

Bayesian Inverse Physics for Neuro-Symbolic Robot Learning

Octavio Arriaga

Robotics Research Group, University of Bremen

ARRIAGAC@UNI-BREMEN.DE

Rebecca Adam

Robotics Innovation Center, DFKI GmbH

REBECCA.ADAM@DFKI.DE

Melvin Laux

Robotics Research Group, University of Bremen

Lisa Gutzeit

Robotics Research Group, University of Bremen; Robotics Innovation Center, DFKI GmbH

Marco Ragni

Technical University of Chemnitz

Jan Peters

Technical University of Darmstadt; Hessian.AI; Systems AI for Robot Learning, DFKI GmbH

Frank Kirchner

Robotics Research Group, University of Bremen; Robotics Innovation Center, DFKI GmbH

Editors: Leilani H. Gilpin, Eleonora Giunchiglia, Pascal Hitzler, and Emile van Krieken

Abstract

Real-world robotic applications, from autonomous exploration to assistive technologies, require adaptive, interpretable, and data-efficient learning paradigms. While deep learning architectures and foundation models have driven significant advances in diverse robotic applications, they remain limited in their ability to operate efficiently and reliably in unknown and dynamic environments. In this position paper, we critically assess these limitations and introduce a conceptual framework for combining data-driven learning with deliberate, structured reasoning. Specifically, we propose leveraging differentiable physics for efficient world modeling, Bayesian inference for uncertainty-aware decision-making, and meta-learning for rapid adaptation to new tasks. By embedding physical symbolic reasoning within neural models, robots could generalize beyond their training data, reason about novel situations, and continuously expand their knowledge. We argue that such hybrid neuro-symbolic architectures are essential for the next generation of autonomous systems, and to this end, we provide a research roadmap to guide and accelerate their development.

1. Introduction

One of the most remarkable features of human cognition is its ability to continuously estimate and interact with the world using limited information (Griffiths et al., 2024; Tenenbaum et al., 2011; Pinker, 1984). Humans can recognize an object after seeing it just once and manipulate objects they have never previously encountered (Gopnik, 2012; Lake et al., 2011). Achieving this level of adaptability is critical across the full range of robotic applications. The next generation of industrial and service robots must interact seamlessly with humans and adapt to dynamic surroundings (Lake et al., 2017; Rabinowitz et al., 2018). In healthcare, for example, assistive robots could infer patient needs with minimal supervision, helping to alleviate the

growing demand for medical care in an aging world population (Łukasik et al., 2020; Macis et al., 2022). The need for autonomy with scarce data is even more pronounced in extreme settings. Disaster response robots must navigate unpredictable environments, underwater systems must deal with shifting currents and poor visibility, and space robots face large communication delays and harsh, unpredictable conditions. All of these domains require the ability to reason under uncertainty and adapt autonomously to new circumstances (Nesnas et al., 2021; Obura et al., 2019). Despite all advances in deep learning, current state-of-the-art robot learning lacks these capabilities. Most state-of-the-art algorithms in computer vision and robot learning rely on methods that require massive training datasets with billions of samples or years of simulated experience to solve a set of constrained predefined tasks (Lake et al., 2017). These models struggle to adapt across tasks and environments, making them inefficient and limiting their real-world applicability (Gigerenzer, 2022; Pearl, 2018; Marcus and Davis, 2019; Sünderhauf et al., 2018). Moreover, relying on data-driven black-box models is not only inefficient when known physical approximations exist, but also raises critical safety concerns when robots operate and collaborate in close proximity to humans (Valdenegro-Toro, 2021). These challenges raise two critical questions:

- Why should we ignore the most accurate predictive models that we have of our world?
- Why should we trust the probabilities of a model that we do not understand?

To answer these fundamental questions, this paper argues for an alternative approach that extends current learning and robot perception algorithms by combining and scaling the following ideas:

1. **Uncertainty-aware models** through Bayesian inference, posterior sampling, and variational inference (Ghahramani, 2015; Jaynes, 2003; Parr et al., 2022).
2. **Physics-informed modeling**, combining differentiable physics and graphics-based simulators with neural networks to enhance generalization, efficiency, and physical consistency (Toussaint et al., 2018; Arriaga et al., 2024; Lutter et al., 2020; Heiden et al., 2021).
3. **Program synthesis**, building higher model abstractions with neural program search and transdimensional sampling (Griffiths et al., 2024; Ellis et al., 2020; Cusumano-Towner et al., 2020).

The synthesis of these components provides a direct path to address our previous questions, paving the way for more capable and safe autonomous systems.

2. Limitations of Foundation Models for Robotics

In the following section, we present three positions regarding the limitations of foundation models applied to robotics. Those limitations include sample and model complexity, safety, and lack of integrated physical knowledge. However, foundation models have significantly contributed to robotics, including perceptual feature extraction (Oquab et al., 2023), zero-shot instance segmentation (Kirillov et al., 2023), and language models (Dubey et al., 2024). Our conclusions ultimately point to an integrative approach combining foundation models, uncertainty quantification, and integrated physical and symbolic knowledge.

2.1. Large Data and Model Complexity

- Performance increases with more data, yet it remains insufficient for robot autonomy.

The transformer models’ performance strongly depends on the number of samples and the model size (Kaplan et al., 2020). Experimental studies have shown that this increase in model performance follows a power-law scaling relationship for both data and the number of parameters (Hoffmann et al., 2022; Kaplan et al., 2020). However, both extensive data and large model sizes constitute two of the biggest challenges in using neural networks for robotics. Specifically, there is no internet-scale repository of human-generated robot data, making large-scale learning challenging. Additionally, achieving full autonomy while preserving user privacy necessitates real-time models that run on-device without reliance on external servers. Moreover, highly dynamic and unpredictable environments, such as those found underwater or in space, often lack any pre-existing data, forcing systems to operate under completely unknown conditions.

Current frontier LLMs are estimated to have increased their training data 100-fold since 2020, up to tens of trillions of tokens in 2024 (Brown et al., 2020; Dubey et al., 2024). Moreover, Cottier et al. (2024) estimated computational costs for training LLMs to have increased between 50- and 100-fold since 2020, from hundreds of thousands of USD up to tens of millions in 2024, ultimately resulting in a total estimated cost of more than 100M USD for current frontier models such as GPT4 and Gemini 1.0 Ultra (Cottier et al., 2024). While current frontier models are closed source, and their exact training data and size remain undisclosed, the latest open-source models, such as Llama 3, were trained using 15.6 trillion tokens using a transformer model with 405 billion parameters (Dubey et al., 2024). Two of the latest open-source vision-language-action (VLA) models for robotics, OpenVLA (Kim et al., 2024) and RT-2-X (Collaboration, 2024), consist of 7B and 55B and were trained on 970K and 1M+ demonstration trajectories, respectively. Although trajectories and tokens are not comparable one-to-one, the order of magnitude between both these types of datasets remains considerably large. Furthermore, these robot datasets ignore torque control and operate only in the task space (Kim et al., 2024; Brohan et al., 2022); thus, limiting the applicability to many other robotics tasks that require force feedback.

The success of foundation models relies on web-scale datasets accumulated over decades. Scaling human-collected datasets to that same order of magnitude remains impractical for robotic applications. The limitations of current robotics datasets become particularly acute when considering the goal of enabling robots to perform multiple tasks across diverse environments, while collaborating with other humans and agents. As illustrated in Figure 1, targeting generalization across these multiple axes at once becomes infeasible due to the combinatorial explosion of required data. Current foundation models tend to succeed on only one axis of generalization at the expense of others. For instance, while a segmentation model like SAM2 shows impressive performance across many environments, it is fundamentally limited to a single task (Kirillov et al., 2023; Ravi et al., 2020). Conversely, a model like the Robot Transformer (RT-1) can perform multiple tasks, but its generalization was demonstrated in only three similar kitchen environments and lacked any interaction with humans or other robots (Brohan et al., 2022). This challenge is most evident with real world applications such as self-driving cars. Despite collecting data for years on a scale far beyond typical robotics datasets, self-driving systems have still not achieved Level 5




Figure 1: Left: Scaling deep learning models (e.g., SAM, fully self-driving cars, robot transformers) across agents, tasks, and environments dramatically increases data demands, which emphasizes the need for hybrid approaches that integrate physics and symbolic reasoning. Right: Combining Bayesian inference, deep learning, and physics to a hybrid approach optimally minimizes uncertainty and yields the best achievable compromise balancing sample and computational complexity.

autonomy, and their most advanced deployments remain constrained to a small number of geographic locations worldwide.

Recent estimates predict that current LLMs will run out of human-generated data within the next five years (Villalobos et al., 2022). An apparent solution to this problem in robotics is to generate vast quantities of synthetic data from physical simulators. However, in the following sections, we argue that using data to optimize an uninterpretable function approximator is inefficient when a known physical approximation exists and could be explicitly embedded in the model. Constraining the model to be physically plausible not only reduces sample and model complexity but also provides the safety guarantees essential for human-robot collaboration, particularly when integrated within a Bayesian framework for uncertainty quantification.

2.2. Trust and Safety

- We cannot trust the probabilities of a model that we do not understand.

Current deep learning (DL) models are known to hallucinate and produce overconfident predictions, often extrapolating incorrect assumptions based solely on correlations in the training dataset (Zhang et al., 2017; Ji et al., 2023). This limitation stems from the fact that DL models typically lack an inherent mechanism to estimate their uncertainty, making them prone to confidently incorrect outputs. Most likely, robot foundation models will inherit these issues in deep learning. This issue is particularly concerning in robot learning, as the robot behavior must generalize to new tasks and environments, i.e., to unseen and out-of-distribution examples. While search and fact-checking can mitigate hallucinations in language models, the consequences in robotics are far more severe. In contrast to errors

in LLMs, where users may correct mistakes afterward, failures in robotic systems occur in real-time and can have irreversible consequences. For example, errors in high-torque systems can lead to immediate fatal consequences, particularly in scenarios where robotic systems operate in direct collaboration with humans.

A significant challenge with deep learning is that it relies on probabilistic outputs from black-box models, which could be wrong in a principled manner by only extrapolating incorrect assumptions found only through correlations in the training dataset. This lack of interpretability raises concerns about safety and the overall trustworthiness of autonomous robotic systems. As these systems become increasingly autonomous through advanced learning, ensuring their safe operation becomes critical. However, autonomy must not come at the cost of human control (Shneiderman, 2020); instead, it should integrate mechanisms that enhance controllability and predictability. This increase in autonomy requires, therefore, coupling with an increase in human control to ensure that robots behave controllably, predictably, and in a comprehensible manner Shneiderman (2020). Evaluating reliability solely based on task success rates is insufficient; instead, trust requires the model’s ability to extrapolate causally and measure its uncertainty. Estimating and communicating uncertainty and learning causal world models are the foundations for safe and trustworthy AI.

2.3. Physics Unawareness

- Learning algorithms shouldn’t ignore the most accurate predictive models of our world.

The most promising large data collection solution for robotics is through simulated data. However, we argue that it is inefficient to use a physics function p to simulate millions of state-action pairs (s_i, u_i) to train a black-box function f to approximate the original function p . In many applications of deep RL (DRL) to robotics (Rudin et al., 2022), as well as current robot foundation models (Brohan et al., 2022, 2023; Kim et al., 2024) the known physics are partially ignored and instead learned in combination or parallel with the task itself. With this approach, the learning challenge is made significantly more difficult and valuable resources such as samples and computation are wasted to approximate already known quantities. Especially the notoriously high sample complexity of DRL could greatly benefit from a more principled integration of known physics and lead to better applicability of DRL methods in the realm of robotics (Dulac-Arnold et al., 2021; Ibarz et al., 2021; Kober et al., 2013). A more reasonable alternative is to directly use the physics function p to solve the task, given that physics is the most accurate predictive model of the world, with a short model description, and a strong generalization across domains (Fetter and Walecka, 2003; Landau and Lifshits, 1960; Feynman et al., 1965). The direct application and approximation of the function p have resulted in optimal control methods and differentiable ray tracers, capable of performing highly athletic behaviors (Kim et al., 2019; Bergonzani et al., 2023; Dai et al., 2014; Javadi et al., 2023; Vyas et al., 2023), and precise scene reconstructions (Nicolet et al., 2021; Loubet et al., 2019). Optimal control models can solve high-level tasks by optimizing at a high frequency the joint torques that control the full rigid body dynamics of the robot, and in some instances elements in the environment (Sleiman et al., 2021). However, there is a detachment between state-of-the-art perception models and action models. Current state-of-the-art computer vision models optimize large-scale neural networks with vast amounts of data, which can be applied to unseen environments and tasks (Bommasani et al., 2021). This

paradigm has provided enormous contributions to the field of computer vision, leading to foundation models in instance segmentation (Kirillov et al., 2023), feature extraction (Oquab et al., 2023), and monocular depth estimation (Yang et al., 2024). However, contrary to the previously described physics-based action models, most computer vision algorithms often ignore the best known physical approximations of the world (Lake et al., 2017). Specifically, most deep learning architectures disregard any information about the interaction between light and matter, even though they are tasked to solve problems that explicitly depend on this physical phenomenon (Spielberg et al., 2023).

3. Potentials of Bayesian Inference, Physics and Program Synthesis

3.1. Differentiable Physics for Perception and Action

The field of computer graphics has researched the appropriate level of approximation of physics-based rendering functions p that balance computational resources and physical realism (Pharr et al., 2023). Recent advances in computer graphics have led to the development of differentiable physics-based graphics engines (Nimier-David et al., 2019; Laine et al., 2020; Ravi et al., 2020) capable of reconstructing objects from multiple views (Liu et al., 2019; Loubet et al., 2019; Nicolet et al., 2021; Vicini et al., 2022; Arriaga et al., 2024). Furthermore, we can use a differentiable physics function to optimize its inputs using gradient descent to find a desired state. Optimizing the control inputs of a robot system leads to a large set of optimal control algorithms, including trajectory optimization (Betts, 2010; Kelly, 2017) and model predictive control (MPC). Moreover, new developments in differentiable hydroelastic contact models (Elandt et al., 2022; Masterjohn et al., 2022) could enhance the optimization of contact-invariant methods (Mordatch et al., 2012b,a). These methods have shown robots performing complex tasks using high-level descriptions. Furthermore, differentiable physics simulators have allowed building robot algorithms that learn to use tools in order to satisfy their objectives (Toussaint et al., 2018). The broad applications of differentiable physics engines have resulted in a large set of available physics simulators (Todorov et al., 2012; Degraeve et al., 2019; Heiden et al., 2021; Geilinger et al., 2020; Freeman et al., 2021). These differentiable physics engines have been combined with neural networks in order to provide more accurate dynamics or close the simulation reality gap (Heiden et al., 2021).

In addition to optimal control, a second promising paradigm for robot control is reinforcement learning (RL) (Sutton and Barto, 2018). Similar to optimal control, RL tackles the problem of sequential decision-making in dynamic systems. A notable difference between RL and optimal control is that optimal control methods assume a known dynamics model, whereas RL aims to learn to make decisions from experienced interactions without making any assumptions about the system dynamics. One promising subfield of RL in particular that could greatly benefit from differentiable physics engines is model-based RL (MBRL) (Moerland et al., 2023). MBRL, either a priori knows or learns a (partial) dynamics model of the environment to speed up learning by allowing the learning algorithm to evaluate long-term effects of the agent’s actions based on its approximation of its environment. By enabling the learning RL agent to plan with a model of its environment, the sample-efficiency of RL algorithms can be significantly improved, however, at the cost higher computational requirements. Using computationally efficient and physically sound differentiable dynamics models could further mitigate the negative effects of this trade-off. Moreover, MBRL algo-

gorithms with access to a differentiable simulator can compute first-order gradients induced by the policy (first-order MBRL (FO-MBRL) (Freeman et al., 2021; Xu et al., 2022; Georgiev et al., 2024)).

Most of the previous models in both differentiable rendering and differentiable robot simulators do not require additional *training* data or black-box functions to perform dynamic high-level tasks or reconstruct new scenes. Current physical descriptions are accurate and fast enough to be optimized to solve a task under reasonable time constraints. While the lack of data can be advantageous, this implies that the model representation of the system and the environment remains close. While tracking algorithms can compensate for partially unknown dynamics, the environment remains completely predefined. To extend models to deal with dynamic environments and physical changes, we need to consider models that are capable of estimating their uncertainty (Sünderhauf et al., 2018) and expand their model description (Tenenbaum et al., 2011). In the following sections, we describe these two elements in more detail.

3.2. Bayesian Inference for Uncertainty-Aware Decision Making

Understanding the physical world requires not only solving inverse problems but also quantifying uncertainty, making Bayesian inference a powerful framework for integrating prior knowledge, refining models with data, and ensuring robust, uncertainty-aware predictions in physics-based applications. A physical simulation p , such as an image renderer or a robotic simulator, uses model parameters to describe its dynamics. In the case of image rendering, those parameters could include the positions of the mesh vertices or the location and intensity of light sources. The *forward* problem consists in predicting a measurement given these parameters. Following our computer graphics example, this would render an image given the mesh vertices and the light parameters. The *inverse* problem consists in predicting the parameters given a measurement; thus, predicting the vertices and the light parameters from an image or a set of images. Inverse problem solving may result in solutions where many model parameters describe the same observation (Tarantola, 2005) and need to be constrained by using their physical description or additional information. The Bayesian workflow (Gelman et al., 2020; Martin et al., 2022) allows us to encode uncertainty and knowledge in our system and solve inverse problems through Bayes’s theorem (Jaynes, 2003). Computing all possible solutions of an inverse problem can be computationally expensive; thus, instead one can resort to computing only the most likely single solution through maximum a posteriori estimation (MAP) (Murphy, 2022).

An often disregarded application of automatic differentiation is its use in probabilistic programming frameworks. Specifically, the availability of the physical simulations derivatives opens the possibility of using efficient inference algorithms such as Hamiltonian Monte Carlo (HMC), which require the gradient of the target distribution (Betancourt, 2017). Furthermore, we can use approximate and variational inference (VI) to minimize the Kullback-Leibler divergence between the approximate and true distributions (Zhang et al., 2018). Other popular posterior approximations, such as Laplace’s approximation, require computing Hessian matrices. Finally, the maximum-entropy principle can be formulated as an optimization problem to build prior distributions using gradient descent (Jaynes, 2003). The widespread adoption and development of automatic differentiation engines in specialized hardware, such

as GPUs, recently bootstrapped the interest in large-scale Bayesian inference frameworks. Some of these new probabilistic programming frameworks include: tensorflow-probability (Dillon et al., 2017), gen (Cusumano-Towner et al., 2019), numpyro (Phan et al., 2019), and blackjax (Cabezas et al., 2024). In general, derivatives are not only useful for computing point estimates of large black-box function approximators, but also to compute the uncertainty of physically grounded models. Uncertainty quantification as well as the direct application of a differentiable physical model can help increase trust and safety in our models, reducing the sample and model complexity of learning algorithms, and building models that can learn continuously Sünderhauf et al. (2018). Bayesian inference, though computationally complex, requires fewer samples and parameters than deep learning by leveraging priors and known models. Integrating it with physics and deep learning offers a balanced tradeoff for uncertainty-aware and complex learning as shown in Fig. 1 on the right.

3.3. Integrating Deep Learning and Program Synthesis to Model Worlds

Most probabilistic programming frameworks focus on static probabilistic graphical models. However, machines that learn and think like humans must be able to learn and expand their model of the world (Lake et al., 2017; Tenenbaum et al., 2011). To facilitate incremental continuous robot learning, continual deep learning strategies tackle catastrophic forgetting, by either replaying trained tasks and classes (Rolnick et al., 2019), or by partially incrementally adapting the last output network layers (Rusu et al., 2016; Fernando et al., 2017). However, incorporating explicit knowledge into neural networks and evolving it remain open challenges for the neuro-symbolic AI community. A viable option for encoding explicit knowledge in combination with neural networks is through the use of programs (Rule et al., 2020) and program synthesis (Ellis et al., 2020). State-of-the-art approaches in program synthesis automate the generation of functional programs by leveraging domain-specific languages (DSLs), a program library, and a search strategy (Ellis et al., 2020; Bowers et al., 2023; Grayeli et al., 2024). A DSL provides structured representations of programs, allowing program synthesis systems to explore hierarchical program trees efficiently. However, brute-force searching through these trees is computationally infeasible. Deep learning can manage the tree search in program synthesis to find new programs by scoring, prioritizing, and pruning paths to improve efficiency (Balog et al., 2016; Silver et al., 2016).

Additionally, probabilistic programming can balance program complexity and accuracy by learning optimal model transitions using Reversible Jump Markov Chain Monte Carlo (RJMCMC) (Green, 1995). RJMCMC enables adaptive model selection by jumping between different parameter dimensionalities, naturally favoring simpler, more efficient models. By dynamically exploring both complex and simple solutions, the system can flexibly synthesize programs that meet the required functional specifications while minimizing unnecessary complexity. Initially, computing transdimensional jumps in RJMCMC required manually implementing the proposal kernels and their respective measure-preserving correction factors. Finally, the combination of program synthesis and (deep) RL provides several promising research avenues, e.g., when applying RL for finetuning program synthesis (Qian et al., 2022; Bunel et al., 2018), guiding RL with program synthesis (Yang et al., 2021), or improving policies with program synthesis (Bhupatiraju et al., 2018).

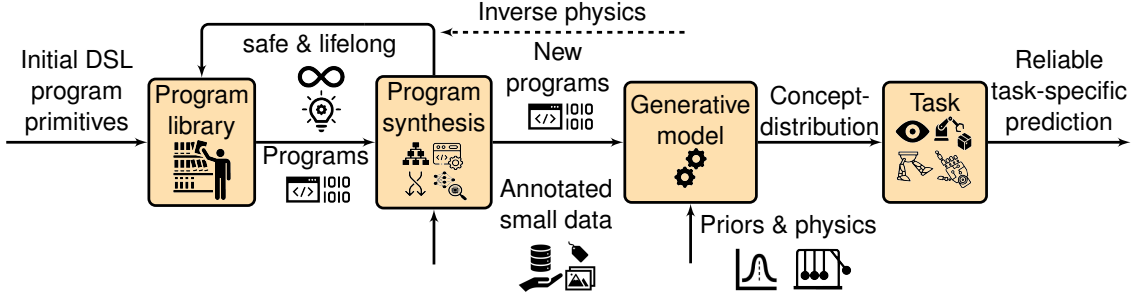


Figure 2: Our neuro-symbolic robot learning framework uses Bayesian inverse physics and program synthesis to solve diverse tasks safely, and physically consistently. Theoretically, this concept continuously evolves knowledge without forgetting.

4. Roadmap for Future Research

Our proposed neuro-symbolic roadmap for robotics prioritizes physically-grounded and uncertainty-aware reasoning for reliable real-world interaction. In this framework, deep learning and foundation models are not the primary tool but rather a complement to physics-based models, and are instead used to learn residual components, find data-driven priors, or initialize optimization problems. This integration is fundamental to lifelong learning systems, where methods like program synthesis can continuously expand the initial robot’s knowledge. Therefore, our short-term roadmap synthesizes physics, Bayesian inference, and program synthesis to build a framework for adaptable intelligent systems (Fig. 2).

- **Domain-specific languages (DSL) and primitives:** Program synthesis techniques require an initial set of reusable and extensible program primitives. To this end, the first step will include defining a versatile DSL and primitives useful for diverse robotic tasks employing both Bayesian inference and differentiable physical simulators. This new language can reuse or draw inspiration from established frameworks, including probabilistic programming systems like Pyro (Bingham et al., 2019), and automatic differentiation libraries like JAX (Frostig et al., 2018).
- **Differentiable physics:** Integrating differentiable physical simulations of perception and action within machine learning models can reduce the number of parameters and training examples used in neural networks. Moreover, physical consistency can limit the search space of program synthesizers by pruning infeasible solutions. Within this context, we need to integrate multiple physical simulators, including rigid body simulators (Todorov et al., 2012), hydroelastic contact models (Elandt et al., 2022), and differentiable renderers (Nimier-David et al., 2019).
- **Bayesian inference:** Robots need to be capable of reasoning under uncertainty. Specifically, they must estimate their posterior distributions. The second step should involve implementing probabilistic models within the DSL, utilizing probabilistic programming frameworks, and applying variational inference to handle computational complexity (Cabezas et al., 2024; Blei et al., 2017).

- **Deep learning integration:** Naive program synthesis involves searching over a large discrete space. Lacking a differentiable continuous search space, one cannot directly utilize gradient-based optimization. Neural search in combination with current LLMs provides a path forward to alleviate these issues (Balog et al., 2016; Shi et al., 2023; Ellis et al., 2020; Guo et al., 2024).
- **Program synthesis for continual library learning:** Program synthesis presents a viable option to represent multiple domains of knowledge (Rule et al., 2020). By iteratively refining and synthesizing symbolic program libraries and updating knowledge, program synthesis retains old knowledge and is inherently adaptive (Bowers et al., 2023; Ellis et al., 2020; Grayeli et al., 2024). Moreover, recent developments in Reversible Jump Markov Chain Monte Carlo (RJMCMC) (Neklyudov et al., 2020; Cusumano-Towner et al., 2020) open the possibility for Bayesian program synthesis (Saad et al., 2019; Matheos et al., 2020).

Bayesian inverse physics demands extensive simulations and high-dimensional sampling, making inference computationally expensive and challenging for real-time robotic decision-making. Efficient approximations, parallelization, or surrogate models will be essential for practical deployment in dynamic environments and require further research. Building Bayesian models requires careful prior and likelihood selection, limiting automation in model construction. However, future work could explore automating prior selection while ensuring physical consistency. Targeting long-term real-world deployability for the hybrid approach requires more research on novel benchmarking methods and real-world testing. Finally, the seamless integration of symbolic program synthesis, neural networks and differentiable physics engines constitutes a huge challenge and requires ensuring compatibility between discrete symbolic representations and continuous robotic actions.

5. Conclusion

Recent advancements in neural networks have attracted significant private investment, fueling the development of generic function approximators. The available vast private data and computational resources risk fostering a research culture focused on ever-larger datasets and black-box models, consuming energy on the scale of nuclear power plants. However, the human brain operates on approximately 20 watts. A more principled approach would be first to encode the physical approximations we already know into our models, serving as a foundation for model-building algorithms inspired by program synthesis and probabilistic programming, enabling more efficient and interpretable learning. While current foundation models are limited in building general-purpose robotics, their specialized skills are valuable for well-defined tasks in closed environments. We propose exploring alternatives that improve generalization, transparency, and safety while reducing data and computational demands.

6. Acknowledgments

This work was funded with the Robotics Institute Germany (RIG) (Grant 16ME1010) and the ActGPT (Grant 01IW25002) projects, through the German Federal Ministry of Education and Research (BMBF).

References

- Octavio Arriaga, Jichen Guo, Rebecca Adam, Sebastian Houben, and Frank Kirchner. Bayesian inverse graphics for few-shot concept learning. In *International Conference on Neural-Symbolic Learning and Reasoning*, pages 141–165. Springer, 2024.
- Matej Balog, Alexander L Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. Deepcoder: Learning to write programs. *arXiv preprint arXiv:1611.01989*, 2016.
- Ivan Bergonzani, Mihaela Popescu, Shivesh Kumar, and Frank Kirchner. Fast dynamic walking with rh5 humanoid robot. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, pages 1–8. IEEE, 2023.
- Michael Betancourt. A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*, 2017.
- John T Betts. *Practical methods for optimal control and estimation using nonlinear programming*. SIAM, 2010.
- Surya Bhupatiraju, Kumar Krishna Agrawal, and Rishabh Singh. Towards Mixed Optimization for Reinforcement Learning with Program Synthesis, July 2018.
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *Journal of machine learning research*, 20(28):1–6, 2019.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Matthew Bowers, Theo X Olausson, Lionel Wong, Gabriel Grand, Joshua B Tenenbaum, Kevin Ellis, and Armando Solar-Lezama. Top-down synthesis for library learning. *Proceedings of the ACM on Programming Languages*, 7(POPL):1182–1213, 2023.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Rudy Bunel, Matthew Hausknecht, Jacob Devlin, Rishabh Singh, and Pushmeet Kohli. Leveraging Grammar and Reinforcement Learning for Neural Program Synthesis, May 2018.
- Alberto Cabezas, Adrien Corenflos, Junpeng Lao, Rémi Louf, Antoine Carnec, Kaustubh Chaudhari, Reuben Cohn-Gordon, Jeremie Coullon, Wei Deng, Sam Duffield, et al. Blackjax: Composable bayesian inference in jax. *arXiv preprint arXiv:2402.10797*, 2024.
- Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. Technical report, arXiv, June 2024.
- Ben Cottier, Robi Rahman, Loredana Fattorini, Nestor Maslej, Tamay Besiroglu, and David Owen. The rising costs of training frontier ai models. *arXiv preprint arXiv:2405.21015*, 2024.
- Marco Cusumano-Towner, Alexander K Lew, and Vikash K Mansinghka. Automating involutive mcmc using probabilistic and differentiable programming. *arXiv preprint arXiv:2007.09871*, 2020.
- Marco F Cusumano-Towner, Feras A Saad, Alexander K Lew, and Vikash K Mansinghka. Gen: a general-purpose probabilistic programming system with programmable inference. In *Proceedings of the 40th acm sigplan conference on programming language design and implementation*, pages 221–236, 2019.
- Hongkai Dai, Andrés Valenzuela, and Russ Tedrake. Whole-body motion planning with centroidal dynamics and full kinematics. In *2014 IEEE-RAS International Conference on Humanoid Robots*, pages 295–302. IEEE, 2014.
- Jonas Degraeve, Michiel Hermans, Joni Dambre, and Francis Wyffels. A differentiable physics engine for deep learning in robotics. *Frontiers in neurorobotics*, 13:6, 2019.
- Joshua V Dillon, Ian Langmore, Dustin Tran, Eugene Brevdo, Srinivas Vasudevan, Dave Moore, Brian Patton, Alex Alemi, Matt Hoffman, and Rif A Saurous. Tensorflow distributions. *arXiv preprint arXiv:1711.10604*, 2017.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Gabriel Dulac-Arnold, Nir Levine, Daniel J. Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis. *Machine Learning*, 110(9):2419–2468, September 2021. ISSN 1573-0565. doi: 10.1007/s10994-021-05961-4.

- Ryan Elandt, Evan Drumwright, Michael Sherman, and Andy Ruina. A pressure field model for fast, robust approximation of net contact force and moment between nominally rigid objects. in 2019 ieee. In *RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8238–8245, 2022.
- Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B Tenenbaum. Dreamcoder: Growing generalizable, interpretable bayesian programs. *arXiv preprint arXiv:2006.08381*, 2020.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017.
- Alexander L Fetter and John Dirk Walecka. *Theoretical mechanics of particles and continua*. Courier Corporation, 2003.
- Richard P Feynman, Robert B Leighton, Matthew Sands, and Everett M Hafner. The feynman lectures on physics; vol. i. *American Journal of Physics*, 33(9):750–752, 1965.
- C Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax—a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281*, 2021.
- Roy Frostig, Matthew James Johnson, and Chris Leary. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning*, 4(9), 2018.
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bäcker, Bernhard Thomaszewski, and Stelian Coros. Add: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020.
- Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles C Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Burkner, and Martin Modrak. Bayesian workflow. *arXiv preprint arXiv:2011.01808*, 2020.
- Ignat Georgiev, Krishnan Srinivasan, Jie Xu, Eric Heiden, and Animesh Garg. Adaptive Horizon Actor-Critic for Policy Learning in Contact-Rich Differentiable Simulation. In *Forty-First International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024.
- Zoubin Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, 2015.
- Gerd Gigerenzer. *How to stay smart in a smart world: Why human intelligence still beats algorithms*. MIT Press, 2022.
- Alison Gopnik. Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102):1623–1627, 2012.
- Arya Grayeli, Atharva Sehgal, Omar Costilla Reyes, Miles Cranmer, and Swarat Chaudhuri. Symbolic regression with a learned concept library. *Advances in Neural Information Processing Systems*, 37:44678–44709, 2024.

- Peter J Green. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732, 1995.
- Thomas L Griffiths, Nick Chater, and Joshua B Tenenbaum. *Bayesian models of cognition: reverse engineering the mind*. MIT Press, 2024.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*, 2024.
- Eric Heiden, David Millard, Erwin Coumans, Yizhou Sheng, and Gaurav S Sukhatme. Neuralsim: Augmenting differentiable simulators with neural networks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9474–9481. IEEE, 2021.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine. How to train your robot with deep reinforcement learning: Lessons we have learned. *The International Journal of Robotics Research*, page 027836492098785, January 2021. ISSN 1741-3176. doi: 10.1177/0278364920987859.
- Mahdi Javadi, Daniel Harnack, Paula Stocco, Shivesh Kumar, Shubham Vyas, Daniel Pizzutilo, and Frank Kirchner. Acromonk: A minimalist underactuated brachiating robot. *IEEE Robotics and Automation Letters*, 8(6):3637–3644, 2023. doi: 10.1109/LRA.2023.3269296.
- Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge university press, 2003.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12), March 2023. ISSN 0360-0300. doi: 10.1145/3571730. URL <https://doi.org/10.1145/3571730>.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Matthew Kelly. An introduction to trajectory optimization: How to do your own direct collocation. *SIAM Review*, 59(4):849–904, 2017.
- Donghyun Kim, Jared Di Carlo, Benjamin Katz, Gerardo Bledt, and Sangbae Kim. Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control. *arXiv preprint arXiv:1909.06586*, 2019.

- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement Learning in Robotics: A Survey. *The International Journal of Robotics Research*, 32(11):1238–1274, August 2013. doi: 10.1177/0278364913495721.
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics (ToG)*, 39(6):1–14, 2020.
- Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.
- Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.
- Lev Davidovich Landau and Lifshits. *Mechanics*. CUP Archive, 1960.
- Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7708–7717, 2019.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019.
- Sylwia Łukasik, Sławomir Tobis, Sylwia Kropińska, and Aleksandra Suwalska. Role of assistive robots in the care of older people: survey study among medical and nursing students. *Journal of medical Internet research*, 22(8):e18003, 2020.
- Michael Lutter, Johannes Silberbauer, Joe Watson, and Jan Peters. A differentiable newton euler algorithm for multi-body model learning. *arXiv preprint arXiv:2010.09802*, 2020.
- Daniel Macis, Sara Perilli, and Cristina Gena. Employing socially assistive robots in elderly care. In *Adjunct proceedings of the 30th ACM conference on user modeling, adaptation and personalization*, pages 130–138, 2022.
- Gary Marcus and Ernest Davis. *Rebooting AI: Building artificial intelligence we can trust*. Vintage, 2019.
- Osvaldo A Martin, Ravin Kumar, and Junpeng Lao. *Bayesian Modeling and Computation in Python*. Chapman and Hall/CRC, 2022.

- Joseph Masterjohn, Damrong Guoy, John Shepherd, and Alejandro Castro. Velocity level approximation of pressure field contact patches. *IEEE Robotics and Automation Letters*, 7(4):11593–11600, 2022.
- George Matheos, Alexander K Lew, Matin Ghavamizadeh, Stuart Russell, Marco Cusumano-Towner, and Vikash Mansinghka. Transforming worlds: Automated involutive mcmc for open-universe probabilistic models. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2020.
- Thomas M. Moerland, Joost Broekens, Aske Plaat, and Catholijn M. Jonker. Model-based Reinforcement Learning: A Survey. *Found. Trends Mach. Learn.*, 16(1):1–118, 2023. doi: 10.1561/22000000086.
- Igor Mordatch, Zoran Popović, and Emanuel Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 137–144, 2012a.
- Igor Mordatch, Emanuel Todorov, and Zoran Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (ToG)*, 31(4):1–8, 2012b.
- Kevin P Murphy. *Probabilistic machine learning: an introduction*. MIT press, 2022.
- Kirill Neklyudov, Max Welling, Evgenii Egorov, and Dmitry Vetrov. Involutive mcmc: a unifying framework. In *International Conference on Machine Learning*, pages 7273–7282. PMLR, 2020.
- Issa AD Nesnas, Lorraine M Fesq, and Richard A Volpe. Autonomy for space robots: Past, present, and future. *Current Robotics Reports*, 2(3):251–263, 2021.
- Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. Large steps in inverse rendering of geometry. *ACM Transactions on Graphics (TOG)*, 40(6):1–13, 2021.
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (ToG)*, 38(6): 1–17, 2019.
- David O. Obura, Greta Aeby, Natchanon Amornthammarong, Ward Appeltans, Nicholas Bax, Joe Bishop, Russell E. Brainard, Samuel Chan, Pamela Fletcher, Timothy A. C. Gordon, Lew Gramer, Mishal Gudka, John Halas, James Hendee, Gregor Hodgson, Danwei Huang, Mike Jankulak, Albert Jones, Tadashi Kimura, Joshua Levy, Patricia Miloslavich, Loke Ming Chou, Frank Muller-Karger, Kennedy Osuka, Melita Samoilys, Stephen D. Simpson, Karenne Tun, and Supin Wongbusarakum. Coral Reef Monitoring, Reef Assessment Technologies, and Ecosystem-Based Management. *Frontiers in Marine Science*, 6, 2019. ISSN 2296-7745. doi: 10.3389/fmars.2019.00580. URL <https://www.frontiersin.org/articles/10.3389/fmars.2019.00580>.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khilodov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2:

- Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- Thomas Parr, Giovanni Pezzulo, and Karl J Friston. *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press, 2022.
- Judea Pearl. Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv preprint arXiv:1801.04016*, 2018.
- Du Phan, Neeraj Pradhan, and Martin Jankowiak. Composable effects for flexible and accelerated probabilistic programming in numpyro. *arXiv preprint arXiv:1912.11554*, 2019.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. MIT Press, 2023.
- Steven Pinker. Visual cognition: An introduction. *Cognition*, 18(1-3):1–63, 1984.
- Yu Qian, Xuegong Zhou, Hao Zhou, and Lingli Wang. Efficient Reinforcement Learning Framework for Automated Logic Synthesis Exploration. In *International Conference on Field-Programmable Technology, (IC)FPT 2022, Hong Kong, December 5-9, 2022*, pages 1–6. IEEE, 2022. doi: 10.1109/ICFPT56656.2022.9974330.
- Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, SM Ali Eslami, and Matthew Botvinick. Machine theory of mind. In *International conference on machine learning*, pages 4218–4227. PMLR, 2018.
- Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in neural information processing systems*, 32, 2019.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on Robot Learning*, pages 91–100. PMLR, 2022.
- Joshua S Rule, Joshua B Tenenbaum, and Steven T Piantadosi. The child as hacker. *Trends in cognitive sciences*, 24(11):900–915, 2020.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- Feras A Saad, Marco F Cusumano-Towner, Ulrich Schaechtle, Martin C Rinard, and Vikash K Mansinghka. Bayesian synthesis of probabilistic programs for automatic data modeling. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–32, 2019.

- Kensen Shi, Hanjun Dai, Wen-Ding Li, Kevin Ellis, and Charles Sutton. Lambdabeam: Neural program search with higher-order functions and lambdas. *Advances in Neural Information Processing Systems*, 36:51327–51346, 2023.
- Ben Shneiderman. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6):495–504, 2020. doi: 10.1080/10447318.2020.1741118. URL <https://doi.org/10.1080/10447318.2020.1741118>.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Jean-Pierre Sleiman, Farbod Farshidian, Maria Vittoria Minniti, and Marco Hutter. A unified mpc framework for whole-body dynamic locomotion and manipulation. *IEEE Robotics and Automation Letters*, 6(3):4688–4695, 2021.
- Andrew Spielberg, Fangcheng Zhong, Konstantinos Rematas, Krishna Murthy Jatavallabhula, Cengiz Oztireli, Tzu-Mao Li, and Derek Nowrouzezahrai. Differentiable visual computing for inverse problems and machine learning. *Nature Machine Intelligence*, 5(11):1189–1199, 2023.
- Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, et al. The limits and potentials of deep learning for robotics. *The International journal of robotics research*, 37(4-5):405–420, 2018.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018. ISBN 978-0-262-03924-6.
- Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.
- Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285, 2011.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.
- Marc A Toussaint, Kelsey Rebecca Allen, Kevin A Smith, and Joshua B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. *Robotics: Science and systems foundation*, 2018.
- Matias Valdenegro-Toro. I find your lack of uncertainty in computer vision disturbing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1263–1272, 2021.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Differentiable signed distance function rendering. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022.

- Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbhahn. Will we run out of data? limits of llm scaling based on human-generated data. *arXiv preprint arXiv:2211.04325*, 2022.
- Shubham Vyas, Lasse Maywald, Shivesh Kumar, Marko Jankovic, Andreas Mueller, and Frank Kirchner. Post-capture detumble trajectory stabilization for robotic active debris removal. *Advances in Space Research*, 72(7):2845–2859, 2023.
- Jie Xu, Viktor Makoviyshuk, Yashraj S. Narang, Fabio Ramos, Wojciech Matusik, Animesh Garg, and Miles Macklin. Accelerated Policy Learning with Parallel Differentiable Simulation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024.
- Yichen David Yang, Jeevana Priya Inala, Osbert Bastani, Yewen Pu, Armando Solar-Lezama, and Martin Rinard. Program Synthesis Guided Reinforcement Learning for Partially Observed Environments, November 2021.
- Cheng Zhang, Judith Bütepage, Hedvig Kjellström, and Stephan Mandt. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):2008–2026, 2018.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.