



Active learning in robotics: A review of control principles

Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey^{*,1}

Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

ARTICLE INFO

Keywords:

Active learning
Robotics
Robot control
Learning theory
Perception and sensing
Artificial intelligence

ABSTRACT

Active learning is a decision-making process. In both abstract and physical settings, active learning demands both analysis and action. This is a review of active learning in robotics, focusing on methods amenable to the demands of embodied learning systems. Robots must be able to learn efficiently and flexibly through continuous online deployment. This poses a distinct set of control-oriented challenges—one must choose suitable measures as objectives, synthesize real-time control, and produce analyses that guarantee performance and safety with limited knowledge of the environment or robot itself. In this work, we survey the fundamental components of robotic active learning systems. We discuss classes of learning tasks that robots typically encounter, measures with which they gauge the information content of observations, and algorithms for generating action plans. Moreover, we provide a variety of examples – from environmental mapping to nonparametric shape estimation – that highlight the qualitative differences between learning tasks, information measures, and control techniques. We conclude with a discussion of control-oriented open challenges, including safety-constrained learning and distributed learning.

1. Introduction

“Perceptual activity is exploratory, probing, searching; percepts do not simply fall onto sensors as rain falls onto ground. We do not just see, we look”. (R. Bajcsy in her 1988 paper *Active Perception* [1]). The difference between seeing and looking is the presence of action—seeing is passive and looking is active. Unfortunately, we do not use distinct words for passive learning and active learning, often leading to confusing the two and unintentionally treating “learning” as passive learning with active learning as an afterthought. Nevertheless, how we acquire data impacts the quality of learning and what is even possible to learn, indicating that control – both analysis and synthesis – in learning will inevitably be important. More than three decades after Bajcsy’s comments, the key elements of how control synthesis and analysis should inform learning remain largely unaddressed, and the vast majority of work in learning still focuses on analysis of passively collected data; this body of work makes up a statistical theory of learning. Still absent is an action-oriented theory of learning—a control theory for learning. How should control synthesis affect learning? What sort of feedback interconnections facilitate learning?

When prior knowledge and existing datasets are widely available, passive learning has proven to be a successful tool for constructing parametric representations of statistical relationships in data. Broadly, passive learning is an optimization process in which the parameters of a model are fit according to data. The last decade has seen major strides in robotics dependent on the advent of modern learning methodologies, particularly variations of deep neural networks [2]. However, in settings where previously existing datasets are unavailable, and where products of human knowledge (e.g., labeled datasets, knowledge graphs) do not exist, a robot will have to engage in unsupervised discovery and acquire the data it needs [3]. We refer to this process as *active learning* (see Fig. 1). In contrast to passive learning, active learning is a decision-making process where agents take actions to gather the data that best realizes a learning objective.

Animals use their bodies to learn. To paraphrase Bajcsy, we do not just passively learn, we actively learn—the pages of a book do not just turn before our eyes while we absorb information. For agents with physical bodies, such as animals or robots, active learning demands understanding and exploiting the role of embodiment and physical

* Corresponding author.

E-mail addresses: annalisa.taylor@u.northwestern.edu (A.T. Taylor), tberrueta@u.northwestern.edu (T.A. Berrueta), t-murphey@northwestern.edu (T.D. Murphey).

URLs: <https://murpheylab.github.io/people/annalisataylor> (A.T. Taylor), <https://murpheylab.github.io/people/thomasberrueta> (T.A. Berrueta), <http://murpheylab.github.io> (T.D. Murphey).

¹ This material is based upon work supported by the United States National Science Foundation under Grant CNS 1837515 and by the United States Army Research Office MURI award W911NF-19-1-0233. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the aforementioned institutions.

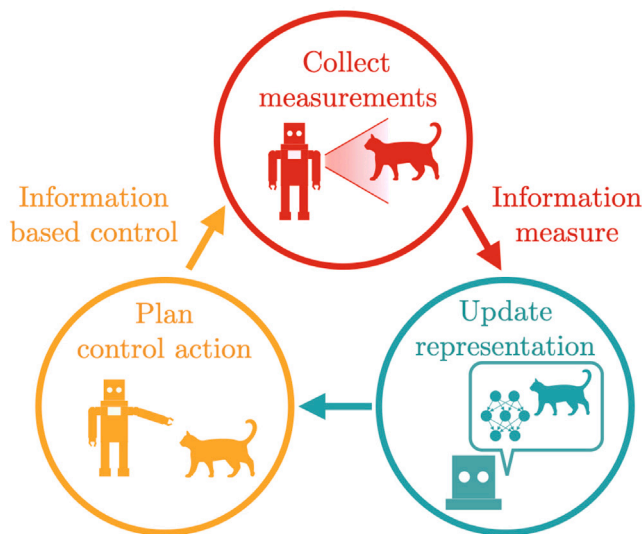


Fig. 1. The active learning process: A learner leverages information measures to formulate actions for collecting relevant or descriptive data. Active learning includes the feedback control of a system for which the internal state is both a learning system and history.

interaction in learning. Insofar as robotics should take inspiration from biology, active learning in robotics will involve the purposeful movement of a robot's body; here, control synthesis tools will connect decision-making to the resulting movement.

There is a rich literature on how animals use their bodies and movements to improve information acquisition [4–12]. For example, in [13] we demonstrated that a variety of animals engage in active information acquisition by exploring their environment in proportion to the local amount of perceived uncertainty. In addition to a medium for embodied movement plans, physical bodies are independently capable of implicit computation [14,15], information storage [16], novelty detection [17], and learning [18]. By harnessing the power of embodiment and morphological computation [19], active learning presents a promising way forward for robotics problems where the outcomes of physical interactions may be unknown *a priori*, such as in soft robotics [20].

Not only is embodiment and movement paramount to information acquisition and active learning, but movements themselves can be informative. Recent work analyzing animal and human movement has begun to interpret physical bodies as information channels and motions as information-carrying signals. This has led to the development of methods that help to understand the pathology of conditions such as autism spectrum disorder [21], schizophrenia [22], and stroke [23,24] through an information-theoretic analysis of movement. More generally, this suggests that in order to realize learning objectives, active learning requires measures that capture the information content of an agent's movements.

Counterintuitively, information-rich movement does not always appear productive, orderly, or carefully planned. A well-studied example of this is the optimality of diffusion in animal foraging—here, purely stochastic motion plans have been shown to be highly informative [25–27]. Another example of interest to researchers for decades is that of playful behavior in animals [28,29]. One may ask why animals would expend significant energy on movement that is not key to survival; for our purposes, we consider these active behaviors as enhancing learning [30]. Hence, to learn through movement, agents must engage in exploratory behaviors that may not always seem useful.

Despite its clear connections to our understanding of learning in animals and humans, the field of active learning finds its origins in theoretical computer science [31]. In this setting, agents are represented

by disembodied algorithms whose actions are limited to making queries about observed data samples. As a result, many modern frameworks for artificial intelligence have tended to neglect the role of physics and embodiment on the learning process. However, adapting to the constraints of the real world is crucial to learning in the wild. Even the most successful traditional machine learning techniques for robot control, such as reinforcement learning, rely on “big data” generated from simulated rollouts. In reality, robot deployment is a time and physically intensive activity, and robots cannot be instantly reset and redeployed at will. To make matters worse, informative data samples are typically sparse. Taken together, these issues highlight the importance of considering sample efficiency and deployment efficiency in robot learning. On the other hand, control theory has a long history of dealing with the constraints imposed by the laws of physics, while simultaneously managing secondary – yet very important – objectives such as safety, robustness, and efficiency.

There are many areas of robotics that will require the type of black-box flexibility of machine learning to make progress. When principled alternatives to modeling the physics of complex interactions between agents and their environments do not exist, machine learning can sometimes be the only way to enable robot control. One area in which flexible learning tools are particularly useful is in high-dimensional nonlinear sensing, where deep convolutional networks excel at integrating potentially hundreds of complex and highly-redundant sensory signals into compressed and informative signals [15,32]. At times, the interactions between a robot and its environment may be infeasible to model either due to properties of the environment (e.g., locomotion in granular media [33]), or the robot itself (e.g., compliant soft robots [34]). Thus, the field of robotic active learning has the potential to overcome the challenges inherent to robot control and machine learning by inheriting the best qualities of both. In this review, we highlight important progress made towards this goal, and motivate future directions for developing an action-oriented theory of embodied learning.

The organization of this review is as follows. First, we cover the history and basic considerations required for an active learning system—what there is to learn, how to measure information in actions, and how to generate such informative actions. Then, we survey key areas of application for active learning and open challenges in the field. The authors' own work plays a role in creating a narrative, but with consistent reference to the broad literature in robotics on related areas. Section 2 covers a brief history of the field of active learning and its origins in the broader field of computer science. In Section 3 we cover different learning goals, increasing in complexity from learning state parameters to abstract features. Then, Section 4 discusses measures of information, focusing on those appropriate to be used as control objectives by synthesis methods such as those in Section 5. In Section 6, we discuss common applications where facets of active learning naturally arise, whether explicitly or implicitly, in problem formulations. Finally, in Section 7 we discuss extensions and open challenges followed by conclusions in Section 8.

2. History of active learning

Since its inception, robotics has been interested in making embodied agents learn and adapt to their surroundings like biological organisms [35]. However, due to fundamental limitations on computing hardware, programming machines [36] and adaptation to external stimuli [37], robot learning was limited to the most rudimentary demonstrations throughout the mid-20th century. After establishing his theory of computation [38], Alan Turing shifted his focus to the question of whether machines could think and learn [39]. Turing's efforts prompted both the philosophical and formal study of artificial intelligence [40].

While hardware posed constraints on applied learning, the second half of the century saw the founding of the field of computational learning theory [41–44]. Analogous to computability theory, computational

learning theory focuses on assessing the “learnability” of concepts under different models of learning, such as inductive inference [45], online learning [46], statistical query learning [47,48], among many others. The diversity of models of learning speaks to the difficulty of capturing what we mean when we say that a concept is learnable. To this day, useful models of learning are being introduced to tackle new problems on learnability [49]. Of the many mathematical frameworks for learning, the most successful and widely used is the Probably Approximately Correct (PAC) learning model [43,50,51]—a particularly important framework because it was the first to bring insights from the theory of computational complexity to the study of learning. Across its many models of learning, computational learning theory forms the primary means through which we mathematically model and formally understand learning as a computational problem.

The influence of computability theory [52] is particularly visible in the field’s focus on automata theory and linguistics [53], where problems are often framed as learning languages or equivalent automata specifying the languages. In contrast, much of robotics is grounded in the history of industrial automation, where mechanical interactions are the fundamental object of interest [54]. As a result, robot learning focuses on the role of physics on sensing, actuation, and mechanical interactions with the environment for the purpose of learning.

One of the most important areas within computational learning theory is that of query learning [47,55,56]. This field is concerned with identifying the classes of functions that a “learner” (e.g., an algorithm) can learn by observing samples of data provided by an “oracle” (e.g., a teacher or an environment) using a given model of learning. At each stage of the learning process, the learner has a “learning hypothesis” about the nature of the function class that it is learning. In the context of query learning, the learner is additionally allowed to ask the oracle for information about the samples it is observing or about its current learning hypothesis [31]. The learner then must make decisions about what queries to present to the oracle in order to advance its learning objective [57]. In this way, learning is no longer framed as a passive process. Instead, it is a decision-theoretic process through which the learner takes actions in order to further its objective—or in other words, *active learning*. By leveraging their decision-making, active learners can almost always achieve the same performance as an equivalent passive learner with exponentially fewer data samples [58]. This framing can be restrictive in a robotics context where actions have the potential to elicit information and affect the environment or learning objective. Despite forming a theory grounded in the decision-making of learning agents, computational learning theory has not concerned itself with these types of practical considerations that embodied robot learning demands.

Another theory of learning largely independent from those discussed above is reinforcement learning (RL), which finds its origins in the study of conditioning in psychology [59]. As originally envisioned, RL refers to the use of external stimuli and incentive structures to elicit desired behavior out of animals or humans [60]. In this sense, RL was established as a theory of learned behavior rather than learning in-itself. However, its mathematical underpinnings were not established until the second half of the 20th century in the work of Richard Sutton and Andrew Barto among others [61–63]. By grounding their work in the theory of dynamic programming [64] and optimal control [65], Sutton and Barto created a rich mathematical theory of learning and control based on the behavioral psychology of reinforcement [66]. Typically, an RL problem is framed as a Markov Decision Process (MDP) where an agent must take actions in order to explore their environment and learn how to maximize their reward signal [67]. When agents are making decisions and taking actions to actively gather data and learn about their objective, we consider RL to be a type of active learning. In contrast, if exploration is being handled passively through naively randomized simulated experience, we do not.

Despite its early uses for optimal control [63], RL has only recently become a primary technique for robot learning due to the many

successes of deep RL in continuous control [68–70]. However, most methods developed for deep RL are ill-suited to robot learning because of their large data requirements, lack of generalizability between tasks, as well as their inability to learn incrementally and guarantee safety [71–73]. While techniques such as Maximum Entropy RL have taken steps to improve data efficiency and generalizability in robot learning settings [74–76], deep RL is still far off from seamless deployment in the real world due to its reliance on simulated experience to make progress on learning and control objectives [77–79]. Moreover, easily specifying and incorporating safety [80], stability [81], controllability [82,83], or reachability [84] remains an open challenge. Taken together, these points highlight that – despite being a theory of active learning based on the behavior of embodied agents – RL is underdeveloped for many robotic applications in its present form.

In this section we have briefly outlined the historical development of active learning as a field. Throughout the literature and across its different subfields, we have found that although researchers have had great interest in applying active learning methods to robotics problems, there is still a need for the development of theories of active learning specifically *for* robotics. Such theories of robot learning should center the properties of the agent as an embodied control system with requirements for stability, safety, sample efficiency, and continuous deployment. To this end, much of the work that we present in this review focuses on aspects of embodiment, and suggests the possibility of developing a control-oriented theory of embodied active learning.

3. What do robots need to learn?

What does a robot need to learn from data? Learning goals can be grouped into problems of increasing sophistication and level of abstraction. Here we will distinguish between learning *parameters*, as a relatively simple starting point, learning *models*, and learning *features*. This division is by no means unique, but provides a useful taxonomy for discussing what learning goals we may have for a robotic system.

3.1. Parameters

Learning parameters is relevant in many settings. For instance, one may wish to determine the location of an object, food, or predators. In this case, the parameters of interest are spatial coordinates that localize the object. If the parameters evolve in time (e.g., a mobile object) they may have dynamical properties that can be exploited or learned. If a model is known, parametric filters [85–88] may be used. When the posterior probabilities of an inference model are not expected to be approximately Gaussian, nonparametric filters, such as Bayesian filters [89,90], histogram filters [91], or particle filters [92–94] are often used instead. Active learning can be critical to overcoming sensor limitations and identifying a wide variety of parameters. A salient setting for active learning is near-field sensing. Near-field sensing includes tactile sensing which requires mechanical contact and electrosense, where close proximity is necessary. Hence, when subject to near-field sensing constraints, robots must leverage their agency for successful parameter identification. In far-field sensing, such as cameras and radar at a distance, actions may play a more limited role in parameter identification because the sensor range automatically provides substantial information without the need for movement.

3.2. Models

Models generalize parameters, and can be models of either the robot itself, such as a model of the dynamics, or the environment, such as a topographical map. The ability of a robot to learn a model of its dynamics is important in rapidly shifting environments where first-principle models struggle to make reliable predictions. The problem of system identification is often parametric, focusing on describing the dynamics using models whose structure and number of parameters are




Fig. 2. Shape reconstruction: This example shows the active identification of an unknown geometry in the environment, using binary contact measurements as the measurement modality [107]. By developing data-driven models of objects, robots can search for and recognize obstacles or tools without needing analytic or CAD models.

fixed *a priori*, such as in neural networks. However, system identification may be *nonparametric* as well, as in Gaussian process regression and other kernel-based methods. Nonparametric models may be particularly useful when robots operate in unstructured or unknown environments. While parametric models have also been successfully used in this context, it is difficult to know ahead of time that a parametrized model will have the representational capacity to characterize the environment. This has led to the use of models with an increasing number of parameters – sometimes on the order of billions of parameters – to ensure that the network can capture the properties of the environment.

3.2.1. Mapping

Mapping is one form of modeling the environment that emphasizes its geometry. Mapping applications often use occupancy grids [95,96], coverage maps [91], and Gaussian process regression to represent spatially-varying phenomena or high-dimensional belief spaces [97–102]. These techniques presume coverage—that data has been taken over a sufficiently varied area to reconstruct and represent the properties of the environment. The active learning approach instead suggests that an agent reacts to data it collects locally and then adjusts its mapping strategy. While environmental mapping in open air is not an application that necessarily demands the use of active learning methods, other types of environments may not be as straightforward. For example, underwater exploration is difficult because robots are subject to stringent constraints on sensing, actuation, and communication. Here, robots often need to operate in environments where light levels prevent long-range visual monitoring, which demands the use of active learning tools in order to construct motion plans that incrementally adapt to the robot's uncertain measurements [103–105]. In [106], the authors use control and Gaussian process regression to model, map, and actively sample the distribution of phytoplankton in the ocean off the coast of Norway, thereby greatly accelerating environmental monitoring and mapping of oceanic resources.

3.2.2. Shape

Similarly to mapping, nonparametric shape estimation is another area of model learning that focuses on the geometric relationship between collected data samples [108]. The shape estimation literature grew from the field of computer vision, and has traditionally focused on static tasks, such as estimating the poses of human bodies [109] or the curvature of roads from image samples [110]. However, as we increasingly deploy autonomy in the real world, determining the shape and material properties of unknown objects may be necessary to interact with them and potentially employ them as tools. To this end, the process of shape estimation may need to be dynamic and probing in nature, requiring that agents leverage their control authority to actively learn the properties of the object.

As an informative example of this kind of learning problem, we share some results from our own work. In [111], we considered nonparametric shape estimation using contact-based sensors to actively learn the shapes of obstacles in the robot's environment, which we then extended towards data-driven mapping and localization [107].

Fig. 2 shows a three dimensional set of objects whose shapes are being reconstructed from binary contact measurements made by a simulated mobile robot. By actively generating trajectories that make contact with the object surfaces, we maximize the Fisher information of the support vector machine (SVM) object model and successfully identify them. The enabling insight is the use of the Fisher information, which we discuss at length in Section 4, to synthesize object–robot interactions that are optimally informative.

3.2.3. Dynamics

One of the most crucial learning tasks is that of identifying the agent's own dynamics. Whether learning the dynamics is necessary due to their intrinsic complexity, or as a result of a sudden malfunction or compliant interaction, there are many scenarios in which it may be impossible or infeasible to have an accurate prior representation of the system's dynamics. Self-identifying dynamics is an active process, where the agent needs to take actions and collect data that explore its different behavioral regimes. In some settings, models that are well-specified in certain behavioral regimes may have to be augmented through data-driven means to work in extraneous conditions. For instance, in aerospace applications data-driven techniques will have excellent data available for nominal conditions but often no data available for specific off-nominal conditions, suggesting the need for active learning outside of the nominal regime [112]. Since the literature on learning dynamics is very diverse, providing a comprehensive survey would require its own review [113–117]. Instead, here we review a few particular representations of dynamics that are of particular interest to the field of robotics.

Deep neural networks (DNNs) are models comprised of many individual units (*i.e.*, computational synthetic neurons) with limited capabilities that together, through their interconnections, are capable of great representational power [2]. As we have discussed earlier in this review, deep networks are not always suited to the demands of robot learning due to their high data and computational requirements. Nonetheless, certain network architectures have been shown to be well-suited to predicting dynamics, such as recurrent neural networks [118], whose capabilities enable them to predict the global structure of temporal dynamics from local measurements. In settings where learning does not need to occur rapidly or incrementally, carefully chosen deep learning architectures have been successful in learning robot dynamics for control [119,120]. While DNNs have been successful in many robotic applications, the online nature of active learning tasks often prevent them from being used in these settings.

A nonparametric alternative to learning dynamics is the use of kernel-based methods [121]. Kernel regression methods frame learning and estimation problems as one of learning functions embedded in high-dimensional – or even infinite-dimensional – spaces defined over the data domain. The properties of the function space are determined by the choice of kernel, which acts as a generalized inner-product that induces a notion of distance between data samples in the function space. These types of methods have been successfully deployed in robotic systems for both dynamics and inverse dynamics learning [122–124]. However, as typically formulated, kernel methods do not have an easy way to model measurement uncertainty and noise in their function spaces. To this end, Bayesian formulations of kernel methods have been developed [125], the most common of which are Gaussian processes. Gaussian processes (GPs) are one of the primary objects of interest in the study of stochastic processes [126]. In GPs, any collection of random variables drawn from the process must be jointly Gaussian. Alternatively, one can insist that functions of the random variables be jointly Gaussian instead, which forms the basis for their application in machine learning [127]. In this context, kernels naturally arise in the specification of the mean and covariance statistics of the Gaussian process in function spaces. Using GPs, researchers have been able to parsimoniously incorporate uncertainty and noise into learning robot dynamics [128]. However, GPs, kernel methods, and nonparametric

learning tools at-large typically have difficulty adapting to online learning settings such as robotic active learning. The primary underlying reason is the fact that nonparametric methods tend to grow in complexity as a function of data. Hence, as a robotic agent acquires more data it becomes more computationally expensive to make predictions with the model.

A promising compromise between the representational capacity of neural networks and the simplicity of kernel methods can be found in techniques like the Koopman operator [129]. The Koopman operator was first introduced in the study of Hamiltonian dynamics and operator theory [130]. Formally, it is an infinite-dimensional, but *linear*, operator that describes the evolution of measure-preserving dynamical systems in a lifted function space. However, to apply Koopman operators numerically they must be approximated in finite dimensions using schemes like Dynamic Mode Decomposition (DMD) [131,132]. Algorithms like DMD use a finite basis for the function space that the Koopman operator acts on to describe the underlying dynamics [133]. Koopman operator theory and its resulting algorithms have been to a large degree developed in the context of dynamics and control, making it an ideal candidate for active learning of dynamics in robotics [134–137]. The linearity of the operator lends itself to the use of canonical control techniques such as linear–quadratic regulators, allowing for computationally-efficient nonlinear optimal control [138]. An important feature of this approach is that it does not scale in complexity with data and allows for adaptable incremental learning. The primary caveat with employing these methods is the difficulty of choosing good basis functions with which to describe the dynamics.

As an illustrative example of learning dynamics in a context that demands rapid adaptation, we compare passive and active learning in the stabilization of a malfunctioning quadrotor vehicle [115]. In this simulation, we equip two quadrotors with a data-driven model of their nominal dynamics that they can use for model-predictive control. However, at the start of the simulation we disable one of the rotors on each robot causing them to free-fall. To recover, each robot must update their internal dynamics model and stabilize themselves using control. Both agents have a single second during which they can collect data to adapt their dynamics models, after which they switch to a stabilizing controller that tries to regain control of the free-fall. Crucially, one agent learns passively and another actively by optimizing the Fisher information with respect to the unknown Koopman operator, which we discuss in the next section. Fig. 3(a) shows snapshots of the different agent trajectories, indicating that the active learning agent is able to stabilize itself much more rapidly than its passive counterpart (see Fig. 3(b) as well), potentially avoiding a crash. The active trajectory greatly exceeds the information gain of the passive approach (Fig. 3(c)) while also achieving lower stabilization error (Fig. 3(d)). Hence, by using control and movement to optimize information measures, robots can learn dynamics faster and more reliably.

3.3. Features

The final and most broad category of learning goals we discuss is that of feature learning. Consider being blindfolded and handed a baseball and a tennis ball in either hand at random. Most people would likely be able to tell them apart with ease. But what is it about either ball that differentiates one from the other? What kinds of properties best represent each ball and its characteristics? Despite having tens of thousands of nerve endings embedded in the palm of our hand, we only need to track a few properties to be able to distinguish between the balls, such as texture and weight. We refer to the general problem of finding informative representations of high-dimensional data that can aid in a task as feature learning.

In the pattern recognition and machine learning literature, features are any measurable characteristics of a phenomenon being observed [139]. Traditional feature learning is the use of machine learning techniques to represent the “intrinsic” structure of data from raw

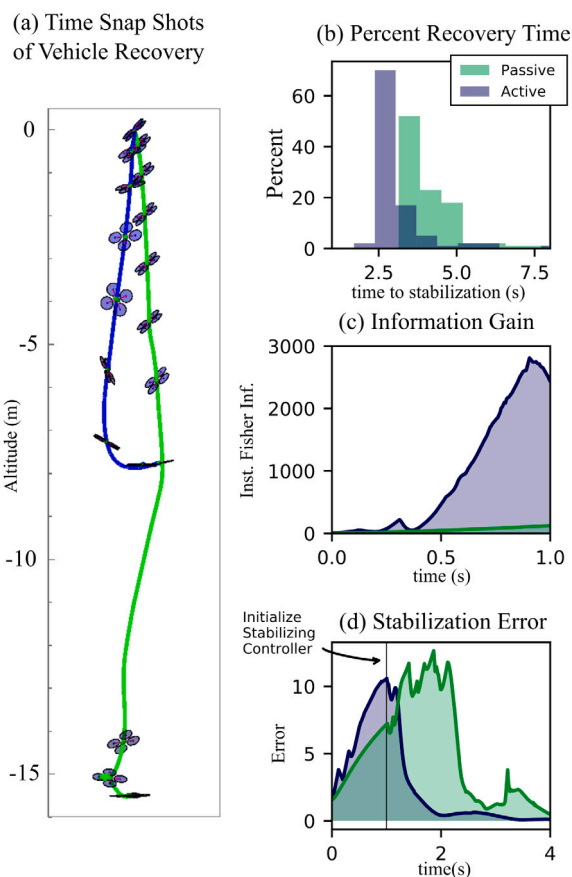


Fig. 3. Online quadrotor recovery: Rotor vehicle recovery using active learning in a real-time single-shot learning context [115]. The rotor vehicle with the blue trajectory uses an actively learned Koopman operator representation of its dynamics. The green trajectory is the result of a passively learned Koopman operator representation of the dynamics. The rotor vehicle with the passively learned representation drops further in altitude – potentially crashing – before recovering after one rotor is disabled. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and possibly highly-redundant measurements [140]. Recent work in this domain has focused on the use of deep learning towards finding succinct representations of human movement [141] and speech [142]. In robotics, tasks are not always well-specified and disentangling the relationship between a robot’s internal state and the intended goal may be difficult. This is primarily a challenge in deep reinforcement learning where problems can become intractable when a naive state representation is used. To this end, feature learning can be leveraged towards making deep RL methods computationally tractable, and to develop schemes that better generalize to the variety of sensory inputs to which an RL agent may be exposed [143].

A simple example of feature learning can be seen in the Koopman literature. As we previously mentioned, finding the correct choice of basis functions for arbitrary dynamical systems can be very difficult. Nonetheless, recent work has been able to construct basis functions that best describe dynamics – also known as the Koopman operator eigenfunctions, or the intrinsic coordinates of the system – using deep learning [114,135,144]. In general, feature learning of this sort will be particularly important for robots with high-dimensional sensing modalities such as e-skins [145], or computer vision [146], and active learning can aid in enhancing rapid identification of intrinsic coordinates.

Our discussion in this review focuses on measures in Section 4 and synthesis tools in Section 5 for active learning using location and other low dimensional learning goals as examples. But the learning goal

can be very high dimensional, as in the case learning dynamics of a vehicle, or in the case of learning representations (e.g., machine vision applications). Regardless of whether a learning goal is low dimensional or high dimensional, the robot still has the same control authority to affect learning—it can move its body and take other physical actions to evoke response and facilitate model updates.

4. Measures for learning

Active learning is rooted in the extraction of information from sensors [1,94,147–152]. Accordingly, measures of information should be expected to play a significant role. The aspects of the objective that can be captured by different information measures as well as how this information can be quantified is key in both control analysis and optimal control synthesis. The approach we discuss here follows this perspective, looking for measures appropriate both for information needs and suitable for numerical synthesis. In this section, we cover three important measures relevant to active learning—entropy, Fisher information, and ergodicity.

4.1. Entropy

Entropy-based measures have been employed in a wide range of action sensing results to calculate the expected information gain for each potential action before collecting measurements [86,93,101,149,153–166,166,167,167–170]. This modern concept of entropy was developed by Claude Shannon for use in the communication and transmission of information [171]. Shannon was concerned with the amount of information necessary to reproduce the content of an information source. To this end, entropy is the expected amount of information or “uncertainty” contained in a random variable. In the case of a discrete random variable X where each x_i is a different outcome of the variable, the amount of information content in a particular event is defined by $I(x_i) = -\log p(x_i)$, referred to as bits when in base 2. The entropy of X is the expected value of the information content of each of the possible events.

$$H(X) = -\sum_i p(x_i) \log p(x_i) \quad (1)$$

The information content of a particular event decreases as the probability of that outcome increases, so low probability events provide more information than high probability events. As entropy is the average value of the information content of a random variable, the maximum value of $H(X)$ for X , would occur when each outcome of the random variable is equiprobable, i.e., when there is maximum uncertainty about a particular outcome. Thus, any particular outcome for a uniformly distributed random variable does not provide much information. In the context of robotics, this is an explanation for why *rare* or *sparse* events are particularly valuable to a robot’s estimation process.

By calculating the Expected Entropy Reduction (EER) of each candidate action, measures of entropy can be readily applied in the context of active sensing. However, exhaustively searching for an optimally informative solution over sensor state space and belief state is a computationally prohibitive process, as it is necessary to calculate an expectation over both the belief and the set of candidate control actions [85,86,101,158,161,172]. Alternatively, the expected information gain can be locally optimized by selecting a control action based on a local estimate of the expected information [88,90,92,94,95,156,162,166,173]. Often times, these methods do not or cannot incorporate general sensor dynamics [88,90,156,166,173] and even the global strategies are likely to suffer when uncertainty is high and information diffuse [91,100,174].

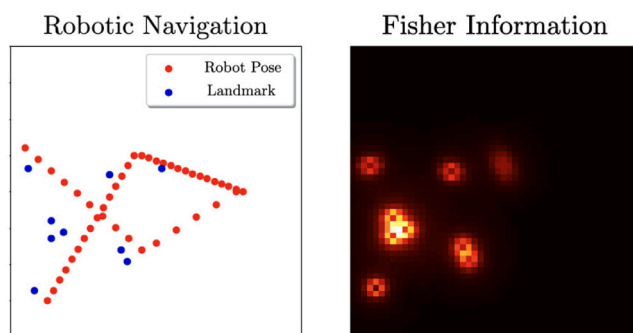


Fig. 4. Fisher information: A measurement model indicates how a sensor will respond to the unknown parameter α , based on the current state. In SLAM applications, the measurement model might be a model of the detection of landmarks in an environment. The Fisher information distribution over landmarks provides a mechanism for determining what states the dynamical system should achieve to provide maximally informative measurements.

4.2. Fisher Information

Active learning relies on collecting informative sensor measurements to support the learning process. In order to do so, there must be a way to locally measure the information contained in sensor readings. Used commonly in maximum likelihood estimation, Fisher information is a method of quantifying the amount of information that a random variable X contains about an estimate of an unknown parameter, or vector of parameters $\alpha \in \mathbb{R}^M$. Using $p(x|\alpha)$, the density function of X parametrized by the value of the vector α , one can determine the likelihood of an observation x given a value of α . The Fisher information is an $M \times M$ matrix that captures the local sensitivity between parameters and observations [170,175]:

$$I(\alpha) = E_X \left[\left(\frac{\partial}{\partial \alpha} \log p(x|\alpha) \right) \left(\frac{\partial}{\partial \alpha} \log p(x|\alpha) \right)^\top \middle| \alpha \right], \quad (2)$$

where the expectation is taken over realizations of X at a given value of the parameter vector α . If $p(x|\alpha)$ is highly sensitive to changes in α —e.g., the distribution of observations exhibits a steep dependence on α —then for a given measurement there will be a small range of highly probable values of α . If $p(x|\alpha)$ is not sensitive to changes in α , then there will be many candidates of comparable likelihood.

In robotics, Fisher information is well suited for measurement models that are naturally parametric (e.g., size, weight, location). Measurement models, sometimes called observation models, are predictions of how unknown variables will impact a sensor reading. This sensor reading can be very sophisticated, like a camera being used in a pixels-to-torque application [176], or very simple, such as a one-bit sensor being used for trajectory tracking [177,178]. The measurement model provides a way of expressing *what* the robot is attempting to learn in terms of its sensing capability and means to adjust its sensors. A commonly used measurement model form is $z = Y(\alpha, x) + \Delta$, where z is the measurement, α is the parameter being estimated, x is the state of the agent, and Δ is (possibly multi-dimensional) zero-mean Gaussian noise. This model is in the form of a sum of a deterministic term—typically modeled by first-principle physics—and a noise term which can be rather challenging to justify, since most robotic applications will not have such convenient additive normal distributions.

For active learning applications, measurement models can play an important role in calculating information measures over a space. To estimate a parameter vector α , the Fisher information matrix has each element (i, j) given by:

$$I_{i,j}(x, \alpha) = \frac{\partial Y(\alpha, x)}{\partial \alpha_i} \Sigma^{-1} \frac{\partial Y(\alpha, x)}{\partial \alpha_j}, \quad (3)$$

where the multi-dimensional noise is assumed to be zero-mean Gaussian with covariance Σ . Intuitively, Fisher information can be expected




Fig. 5. Ergodicity: For an agent to be ergodic with respect to a target distribution, the spatial statistics of the agent’s trajectory must match the statistics of the target distribution. This means that the time spent in a particular area is proportional to the density of the target distribution in that area. Here an agent traverses a bimodal distribution. The size of each waypoint is proportional to the time spent at that location.

to be higher where the expected measurement signal is greater than that of the noise. The expected information density $EID(x)$ over a search space can be constructed by computing the *expected* Fisher information with respect to a probability distribution representing an estimate of a parameter $p(\alpha)$. This $EID(x)$ would then form the information landscape against which active learning decisions are made and then executed.

As an example, we consider the use of the Fisher information in Simultaneous Localization and Mapping (SLAM) problems subject to measurement models of the form discussed above. While the SLAM literature in robotics is diverse and well-established, the more recent field of *active* SLAM has seen much growth [179]. Active SLAM makes use of representations of uncertainty and information to generate exploration plans. In active SLAM, different information measures can capture different features of an environment. In Fig. 4, measurement models for landmark detection are used to provide a basis for calculating information measures to inform the agent’s exploration plan. In this case, the Fisher information over each landmark attracts the robot to landmarks with lower uncertainty, thereby enabling efficient loop closure. This allows an agent to discern an ensemble of locations that are expected to provide more informative measurements.

4.3. Ergodicity

Ergodicity is a fundamental property of dynamical systems and stochastic processes. Formally, achieving ergodicity implies that the dynamical system uniformly visits all parts of the space in which it exists [180]. However, more often what we mean when we say that a system is “ergodic” is whether or not it satisfies the pointwise ergodic theorem [181]. In this sense, being ergodic requires that the system spend time in regions of space in proportion to the measure of said regions. The specific measure used can vary with context, but very often probability measures are used.

In engineering contexts such as active learning, we are free to choose or construct the spatial measure. Particularly, when a system is ergodic with respect to measures representing an information distribution over the space, ergodicity demands perfect asymptotic sampling of informative states. As a simple example, consider a system trajectory $x(t) \in \mathcal{X}$ and a probability density function (PDF) capturing the expected distribution of information over the space. If the trajectory is ergodic, then the amount of time the agent spends in each




Fig. 6. Fourier transform of a trajectory: The Fourier transform of a constant speed trajectory represents the trajectory in the form of a spatial distribution. The representation of the trajectory by its transform changes in granularity for $k = 1, 3, 5, 10, 20, 50$ Fourier coefficients.

neighborhood $\mathcal{N} \subset \mathcal{X}$ is going to be proportional to the amount of information in \mathcal{N} as measured by the PDF (see Fig. 5). Hence, designing ergodic dynamics with respect to desired measures is of interest to active learning [182]. However, in order to do so we need a metric that captures how “ergodic” our trajectories are.

Because perfect ergodicity is only possible on infinite time horizons, we require a metric that can be maximized over finite-horizons through decision-making—such a metric was developed in [183]. Metrics on ergodicity provide a principle of motion [13,24] similar to energy minimization and error minimization, and can be used to synthesize automated exploration for learning, as we will see in Section 5. The ergodic metric in [183] provides a method for comparing a trajectory $x(t)$ —a singleton at any given time t —to a distribution $\Phi(x)$ through their *spatial Fourier transforms*. This suggests that one can compare the coefficients c_k of $x(t)$ and ϕ_k of $\Phi(x)$ respectively and measure a distance between the two. In general, it is not obvious how one might do this otherwise since information content between dimensionally different objects is typically not well defined.

Comparing how “close” two quantities are to each other is imperative for control when using optimization-based methods. To compute the Fourier coefficients ϕ_k of a distribution $\Phi(x)$, we use the inner product

$$\phi_k = \int_{\mathcal{X}} \phi(x) F_k(x) dx, \tag{4}$$

where F_k ’s represent the choice of Fourier basis functions. For trajectories, we begin by interpreting them as distributions comprised of sequences of impulses:

$$C(x) = \frac{1}{T} \int_0^T \delta[x - x(t)] dt, \tag{5}$$

where δ is the Dirac delta [183]. Then from the properties of the Dirac delta function, we can calculate the Fourier coefficients

$$c_k = \frac{1}{T} \int_0^T F_k(x(t)) dt, \tag{6}$$

where the coefficients take on the value of the basis functions averaged over a time window of duration T . An example of such a spatial representation is shown in Fig. 6, where a trajectory along with its Fourier decomposition is shown for different numbers of coefficients c_k . As the number of coefficients k increases the spatial resolution of the trajectory improves, showing how the statistics of a trajectory may be represented as a spatial distribution.

With this in mind, the ergodic metric represents a *distance from ergodicity* that is measured from a time-averaged trajectory $x(t)$ with

respect to a distribution $\Phi(x)$. This distance is calculated by imposing a norm on the difference between the trajectory's c_k and the distribution's ϕ_k coefficients. Particularly, we take the Sobolev norm between the coefficients by using the sum of the weighted squared distance between them:

$$\mathcal{E}(x(t)) = \sum_{k_1=0}^K \dots \sum_{k_n=0}^K \Lambda_k |c_k - \phi_k|^2 \quad (7)$$

where K is number of Fourier coefficients used for each of the n dimensions, and k is a multi-index (k_1, \dots, k_n) . The coefficient $\Lambda_k = (1 + \|k\|^2)^{-s}$ is a weight where $s = \frac{n+1}{2}$, which places larger weight on lower frequency information, ensuring convergence [183]. It is worth noting that spectral methods, and the ability to generate a norm on a trajectory $x(t)$ using them, offer opportunities in measuring entropy as well. The entropy of a distribution could also be measured in the Fourier domain—yielding an objective function that is differentiable and amenable to control synthesis, enabling one to avoid approximating entropy in an optimization with Fisher information (e.g., as in [184]).

The measures discussed in this section form the basis for how we measure performance of an active learning system. The next section focuses on synthesizing decisions that optimize, or at least improve, those measures.

5. Control synthesis for active learning

Active learning has a wide range of applications in robotics including prioritized decision making [159,185], inspection [165], mine detection [186], object recognition or classification [155,160,187], next-best-view problems [156,157,188], and environmental modeling [97, 98,189]. As a result, particular controller architectures may be advantageous for different environments, tasks, and constraints. Here, we survey several model-based optimal control methods that provide distinct advantages for active learning.

Model-predictive control (MPC) is an optimal control framework that optimizes current actions with respect to an objective while taking into account the future behavior of the system over a finite time horizon. Once the current action is taken, MPCs reoptimize from the new starting point and continually plan actions throughout the receding horizon. MPCs are particularly suited to active learning because receding horizon planning lends itself to continuous incremental learning, while simultaneously enabling assessments of the safety and stability of trajectories. In contrast, other optimal control approaches such as the linear–quadratic regulator (LQR) must solve the entire control problem without replanning.

One of the primary optimal control algorithms is Differential Dynamic Programming (DDP) [190], which is an extension of the seminal work by Bellman [191]. DDP is a model-predictive method requiring second derivatives of the dynamics to realize quadratic convergence to the optimal solution. While DDP has fast convergence guarantees, calculating the Hessian of the dynamics can be computationally intractable. If one is willing to forego the fast convergence rate by disregarding the second order terms of the control solution, DDP becomes equivalent to the first order iterative LQR (iLQR) method. DDP and iLQR have both been shown to be effective in the context of robot control in a variety of applications [192]. For example, in [193] the authors use local trajectory optimization methods in combination with RL to learn policies for dexterous manipulation with a five-fingered robotic hand. In scenarios where the dynamics are known or easily modeled, and their Jacobians and Hessians are inexpensive to compute, DDP and iLQR may be well-suited to active learning applications.

A method that generalizes MPC to both convex and nonconvex objectives is the sampling-based Model Predictive Path Integral (MPPI) control algorithm [194]. In MPPI, Monte Carlo sampled trajectories are used to approximately extremize a free energy objective [195]. These types of objectives are designed in analogy to thermodynamic

free energy from the statistical mechanics literature and can be used to synthesize control [196]. Moreover, the synthesized control actions are formally equivalent to Bellman optimal control without the need for computing derivatives, and their computation can be easily parallelized [197]. As a result, MPPI is particularly well-suited for use in learning problems where the dynamics of the agent are non-differentiable or too complex to differentiate in a computationally-efficient way as with neural network models. For example, in [194] the authors use MPPI to learn a neural network model of the dynamics of an auto-rally autonomous race car. However, depending on the structure of the task, generating enough simulated trajectories to sufficiently sample a learning objective may become prohibitive.

Another model-based control synthesis method is Sequential Action Control (SAC), which is inspired by hybrid systems theory [198]. Unlike other MPC techniques, SAC explicitly tries to expend the least control effort possible in generating actions by taking into account the benefits of taking optimal actions as opposed to alternative policies or doing nothing. SAC simultaneously finds the actions that optimize an objective, the best time to apply said actions, and the application duration. Due to its hybrid specification, SAC can naturally handle non-smooth dynamics, and can also be easily wrapped around other controllers to enable more exotic control architectures [199]. In [200], SAC was used for active parameter estimation with a robotic system. This work uses SAC to control a robot to determine the length of a pendulum by maximizing the Fisher information with respect to the pendulum parameters. The SAC control actions sequence is piecewise continuous, with generally short application durations for any control. This allows a robot to reactively generate motions towards information dense regions. However, like most MPC techniques, it requires having access to the derivatives of the objective and dynamics, which can constrain its usage in learning scenarios as previously discussed.

An important consideration when choosing a controller for active learning is the global characteristics of the search process. Depending on the structure of the learning task, there may be a single optimum that represents the true parameter value that is being estimated. Other learning tasks require that the agent avoid fixating on a single information source and instead visit many sources. We distinguish between these approaches by referring to them as myopic and non-myopic respectively. Myopic learning uses local algorithms that greedily take actions over short horizons that optimize the immediate learning objective. While these methods are prone to getting trapped in local minima, they have lower computational overhead than non-myopic learning methods. Non-myopic approaches plan control actions over long time horizons so as to produce coverage over distinct information sources. These are often used to avoid local minima associated with fixation [98,100,101], and can take advantage of approximate solutions [86,87,91,98–100,164,186,201,202].

Choosing a mechanism by which one can avoid myopic learning is critical to operating in environments that have unmodeled effects, such as visual occlusion, where the *expected* most informative state may not provide information. For example, a camera taking a picture of a person behind and a piece of furniture does not benefit from multiple pictures taken from the same state. As a result, dynamic coverage of high information density areas can keep a robot collecting good data while acknowledging unmodeled uncertainty effects through decision-making. Taking these factors into account can be critical to the success of the active learning process. Next, we will examine two approaches to active learning and exploration – infotaxis and ergodic control – that take opposing attitudes towards this question.

5.1. Infotaxis

Inspired by animals' search for chemical sources in a fluid such as air or water, infotaxis is an information-maximizing search strategy using entropy reduction as an information criteria [153]. This technique was developed to show that a search plan does not need to depend on




Fig. 7. Infotaxis: *Upper panel:* A cat searches for the two yellow mice. The distribution around the mice represents the probability of detecting the mouse at that location where white is high probability. The gradient in the cone around the cat represents the cat's measurement model. The cat has a mental image of the expected local reduction in entropy for moving up, down, left or right. *Lower panel:* This is an example of an infotactic trajectory of an agent searching for target locations, represented by the yellow stars. The likelihood of detecting the target is represented by the distribution around the target locations, becoming more dense closer to the target. The measurement model encodes the ability of a camera to detect an object at a particular range and angle of attack. The search strategy selects the direction of movement that maximizes the expected entropy reduction at each time step. The infotaxis strategy succeeds at finding only one of the two targets in the environment and stops searching. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

environmental gradients, such as the concentration of a scent smoothly increasing in proximity to a flower. Instead, animals may sense traces of a source dispersed by wind or currents and formulate a movement strategy based on infrequent detections.

In this work, an agent attempts to localize a target or source in a 2D environment based on detections of the target. To generate an infotactic trajectory, the searching agent chooses a control action at each time step that locally maximizes the expected reduction in entropy, thereby maximizing expected information gain. Concretely, the agent considers moving to adjacent positions on a lattice, or staying in the same location to take more measurements.

To determine an action, it is necessary to have a probability distribution $p(r)$ representing the unknown location r of the source. The probability of detecting the source at a given location is dependent on the distance from the source, meaning that the record of detections along the trajectory of the searcher, $x(t)$, carries information about the source location. When a detection event occurs, the times and coordinates are stored in the random variable \mathcal{T}_t . From this record of detections, the searcher is able to represent the location of the source as a posterior probability distribution that is updated based on the measurement taken at each time step.

$$p_t(r_0) = \frac{\mathcal{L}_{r_0}(\mathcal{T}_t)}{\int \mathcal{L}_x(\mathcal{T}_t) dx} \quad (8)$$

Here, \mathcal{L}_{r_0} is the likelihood of detections \mathcal{T}_t for a source located at r_0 . From the posterior distribution one can calculate a control action that minimizes the expected entropy at the next time step by selecting a set of potential actions, computing the EER given the current $p(r)$, and then selecting the action that provides the minimal EER. This strategy can be computationally prohibitive for many systems.

The trajectories produced by infotaxis exhibit similarities to biological organisms such as moths or bacteria that engage in olfactory search [203]. However, infotaxis-type approaches can fail when there are distractors – states that appear similar to the target but are not the

target – in the environment [13]. The searcher may conflate the actual target with the distractor and then ignore the intended target. Practically, infotaxis can only be implemented using short time horizons as the computational requirements of predicting for longer horizons are significant. For each control action considered, the expected entropy reduction must be calculated, including calculating a posterior for each possible outcome of the measurement random variable. Fig. 7 provides an example of an infotactic search with two target locations. Here, the agent successfully determines the location of one source and stops searching. This strategy is purposefully ignorant of a signature that may conflict with the perceived location of the target in favor of detecting the same target to increase its certainty. This example illustrates that the infotactic strategy is myopic when confronted with multiple sources or environments with convincing distractors.

While an infotactic search strategy can experience difficulties when there are multiple targets in the environment that require persistent monitoring, it is well suited to react to sporadic cues and requires only local information. Infotaxis represents one of the most straightforward examples of active learning in which an agent acts greedily to maximize expected entropy reduction.

5.2. Ergodic control

Recent work by the authors and colleagues has analyzed biological motion by introducing energy constrained proportional betting [13, 204], where the energetic cost of movement is balanced against the desire to gain sensory information about a source. This approach uses the ergodic metric, discussed in Section 4.3, to quantify how well a trajectory covers a distribution of expected entropy reduction. The resulting algorithm produces trajectories that balance informative sampling – collecting many samples in high information areas – with the amount of energy expended from motion. These types of trajectories were observed in the behavior of electric fish, moles, and cockroaches. This suggests that the strategy of energy constrained proportional betting provides a competitive hypothesis for the ways in which living creatures collect information about their surroundings, and may be a robust approach for robotic systems to acquire information. Extensions and variations of this idea now arise in many robotic applications [205–213].

If the goal of infotaxis is to maximize the information content of a series of measurements collected along a trajectory, the goal of ergodic control – first developed in [183] – is to control the spatial statistics of a trajectory $x(t)$ to match those of an expected information density distribution $EID(x)$. This requires the choice of a norm on the difference between the distributions $EID(x)$ and the trajectory $x(t)$ interpreted as a distribution $C(x)$, defined in Eq. (5). To this end, we use the ergodic metric from Section 4.3 as an objective to synthesize maximally ergodic trajectories for general nonlinear systems using tools from model-predictive control [204]. However, we note that any trajectory optimization tools or direct optimization tools could be used; we use the results from [204] primarily because they are amenable to real-time computation [214].

The first thing to note is that the ergodic metric \mathcal{E} in Eq. (7) is not of the form of a running cost—as a result it is not a Bolza problem (although one can turn it into a Bolza problem by appending the Fourier coefficients to the state vector [215], creating an infinite dimensional state space). Nevertheless, one can calculate the adjoint variable ρ of the ergodic metric function:

$$\dot{\rho} = -\frac{2}{T} \sum_k A_k (c_k - \phi_k) \frac{\partial F_k}{\partial x} - \frac{\partial f}{\partial x}^\top \rho \quad (9)$$

where the dynamics are represented by $\dot{x} = f(x, u)$, and get a descent direction for locally minimizing the ergodic metric [216]. Other approaches can be used that lead to slightly different solutions (e.g., the projection-based trajectory optimization method for ergodic control in [204,217,218], where higher-order convergence properties come at




Fig. 8. Ergodic control with respect to the expected entropy reduction over the search space: *Upper panel:* A cat searches for two targets represented by the yellow mice. Here, the cat has access to its past trajectory and has a mental image of the expected entropy reduction over the whole search space. *Middle panel:* An agent is searching for two targets located at the yellow stars. The information distribution becomes more dense closer to the targets. Here the agent takes a measurement and updates its belief using the same measurement model and likelihood function as in the infotaxis implementation. *Lower panel:* The agent chooses its next control action based on the global expected entropy reduction. This is determined from its belief of the information content in a particular location. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the expense of high computational cost). A key property of the metric \mathcal{E} is that it is differentiable with respect to $x(t)$, so most optimal control techniques can be easily applied.

An example of an ergodic trajectory can be seen in Fig. 8, where the agent is exploring with respect to the expected entropy reduction distribution over the whole environment. The agent is able to successfully locate both target locations in this scenario because the ergodic control strategy is amenable to persistent monitoring of multiple targets. As perfect ergodicity can only be realized as time goes to infinity, the agent will continue to explore the space. Using infotaxis, the agent would conclude its exploration once a target has been detected. Here, we make use of global information to plan control actions over longer time horizons.

With both these local and global information-based synthesis techniques in mind, we next move on to applications in robotics that will depend on active learning strategies.

6. Applications in robotics

While the landscape of applications for active learning is almost as broad as that of machine learning itself, here we will focus on settings where datasets are rarely available ahead of time. Active exploration applications such as search and rescue or mapping are particularly relevant in this class of problems, especially when the environments are dynamic and hard to predict. We also discuss applications in which system models are either unknown or difficult to parametrize, as is the case for soft robotics and for many of the areas of application of imitation learning.

6.1. Soft robotics

Soft robots are made from compliant materials, enabling them to be well suited for delicate tasks and environmental adaptation [219–221]. Unfortunately, precise modeling and control of soft robots poses

challenges because soft materials are continuously deformable and thus nominally have infinite degrees of freedom. There is no clear method of representing the geometry of such a robot without making significant simplifications [20]. The most important functional property of a soft robotic system – deformation in response to the environment – makes soft robotic systems practically impossible to meaningfully model for control based on first-principles (*e.g.*, partial differential equations based on elastic body mechanics). Data-driven modeling is a natural alternative when first principle arguments are either not tractable or do not involve the use of a state space.

Learned representations, such as those constructed by DNNs, have been shown to find input to output mappings that predict the behavior of soft robots [222]. However, these models are difficult to apply using known model-based control techniques. Alternatively, the Koopman operator has also been used for modeling and control of soft robots [137]. Described earlier in Section 3, Koopman operators provide a linear representation for nonlinear dynamical systems that is compatible with linear control methods such as LQR synthesis. In practice, a data-driven approximation is adopted. As an example, [137] develops a model predictive controller with a Koopman operator representation of a soft robotic arm for tracing reference trajectories. The data collection strategy for soft systems plays an important role in determining a model. For instance, though obvious, data collected while an end-effector is out of contact with the environment cannot provide useful modeling data. In prior work we showed that a Koopman operator representation of a robotic system can be actively learned using information-theoretic strategies [114].

Despite the complexity that soft elastic structures introduce to the analysis of robotic motion, soft robots can beneficially exploit these physical properties. For example, soft structures can be leveraged as a computational resource, sometimes called morphological computation or embodied intelligence [223]. A soft body that deforms around an object, in principle, will make manipulation easier, and will imply that the

amount of explicit computation needed will be lower in exchange for the implicit computation afforded by the soft body. For instance, [224] shows that stable hopping behavior of a soft underwater robot can be achieved experimentally by dynamically changing the size of its body. Moreover, with actuator saturation, adapting the morphology of the robot's body was the only route to achieve stable behavior, implying that control over the continuous shape properties of the robot was key to task success.

In addition to articulation, sensory acquisition via morphological computation is connected to biological systems and present in structures such as the cochlea of the human ear [225] and the bodies of octopuses [226]. While data can be passively collected through the physical structure, active sensing is a biologically motivated extension. In [120], the authors build a perception system to learn the kinematics of a soft actuator and estimate interaction forces with embedded sensors and recurrent neural networks. In their approach, the authors consider the relationship between action and perception in the learning process by quantifying sensor information as a result of commanded actuation information. Work in [227] uses a soft robotic probe to palpate imitation tissue to determine the location of a hard tumor-like nodule. The soft robot was able to adjust its stiffness across iterations of the palpation task based on information metrics calculated from human test subjects. These findings suggest that active haptic perception through physical changes to the probe improves estimation accuracy, motivating active learning techniques that could automate learning for this and other soft systems.

6.2. Search and rescue

Prevention, response, and recovery from disasters can be dangerous for emergency professionals who may need to interact with areas affected by events such as hurricanes, oil spills, and earthquakes. Disaster robotics is an area that works to augment the capabilities of workers by delivering real-time data to experts and intervening in the environment [228]. The need to efficiently search an environment is an issue at the core of disaster robotics. One of the most visible examples of the need to search an extremely large, dynamic environment in recent years is the investigation of the crash site of Malaysia Airlines Flight 370 (MH370) in March of 2014. In the first 52 days after the crash, the Australian government reported that air crafts and surface vessels covered an area of over 1.6 million square miles. By June of 2018, the final search effort was suspended without success. Although there may be many points of failure in this search effort, one dimension involved robotic technologies that scanned the bottom of the ocean that were incapable of reasoning about potential debris signatures, the dynamic environment, and their own capabilities.

When searching large areas where information is sparse, active coverage algorithms are important in determining important areas of a search region and the schedule to visit these regions. Coverage algorithms are used in many robotic applications such as underwater exploration [229], agriculture [230], and inspection [231]. The goal of coverage algorithms is to visit all points in an area or volume while avoiding obstacles [232]. Commonly used approaches for coverage, a taxonomy of which is included in [233], include cellular decomposition or grid-based methods to divide the area into manageable sections [234–237]. However, as the complexity of the environment increases, the number of cells necessary to represent the environment increases. These methods typically do not take into account the physical properties of sensing capabilities of the robots or the dynamics of the environment. As a result, coverage is treated as both necessary and sufficient for capturing needed data. This attitude about coverage can be seen in the search strategy of the MH370 investigation which focused on area coverage, neglecting factors such as how the ocean currents might pull debris away from the site [238].

6.3. Localization and mapping

SLAM algorithms create a map of an unknown environment while simultaneously estimating the state of the robot within that environment. This is a major success story in robotics, with the current flood of driverless car technologies all dependent upon SLAM algorithms. When navigating an unknown environment, a robot may lose its ability to localize itself due to accumulated small errors in sensors and actuators, known as representation drift. To correct for this drift, SLAM algorithms use loop closure – the task of identifying whether an agent has returned to a previously visited location – to maintain an accurate representation of the location of the robot relative to environmental features. To maintain loop closure, the robot revisits regions with low estimation uncertainty or informative features to combat representation drift. Beyond localization, loop closure allows the robot to represent the topology of the environment, instead of simply a record of where it has been.

In passive approaches, a robot performs SLAM with sensor information provided to it. For instance a lidar sensor collects data while driving down a road. In contrast, active SLAM leverages the actions of the robot to seek out informative measurements that efficiently decrease localization and mapping uncertainty. Fig. 9 illustrates the flow of information in passive versus active SLAM. Active SLAM generates controls based on the current state of both the map estimate and robot states. The review paper [179] summarizes methods that have been employed in the development of active SLAM including the theory of optimal experimental design [239], information theoretic approaches [240–242] and control theoretic approaches [243,244]. Active SLAM can also be formulated as a Partially Observable MDP (POMDP) and approximated using Bayesian optimization or Gaussian belief propagation to attain computational tractability. Belief space planning entails planning in the space of probabilistic estimates of a robot's state and additional variables of interest [245,246]. This method has also been used in combination with navigation error [247–249].

Using planning algorithms in SLAM is challenging because SLAM is generally executed on a pre-planned trajectory. This trajectory can greatly affect the quality of performance. Conversely, path planning algorithms typically assume a given map. Hence, planning and SLAM are nontrivially interdependent. Work in [250] attempts to integrate SLAM with a coverage path planning problem by developing a movement strategy they call perception-driven navigation. The authors use a cost function that weights navigation uncertainty, evaluated using the Fisher information matrix described in Section 4, with the ratio of unexplored to total coverage area. This method plans paths between waypoints that are selected based on a measure of visual saliency, prioritizing areas in which notable environmental features have been detected. The integration of perception based navigation in the SLAM framework is key to balancing effective mapping alongside exploration as the distribution of features in an environment is often highly uneven. It also allows for operating in limited field of view environments, such as underwater inspection tasks.

Developing a method to determine informative features from images is an important aspect of visual SLAM, in which SLAM is performed using only camera inputs [251]. Image pre-processing with feature selection reduces the computational burden of scanning all the pixels in images, leading to many active feature selection algorithms [252]. One can also selectively process informative regions of images or videos using a recurrent neural network [253]. Lastly, visual-inertial navigation – where a robot must estimate its state using only a camera and inertial sensors – can supplement the visual SLAM process. In [254] visual-inertial navigation selects features based on the state of the observer and the context of the scene, using information theoretic constructions as a basis for prioritizing features to be used in state estimation.

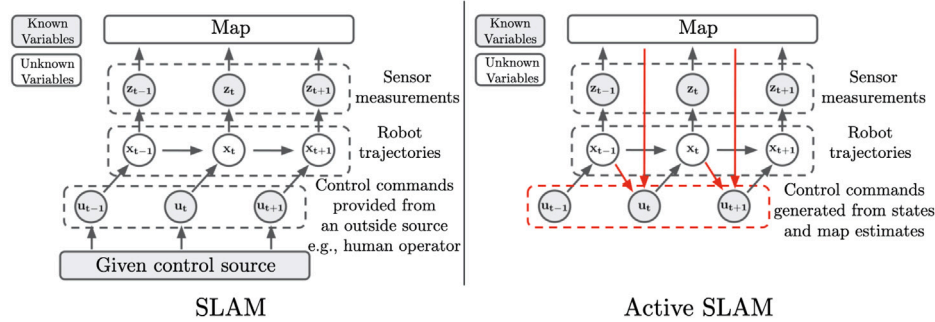


Fig. 9. **SLAM versus Active SLAM:** Active SLAM uses control commands generated to decrease localization and mapping uncertainty. In traditional SLAM, the control signal is given in the problem statement.

6.4. Imitation learning

Imitation learning is a widely used and effective method of imparting human skills to robots by learning desired behaviors from demonstrations. To transfer knowledge about a task through imitation, it is important to capture salient features of a demonstration in efficient and generalizable representations of a skill. Here, active learning can play an important role in capturing knowledge from a demonstration.

The field of imitation learning is expansive [255–257] and has been used in numerous settings including autonomous driving [258], virtual games [259], and replicating human motion in robots [260]. Capturing knowledge about a task from human experts is especially applicable to robotics, where autonomous systems are charged with operating in complex and unstructured environments. In these situations it can be difficult to manually program specific behaviors and engineer reward functions to suit a task. Imitation learning is commonly tied to deep neural networks to take state/action pairs from demonstrations and learn a policy for a skill. This can often require large amounts of data, leading to questions about what aspects of demonstrations are particularly useful to impart a skill to an autonomous system.

When transferring skills from a human operator to a robot, active learning occurs when a human operator is queried for information. For instance, work in [261] considers two approaches to active learning from demonstration in the context of autonomous navigation. A learner, such as a robot, selects expert demonstrations that they believe to be informative based on either novelty or uncertainty reduction criteria. In novelty management, demonstrations are selected based on a density model from which a test feature vector can be compared to demonstrations previously seen in training to provide exposure to unobserved or anomalous data. For uncertainty reduction based active learning, the authors used the Query Bagging Method [262], in which training data is partitioned into multiple subsets. A demonstration would be deemed to have high uncertainty if the variance over these subsets for the demonstration was high.

Inverse reinforcement learning (IRL), also called inverse optimal control, is a method of determining the goals of desired behavior from trajectories executing a policy [263]. The aim of IRL is to find a reward function that describes the desired task from expert demonstrations. When a task is well suited to be described by a single reward function, IRL is most applicable. However, a policy may be optimal for multiple reward functions, making it difficult to discern intent. In response, it may be necessary to include other objectives. Work in [264,265] focuses on active learning in the context of IRL, which seeks to reduce the demonstrations from full trajectories to particularly useful states. In this case, active learning means selecting particularly informative samples to be labeled by an oracle. In [264], a robot learns a reward function and movement policy for a grasping task. The reward function is in the form of a Gaussian process model and is based on human evaluations of the quality of the grasp. In this method, the learning agent is able to impact the demonstrations it sees by choosing to query

human expert ratings based on acquisition functions from the Bayesian optimization literature.

Generative adversarial imitation learning (GAIL) is a model-free imitation learning approach that scales well to high dimensional environments [266]. Inspired by generative adversarial networks, GAIL produces behaviors similar to demonstrated behaviors while training a discriminator to differentiate expert attempts with generated attempts. An extension of GAIL, called InfoGAIL, attempts to find latent structure across human demonstrations – that can be highly variable – to describe interpretable concepts [267]. Related to techniques that train a discriminator to differentiate between expert and learned policies (such as InfoGAN [268]), InfoGAIL approximately maximizes mutual information between latent space and trajectories to deduce meaningful latent variables. In this way, it is possible to produce semantically meaningful or informative data that pertains to a particular task.

Imitation learning, and the other applications mentioned above, stand to benefit from robots that physically manipulate when and how they learn, rather than relying on visual and aural requests for more or better data, which is one of the principal goals of active learning in robotics.

7. Open challenges

Closed-loop active learning presents a key opportunity for improving the quality and rate of learning. In this section, we focus on specific challenges in both the near and far term, such as safety and distributability. These challenges are specific to the expertise of the controls community—e.g., analyzing properties like complexity, convergence, and motion feasibility. We end with a broader discussion of questions such as how can one assess the sufficiency of a learning model for a given task? These challenges, among others that we may not yet understand, are at the core of what it means to construct a robotic theory of active learning.

7.1. Distributability

Distributability has become a widely studied and often implemented goal for control systems, enabling a swarm of robots to accomplish what an individual robot cannot. In the context of control-driven tasks such as exploration or search, the benefits of distributability are immediately apparent—multiple robots will be able to cover an area more efficiently than a single robot could. Distributed data collection of this kind has been widely and successfully applied in a variety of contexts, such as environmental monitoring [236,269]. The key feature underlying the success of these distributed control applications is that the dynamics of the robot collective are factorable into a block-diagonal representation—the dynamics of each robot agent are independent from one another [214,270]. However, can we expect this to be the case across active learning applications?




Fig. 10. Safe decentralized ergodic control: *Upper panel:* Fire trucks attempt to reach the site of a fire guided by helicopters above. The firetrucks are able to explore in the areas that have already been explored by the helicopters. At the same time, the helicopters must maintain the ability to return to a refueling station. *Lower panel:* Here are three time snapshots of an ensemble of six robots – three purple and three blue – explore an environment subject to the condition that blue robots can only go some place purple robots have already visited. The purple robots are tasked with exploring the purple building while the blue robots are tasked with exploring the blue building. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

While independent robots can easily coordinate to collect measurements and effectively augment their perception [271], learning collectively may prove to be much more challenging for a variety of reasons. For one, when robots are not just collecting data but also using it to learn as a group, they must be in constant communication and sharing data samples with one another. Another important challenge is that the data samples that each agent is locally exposed to may be statistically distinct. Moreover, the noise and disturbances that robots are exposed to may be heterogeneous across agents as well. Taken together, these observations suggest that during distributed learning the samples that a swarm collects may not be independent and identically distributed, which is a key assumption underlying most learning methods and can create issues with fundamental properties of the learning process (e.g., convergence). Most of the difficulties outlined so far have been described by the fields of distributed [272] and federated [273] machine learning. Hence, the success of distributed active learning is in part tied to the challenges of distributed learning generally.

Nonetheless, some challenges in distributability will be unique to active learning. As we have discussed, when the dynamics of robotic agents are left uncoupled making control decisions may be simple. However, active learning in robotics precisely requires a coupling between learning and taking actions. Then, when agents share a common distributed learning objective, their dynamics may become effectively coupled through the contingent relationship between acting and learning. As a best-case scenario, this can lead to redundant data collection and learning, but in the worst-case this can create stability issues in the learning process. Highly-coupled dynamics, along with extended network dropouts, will generate high degrees of disagreement between agents, making both analysis and prediction more difficult. Thus, eliciting useful collective behavior from decentralized systems based on local decisions is still an open challenge.

7.2. Safe active learning

Safety is a problem of both specification and prediction—one needs to specify what is meant by safety and be able to predict that the

specification will be satisfied. Imposing safety enables learning in high-consequence environments with continuous deployment, making reliance on models and prior experience less risky.

Common tools available for imposing safety constraints often depend on Control Lyapunov Functions (CLFs) [274,275]. These control approaches enforce stability properties of a robotic system through a feedback stabilizing control law that drives a positive-definite differentiable function to zero over time. In the context of active learning, one may desire to have a CLF for ergodic control, using the ergodic metric as the candidate Lyapunov function [216]. One can use Control Barrier Functions (CBFs) [276–278] that encode safety constraints, such as in Fig. 10 where we impose the constraint that one set of vehicles can only enter a region after another set of vehicles has explored it. Both CLFs and CBFs can be combined with other objective functions that are task-oriented rather than safety oriented; these often then involve solving quadratic programs to satisfy safety constraints [275,277,279,280]. The CLF/CBF approach is the most amenable to computation in high dimensional spaces, but in lower dimensional spaces one can directly solve for safety sets using reachability analysis, which depends on solving a Hamilton–Jacobi–Isaacs partial differential equation [281,282]. Though not necessarily practical for high dimensional systems, this guarantees an optimal trade-off between safety and performance.

An important challenge in these safe learning techniques is that they are model-based. They require a model to evaluate the monotonic decrease of the CLFs/CBFs or to evaluate reachability conditions. Since a robot will typically be learning something about the environment relevant to its evolution, its own dynamics, or its interactions with the environment, all these techniques will rely on model updates of some form along with real-time updates to statistical analysis. A key question is how should a robot stay safe during this process, and what should safety mean when representations critical to safety are not known?

In recent work – following the CLF/CBF viewpoint of safety – we showed that one can use hybrid control methods to schedule switching between a safe controller and a learning controller, while maintaining the asymptotic properties of the safe controller [283]. The critical assumption in that work is that there is an operating point where stability of the robot–environment combination is already established and using the safety of that state as a starting point for safe learning. This is often a reasonable assumption; for instance, one might have an empirically safe PID controller for a humanoid robot near upright posture without having model-based safety analysis. Additionally, CBFs have been used to guide the learning process in reinforcement learning [284]. In this work, the CBFs restrict exploration to safe policies and become less conservative as an online learning process learns a model of the dynamical system. This makes the learning process more efficient while guaranteeing safety. This method incorporates online measurements to improve the CBF-RL controller, providing an opportunity for active learning approaches such as those discussed here to facilitate information gathering. Other approaches to simultaneously satisfying safety guarantees with *a priori* unknown dynamics and/or unknown environments need control formalisms that enforce safety criteria in the absence of any certainty.

7.3. Stability, invariance, and specification

Another concern critical to learning is how to impose prior knowledge on learned models. Particularly in the context of physical learning, where a model does not need to be an ordinary differential equation or a statistical pattern, but can instead be a principle (such as a motion symmetry [285] or energetic dissipation). Among these principled statements of modeling assumptions, stability, the property that the unforced system asymptotically converges to an equilibrium, may be the most common property in a physical system that we may wish to insist upon [286]. In [287]—following [81,288–292]—we used recent results in linear algebra to project linear operators (such as the Koopman representations discussed earlier) onto the closest stable

linear operators. Moreover, in [293] we applied these techniques to robotic manipulation examples, where notably the experiments were implausible without constraining the learning to stable models.

There is a wide range of potential specifications one may wish to impose on a learning system. How would one specify that a learned model must satisfy a linear temporal logic (LTL) constraint such as those described in [294]? What about symmetries in time and space, implying conservation of energy and momentum? Developing formally correct methods for combining learning tools with these specifications is a key step forward towards robot learning under user-generated constraints on what should be learned.

7.4. Actionable learning

A key property of linear control systems is the separation principle. This principle asserts that an optimal estimator can be designed independently from the optimal control. A consequence of the separation principle is that as soon as a measurement has been taken, one knows that the automation system can start to productively take actions. That is, every measurement is *actionable* for the control system. A generalization of the separation principle is to ask whether designing a learning algorithm can be done independently from designing its control system. In general, this review assumes that this is not possible—the learning and control goals are mutually dependent. However, in some learning cases the relationship between what is being learned and when or how soon one can take action may be important. For instance, in the case of shape recognition in Fig. 2, exploring an object to determine its shape properties must happen prior to exploring an unknown environment in search of that shape. This transition is an example of the representation (in this case the abstraction’s “shape”) becoming actionable to the control system. As far as the authors are aware, this topic is little studied in control, but has a long history in psychological study of decision making (e.g., see the many books on this topic by Alain Berthoz [295]).

When a control system becomes actionable is particularly important when distinguishing between active learning and passive learning. During the active learning phase, learning may be the primary goal of the control system. During the passive learning phase the robotic system (or animal) may transition to attempting its ultimate task while continuing to run online passive learning updates. In single-shot learning, where the learner only has one trajectory to exploit for the purpose of learning, being able to robustly detect when learning has become sufficient to take action is a critical part of the path to task success.

Analysis methods are needed for describing conditions under which learned models are *sufficient* for making a decision to combine the estimation aspects of learning with the control aspects of learning. This transition is often characterized in terms of exploration/exploitation trade-offs [296] in the context of sampling-based learning. In the context of a physical system, exploration and exploitation depend on the physics of the learner and environment, and the transition between them will be regulated by the control system. In the case of the example in Fig. 3, this would be a safety-critical decision—devoting inadequate time for active learning yields an insufficient model for recovery prior to the vehicle hitting the ground, while engaging in active learning too long will lead to a catastrophic failure. This particular example would likely yield a convex function that represents safety as a function of transition time. However, how to analyze and compute this transition in general is unknown.

Efficiently forming representations relevant to task completion is part of the challenge in forming actionable representations. When a representation becomes actionable, we capture particular elements of the underlying object or task relevant for decision making while ignoring irrelevant sensory data. The question of determining functionally applicable representations has been explored in [297]. The authors claim that the structure of the environment can be modeled with a known goal-conditioned policy—a policy that can achieve a goal state

from a given state. The authors refine this policy by differentiating states using the actions necessary to reach them. Thus, states that are functionally similar are closer to each other in the representation than they would be when representing their location with an Euclidean distance. This method could benefit from active learning. For instance, one may use the entropy of the representation rather than the entropy of the input or entropy of the physical states, as the information quantity to force active learning capabilities. However one constructs representations from data-driven experience, an important question will be how to synthesize active learning to close the loop on representation generation.

8. Conclusion

Active learning and data-driven control will play a major role in future robotic systems operating without access to reliable analytic models or prior datasets in uncertain environments. Robots will need to become fluent learners—routinely investing time and energy in single-shot learning through purposeful data collection and interpretation. This high level goal transcends the capabilities currently available for robotics in machine learning, both in terms of specifying behavior and representing learning goals. Machine intelligence in general has almost entirely been viewed as an extension of estimation theory, focusing on the processing of data. Even reinforcement learning assumes that the data needed for updating a policy is available or that it can be created in simulation. Here we view learning, in part, as an extension of control theory, focusing on how decisions impact learning outcomes. Before these two views can be synthesized into a single coherent theory, many challenges need to be addressed including those mentioned earlier and many not yet understood.

Expanding our notion of a *model* becomes a key effort moving forward. Models should no longer be solely defined by an ordinary differential equation, though ordinary differential equations may still play critical roles during analysis and computation. Instead, a theme in this review is that model-based reasoning needs to admit any set of meta-principles one asserts, such as symmetries in the system, its stability properties, what equilibria are expected, or its logical structure. These assertions will constrain numerical inference, thereby improving learning by reducing the classes of admissible models.

We have outlined and argued for the development of a theory of robot learning—one that deals with the difficulties and constraints that an embodied learning agent would face in the physical world. While much of machine learning has neglected the challenges that physical embodiment brings, this presents a great opportunity for control theorists at-large. The historical arc of robot control has retained a clear focus on the physical properties that ensure safe, robust, and reliable performance. By merging our understanding of controllability, stability, and compliance, with the flexibility of black-box learning, an action-oriented theory of learning will be key to enable future robot technologies.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank Muchen Sun, Ana Pervan, Kyra Rudy, Frank Park, and the anonymous reviewers of the first draft for their many helpful comments on this manuscript.

References

- [1] Bajcsy R. Active perception. *Proc IEEE* 1988;76(8):996–1005.
- [2] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015;521(7553):436–44.
- [3] Stanton SC. Situated experimental agents for scientific discovery. *Science Robotics* 2018;3(24):4978.
- [4] Martin H. Osmotropotaxis in the honey-bee. *Nature* 1965;208(5005):59–63.
- [5] Basil JA, Hanlon RT, Sheikh SI, Atema J. Three-dimensional odor tracking by nautilus pompilius. *J Exp Biol* 2000;203(9):1409–14.
- [6] Yovel Y, Falk B, Moss CF, Ulanovsky N. Optimal localization by pointing off axis. *Science* 2010;327(5966):701–4.
- [7] Webb B, Harrison RR, Willis MA. Sensorimotor control of navigation in arthropod and artificial systems. *Arthropod Struct Develop* 2004;33(3):301–29.
- [8] Khan AG, Sarangi M, Bhalla US. Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. *Nature Commun* 2012;3(1):1–10.
- [9] Stamper SA, Roth E, Cowan NJ, Fortune ES. Active sensing via movement shapes spatiotemporal patterns of sensory feedback. *J Exp Biol* 2012;215(9):1567–74.
- [10] Catania KC. Stereo and serial sniffing guide navigation to an odour source in a mammal. *Nature Commun* 2013;4(1):1–8.
- [11] Hartmann MJ. Active sensing capabilities of the rat whisker system. *Auton Robots* 2001;11(3):249–54.
- [12] Nelson ME, MacIver MA. Sensory acquisition in active sensing systems. *J Comp Physiol [A]* 2006;192(6):573–86.
- [13] Chen C, Murphey TD, MacIver MA. Tuning movement for sensing in an uncertain world. *ELife* 2020;9:e52371.
- [14] Nakajima K, Hauser H, Li T, Pfeifer R. Information processing via physical soft body. *Sci Rep* 2015;5(1):10487.
- [15] Yin X, Müller R. Integration of deep learning and soft robotics for a biomimetic approach to nonlinear sensing. *Nat Mach Intell* 2021.
- [16] Chen T, Pauly M, Reis PM. A reprogrammable mechanical metamaterial with stable memory. *Nature* 2021;589(7842):386–90.
- [17] Gold JM, England JL. Self-organized novelty detection in driven spin glasses. 2019.
- [18] Zhong W, Gold JM, Marzen S, England JL, Halpern NY. Learning about learning by many-body systems. 2020.
- [19] Pfeifer R, Gómez G. Morphological computation—connecting brain, body, and environment. In: *Creating brain-like intelligence*. Springer; 2009, p. 66–83.
- [20] Rus D, Tolley MT. Design, fabrication and control of soft robots. *Nature* 2015;521(7553):467–75.
- [21] Furutani N, Takahashi T, Naito N, Maruishi T, Yoshimura Y, Hasegawa C, et al. Complexity of body movements during sleep in children with autism spectrum disorder. *Entropy* 2021;23(4).
- [22] Osipov M, Behzadi Y, Kane JM, Petrides G, Clifford GD. Objective identification and analysis of physiological and behavioral signs of schizophrenia. *J Ment Health* 2015;24(5):276–82.
- [23] Berrueta T, Pervan A, Fitzsimons K, Murphey T. Dynamical system segmentation for information measures in motion. *IEEE Robot Autom Lett* 2019;4(1):169–76.
- [24] Fitzsimons K, Acosta AM, Dewald J, Murphey TD. Ergodicity reveals assistance and learning in physical human robot interaction. *Sci Robot* 2019;4(29):6079.
- [25] Viswanathan GM, Buldyrev SV, Havlin S, da Luz MGE, Raposo EP, Stanley HE. Optimizing the success of random searches. *Nature* 1999;401(6756):911–4.
- [26] Bartumeus F, Catalan J. Optimal search behavior and classic foraging theory. *J Phys A* 2009;42(43):434002.
- [27] Baddeley RJ, Franks NR, Hunt ER. Optimal foraging and the information theory of gambling. *J R Soc Interface* 2019;16(157):20190162.
- [28] Bekoff M. Animal play: Problems and perspectives. In: Bateson PPG, Klopfer PH, editors. *Perspectives in ethology: Volume 2*. Springer US; 1976, p. 165–88.
- [29] Reinhold AS, Sanguinetti-Scheck JI, Hartmann K, Brecht M. Behavioral and neural correlates of hide-and-seek in rats. *Science* 2019;365(6458):1180–3.
- [30] Smith PK. Does play matter? Functional and evolutionary aspects of animal and human play. *Behav Brain Sci* 1982;5(1):139–55.
- [31] Settles B. Active learning literature survey. Technical Report 1648, University of Wisconsin–Madison; 2009.
- [32] Gao Y, Hendricks LA, Kuchenbecker KJ, Darrell T. Deep learning for tactile understanding from visual and haptic data. In: *2016 IEEE international conference on robotics and automation (ICRA)*; 2016, p. 536–43.
- [33] Li C, Zhang T, Goldman DI. A terradynamics of legged locomotion on granular media. *Science* 2013;339(6126):1408–12.
- [34] Laschi C, Mazzolai B, Cianchetti M. Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Science Robotics* 2016;1(1).
- [35] Merlet J-P. A historical perspective of robotics. In: Ceccarelli M, editor. *International symposium on history of machines and mechanisms proceedings HMM 2000*. Dordrecht: Springer Netherlands; 2000, p. 379–86.
- [36] Devol GC. Programmable article transfer. Google Patents; 1954, U.S. Patent 2,988,237.
- [37] Walter W. A machine that learns. *Sci Am* 1951;185:60–3.
- [38] Turing AM. On computable numbers, with an application to the entscheidungsproblem. *Proc Lond Math Soc* 1937;52-42(1):230–65.
- [39] Turing AM. Computing machinery and intelligence. *Mind* 1950;59(236):433–60.
- [40] McCarthy J, Hayes PJ. Some philosophical problems from the standpoint of artificial intelligence. In: *Readings in artificial intelligence*. Elsevier; 1969, p. 431–50.
- [41] Gold EM. Language identification in the limit. *Inf Control* 1967;10(5):447–74.
- [42] Angluin D. Inductive inference of formal languages from positive data. *Inf Control* 1980;45(2):117–35.
- [43] Valiant LG. A theory of the learnable. *Commun ACM* 1984;27(11):1134–42.
- [44] Rosenblatt F. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychol Rev* 1958;65(6):386–408.
- [45] Angluin D, Smith CH. Inductive inference: Theory and methods. *ACM Comput Surv* 1983;15(3):237–69.
- [46] Littlestone N. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Mach Learn* 1988;2(4):285–318.
- [47] Kearns M. Efficient noise-tolerant learning from statistical queries. *J ACM* 1998;45(6):983–1006.
- [48] Ben-David S, Itai A, Kushilevitz E. Learning by distances. *Inform and Comput* 1995;117(2):240–50.
- [49] Ben-David S, Hrubeš P, Moran S, Shpilka A, Yehudayoff A. Learnability can be undecidable. *Nat Mach Intell* 2019;1(1):44–8.
- [50] Blumer A, Ehrenfeucht A, Haussler D, Warmuth MK. Occam's razor. *Inform Process Lett* 1987;24(6):377–80.
- [51] Blumer A, Ehrenfeucht A, Haussler D, Warmuth MK. Learnability and the vapnik-chervonenkis dimension. *J ACM* 1989;36(4):929–65.
- [52] Cooper SB. *Computability Theory*. CRC Press; 2003.
- [53] Angluin D. Inference of reversible languages. *J ACM* 1982;29(3):741–65.
- [54] Nocks L. *The robot: The life story of a technology*. Greenwood technographies, Greenwood Press; 2007.
- [55] Angluin D. Queries and concept learning. *Mach Learn* 1988;2(4):319–42.
- [56] Cohn DA, Ghahramani Z, Jordan MI. Active learning with statistical models. *J Artificial Intelligence Res* 1996;4:129–45.
- [57] Balcan M-FF, Feldman V. Statistical active learning algorithms. In: *Advances in neural information processing systems (NeurIPS)*. Vol. 26. 2013.
- [58] Balcan M, Hanneke S, Vaughan JW. The true sample complexity of active learning. *Mach Learn* 2010;80(2):111–39.
- [59] Watson JB. Psychology as the behaviorist views it. *Psychol Rev* 1913;20(2):158–77.
- [60] Skinner BF. *The behavior of organisms: An experimental analysis*. Appleton-Century-Crofts; 1938.
- [61] Barto AG, Sutton RS, Brouwer PS. Associative search network: A reinforcement learning associative memory. *Biol Cybernet* 1981;40(3):201–11.
- [62] Sutton RS, Barto AG. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychol Rev* 1981;88(2):135–70.
- [63] Barto AG, Sutton RS, Anderson CW. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans Syst Man Cybern* 1983;SMC-13(5):834–46.
- [64] Bellman R. Dynamic programming. *Science* 1966;153(3731):34–7.
- [65] Sutton RS, Barto AG, Williams RJ. Reinforcement learning is direct adaptive optimal control. *IEEE Control Syst Mag* 1992;12(2):19–22.
- [66] Thorndike EL. The law of effect. *Am J Psychol* 1927;39(1):212–22.
- [67] Sutton RS, Barto AG. *Reinforcement learning: An introduction*. MIT press; 2018.
- [68] Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, et al. Continuous control with deep reinforcement learning. In: *Proceedings of the international conference on learning representations (ICLR)*; 2016.
- [69] Duan Y, Chen X, Houthoofd R, Schulman J, Abbeel P. Benchmarking deep reinforcement learning for continuous control. In: *Proceedings of the International Conference on Machine Learning (ICML)*. Vol. 48; 2016, p. 1329–38.
- [70] Gu S, Holly E, Lillicrap T, Levine S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE; 2017, p. 3389–96.
- [71] Kaelbling LP. The foundation of efficient robot learning. *Science* 2020;369(6506):915–6.
- [72] Ibarz J, Tan J, Finn C, Kalakrishnan M, Pastor P, Levine S. How to train your robot with deep reinforcement learning: lessons we have learned. *Int J Robot Res* 2021.
- [73] Sünderhauf N, Brock O, Scheirer W, Hadsell R, Fox D, Leitner J, et al. The limits and potentials of deep learning for robotics. *Int J Robot Res* 2018;37(4–5):405–20.
- [74] Haarnoja T, Tang H, Abbeel P, Levine S. Reinforcement learning with deep energy-based policies. In: *Proceedings of the international conference on machine learning (ICML)*. Vol. 70; 2017, p. 1352–61.
- [75] Haarnoja T, Zhou A, Abbeel P, Levine S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: *Proceedings of the international conference on machine learning (ICML)*. Vol. 80; 2018, p. 1861–70.
- [76] Eysenbach B, Levine S. Maximum entropy RL (provably) solves some robust RL problems. 2021.
- [77] Peng XB, Andrychowicz M, Zaremba W, Abbeel P. Sim-to-real transfer of robotic control with dynamics randomization. In: *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE; 2018, p. 3803–10.

- [78] Rusu AA, Večerík M, Rothörl T, Heess N, Pascanu R, Hadsell R. Sim-to-real robot learning from pixels with progressive nets. In: Proceedings of the 1st annual conference on robot learning. Proceedings of machine learning research, 78, PMLR; 2017, p. 262–70.
- [79] James S, Wohlhart P, Kalakrishnan M, Kalashnikov D, Irpan A, Ibarz J, et al. Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR); 2019.
- [80] García J, Fernández F. A comprehensive survey on safe reinforcement learning. *J Mach Learn Res* 2015;16(42):1437–80.
- [81] Kolter JZ, Manek G. Learning stable deep dynamics models. In: Advances in neural information processing systems. 2019, p. 11126–34.
- [82] Gehring C, Precup D. Smart exploration in reinforcement learning using absolute temporal difference errors. In: Proceedings of the 2013 international conference on autonomous agents and multi-agent systems; 2013, p. 1037–44.
- [83] Tsiamis A, Pappas GJ. Linear systems can be hard to learn. 2021.
- [84] Akametalu AK, Fisac JF, Gillula JH, Kaynama S, Zeilinger MN, Tomlin CJ. Reachability-based safe learning with Gaussian processes. In: 53rd IEEE conference on decision and control (CDC); 2014, p. 1424–31.
- [85] Feder HJS, Leonard JJ, Smith CM. Adaptive mobile robot navigation and mapping. *Int J Robot Res* 1999;18(7):650–68.
- [86] Leung C, Huang S, Kwok N, Dissanayake G. Planning under uncertainty using model predictive control for information gathering. *Robot Auton Syst* 2006;54(11):898–910.
- [87] Sim R, Roy N. Global A-optimal robot exploration in SLAM. In: IEEE int. conf. on robotics and automation (ICRA); 2005, p. 661–6.
- [88] Vander Hook J, Tokekar P, Isler V. Cautious greedy strategy for bearing-based active localization: Experiments and theoretical analysis. In: IEEE international conference on robotics and automation (ICRA); 2012, p. 1787–92.
- [89] Marchant R, Ramos F. Bayesian optimisation for intelligent environmental monitoring. In: IEEE int. conf. on intelligent robots and systems (IROS); 2012, p. 2242–49.
- [90] Wong E-M, Bourgault F, Furukawa T. Multi-vehicle Bayesian search for multiple lost targets. In: IEEE int. conf. on robotics and automation (ICRA); 2005, p. 3169–74.
- [91] Stachniss C, Burgard W. Exploring unknown environments with mobile robots using coverage maps. In: International joint conference on artificial intelligence; 2003, p. 1127–34.
- [92] Kreucher C, Wegrzyn J, Beauvais M, Conti R. Multiplatform information-based sensor management: an inverted UAV demonstration. In: SPIE defense transformation and network-centric systems. Vol. 6578; 2007, p. 65780Y–1–11.
- [93] Roy N, Earnest C. Dynamic action spaces for information gain maximization in search and exploration. In: American controls conf. (ACC); 2006, p. 1631–36.
- [94] Lu W, Zhang G, Ferrari S, Fierro R, Palunko I. An information potential approach for tracking and surveilling multiple moving targets using mobile sensor agents. In: SPIE Unmanned Systems Technology. Vol. 8045. 2011, 80450T–1–80450T–13.
- [95] Bourgault F, Makarenko AA, Williams S, Grocholsky B, Durrant-Whyte H. Information based adaptive robotic exploration. In: IEEE Int. Conf. on Intelligent Robots and Systems (IROS). Vol. 1; 2002, p. 540–5.
- [96] Elfes A. Using occupancy grids for mobile robot perception and navigation. *Computer* 1989;22(6):46–57.
- [97] Bender A, Williams SB, Pizarro O. Autonomous exploration of large-scale benthic environments. In: IEEE Int. Conf. on Robotics and Automation (ICRA); 2013, p. 390–6.
- [98] Cao N, Low KH, Dolan JM. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In: International conference on autonomous agents and multi-agent systems; 2013, p. 7–14.
- [99] Hoang TN, Low KH, Jaillet P, Kankanhalli M. Nonmyopic ϵ -Bayes-optimal active learning of Gaussian processes. In: International conference on machine learning; 2014, p. 739–47.
- [100] Low KH, Dolan JM, Khosla P. Adaptive multi-robot wide-area exploration and mapping. In: Conference on Autonomous Agents and Multiagent Systems. 2008, p. 23–30.
- [101] Singh A, Krause A, Guestrin C, Kaiser WJ. Efficient informative sensing using multiple robots. *J Artif Intell Res (JAIR)* 2009;34:707–55.
- [102] Souza J, Marchant R, Ott L, Wolf D, Ramos F. Bayesian optimisation for active perception and smooth navigation. In: IEEE Int. Conf. on Robotics and Automation (ICRA); 2014, p. 4081–87.
- [103] Picardi G, Chellapurath M, Iacoponi S, Stefanni S, Laschi C, Calisti M. Bio-inspired underwater legged robot for seabed exploration with low environmental disturbance. *Science Robotics* 2020;5(42).
- [104] Breier JA, Jakuba MV, Saito MA, Dick GJ, Grim SL, Chan EW, et al. Revealing ocean-scale biochemical structure with a deep-diving vertical profiling autonomous vehicle. *Science Robotics* 2020;5(48).
- [105] Zhang Y, Ryan JP, Hobson BW, Kieft B, Romano A, Barone B, et al. A system of coordinated autonomous robots for Lagrangian studies of microbes in the oceanic deep chlorophyll maximum. *Science Robotics* 2021;6(50).
- [106] Fossum TO, Fragoso GM, Davies EJ, Ullgren JE, Mendes R, Johnsen G, et al. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics* 2019;4(27).
- [107] Abraham I, Mavrommati A, Murphey TD. Data-driven measurement models for active localization in sparse environments. In: Robotics: Science and systems proceedings. 2018.
- [108] Guntuboyina A, Sen B, et al. Nonparametric shape-restricted regression. *Statist Sci* 2018;33(4):568–94.
- [109] Hasler N, Ackermann H, Rosenhahn B, Thormählen T, Seidel H-P. Multilinear pose and body shape estimation of dressed subjects from image sets. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE; 2010, p. 1823–30.
- [110] Southall B, Taylor C. Stochastic road shape estimation. In: Proceedings of the 8th IEEE international conference on computer vision (ICCV). Vol. 1; 2001, p. 205–212.
- [111] Abraham I, Prabhakar A, Hartmann M, Murphey T. Ergodic exploration using binary sensing for non-parametric shape estimation. *IEEE Robot Autom Lett* 2017;2(2):827–34.
- [112] Bayen AM, Mitchell IM, Oishi MM, Tomlin CJ. Aircraft autolander safety analysis through optimal control-based reach set computation. *J Guid Control Dyn* 2007;30(1):68–77.
- [113] Schmidt M, Lipson H. Distilling free-form natural laws from experimental data. *Science* 2009;324(5923):81–5.
- [114] Abraham I, Murphey TD. Active learning of dynamics for data-driven control using Koopman operators. *IEEE Trans Robot* 2019;35(5):1071–83.
- [115] Berrueta TA, Abraham I, Murphey T. Experimental applications of the Koopman operator in active learning for control. In: The Koopman operator in systems and control: Concepts, methodologies, and applications. Springer International Publishing; 2020, p. 421–50.
- [116] Oubbati M, Palm G. A neural framework for adaptive robot control. *Neural Comput Appl* 2010;19(1):103–14.
- [117] Nguyen-Tuong D, Seeger M, Peters J. Model learning with local Gaussian process regression. *Adv Robot* 2009;23(15):2015–34.
- [118] Kim JZ, Lu Z, Nozari E, Pappas GJ, Bassett DS. Teaching recurrent neural networks to infer global temporal structure from local examples. *Nat Mach Intell* 2021;3(4):316–23.
- [119] Karkus P, Ma X, Hsu D, Kaelbling LP, Lee WS, Lozano-Pérez T. Differentiable algorithm networks for composable robot learning. *Robot Sci Syst* 2019.
- [120] Thuruethel TG, Shih B, Laschi C, Tolley MT. Soft robot perception using embedded soft sensors and recurrent neural networks. *Science Robotics* 2019;4(26).
- [121] Hofmann T, Schölkopf B, Smola AJ. Kernel methods in machine learning. *Ann Statist* 2008;1171–220.
- [122] Schaal S, Atkeson C, Vijayakumar S. Real-time robot learning with locally weighted statistical learning. In: Proceedings of the IEEE international conference on robotics and automation (ICRA). Vol. 1; 2000, p. 288–93.
- [123] Cheng C-A, Huang H-P, Hsu H-K, Lai W-Z, Cheng C-C. Learning the inverse dynamics of robotic manipulators in structured reproducing kernel Hilbert space. *IEEE Trans Cybern* 2016;46(7):1691–703.
- [124] Dalla Libera A, Carli R. A data-efficient geometrically inspired polynomial kernel for robot inverse dynamic. *IEEE Robot Autom Lett* 2019;5(1):24–31.
- [125] Smola AJ, Schölkopf B. Bayesian kernel methods. In: Advanced lectures on machine learning. Springer; 2003, p. 65–117.
- [126] Gallager RG. Stochastic processes: Theory for applications. Cambridge University Press; 2013.
- [127] Rasmussen CE, Williams CKI. Gaussian processes for machine learning. The MIT Press; 2005.
- [128] Deisenroth M, Fox D, Rasmussen CE. Gaussian Processes for data-efficient learning in robotics and control. *IEEE Trans Pattern Anal Mach Intell* 2015;37(2):408–23.
- [129] Otto SE, Rowley CW. Koopman operators for estimation and control of dynamical systems. *Annu Rev Control Robot Auton Syst* 2021;4(1).
- [130] Koopman BO. Hamiltonian systems and transformation in Hilbert space. *Proc Natl Acad Sci* 1931;17(5):315–8.
- [131] Tu JH, Rowley CW, Luchtenburg DM, Brunton SL, Kutz JN. On dynamic mode decomposition: Theory and applications. *J Comput Dyn* 2014;1:391.
- [132] Williams M, Kevrekidis IG, Rowley CW. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *J Nonlinear Sci* 2015;25:1307–46.
- [133] Brunton S, Brunton B, Proctor J, Kutz J. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE* 2016;11(2):1–19.
- [134] Proctor JL, Brunton SL, Kutz JN. Generalizing Koopman theory to allow for inputs and control. *SIAM J Appl Dyn Syst* 2018;17(1):909–30.
- [135] Kaiser E, Kutz JN, Brunton SL. Data-driven discovery of Koopman eigenfunctions for control. 2017, ArXiv.
- [136] Abraham I, de la Torre G, Murphey T. Model-based control using Koopman operators. In: Robotics: Science and systems proceedings. 2017.
- [137] Bruder D, Gillespie B, Remy CD, Vasudevan R. Modeling and control of soft robots using the Koopman operator and model predictive control. In: Robotics: Science and systems. 2019.
- [138] Mamakoukas G, Castaño ML, Tan X, Murphey TD. Local Koopman operators for data-driven control of robotic systems. In: Robotics: Science and Systems. 2019.

- [139] Bishop CM. Pattern recognition and machine learning (Information science and statistics). Springer; 2006.
- [140] Zhong G, Wang L-N, Ling X, Dong J. An overview on data representation learning: From traditional feature learning to recent deep learning. *J Financ Data Sci* 2016;2(4):265–78.
- [141] Butepage J, Black MJ, Kragic D, Kjellstrom H. Deep representation learning for human motion prediction and classification. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR); 2017. p. 6158–66.
- [142] Liu J, Liu Z, Wang L, Guo L, Dang J. Time-frequency deep representation learning for speech emotion recognition integrating self-attention. In: Gedeon T, Wong KW, Lee M, editors. Neural information processing. Springer International Publishing; 2019, p. 681–9.
- [143] de Bruin T, Kober J, Tuyls K, Babu ska R. Integrating state representation learning into deep reinforcement learning. *IEEE Robot Autom Lett* 2018;3(3):1394–401.
- [144] Lusch B, Kutz JN, Brunton SL. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Commun* 2018;9(1):4950.
- [145] Shih B, Shah D, Li J, Thuruthel TG, Park Y-L, Iida F, et al. Electronic skins and machine learning for intelligent soft robots. *Science Robotics* 2020;5(41).
- [146] Madokoro H, Sato K, Shimoi N. Adaptive category mapping networks for all-mode topological feature learning used for mobile robot vision. In: Proceedings of the 23rd IEEE International Symposium on Robot and Human Interactive Communication. IEEE; 2014, p. 678–83.
- [147] Spletzer JR, Taylor CJ. Dynamic sensor planning and control for optimally tracking targets. *Int J Robot Res* 2003;22(1):7–20.
- [148] DasGupta B, Hespanha JP, Riehl J, Sontag E. Honey-pot constrained searching with local sensory information. *Nonlinear Anal TMA* 2006;65(9):1773–93.
- [149] Zhang G, Ferrari S. An adaptive artificial potential function approach for geometric sensing. In: IEEE int. conf. on decision and control (CDC); 2009. p. 7903–10.
- [150] Hager G, Mintz M. Computational methods for task-directed sensor data fusion and sensor planning. *Int J Robot Res* 1991;10(4):285–313.
- [151] Benet G, Blanes F, Simó J, Pérez P. Using infrared sensors for distance measurement in mobile robots. *Robot Auton Syst* 2002;40(4):255–66.
- [152] Denzler J, Zobel M, Niemann H. Information theoretic focal length selection for real-time active 3D object tracking. In: IEEE int. conf. on computer vision, 2003. p. 400–7.
- [153] Vergassola M, Villermaux E, Shraiman BI. Infotaxis as a strategy for searching without gradients. *Nature* 2007;445(7126):406.
- [154] Fox D, Burgard W, Thrun S. Active Markov localization for mobile robots. *Robot Auton Syst* 1998;25(3–4):195–207.
- [155] Arbel T, Ferrie F. Viewpoint selection by navigation through entropy maps. In: IEEE int. conf. on computer vision, 1999. p. 248–54.
- [156] Vázquez P-P, Feixas M, Sbert M, Heidrich W. Viewpoint selection using viewpoint entropy. In: Vision modeling and visualization conference. vol. 1, 2001, p. 273–80.
- [157] Takeuchi Y, Ohnishi N, Sugie N. Active vision system based on information theory. *Syst Comput Japan* 1998;29(11):31–9.
- [158] Kreucher C, Kastella K, Hero AO. Sensor management using an active sensing approach. *Signal Process* 2005;85(3):607–24.
- [159] Toh J, Sukkarieh S. A Bayesian formulation for the prioritized search of moving objects. In: IEEE int. conf. on robotics and automation (ICRA), 2006. p. 219–24.
- [160] Denzler J, Brown C. Information theoretic sensor data selection for active object recognition and state estimation. *IEEE Trans Pattern Anal Mach Intell* 2002;24(2):145–57.
- [161] Tisdale J, Kim Z, Hedrick JK. Autonomous UAV path planning and estimation. *IEEE Robot Autom Magaz* 2009;16(2):35–42.
- [162] Grocholsky B, Keller J, Kumar V, Pappas G. Cooperative air and ground surveillance. *IEEE Robot Autom Magaz* 2006;13(3):16–25.
- [163] Lu W, Zhang G, Ferrari S. An information potential approach to integrated sensor path planning and control. *IEEE Trans Robot* 2014;30(4):919–34.
- [164] Zhang G, Ferrari S, Qian M. An information roadmap method for robotic sensor path planning. *J Intell Robot Syst* 2009;56(1–2):69–98.
- [165] Hollinger GA, Englot B, Hover FS, Mitra U, Sukhatme GS. Active planning for underwater inspection and the benefit of adaptivity. *Int J Robot Res* 2013;32(1):3–18.
- [166] Liao X, Carin L. Application of the theory of optimal experiments to adaptive electromagnetic-induction sensing of buried targets. *IEEE Trans Pattern Anal Mach Intell* 2004;26(8):961–72.
- [167] Emery A, Nenarokomov AV. Optimal experiment design. *Meas Sci Technol* 1998;9(6):864.
- [168] Ucinski D, Korbicz J. Path planning for moving sensors in parameter estimation of distributed systems. In: Workshop on robot motion and control (RoMoCo), 1999. p. 273–8.
- [169] Ucinski D. Optimal sensor location for parameter estimation of distributed processes. *Internat J Control* 2000;73(13):1235–48.
- [170] Frieden RB. Science from fisher information: A unification. Cambridge University Press; 2004.
- [171] Shannon C. A mathematical theory of communication. *Bell Syst Tech J* 1948;27:379–423.
- [172] Atanasov N, Sankaran B, Le Ny J, Pappas G, Daniilidis K. Nonmyopic view planning for active object classification and pose estimation. *IEEE Trans Robot* 2014;30(5):1078–90.
- [173] Li Y, Liu ZG. Information entropy based viewpoint planning for 3-D object reconstruction. *IEEE Trans Robot* 2005;21(3):324–7.
- [174] Rahimi M, Hansen M, Kaiser W, Sukhatme G, Estrin D. Adaptive sampling for environmental field estimation using robotic sensors. In: IEEE int. conf. on intelligent robots and systems (IROS), 2005. p. 3692–8.
- [175] Cover TM, Thomas JA. Elements of information theory. John Wiley & Sons; 2012.
- [176] Wahlström N, Schön TB, Deisenroth MP. From pixels to torques: Policy learning with deep dynamical models. 2015.
- [177] Tovar B, Munoz-Gomez L, Murrieta-Cid R, Alencastre-Miranda M, Monroy R, Hutchinson S. Planning exploration strategies for simultaneous localization and mapping. *Robot Auton Syst* 2006;54(4):314–31.
- [178] Tovar B, Murphey TD. Trajectory tracking among landmarks and binary sensor beams. In: IEEE int. conf. on robotics and automation (ICRA), 2012. p. 2121–7.
- [179] Cadena C, Carlone L, Carrillo H, Latif Y, Scaramuzza D, Neira J, et al. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans Robot* 2016;32(6):1309–32.
- [180] Neumann Jv. Proof of the quasi-ergodic hypothesis. *Proc Natl Acad Sci* 1932;18(1):70–82.
- [181] Krengel U. Ergodic theorems. de Gruyter; 1985.
- [182] Shell DA, Jones CV, Mataric MJ. Ergodic dynamics by design: A route to predictable multi-robot systems. In: Multi-robot systems. from swarms to intelligent automata. Springer; 2005. p. 291–7.
- [183] Mathew G, Mezić I. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D* 2011;240(4):432–42.
- [184] Wilson A, Schultz J, Murphey TD. Trajectory synthesis for Fisher information maximization. *IEEE Trans Robot* 2014;30(6):1358–70.
- [185] Cooper J, Goodrich M. Towards combining UAV and sensor operator roles in UAV-enabled visual search. In: IEEE int. conf. on human robot interaction (HRI), 2008. p. 351–8.
- [186] Cai C, Ferrari S. Information-driven sensor path planning by approximate cell decomposition. *IEEE Trans Syst Man Cybern* 2009;39(3):672–89.
- [187] Ye Y, Tsotsos JK. Sensor planning for 3D object search. *Comput Vis Image Underst* 1999;73(2):145–68.
- [188] Massios NA, Fisher RB. A best next view selection algorithm incorporating a quality criterion. In: British machine vision conference, 1998. p. 78.1–78.10.
- [189] Marchant R, Ramos F. Bayesian Optimisation for informative continuous path planning. In: IEEE int. conf. on robotics and automation (ICRA), 2014. p. 6136–43.
- [190] Mayne D. A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems. *Internat J Control* 1966;3(1):85–95.
- [191] Bellman R. On the theory of dynamic programming. In: Proceedings of the national academy. vol. 38.8, 1952. p. 716.
- [192] Tassa Y, Mansard N, Todorov E. Control-limited differential dynamic programming. In: 2014 IEEE international conference on robotics and automation (ICRA). IEEE; 2014, p. 1168–75.
- [193] Kumar V, Gupta A, Todorov E, Levine S. Learning dexterous manipulation policies from experience and imitation. 2016, ArXiv.
- [194] Williams G, Wagener N, Goldfain B, Drews P, Rehg JM, Boots B, et al. Information theoretic MPC for model-based reinforcement learning. In: International conference on robotics and automation (ICRA). 2017.
- [195] Theodorou EA, Todorov E. Relative entropy and free energy dualities: Connections to Path Integral and KL control. In: 2012 IEEE 51st IEEE conference on decision and control (CDC), 2012. p. 1466–73.
- [196] Kappen HJ. Path integrals and symmetry breaking for optimal control theory. *J Stat Mech Theory Exp* 2005;2005(11):P11011.
- [197] Williams G, Aldrich A, Theodorou EA. Model predictive path integral control: From theory to parallel computation. *J Guid Control Dyn* 2017;40(2):344–57.
- [198] Ansari A, Murphey TD. Sequential action control: Closed-form optimal control for nonlinear and nonsmooth systems. *IEEE Trans Rob* 2017;32.
- [199] Abraham I, Broad A, Pinosky A, Argall B, Murphey TD. Hybrid control for learning motor skills. In: Workshop on the algorithmic foundations of robotics (WAFR). 2020.
- [200] Wilson A, Schultz J, Ansari A, Murphey TD. Real-time trajectory synthesis for information maximization using sequential action control and least-squares estimation. In: IEEE int. conf. on intelligent robots and systems (IROS), 2015. p. 4935–40.
- [201] Hollinger GA, Sukhatme GS. Sampling-based robotic information gathering algorithms. *Int J Robot Res* 2014;33(9):1271–87.
- [202] Ryan A, Hedrick JK. Particle filter based information-theoretic active sensing. *Robot Auton Syst* 2010;58(5):574–84.
- [203] Vickers NJ. Mechanisms of animal navigation in odor plumes. *Biol Bull* 2000;198(2):203–12.
- [204] Miller L, Silverman Y, MacIver MA, Murphey T. Ergodic exploration of distributed information. *IEEE Trans Robot* 2016;32(1):36–52.

- [205] Nishimura H, Schwager M. SACBP: belief space planning for continuous-time dynamical systems via stochastic sequential action control. In: Workshop on the algorithmic foundations of robotics, 2018. p. 267–83.
- [206] Dressel L, Kochenderfer MJ. Tutorial on the generation of ergodic trajectories with projection-based gradient descent. *IET Cyber-Phys Syst Theory Appl* 2019;4(2):89–100.
- [207] Dressel L, Kochenderfer MJ. Using neural networks to generate information maps for mobile sensors. In: 2018 IEEE conference on decision and control (CDC), 2018. p. 2555–60.
- [208] Paley DA, Wolek A. Mobile sensor networks and control: Adaptive sampling of spatiotemporal processes. *Annu Rev Control Robot Auton Syst* 2020;3:91–114.
- [209] Chen Z, Xiao L, Wang Q, Wang Z, Sun Z. Coverage control of multi-agent systems for ergodic exploration. In: 2020 39th Chinese control conference (CCC), 2020. p. 4947–52.
- [210] Khodayi-mehr R, Aquino W, Zavlanos MM. Model-based active source identification in complex environments. *IEEE Trans Robot* 2019;35(3):633–52.
- [211] Veitch C, Render D, Aravind A. Ergodic flocking. In: 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS). IEEE; 2019, p. 6957–62.
- [212] Salman H, Ayvali E, Choset H. Multi-agent ergodic coverage with obstacle avoidance. In: International conference on automated planning and scheduling. 2017.
- [213] Ayvali E, Salman H, Choset H. Ergodic coverage in constrained environments using stochastic trajectory optimization. In: IEEE/RSJ international conference on intelligent robots and systems (IROS), 2017. p. 5204–10.
- [214] Prabhakar A, Abraham I, Schlafly M, Taylor A, Popovic K, Diniz G, et al. Ergodic specifications for flexible swarm control: from user commands to persistent adaptation. In: Robotics: science and systems proceedings, 2020.
- [215] De La Torre G, Flaßkamp K, Prabhakar A, Murphey TD. Ergodic exploration with stochastic sensor dynamics. In: American controls conf. (ACC), 2016. p. 2971–6.
- [216] Mavrommati A, Tzorakoleftherakis E, Abraham I, Murphey TD. Real-time area coverage and target localization using receding-horizon ergodic exploration. *IEEE Trans Robot* 2018;34(1):62–80.
- [217] Miller LM, Murphey TD. Trajectory optimization for continuous ergodic exploration. In: American controls conf. (ACC), 2013. p. 4196–201.
- [218] Miller L, Murphey TD. Optimal planning for target localization and coverage using range sensing. In: IEEE int. conf. on automation science and engineering (CASE), 2015. p. 501–8.
- [219] Agharese N, Cloyd T, Blumenschein LH, Raitor M, Hawkes EW, Culbertson H, et al. HapWRAP: soft growing wearable haptic device. In: IEEE international conference on robotics and automation (ICRA), 2018. p. 5466–72.
- [220] Galloway KC, Becker KP, Phillips B, Kirby J, Licht S, Tchernov D, et al. Soft robotic grippers for biological sampling on deep reefs. *Soft Robot* 2016;3(1):23–33.
- [221] Tolley MT, Shepherd RF, Mosadegh B, Galloway KC, Wehner M, Karpelson M, et al. A resilient, untethered soft robot. *Soft Robot* 2014;1(3):213–25.
- [222] Gillespie MT, Best CM, Townsend EC, Wingate D, Killpack MD. Learning nonlinear dynamic models of soft robots for model predictive control with neural networks. In: 2018 IEEE international conference on soft robotics (RoboSoft), 2018. p. 39–45.
- [223] Laschi C, Cianchetti M. Soft robotics: new perspectives for robot bodyware and control. *Front Bioeng Biotechnol* 2014;2:3.
- [224] Picardi G, Hauser H, Laschi C, Calisti M. Morphologically induced stability on an underwater legged robot with a deformable body. *Int J Robot Res* 2019;40:435–48.
- [225] Mammano F, Nobili R. Biophysics of the cochlea: linear approximation. *J Acoust Soc Am* 1993;93(6):3320–32.
- [226] Sumbre G, Fiorito G, Flash T, Hochner B. Motor control of flexible octopus arms. *Nature* 2005;433(7026):595–6.
- [227] Sornkarn N, Nanayakkara T. Can a soft robotic probe use stiffness control like a human finger to improve efficacy of haptic perception? *IEEE Trans Haptics* 2017;10(2):183–95.
- [228] Murphy RR, Tadokoro S, Kleiner A. Disaster robotics. In: Springer handbook of robotics. Springer; 2016. p. 1577–604.
- [229] Paull L, Saeedi S, Seto M, Li H. Sensor-driven online coverage planning for autonomous underwater vehicles. *IEEE/ASME Trans Mechatronics* 2013;18(6):1827–38.
- [230] Oksanen T, Visala A. Coverage path planning algorithms for agricultural field machines. *J Field Robotics* 2009;26(8):651–68.
- [231] Englot B, Hover FS. Sampling-based coverage path planning for inspection of complex structures. In: International conference on automated planning and scheduling (ICAPS). 2012.
- [232] Galceran E, Carreras M. A survey on coverage path planning for robotics. *Robot Auton Syst* 2013;61(12):1258–76.
- [233] Choset H. Coverage for robotics—a survey of recent results. *Ann Math Artif Intell* 2001;31(1–4):113–26.
- [234] Karapetyan N, Benson K, McKinney C, Taslakian P, Rekleitis I. Efficient multi-robot coverage of a known environment. In: 2017 IEEE/RSJ international conference on intelligent robots and systems (IROS), 2017. p. 1846–52.
- [235] Jan GE, Luo C, Hung L, Shih S. A computationally efficient complete area coverage algorithm for intelligent mobile robot navigation. In: 2014 international joint conference on neural networks, (IJCNN), 2014. p. 961–6.
- [236] Schwager M, Rus D, Slotine J-J. Decentralized, adaptive coverage control for networked robots. *Int J Robot Res* 2009;28(3):357–75.
- [237] Stergiopoulos Y, Tzes A. Spatially distributed area coverage optimisation in mobile robotic networks with arbitrary convex anisotropic patterns. *Automatica* 2013;49(1):232–7.
- [238] García-Garrido V, Mancho A, Wiggins S, Mendoza C. A dynamical systems approach to the surface search for debris associated with the disappearance of flight MH370. *Nonlin Process Geophys* 2015;22:701–12.
- [239] Rodríguez-Arévalo ML, Neira J, Castellanos JA. On the importance of uncertainty representation in active SLAM. *IEEE Trans Robot* 2018;34(3):829–34.
- [240] Stachniss C, Grisetti G, Burgard W. Information gain-based exploration using rao-blackwellized particle filters. In: Robotics: science and systems. vol. 2, 2005, p. 65–72.
- [241] Carrillo H, Dames P, Kumar V, Castellanos JA. Autonomous robotic exploration using occupancy grid maps and graph SLAM based on Shannon and Rényi entropy. In: IEEE international conference on robotics and automation (ICRA), 2015. p. 487–94.
- [242] Carlone L, Du J, Ng MK, Bona B, Indri M. Active SLAM and exploration with particle filters using Kullback-Leibler divergence. *J Intell Robot Syst* 2014;75(2):291–311.
- [243] Leung C, Huang S, Dissanayake G. Active SLAM using model predictive control and attractor based exploration. In: 2006 IEEE/RSJ international conference on intelligent robots and systems, 2006. p. 5026–5031.
- [244] Atanasov N, Le Ny J, Daniilidis K, Pappas GJ. Decentralized active information acquisition: Theory and application to multi-robot SLAM. In: 2015 IEEE international conference on robotics and automation (ICRA), 2015. p. 4775–82.
- [245] Bonet B, Geffner H. Planning with incomplete information as heuristic search in belief space. In: Proceedings of the fifth international conference on artificial intelligence planning systems. AAAI Press; 2000. p. 52–61.
- [246] Platt Jr R, Tedrake R, Kaelbling L, Lozano-Perez T. Belief space planning assuming maximum likelihood observations. In: Robotics: science and systems. 2006.
- [247] Prentice S, Roy N. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Int J Robot Res* 2009;28(11–12):1448–65.
- [248] Valencia R, Morta M, Andrade-Cetto J, Porta JM. Planning reliable paths with pose SLAM. *IEEE Trans Robot* 2013;29(4):1050–9.
- [249] Patil S, Kahn G, Laskey M, Schulman J, Goldberg K, Abbeel P. Scaling up Gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In: Workshop on the algorithmic foundations of robotics. Springer; 2015. p. 515–33.
- [250] Kim A, Eustice RM. Active visual SLAM for robotic area coverage: Theory and experiment. *Int J Robot Res* 2015;34(4–5):457–75.
- [251] Taketomi T, Uchiyama H, Ikeda S. Visual SLAM algorithms: a survey from 2010 to 2016. *IPSN Trans Comput Vision Appl* 2017;9(1):16.
- [252] Chen S, Li Y, Kwok NM. Active vision in robotic systems: A survey of recent developments. *Int J Robot Res* 2011;30(11):1343–77.
- [253] Mnih V, Heess N, Graves A, et al. Recurrent models of visual attention. *Adv Neural Inf Process Syst* 2014;27:2204–12.
- [254] Carlone L, Karaman S. Attention and anticipation in fast visual-inertial navigation. *IEEE Trans Robot* 2018;35(1):1–20.
- [255] Hussein A, Gaber MM, Elyan E, Jayne C. Imitation learning: A survey of learning methods. *ACM Comput Surv* 2017;50(2).
- [256] Argall BD, Chernova S, Veloso M, Browning B. A survey of robot learning from demonstration. *Robot Auton Syst* 2009;57(5):469–83.
- [257] Osa T, Pajarinen J, Neumann G, Bagnell JA, Abbeel P, Peters J. An algorithmic perspective on imitation learning. 2018, ArXiv.
- [258] Codevilla F, Müller M, López A, Koltun V, Dosovitskiy A. End-to-end driving via conditional imitation learning. In: 2018 IEEE international conference on robotics and automation (ICRA). IEEE; 2018, p. 1–9.
- [259] Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, et al. Mastering the game of go with deep neural networks and tree search. *Nature* 2016;529(7587):484–9.
- [260] Ijspeert AJ, Nakanishi J, Schaal S. Movement imitation with nonlinear dynamical systems in humanoid robots. In: Proceedings 2002 IEEE international conference on robotics and automation (Cat. No.02CH37292). vol. 2, 2002, p. 1398–403 vol.2.
- [261] Silver D, Bagnell JA, Stentz A. Active learning from demonstration for robust autonomous navigation. In: 2012 IEEE international conference on robotics and automation. IEEE; 2012, p. 200–7.
- [262] Dima C, Hebert M. Active learning for outdoor obstacle detection. In: Robotics: science and systems. 2005, p. 9–16.
- [263] Ab Aza N, Shahmansoorian A, Davoudi M. From inverse optimal control to inverse reinforcement learning: A historical review. *Annu Rev Control* 2020.
- [264] Daniel C, Kroemer O, Viering M, Metz J, Peters J. Active reward learning with a novel acquisition function. *Auton Robots* 2015;39(3):389–405.
- [265] Judah K, Fern A, Dieterich TG. Active imitation learning via reduction to iid active learning. 2012, ArXiv.

- [266] Ho J, Ermon S. Generative adversarial imitation learning. 2016, ArXiv.
- [267] Li Y, Song J, Ermon S. Infogail: Interpretable imitation learning from visual demonstrations. In: *Advances in neural information processing systems*. 2017, p. 3812–22.
- [268] Chen X, Duan Y, Houthoofd R, Schulman J, Sutskever I, Abbeel P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Adv Neural Inf Process Syst* 2016;29:2172–80.
- [269] Low KH, Chen J, Dolan JM, Chien S, Thompson DR. Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing. In: *Proceedings of the 11th international conference on autonomous agents and multiagent systems*, 2012. p. 105–12.
- [270] Abraham I, Murphey T. Decentralized ergodic control: Distribution-driven sensing and exploration for multi-agent systems. *IEEE Robot Autom Lett* 2018;3(4):2987–94.
- [271] Best G, Cliff OM, Patten T, Mettu RR, Fitch R. Dec-MCTS: Decentralized planning for multi-robot active perception. *Int J Robot Res* 2019;38(2–3):316–37.
- [272] Verbraeken J, Wolting M, Katzy J, Kloppenburg J, Verbelen T, Rellermeyer JS. A survey on distributed machine learning. *ACM Comput Surv* 2020;53(2):1–33.
- [273] Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process Mag* 2020;37(3):50–60.
- [274] Ames AD, Galloway K, Sreenath K, Grizzle JW. Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics. *IEEE Trans Automat Control* 2014;59(4):876–91.
- [275] Ames AD, Powell M. Towards the unification of locomotion and manipulation through control Lyapunov functions and quadratic programs. In: *Control of cyber-physical systems*. Springer; 2013, p. 219–40.
- [276] Ames AD, Coogan S, Egerstedt M, Notomista G, Sreenath K, Tabuada P. Control barrier functions: Theory and applications. In: *2019 18th european control conference (ECC)*, 2019. p. 3420–3431.
- [277] Ames AD, Xu X, Grizzle JW, Tabuada P. Control barrier function based quadratic programs for safety critical systems. *IEEE Trans Automat Control* 2016;62(8):3861–76.
- [278] Wang L, Ames AD, Egerstedt M. Safety barrier certificates for collisions-free multirobot systems. *IEEE Trans Robot* 2017;33(3):661–74.
- [279] Berkenkamp F, Turchetta M, Schoellig A, Krause A. Safe model-based reinforcement learning with stability guarantees. In: *Advances in neural information processing systems*. 2017, p. 908–18.
- [280] Choi J, Castañeda F, Tomlin CJ, Sreenath K. Reinforcement learning for safety-critical control under model uncertainty, using control Lyapunov functions and control barrier functions. 2020, ArXiv.
- [281] Akametalu AK, Fisac JF, Gillula JH, Kaynama S, Zeilinger MN, Tomlin CJ. Reachability-based safe learning with Gaussian processes. In: *53rd IEEE conference on decision and control*, 2014. p. 1424–31.
- [282] Bansal S, Chen M, Herbert S, Tomlin CJ. Hamilton-Jacobi reachability: A brief overview and recent advances. In: *IEEE 56th annual conference on decision and control*, 2017. p. 2242–53.
- [283] Abraham I, Prabhakar A, Murphey TD. An ergodic measure for active learning from equilibrium. *IEEE Trans Autom Sci Eng* 2021. <http://dx.doi.org/10.1109/TASE.2020.3043636>.
- [284] Cheng R, Orosz G, Murray RM, Burdick JW. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, 2019. p. 3387–95.
- [285] Marsden JE, Ratiu TS. *Introduction to mechanics and symmetry: A basic exposition of classical mechanical systems*. vol. 17, Springer Science & Business Media; 2013.
- [286] Chui NLC, Maciejowski JM. Realization of stable models with subspace methods. *Automatica* 1996;32(11):1587–95.
- [287] Mamakoukas G, Xherija O, Murphey TD. Learning memory-efficient stable linear dynamical systems for prediction and control. In: *Conference on neural information processing systems (NeurIPS)*. 2020.
- [288] Boots B, Gordon GJ, Siddiqi SM. A constraint generation approach to learning stable linear dynamical systems. In: *Advances in neural information processing systems*. 2008, p. 1329–36.
- [289] Huang W-B, Cao L-L, Sun F, Zhao D, Liu H, Yu S. Learning stable linear dynamical systems with the weighted least square method. In: *International joint conference on artificial intelligence (IJCAI)*, 2016. p. 1599–605.
- [290] Erichson NB, Muehlebach M, Mahoney MW. Physics-informed autoencoders for Lyapunov-stable fluid flow prediction. ArXiv 2019.
- [291] Boffi NM, Tu S, Matni N, Slotine J-JE, Sindhvani V. Learning stability certificates from data. 2020, ArXiv.
- [292] Richards SM, Berkenkamp F, Krause A. The Lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems. In: Billard A, Dragan A, Peters J, Morimoto J, editors. *Proceedings of the 2nd conference on robot learning*. *Proceedings of machine learning research*, vol. 87, PMLR; 2018, p. 466–76.
- [293] Mamakoukas G, Abraham I, Murphey TD. Learning stable models for prediction and control. *IEEE Trans Robot* 2020. submitted for publication. <https://arxiv.org/abs/2005.04291>.
- [294] Mehta AM, DelPreto J, Wong KW, Hamill S, Kress-Gazit H, Rus D. Robot creation from functional specifications. In: *Robotics research*. 2018, p. 631–48.
- [295] Berthoz A, Petit J. *The physiology of action and phenomenology*. Oxford University Press; 2008.
- [296] Audibert J-Y, Munos R, Szepesvári C. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoret Comput Sci* 2009;410(19):1876–902.
- [297] Ghosh D, Gupta A, Levine S. Learning actionable representations with goal-conditioned policies. 2018, ArXiv.



Annalisa T. Taylor received her B.S. in Mechanical Engineering from Stanford University in 2017. She is currently a graduate student at Northwestern University in the Center for Robotics and Biosystems. Her research interests include collective behavior, swarm robotics, and safety specifications for robotic systems.



Thomas A. Berrueta received his B.S. degree in engineering from Harvey Mudd College in 2017, and his M.S. in Mechanical Engineering from Northwestern University in 2020. He is currently a doctoral candidate at Northwestern University where he analyzes and develops algorithmic tools for exploiting emergent behavior in complex robotic systems. His research interests include robotics, control, information theory, statistical physics, and their intersections.



Todd D. Murphey received his B.S. degree in mathematics from the University of Arizona and the Ph.D. degree in Control and Dynamical Systems from the California Institute of Technology. He is a Professor of Mechanical Engineering at Northwestern University. His laboratory is part of the Neuroscience and Robotics Laboratory, and his research interests include robotics, control, computational methods for biomechanical systems, and computational neuroscience. Honors include the National Science Foundation CAREER award in 2006, membership in the 2014–2015 DARPA/IDA Defense Science Study Group, and Northwestern's Professorship of Teaching Excellence. He was a Senior Editor of the *IEEE Transactions on Robotics*.