Lessons from the Past for the Future of Robotics

Mohan Sridharan
Institute of Perception, Action, and Behavior
School of Informatics, University of Edinburgh, UK
m.sridharan@ed.ac.uk

May 6, 2026

1 Motivation

Robots are increasingly being used in practical applications due to advancements in AI methods. Modern AI frameworks based on deep networks and foundation models such as Large Language Models (LLMs) and Vision Language Action Models (VLAs) are considered state of the art for perception, reasoning, manipulation, and interaction problems [4, 9, 43]. These methods and models are excellent statistical predictors for well-defined tasks, but they make arbitrary decisions in truly novel situations [19, 26]. Also, the many network architectures being developed are based on a narrow set of representations and update processes. In addition, they are resource-hungry batch-learning systems whose internal operation remains opaque. Furthermore, they are having a negative impact on the reproducibility and sustainability of robotics research.

Robot systems that sense and (inter)act in the physical world are, on the other hand, required to make decisions at different spatiotemporal abstractions based on multimodal (e.g., visual, verbal, tactile) inputs. They have to operate under strict resource constraints and open world uncertainty, where optimal decisions are unknowable and probabilities may not always meaningfully model the uncertainty. Also, they have to rapidly revise models for perception, planning, and navigation, and they may be required to express their decisions in terms of human concepts to promote understanding in critical applications. There is thus a fundamental mismatch between the requirements of robot systems and the characteristics of modern AI methods.

The mismatch identified above can be addressed by revisiting key principles that can be traced back to the pioneers of AI but are not fully leveraged in modern robotics/AI research. These pioneers had a deep understanding of cognition and control in humans (i.e., *natural intelligence*). They knew that human behavior is determined by internal cognitive processes and the environment, and that we maintain and automatically direct *attention* [6] to relevant *representations and processes at different abstractions* [35, 40]. They understood that we acquire skills *incrementally*, *interactively*, and *compositionally*, making *rational* decisions under resource constraints and open world uncertainty [33]. They also appreciated that replicating some of our *hardware* on robots will not equip them with the sensorimotor capabilities that have evolved in humans over a long time [27, 41]. Based on these insights, I would like to outline some of these principles and draw on my work over the last 5-10 years to highlight the benefits of embedding these principles in robot architectures [37].

2 Key Principles and Examples

I would like to draw attention to three sets of principles that I think are particularly useful in robotics.

Refinement, Compositionality, Attention. The first set of principles represent a domain's actions and changes as transition diagrams at different abstractions, with the fine(r)-granularity diagram being a *refinement* of the coarse(r)-granularity diagram. Refinement has been explored in computing over many decades; it is also related to *compositionality*, the hierarchical representation of knowledge at different resolutions. These principles have played a key role in computing and other disciplines over many decades [12, 7]. Research has also identified that these principles and related representations lead to a good computational model for human cognition [22, 30], and for computer vision and robotics tasks [10, 42].

These principles can be applied to different robotics problems by establishing a suitable representation (i.e., *vocabulary*) and update processes at each level of abstraction, and defining a formal relationship between these levels. The relevant information can then be chosen automatically for any given domain and tasks using the principle of *selective attention* [6]; a limited exploration of attention has provided impressive results with deep networks too [9]. Refinement of an agent's action theories has been defined using situation calculus, but with limited expressivity due to the assumption of a strong bisimulation relation between these theories [2] or the assumption of deterministic causal models [3]. There has also been related work on combining discrete and continuous planning at different resolutions for task and motion planning in robotics [13, 23], but existing methods do not fully support bidirectional flow of relevant information, or handle uncertainty and discontinuous interaction dynamics.

Toward addressing these limitations, our refinement-based architecture supports different representations (logics, probabilities) and processes (non-monotonic logical reasoning, probabilistic sequential decision making) at two abstractions [38]. It acknowledges that uncertainty does not always have to be modeled using probabilities [17]; embeds cognitive theories of intention, affordance, and explainable agency [16, 36]; and enables the robot to identify and reason with relevant information using the other principles described below. We demonstrated better performance than knowledge-based or data-driven baselines, with the robot using the representations to describe its decisions such that they make contact with human concepts.

Ecological Rationality (ER) and Decision Heuristics. The second set of principles build on Herb Simon's definition of *Bounded Rationality* [33] and the algorithmic theory of heuristics [14]. ER explores decision making under *open world uncertainty*, i.e., when the space of possible scenarios is not known. It views an agent's behavior as a joint function of its cognitive processes and the environment, making decisions based on *adaptive satisficing*. It uses *decision heuristics* such as tallying, sequential search, and fast and frugal (FF) trees, to ignore some information and make decisions more quickly, frugally, and accurately than complex methods with many free parameters [15]. The associated models are both *prescriptive* and *descriptive*, automatically supporting process-level explanations of the decisions. Such an approach is complementary to theories and methods that define and use causal relationships to provide explanations [18], or categorize methods in terms of their ability to extract causal relationships and provide explanation [29].

Although ER and decision heuristics have provided good performance in different domains [5, 21], there is very limited use of these methods in robotics [25] because their inherent simplicity makes researchers doubt their suitability. We have used decision heuristics for multiagent collaboration without prior coordination, i.e., *ad hoc teamwork* (AHT) [28]. Methods considered state of the art for AHT use a large labeled dataset of prior observations to model the behavior of other agent types and determine the ad hoc agent's behavior [31]. We instead adapted our refinement-based architecture, with an ad hoc agent choosing its actions based on non-monotonic logical reasoning with prior knowledge (at different abstractions) and an ensemble of FF trees learned rapidly to predict the behavior of other agents. We experimentally demonstrated collaboration with other agents in complex and novel scenarios, providing better performance than data-driven baselines while using orders of magnitude fewer resources [8].

We also used these principles on a robot manipulator making and breaking contacts with objects and surfaces. These changing-contact manipulation tasks are characterized by abrupt transitions (in force, acceleration) that can damage the robot or the objects. Unlike data-driven methods that explore different transitions in advance, posing smooth motion as an optimization or learning problem, we drew inspiration from human motor control [11]. We enabled the robot to rapidly learn and revise simple models predicting the end-effector sensor observations from a single demonstration. During run-time, any mismatch between predicted and actual measurements revises these predictive *forward models* and a simple control law, leading to smooth changing-contact manipulation in novel situations [32].

Interactive learning and Memory Consolidation. The third set of principles refer to different types of learning such as supervised (or unsupervised) learning and learning from reinforcement [24]. Modern AI systems focus on learning a single (foundation) model or policy for different categories, situations, platforms, and domains. This approach is considered essential for *generalization*, but the corresponding design choices make the learned model (policy) hard to understand, explain, or revise meaningfully. These approaches are appropriate when the space of options or situations is known a priori and there are no strict resource constraints; they are unsuitable for decision making *in the wild* [20]. Interactive learning, on the other hand, uses reasoning and decision heuristics to trigger (and guide) online learning as needed to adapt to any given domain and set of tasks, revising the learned knowledge offline through memory consolidation [39, 41].

We used these principles for vision-based scene understanding, planning, and question answering. Unlike methods that train deep networks on a large dataset of images, questions, and answers, we adapted our refinement-based architecture to determine the occlusion and stability of objects, achieve desired object configurations, and answer questions about the decisions. This architecture reasoned with prior domain knowledge to make the target decisions if possible. For examples that could not be handled through reasoning, the robot automatically identified the relevant regions in images to be used for learning. In addition, these examples were processed using decision heuristics to induce new actions and axioms that revised existing knowledge. We experimentally demonstrated how leveraging the interplay between reasoning and learning resulted in better performance than deep network baselines while using orders of magnitude fewer examples, and in on-demand relational descriptions as explanations in response to different types of questions (causal, contrastive, counterfactual) [36].

We also used these principles to enable AI agents to collaborate with humans to complete daily living tasks in a 3D simulation environment. Instead of making dubious claims about the "planning" or "general intelligence" capabilities of LLMs, our architecture leveraged their predictive abilities to anticipate the upcoming sequence of high-level tasks. The AI agent then jointly planned for the current and anticipated tasks, using decision heuristics to collaborate effectively with humans. We demonstrated substantial improvement in reliability and computational efficiency compared with LLM baselines [1, 34].

Summary: Although I unfortunately never had the opportunity to work with Joseph Halpern, my work (like that of many other researchers around the world) has been influenced by his papers and books. I would like to use this workshop as an opportunity to express my gratitude, promote appreciation for some key principles that are not being leveraged in modern robotics research, and illustrate how embedding these principles in architectures is an important step toward enabling robust decision making in robots.

Acknowledgements

The ideas described in this paper were informed by research threads pursued in collaboration with Hasra Dodampegama, Michael Gelfond, Gerd Gigerenzer, Rocio Gomez, Konstantinos Katsikopoulos, Pat Langley, Ales Leonardis, Ben Meadows, Tiago Mota, Heather Riley, Saif Sidhik, Özgür Şimşek, Jeremy Wyatt, and Shiqi Zhang. This work was supported in part by the U.S. ONR Awards N00014-13-1-0766, N00014-17-1-2434, N00014-20-1-2390, and N00014-24-1-2737, AOARD award FA2386-16-1-4071, and U.K. EPSRC award EP/S032487/1. All conclusions described are those of the author alone.

References

[1] Raghav Arora, Shivam Singh, Karthik Swaminathan, Snehasis Banerjee, Brojeshwar Bhowmick, Krishna Murthy Jatavallabhula, Mohan Sridharan, and Madhava Krishna. Anticipate & Act: Integrating LLMs and Classical Planning for Efficient Task Execution in Household Environments. In *International Conference on Robotics and Automation (ICRA)*, Yokohoma, Japan, May 13-17, 2024.

[2] Bita Banihashemi, Giuseppe De Giacomo, and Yves Lesperance. Abstraction of Agents Executing Online and their Abilities in Situation Calculus. In *International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, July 13-19, 2018.

[3] Sander Beckers and Joseph Halpern. Abstracting Causal Models. In *AAAI Conference on Artificial Intelligence*, Honolulu, Hawaii, January 27-February 1, 2019.

[4] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A Vision-Language-Action Flow Model for General Robot Control. In *Robotics: Science and Systems Conference*, Los Angeles, USA, June 21-25, 2025.

[5] Henry Brighton and Gerd Gigerenzer. How Heuristics Handle Uncertainty? In *Ecological Rationality: Intelligence in the World*. Oxford University Press, New York, 2012.

[6] Donald E. Broadbent. A Mechanical Model for Human Attention and Immediate Memory. *Psychology Review*, 64:205–215, 1957.

[7] T. Dietterich. The MAXQ Method for Hierarchical Reinforcement Learning. In *International Conference on Machine Learning (ICML)*, 1998.

[8] Hasra Dodampegama and Mohan Sridharan. Knowledge-based Reasoning and Learning under Partial Observability in Ad Hoc Teamwork. *Theory and Practice of Logic Programming*, 23(4):696–714, 2023.

[9] Ria Doshi, Homer Walke, Oier Mees, Sudeeo Dasai, and Sergey Levine. Scaling Cross-Embodied Learning: One Policy for Manipulation, Navigation, Locomotion, and Aviation. In *International Conference on Robot Learning*, Munich, Germany, November 6-9, 2024.

[10] Sanja Fidler and Ales Leonardis. Towards Scalable Representations of Object Categories: Learning a Hierarchy of Parts. In *International Conference on Computer Vision and Pattern Recognition*, 2007.

[11] J. Randall Flanagan, Philipp Vetter, Roland S. Johansson, and Daniel M. Wolpert. Prediction Precedes Control in Motor Learning. *Current Biology*, 13:146–150, 2003.

[12] Jerry A. Fodor. *The Language of Thought*. Harvard University Press, 1975.

[13] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Integrated Task and Motion Planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:265–293, 2021.

[14] Gerd Gigerenzer. What is Bounded Rationality? In *Routledge Handbook of Bounded Rationality*. Routledge, 2020.

[15] Gerd Gigerenzer and Wolfgang Gaissmaier. Heuristic Decision Making. *Annual Review of Psychology*, 62:451–482, 2011.

[16] Rocio Gomez, Mohan Sridharan, and Heather Riley. What do you really want to do? Towards a Theory of Intentions for Human-Robot Collaboration. *Annals of Mathematics and Artificial Intelligence, special issue on commonsense reasoning*, 89:179–208, February 2021.

[17] Joseph Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.

[18] Joseph Y. Halpern and Judea Pearl. Causes and Explanations: A Structural-Model Approach. Part I: Causes. *British Journal for the Philosophy of Science*, 56(4):843–887, 2005.

[19] Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Mudit Verma, Kaya Stechly, Siddhant Bhambri, Lucas Paul Saldyt, and Anil Murthy. Position: LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks. In *International Conference on Machine Learning*, Vienna, Austria, July 2024.

[20] Konstantinos Katsikopoulos, Ozgur Simsek, Marcus Buckmann, and Gerd Gigerenzer. *Classification in the Wild: The Science and Art of Transparent Decision Making*. MIT Press, February 2021.

[21] Konstantinos Katsikopoulos, Ozgur Simsek, Marcus Buckmann, and Gerd Gigerenzer. Transparent modeling of influenza incidence: Big data or a single data point from psychological theory? *International Journal of Forecasting*, 2021.

[22] G. Knoblich and R. Flach. Predicting the Effects of Actions: Interactions of Perception and Action. *Psychological Sciences*, 12(6):467–472, 2001.

[23] Harsha Kokel, Sriraam Natarajan, Balaraman Ravindran, and Prasad Tadepalli. RePReL: A Unified Framework for Integrating Relational Planning and Reinforcement Learning for Effective Abstraction in Discrete and Continuous Domains. *Neural Computing and Applications*, 35:16877–16892, 2023.

[24] John E. Laird, Kevin Gluck, John Anderson, Kenneth D. Forbus, Odest Chadwicke Jenkins, Christian Lebiere, Dario Salvucci, Matthias Scheutz, Andrea Thomaz, Greg Trafton, Robert E. Wray, Shiwali Mohan, and James R. Kirk. Interactive Task Learning. *IEEE Intelligent Systems*, 32(4):6–21, 2017.

[25] Pat Langley and Edward P. Katz. Motion Planning and Continuous Control in a Unified Cognitive Architecture. In *Annual Conference on Advances in Cognitive Systems*, Arlington, VA, August 2022.

[26] Sheng Lu, Irina Bigoulaeva, Rachneet Sachdeva, Harish Tayyar Madabushi, and Iryna Gurevych. Are Emergent Abilities in Large Language Models just In-Context Learning? In *Annual Meeting of the Association for Computational Linguistics*, pages 5098–5139, Bangkok, Thailand, August 11-16, 2024.

[27] Marvin L. Minsky. *The Society of Mind*. Simon and Schuster, 1986.

[28] Reuth Mirsky, Ignacio Carlucho, Arrasy Rahman, Elliot Fosong, William Macke, Mohan Sridharan, Peter Stone, and Stefano Albrecht. A Survey of Ad Hoc Teamwork: Definitions, Methods, and Open Problems. In *European Conference on Multiagent Systems*, 2022.

[29] Judea Pearl and Dana Mackenzie. *The book of why : the new science of cause and effect*. Allen Lane, 2018.

[30] Steven T. Piantadosi, Joshua B. Tenenbaum, and Noah D. Goodman. The Logical Primitives of Thought: Empirical Foundations for Compositional Cognitive Models. *Psychological Review*, 123(4):392–424, 2016.

[31] Muhammad A Rahman, Niklas Hopner, Filippos Christianos, and Stefano V Albrecht. Towards open ad hoc teamwork using graph-based policy learning. In *International Conference on Machine Learning*, pages 8776–8786, 18–24 Jul 2021.

[32] Saif Sidhik, Mohan Sridharan, and Dirk Ruiken. An Adaptive Framework for Trajectory Following in Changing-Contact Robot Manipulation Tasks. *Robotics and Autonomous Systems*, 181:1–21, 2024.

[33] Herbert A. Simon. Rational Choice and the Structure of the Environment. *Psychological Review*, 63:129–138, 1956.

[34] Shivam Singh, Karthik Swaminathan, Nabanita Dash, Ramandeep Singh, Snehasis Banerjee, Mohan Sridharan, and Madhava Krishna. AdaptBot: Combining LLM with Knowledge Graphs and Human Input for Generic-to-Specific Task Decomposition and Knowledge Refinement. In *International Conference on Robotics and Automation (ICRA)*, Atlanta, USA, May 19-23, 2025.

[35] Aaron Sloman. Meta-morphogenesis and the Creativity of Evolution. In *Workshop on Computational Creativity, Concept Invention, and General Intelligence at ECAI*, Montpellier, France, 2012.

[36] Mohan Sridharan. Integrated Knowledge-based Reasoning and Data-driven Learning for Explainable Agency in Robotics. In *Explainable Agency in Artificial Intelligence: Research and Practice*. CRC Press, 2024.

[37] Mohan Sridharan. Back to the Future of Integrated Robot Systems. In *AAAI Conference on Artificial Intelligence*, Philadelphia, US, February 25-March 4, 2025.

[38] Mohan Sridharan, Michael Gelfond, Shiqi Zhang, and Jeremy Wyatt. REBA: A Refinement-Based Architecture for Knowledge Representation and Reasoning in Robotics. *Journal of Artificial Intelligence Research*, 65:87–180, May 2019.

[39] Robert Stickgold. Sleep-dependent Memory Consolidation. *Nature*, 437:1272–1278, 2005.

[40] Alan Turing. The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society of London B*, 237(641):37–72, 1952.

[41] Daniel M. Wolpert, Jorn Diedrichsen, and J. Randall Flanagan. Principles of Sensorimotor Learning. *Nature Reviews Neuroscience*, 12:739–751, 2011.

[42] Jure Zabkar and Ales Leonardis. Motor Memory: Representation, Learning, and Consolidation. *Biologically Inspired Cognitive Architectures*, 16:64–74, 2016.

[43] Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B. Tenenbaum, Tuamin Shu, and Chuang Gan. Building Cooperative Embodied Agents Modularly with Large Language Models. In *International Conference on Learning Representations*, May 7-11, 2024.