP to practical KG systems through shortest-path distance, providing theoretical foundations for graph neural networks [19, 20, 21] and model-based reinforcement learning [22, 23].

NORA'25: 1st Workshop on Knowledge Graphs & Agentic Systems Interplay co-located with NeurIPS, Dec. 1, 2025, Mexico City, Mexico

✉ gjhajj1@learn.athabascau.ca (G. Jhajj); oscarl@athabascau.ca (F. Lin)

>ID 0000-0001-5817-0297 (G. Jhajj); 0000-0002-5876-093X (F. Lin)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. From Syntax to Semantics

Murphy et al. [2] quantified syntactic surprise via tree depth. We extend this to arbitrary directed graphs with cycles. Given a KG $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ with entities \mathcal{E} , relations \mathcal{R} , and triples $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, geometric surprise is:

$$S_{\text{geo}}(e | C) = \begin{cases} \min_{c \in C} d_{\mathcal{G}}(c, e) & \text{if path exists} \\ \alpha & \text{otherwise} \end{cases} \quad (1)$$

where $d_{\mathcal{G}}(c, e)$ is the shortest directed path length from context $c \in C$ to entity e (computed via BFS, Appendix B), and α is a hyperparameter penalizing disconnection. In our worked example, we set $\alpha = 5$; in general, α should exceed the graph's diameter (longest shortest-path distance) to ensure disconnected entities always have higher surprise than any connected entity. Combined with algorithmic complexity [2]:

$$F(e | C) = S_{\text{geo}}(e | C) + \lambda K(\pi_{C \rightarrow e}) \quad (2)$$

where $K(\pi_{C \rightarrow e})$ is Kolmogorov complexity of the relation path, approximated via Lempel-Ziv compression, and λ weights the components. For trees, this recovers Murphy's tree depth; for general graphs, it handles cycles naturally.

Connection to FEP: Under FEP, agents minimize $F = -\log P(o, s) - H[Q(s)]$ [1]. Interpreting the KG as the agent's generative model, we posit $-\log P(e | C) \propto d_{\mathcal{G}}(C, e)$: shorter distances indicate higher probability. Thus S_{geo} implements the surprise term, while $K(\pi)$ approximates $H[Q(s)]$. Figure 1 illustrates this with a political KG example (detailed calculations in Appendix A).

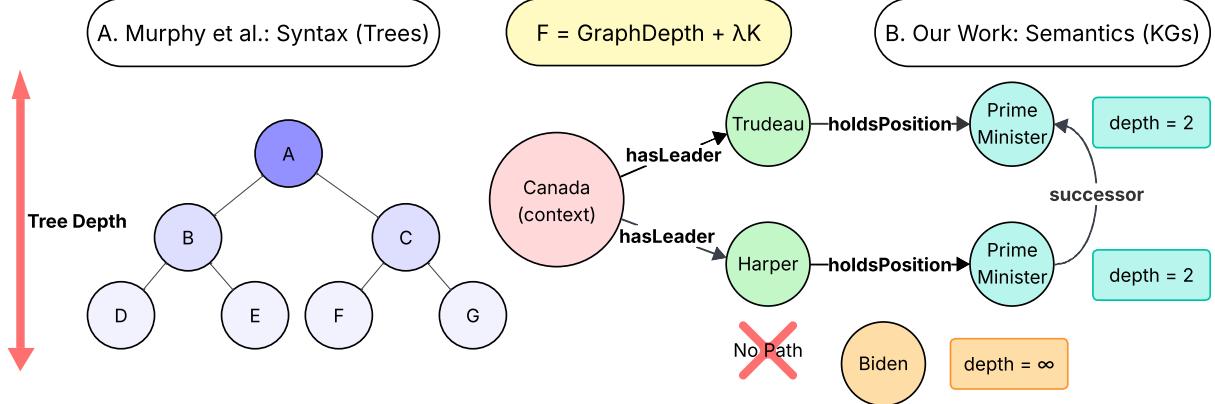


Figure 1: Extending surprise from trees to knowledge graphs. Following standard KG design (e.g., Wikidata), we model “Prime Minister” as a position node. Given context “Canada”, leaders (Trudeau, Harper) are at distance 1, the position node at distance 2, while disconnected entities (Biden) have distance ∞ . The successor relation demonstrates cycle handling.

3. Theoretical Justification

Three principles justify the shortest-path distance: (1) **Proper generalization**: For trees, it recovers Murphy's tree depth. (2) **Least-action**: Shortest paths minimize cumulative cost, aligning with active inference where agents minimize expected free energy [3]. (3) **Computational grounding**: In GNNs, k message-passing iterations aggregate k -hop neighborhoods [19, 21]; minimizing iterations minimizes distance and surprise. Cycles pose no issue: FEP accommodates circular causality [24], and BFS handles cycles via visited sets (Appendix B).

4. Implications and Future Work

This work-in-progress connects FEP from neuroscience to KG reasoning in AI systems. The presented framework offers practical implications: (1) **Entity grounding**: LLM-KG systems could rank candidate entity groundings by computing S_{geo} via BFS from discourse context entities, preferring groundings with lower free energy [10, 9]; (2) **KG embeddings**: embedding methods could preserve distance-based surprise structure [25]; (3) **GNN architecture**: depth could be selected to balance computational cost against the surprise horizon needed for a task.

Future work includes empirical validation on benchmark KG datasets (FB15k-237 [26], YAGO [27]), comparison with human semantic similarity judgments, integration with existing KG reasoning systems [10, 28, 9], and extension to temporal KGs.

This work represents an early-stage exploration of applying FEP to knowledge graph reasoning. While we proposed the shortest-path distance as a principled formalization of surprise, other formulations may be more elegant or practical.

We aim to present this contribution as an initial research direction rather than a definitive solution. We also encourage the community to develop complementary or improved approaches to connecting FEP principles with graph-based reasoning.

Acknowledgments

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), Alberta Innovates, Alberta Advanced Education, and Athabasca University, Canada. We would also like to thank the reviewers for their suggestions on how to improve this work.

Declaration on Generative AI

During the preparation of this work, the author(s) used Grammarly and Claude (Anthropic) for Grammar and spelling checks.

References

- [1] K. Friston, The free-energy principle: a unified brain theory?, *Nature Reviews Neuroscience* 11 (2010) 127–138. URL: <https://www.nature.com/articles/nrn2787>. doi:10.1038/nrn2787.
- [2] E. Murphy, E. Holmes, K. Friston, Natural language syntax complies with the free-energy principle, *Synthese* 203 (2024) 154. URL: <https://link.springer.com/10.1007/s11229-024-04566-3>. doi:10.1007/s11229-024-04566-3.
- [3] K. Friston, T. Fitzgerald, F. Rigoli, P. Schwartenbeck, G. Pezzulo, Active Inference: A Process Theory, *Neural Computation* 29 (2017) 1–49. URL: <https://direct.mit.edu/neco/article/29/1/1-49/8207>. doi:10.1162/NECO_a_00912.
- [4] T. Parr, G. Pezzulo, K. J. Friston, Active Inference: The Free Energy Principle in Mind, Brain, and Behavior, The MIT Press, 2022. URL: <https://direct.mit.edu/books/book/5299/Active-InferenceThe-Free-Energy-Principle-in-Mind>. doi:10.7551/mitpress/12441.001.0001.
- [5] M. Li, P. Vitányi, An Introduction to Kolmogorov Complexity and Its Applications, Texts in Computer Science, Springer New York, New York, NY, 2008. URL: <http://link.springer.com/10.1007/978-0-387-49820-1>. doi:10.1007/978-0-387-49820-1.
- [6] J. Ziv, A. Lempel, A universal algorithm for sequential data compression, *IEEE Transactions on Information Theory* 23 (1977) 337–343. URL: <https://ieeexplore.ieee.org/document/1055714>. doi:10.1109/TIT.1977.1055714.
- [7] L. Chen, P. Tong, Z. Jin, Y. Sun, J. Ye, H. Xiong, Plan-on-graph: self-correcting adaptive planning of large language model on knowledge graphs, in: Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24, Curran Associates Inc., Red Hook, NY, USA, 2025. Event-place: Vancouver, BC, Canada.
- [8] Y. Cui, Z. Sun, W. Hu, A prompt-based knowledge graph foundation model for universal in-context reasoning, in: Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24, Curran Associates Inc., Red Hook, NY, USA, 2025. Event-place: Vancouver, BC, Canada.
- [9] X. He, Y. Tian, Y. Sun, N. V. Chawla, T. Laurent, Y. LeCun, X. Bresson, B. Hooi, G-Retriever: Retrieval-Augmented Generation for Textual Graph Understanding and Question Answering, in: The Thirty-eighth Annual Conference on Neural Information Processing Systems, 2024. URL: <https://openreview.net/forum?id=MPJ3oXtTZl>.
- [10] G. Jhajj, X. Zhang, J. R. Gustafson, F. Lin, M. P.-C. Lin, Educational Knowledge Graph Creation and Augmentation via LLMs, in: A. Sifaleras, F. Lin (Eds.), Generative Intelligence and Intelligent Tutoring Systems, Springer Nature Switzerland, Cham, 2024, pp. 292–304. doi:10.1007/978-3-031-63031-6_25.
- [11] G. Jhajj, J. R. D. Gustafson, R. Morland, C. E. Gutierrez, M. P.-C. Lin, M. A. A. Dewan, F. Lin, Neuromorphic Knowledge Representation: SNN-Based Relational Inference and Explainability in Knowledge Graphs, in: S. Graf, A. Markos (Eds.), Generative Systems and Intelligent Tutoring Systems, volume 15724, Springer Nature Switzerland, Cham, 2026, pp. 159–165. URL: https://link.springer.com/10.1007/978-3-031-98284-2_13. doi:10.1007/978-3-031-98284-2_13.
- [12] G. Jhajj, Y. Nomura, Jack and the beansTALK: Towards question answering in plant biology, in: Eighth Widening NLP Workshop (WiNLP 2024) Phase II, 2024. URL: <https://openreview.net/forum?id=0DlJEPHHKe>.
- [13] R. D. Morland, F. Lin, An adaptable client-server architecture for generating educational content using large language models, *Bulletin of the Technical Committee on Learning Technology (ISSN: 2306-0212)* 25 (2025) 42–49.
- [14] J. R. D. Gustafson, G. Jhajj, X. Zhang, F. O. Lin, Enhancing project-based learning with a genai tool based on retrieval: Augmented generation and knowledge graphs, in: AI Applications and Strategies in Teacher Education, IGI Global, 2025, pp. 161–194.
- [15] G. Jhajj, F. Lin, Augmenting Japanese language acquisition via LLMs and ASR, in: IEEE Smart World Congress 2025 (IEEE SWC'25), Calgary, Canada, 2025, p. 3.84.

- [16] M. R. Kabir, F. Lin, An llm-powered adaptive practicing system., in: LLM@ AIED, 2023, pp. 43–52.
- [17] D. Pathak, P. Agrawal, A. A. Efros, T. Darrell, Curiosity-driven exploration by self-supervised prediction, in: Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17, JMLR.org, 2017, p. 2778–2787.
- [18] T. Rakotoaritina, G. Jhajj, C. Reinke, K. Doya, Information-theoretic formulation and combination of intrinsic rewards: Novelty, surprise and empowerment, in: Seventh International Workshop on Intrinsically Motivated Open-ended Learning, 2025. URL: <https://openreview.net/forum?id=WN7ofwXNvv>.
- [19] T. N. Kipf, M. Welling, Semi-Supervised Classification with Graph Convolutional Networks, in: International Conference on Learning Representations, 2017. URL: <https://openreview.net/forum?id=SJU4ayYgl>.
- [20] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, M. Welling, Modeling Relational Data with Graph Convolutional Networks, in: A. Gangemi, R. Navigli, M.-E. Vidal, P. Hitzler, R. Troncy, L. Hollink, A. Tordai, M. Alam (Eds.), The Semantic Web, volume 10843, Springer International Publishing, Cham, 2018, pp. 593–607. URL: https://link.springer.com/10.1007/978-3-319-93417-4_38. doi:10.1007/978-3-319-93417-4_38.
- [21] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. Dahl, A. Vaswani, K. Allen, C. Nash, V. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, R. Pascanu, Relational inductive biases, deep learning, and graph networks, 2018. URL: <http://arxiv.org/abs/1806.01261>. doi:10.48550/arXiv.1806.01261, arXiv:1806.01261.
- [22] R. S. Sutton, A. G. Barto, Reinforcement Learning: An Introduction, A Bradford Book, Cambridge, MA, USA, 2018.
- [23] B. Millidge, A. Tschantz, A. K. Seth, C. L. Buckley, On the Relationship Between Active Inference and Control as Inference, in: T. Verbelen, P. Lanillos, C. L. Buckley, C. De Boom (Eds.), Active Inference, volume 1326, Springer International Publishing, Cham, 2020, pp. 3–11. URL: https://link.springer.com/10.1007/978-3-030-64919-7_1. doi:10.1007/978-3-030-64919-7_1.
- [24] K. Friston, M. Levin, B. Sengupta, G. Pezzulo, Knowing one’s place: a free-energy approach to pattern regulation, Journal of The Royal Society Interface 12 (2015) 20141383. URL: <https://royalsocietypublishing.org/doi/10.1098/rsif.2014.1383>. doi:10.1098/rsif.2014.1383.
- [25] A. Bordes, N. Usunier, A. Garcia-Durán, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13, Curran Associates Inc., Red Hook, NY, USA, 2013, pp. 2787–2795. Event-place: Lake Tahoe, Nevada.
- [26] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. v. d. Berg, I. Titov, M. Welling, Modeling relational data with graph convolutional networks, in: European semantic web conference, Springer, 2018, pp. 593–607.
- [27] F. M. Suchanek, G. Kasneci, G. Weikum, Yago: a core of semantic knowledge, in: Proceedings of the 16th International Conference on World Wide Web, WWW ’07, Association for Computing Machinery, New York, NY, USA, 2007, p. 697–706. URL: <https://doi.org/10.1145/1242572.1242667>. doi:10.1145/1242572.1242667.
- [28] Y. Deng, C. Ye, Z. Huang, M. D. Ma, Y. Kou, W. Wang, GraphVis: Boosting LLMs with Visual Knowledge Graph Integration, in: The Thirty-eighth Annual Conference on Neural Information Processing Systems, 2024. URL: <https://openreview.net/forum?id=haVPmN8UGi>.

A. Worked Example: Free Energy Calculations

We demonstrate free energy calculations using the Canadian Prime Minister knowledge graph from Figure 1.

A.1. Scenario and Knowledge Graph

Consider query “Who is the Prime Minister?” with context $C = \{\text{Canada}\}$. The knowledge graph contains:

Entities: $\mathcal{E} = \{\text{Canada}, \text{Trudeau}, \text{Harper}, \text{PrimeMinister}, \text{Biden}\}$

Relations: (Canada, hasLeader, Trudeau), (Canada, hasLeader, Harper), (Trudeau, holdsPosition, PrimeMinister), (Harper, holdsPosition, PrimeMinister), (Trudeau, successor, Harper), (Harper, predecessor, Trudeau)

The successor/predecessor relations form a cycle: Trudeau \leftrightarrow Harper. Importantly, Biden has no directed path from Canada (separate subgraph).

A.2. Computing Geometric Surprise

Using BFS from Canada, we compute shortest directed paths:

- $d(\text{Canada}, \text{Trudeau}) = 1$ (direct via hasLeader)
- $d(\text{Canada}, \text{Harper}) = 1$ (direct via hasLeader)
- $d(\text{Canada}, \text{PrimeMinister}) = 2$ (via hasLeader then holdsPosition)
- $d(\text{Canada}, \text{Biden}) = \infty$ (no path)

Therefore: $S_{\text{geo}}(\text{Trudeau}) = S_{\text{geo}}(\text{Harper}) = 1$, $S_{\text{geo}}(\text{PrimeMinister}) = 2$, and $S_{\text{geo}}(\text{Biden}) = \alpha = 5$.

The cycle between Trudeau and Harper does not affect distances: BFS selects the shortest path (direct edge) and handles cycles via visited set (Appendix B).

A.3. Computing Algorithmic Complexity

For each grounding, we estimate Kolmogorov complexity via relation path patterns:

Trudeau & Harper: Paths $\pi = [\text{hasLeader}]$ use frequent relations, yielding high compression (low $K(\pi)$).

PrimeMinister node: Path $\pi = [\text{hasLeader}, \text{holdsPosition}]$ uses standard role-modeling patterns, also yielding low $K(\pi)$.

Biden: No path from Canada. The grounding requires irregular cross-country reasoning not represented in the graph (high $K(\pi)$).

A.4. Free Energy Results

Combining components with $\lambda = 1$:

Entity	S_{geo}	$K(\pi)$	F
Trudeau	1	Low	~ 1.3
Harper	1	Low	~ 1.3
Biden	5	High	~ 5.5

Interpretation: Real groundings (Trudeau, Harper) exhibit low free energy: (1) short distance (1 hop), (2) regular relation patterns. The impossible grounding (Biden) exhibits high free energy: (1) disconnection (no path), (2) irregular pattern. The framework correctly identifies both Trudeau and Harper as plausible (both were Canadian PMs) while rejecting Biden (US president).

We focus on entity groundings (Trudeau, Harper, Biden) rather than the position node itself, as queries about leadership typically seek individuals rather than abstract roles. The PrimeMinister node, at distance 2, would have intermediate surprise ($S_{\text{geo}} = 2, F \approx 2.3$), but is not a direct answer to “Who is the Prime Minister?” This demonstrates how our framework naturally distinguishes between entities at different levels of abstraction in reified KG schemas.

This demonstrates three key properties: **(1)** cycles handled naturally, **(2)** multiple valid answers coexist with equal surprise, **(3)** disconnected entities correctly penalized.

B. Mathematical Details

B.1. Breadth-First Search Algorithm

Given directed graph $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ and context $C \subseteq \mathcal{E}$, we compute $S_{\text{geo}}(e | C)$ via BFS:

Algorithm 1 Compute Geometric Surprise

Require: Knowledge graph \mathcal{G} , context C , target entity e

Ensure: Geometric surprise $S_{\text{geo}}(e | C)$

```

1: Initialize:  $d(c) \leftarrow 0$  for all  $c \in C$ ;  $d(v) \leftarrow \infty$  for  $v \notin C$ 
2:  $Q \leftarrow C$  (queue),  $V \leftarrow C$  (visited set)
3: while  $Q \neq \emptyset$  do
4:    $u \leftarrow \text{dequeue from } Q$ 
5:   for each outgoing edge  $(u, r, v) \in \mathcal{T}$  do
6:     if  $v \notin V$  then
7:        $d(v) \leftarrow d(u) + 1$ 
8:        $V \leftarrow V \cup \{v\}$ , enqueue  $v$  to  $Q$ 
9:     end if
10:   end for
11: end while
12: return  $d(e)$  if  $d(e) < \infty$ , else  $\alpha$ 

```

Properties: (1) *Correctness*: BFS finds shortest paths in $O(|\mathcal{E}| + |\mathcal{T}|)$ time. (2) *Cycle handling*: Visited set V prevents re-visiting nodes, ensuring termination. (3) *Directionality*: Only outgoing edges followed, respecting direction.

B.2. Kolmogorov Complexity Approximation

We approximate $K(\pi_{C \rightarrow e})$ via Lempel-Ziv compression: (1) Extract relation sequence $\pi = [r_1, \dots, r_k]$ from shortest path. (2) Encode as string (e.g., “pm|successor”). (3) Compress with LZ77. (4) Compute ratio $K(\pi) = \text{compressed}/\text{original}$.

Interpretation: Regular patterns (frequent relations, short sequences) achieve high compression (low K). Irregular patterns (rare relations, long sequences) achieve low compression (high K). This approximates Kolmogorov complexity, which is uncomputable [5]. Murphy et al. [2] use the same approximation for syntactic patterns.

B.3. Connection to Active Inference

In active inference, agents minimize expected free energy $G(\pi)$ [3, 4]:

$$G(\pi) = \underbrace{D_{KL}[Q(o|\pi) \| P(o)]}_{\text{Pragmatic}} + \underbrace{\mathbb{E}_{Q(o|\pi)}[H[P(s|o)]]}_{\text{Epistemic}} \quad (3)$$

balancing pragmatic value (exploitation) and epistemic value (exploration).

Pragmatic value: Entities at shorter distances are more likely: $P(\text{observe } e | C)$ increases as S_{geo} decreases, making low-distance entities preferred for goal-directed actions.

Epistemic value: Entities at longer distances provide higher information gain: observing distant entities reduces uncertainty about unexplored graph regions, making high-distance entities preferred for exploration.

Our S_{geo} implements pragmatic value: low surprise entities preferred for exploitation. Extensions could weight distance inversely for epistemic value, valuing high-surprise entities for exploration.