Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

b-it
Bonn-Aachen
International Center for
Information Technology

Institute for AI and
Autonomous Systems

Active Learning

An Overview for Cognitive Robotics

Dr. Alex Mitrevski
Master of Autonomous Systems

Structure

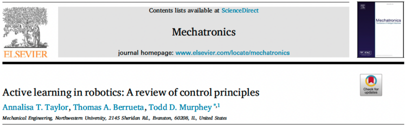
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- Overview of active learning
- Uses of active learning
- Information metrics used in active learning

Motivating Problem: Skill Model Learning

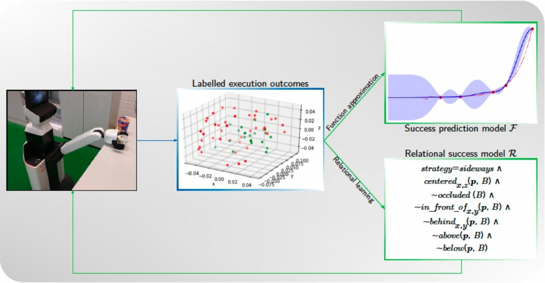

A. Mitrevski, P. G. Plöger, and G. Lakemeyer, "A Hybrid Skill Parameterisation Model Combining Symbolic and Subsymbolic Elements for Introspective Robots," *Robotics and Autonomous Systems*, vol. 161, p. 104350:1–22, Mar. 2023. Available: https://doi.org/10.1016/j.robot.2022.104350

- While learning a skill model, a robot needs to collect experiences by interacting with the environment

Motivating Problem: Skill Model Learning

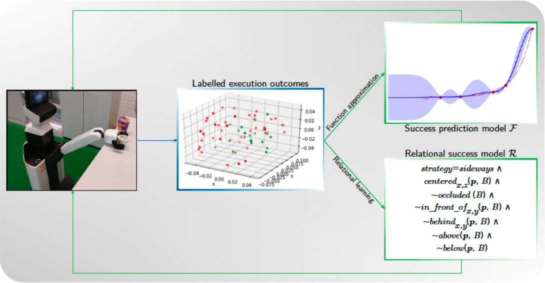

A. Mitrevski, P. G. Plöger, and G. Lakemeyer, "A Hybrid Skill Parameterisation Model Combining Symbolic and Subsymbolic Elements for Introspective Robots," *Robotics and Autonomous Systems*, vol. 161, p. 104350:1–22, Mar. 2023. Available: https://doi.org/10.1016/j.robot.2022.104350

- While learning a skill model, a robot needs to collect experiences by interacting with the environment
- Randomly selecting examples (without external guidance) may result in long learning time

Motivating Problem: Skill Model Learning

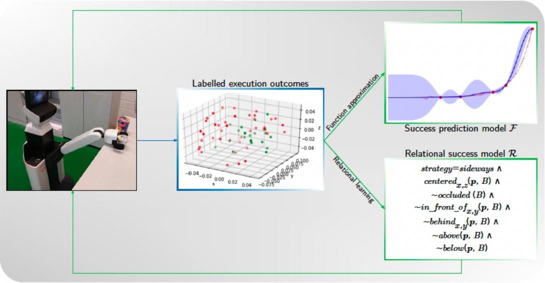

A. Mitrevski, P. G. Plöger, and G. Lakemeyer, "A Hybrid Skill Parameterisation Model Combining Symbolic and Subsymbolic Elements for Introspective Robots," *Robotics and Autonomous Systems*, vol. 161, p. 104350:1–22, Mar. 2023. Available: https://doi.org/10.1016/j.robot.2022.104350

- While learning a skill model, a robot needs to collect experiences by interacting with the environment
- Randomly selecting examples (without external guidance) may result in long learning time
- How can the robot **select informative experiences** so that it can **minimise the number of experiences it has to collect**?

Overview of Active Learning

Active Learning

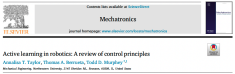
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

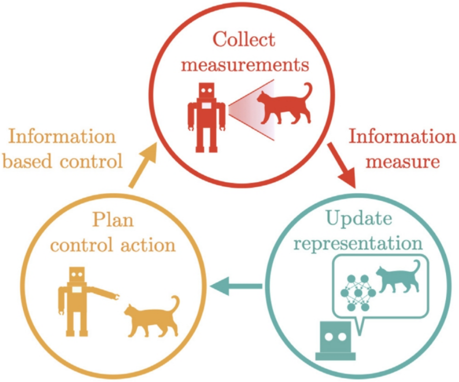

- Enables a learning process during which a robot can actively **focus on relevant learning experiences**

Active Learning

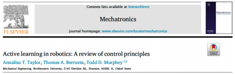
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

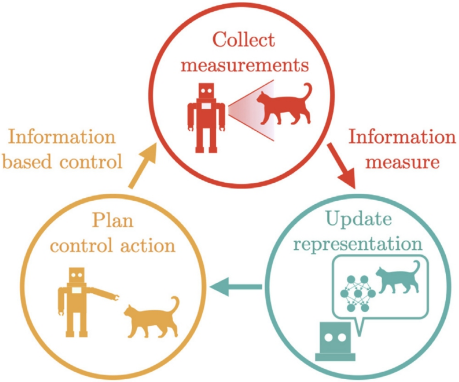

- Enables a learning process during which a robot can actively **focus on relevant learning experiences**
- Experiences are typically selected by **maximising an information measure**

Active Learning

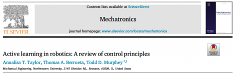
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey

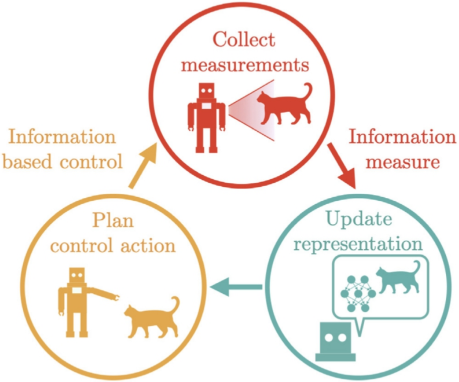

- Enables a learning process during which a robot can actively **focus on relevant learning experiences**
- Experiences are typically selected by **maximising an information measure**
- Aims to **reduce the data requirements of the learning procedure** since the robot does not explore randomly

Active Learning

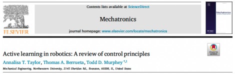
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

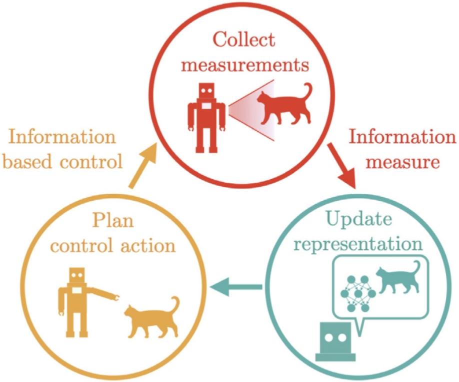

- Enables a learning process during which a robot can actively **focus on relevant learning experiences**
- Experiences are typically selected by **maximising an information measure**
- Aims to **reduce the data requirements of the learning procedure** since the robot does not explore randomly

"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." (Taylor et al., 2021)

Why Active Learning?

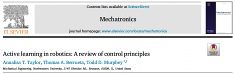
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey *,1
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- Applicable when a robot needs to learn by exploration due to the lack of availability of data

Why Active Learning?

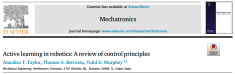
ELSEVIER

Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- Applicable when a robot needs to learn by exploration due to the lack of availability of data
- The general principle is inspired by how people learn optimally — learning is most effective and efficient when examined problems are sufficiently different from what we know so far

Why Active Learning?

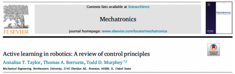
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey *,1
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- Applicable when a robot needs to learn by exploration due to the lack of availability of data
- The general principle is inspired by how people learn optimally — learning is most effective and efficient when examined problems are sufficiently different from what we know so far
- Can be used for a variety of learning problems:

Why Active Learning?

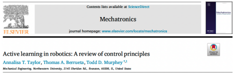
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- Applicable when a robot needs to learn by exploration due to the lack of availability of data
- The general principle is inspired by how people learn optimally — learning is most effective and efficient when examined problems are sufficiently different from what we know so far
- Can be used for a variety of learning problems:
  - **Parameter learning** (e.g. when trying to find optimal parameters for cutting vegetables)

Why Active Learning?

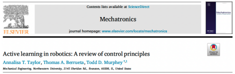
ELSEVIER

Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- Applicable when a robot needs to learn by exploration due to the lack of availability of data
- The general principle is inspired by how people learn optimally — learning is most effective and efficient when examined problems are sufficiently different from what we know so far
- Can be used for a variety of learning problems:
  - **Parameter learning** (e.g. when trying to find optimal parameters for cutting vegetables)
  - **Model learning** (e.g. our motivating skill learning problem, where a success model is represented as a Gaussian process)

Why Active Learning?

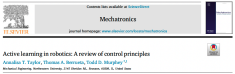
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey [*,1]
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- Applicable when a robot needs to learn by exploration due to the lack of availability of data
- The general principle is inspired by how people learn optimally — learning is most effective and efficient when examined problems are sufficiently different from what we know so far
- Can be used for a variety of learning problems:
  - **Parameter learning** (e.g. when trying to find optimal parameters for cutting vegetables)
  - **Model learning** (e.g. our motivating skill learning problem, where a success model is represented as a Gaussian process)
  - **Feature learning** (finding a suitable low-level representation of high-dimensional data, which can be used for subsequent learning)

Active vs. Reinforcement Learning

- Both active and reinforcement learning enable a robot to learn from exploration by interacting with the environment
- There are some differences between the two though:

Active vs. Reinforcement Learning

- Both active and reinforcement learning enable a robot to learn from exploration by interacting with the environment
- There are some differences between the two though:

Reinforcement learning

- The objective is to learn a policy so that the robot can act accordingly — learning is performed from interaction sequences
- Exploration is done partially randomly (in off-policy methods) or by following the robot's current policy (in on-policy methods)

Active vs. Reinforcement Learning

- Both active and reinforcement learning enable a robot to learn from exploration by interacting with the environment
- There are some differences between the two though:

Reinforcement learning

- The objective is to learn a policy so that the robot can act accordingly — learning is performed from interaction sequences
- Exploration is done partially randomly (in off-policy methods) or by following the robot's current policy (in on-policy methods)

Active learning

- A policy can be learned, but so can other models (as mentioned later)
- The selection of exploration experiences is guided by an information model, such that the most informative examples are selected to be explored

Active Learning and Deep Learning

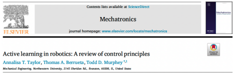
ELSEVIER
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey *,1
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- **Deep learning is generally not suitable for being used in an active learning framework**:

---

[1] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. 34th Int. Conf. Machine Learning*, 2017, pp. 1183–1192. Available: https://dl.acm.org/doi/10.5555/3305381.3305504

Active Learning and Deep Learning

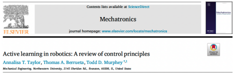
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- **Deep learning is generally not suitable for being used in an active learning framework**:
  - Deep learning models are trained on large datasets; active learning is performed when a large dataset is not available or is costly to collect (as is often the case in robotics)

[1] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. 34th Int. Conf. Machine Learning*, 2017, pp. 1183–1192. Available: https://dl.acm.org/doi/10.5555/3305381.3305504

Active Learning and Deep Learning

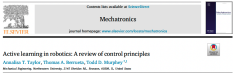
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey *,1
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- **Deep learning is generally not suitable for being used in an active learning framework**:
  - Deep learning models are trained on large datasets; active learning is performed when a large dataset is not available or is costly to collect (as is often the case in robotics)
  - Active learning requires a way of estimating the quality / informativeness of examples so that it can select the most informative one; deep networks are often overconfident about their estimates

[1] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. 34th Int. Conf. Machine Learning*, 2017, pp. 1183–1192. Available: https://dl.acm.org/doi/10.5555/3305381.3305504

Active Learning and Deep Learning

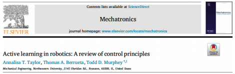
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- **Deep learning is generally not suitable for being used in an active learning framework**:
  - Deep learning models are trained on large datasets; active learning is performed when a large dataset is not available or is costly to collect (as is often the case in robotics)
  - Active learning requires a way of estimating the quality / informativeness of examples so that it can select the most informative one; deep networks are often overconfident about their estimates
- While traditional deep neural networks are not particularly compatible with active learning, some attempts at doing active learning with Bayesian neural networks exist[1]

[1] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. 34th Int. Conf. Machine Learning*, 2017, pp. 1183–1192. Available: https://dl.acm.org/doi/10.5555/3305381.3305504

Uses of Active Learning

Active Perception

Active Perception

RUZENA BAJCSY, MEMBER, IEEE

Invited Paper

- The problem of active learning is closely related to that of active perception, where a robot attempts to **find the best next viewpoint so that it can maximise the information of incoming perceptual information**

Active Perception

Active Perception

RUZENA BAJCSY, MEMBER, IEEE

Invited Paper

- The problem of active learning is closely related to that of active perception, where a robot attempts to **find the best next viewpoint so that it can maximise the information of incoming perceptual information**
- Active perception is also a very biologically inspired approach to perception (visual or otherwise), for instance:
  - When observing a scene visually, we change our viewpoint so that we get a more complete view of the scene
  - When exploring an object using tactile information, we move along informative parts of the object, such as edges

Active Perception

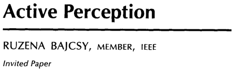

- The problem of active learning is closely related to that of active perception, where a robot attempts to **find the best next viewpoint so that it can maximise the information of incoming perceptual information**
- Active perception is also a very biologically inspired approach to perception (visual or otherwise), for instance:
  - When observing a scene visually, we change our viewpoint so that we get a more complete view of the scene
  - When exploring an object using tactile information, we move along informative parts of the object, such as edges

"The problem of Active Sensing can be stated as a problem of controlling strategies applied to the data acquisition process which will depend on the current state of the data interpretation and the goal or the task of the process." (Bajcsy 1988)

Active Learning and Sim-to-Real Transfer

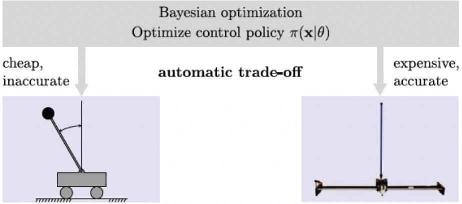

A. Marco et al., "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with Bayesian optimization," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 1557–1563. Available: https://doi.org/10.1109/ICRA.2017.7989186

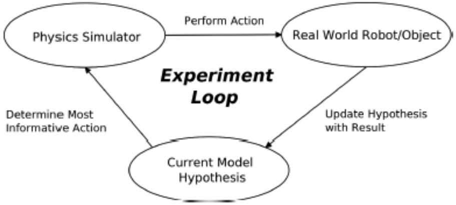

O. O. Sushkov and C. Sammut, "Active robot learning of object properties," in *Proc. IEEE/ RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2012, pp. 2621–2628. Available: https://doi.org/10.1109/IROS.2012.6385717

- A typical use of active learning is in the process of **learning a controller when both a simulation and a real system are available**

Active Learning and Sim-to-Real Transfer

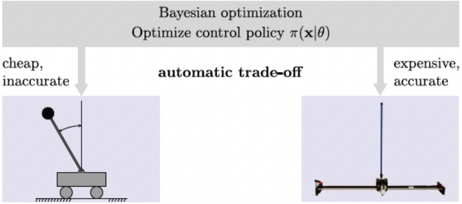

A. Marco et al., "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with Bayesian optimization," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 1557–1563. Available: https://doi.org/10.1109/ICRA.2017.7989186

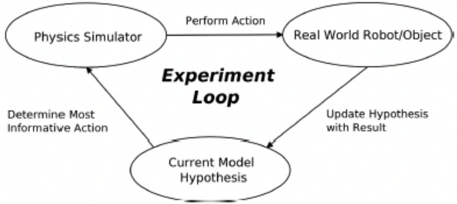

O. O. Sushkov and C. Sammut, "Active robot learning of object properties," in *Proc. IEEE/ RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2012, pp. 2621–2628. Available: https://doi.org/10.1109/IROS.2012.6385717

- A typical use of active learning is in the process of **learning a controller when both a simulation and a real system are available**
  - Simulations are cheap to perform compared to real-robot experiments; real-robot trials should be carefully chosen so that they provide as much information as possible

Active Learning and Sim-to-Real Transfer

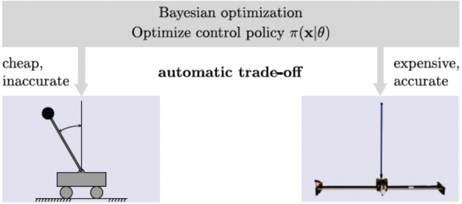

A. Marco et al., "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with Bayesian optimization," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 1557–1563. Available: https://doi.org/10.1109/ICRA.2017.7989186

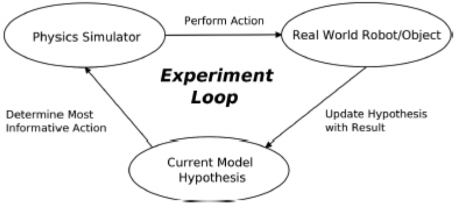

O. O. Sushkov and C. Sammut, "Active robot learning of object properties," in *Proc. IEEE/ RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2012, pp. 2621–2628. Available: https://doi.org/10.1109/IROS.2012.6385717

- A typical use of active learning is in the process of **learning a controller when both a simulation and a real system are available**
  - Simulations are cheap to perform compared to real-robot experiments; real-robot trials should be carefully chosen so that they provide as much information as possible

- In this case, the simulation can be used to **perform predictive experiments**, **identify an experiment that is most ambiguous**, and then **execute that one on the real system**

Active Learning and Human-Robot Interaction

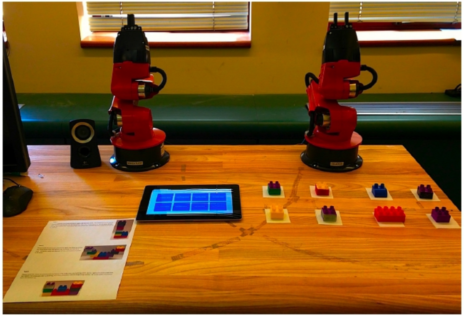

B. Hayes and B. Scassellati, "Discovering task constraints through observation and active learning," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2014, pp. 4442–4449. Available: https://doi.org/10.1109/IROS.2014.6943191

- Active learning can be used to improve the efficiency of autonomous learning, but can **also be used to enhance the process of interactive learning**

Active Learning and Human-Robot Interaction

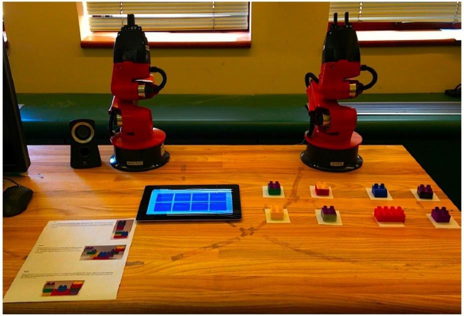

B. Hayes and B. Scassellati, "Discovering task constraints through observation and active learning," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2014, pp. 4442–4449. Available: https://doi.org/10.1109/IROS.2014.6943191

- Active learning can be used to improve the efficiency of autonomous learning, but can **also be used to enhance the process of interactive learning**
- During interactive learning, active learning can be used to **pose informative queries to a teacher**, as opposed to posing undirected questions

Active Learning and Human-Robot Interaction

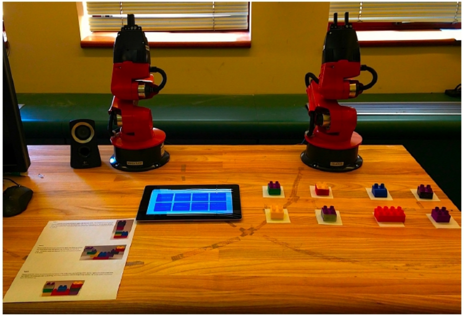

B. Hayes and B. Scassellati, "Discovering task constraints through observation and active learning," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2014, pp. 4442–4449. Available: https://doi.org/10.1109/IROS.2014.6943191

- Active learning can be used to improve the efficiency of autonomous learning, but can **also be used to enhance the process of interactive learning**
- During interactive learning, active learning can be used to **pose informative queries to a teacher**, as opposed to posing undirected questions
- **Human feedback is precious, but expensive** — too many questions or too undirected questions can frustrate a user; a robot that models a good student is likely to be most pleasant to use and teach

Information Metrics Used in Active Learning

Information Metrics

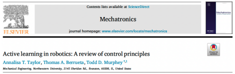
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- To select experiences during active learning, a robot needs to be able to **calculate the value of experiences and select experiences that provide the highest value** (the largest amount of information)
- A variety of metrics can be used for this purpose, but we will mention two that are commonly seen in practical implementation:
  - **Information gain**
  - **Kullback-Leibler divergence (relative entropy)**
- The used information metric is a design choice (just as in the context of loss functions in deep learning)

Information Gain

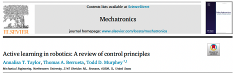
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

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

- **The entropy of an event measures the information that is provided by the event**

Information Gain

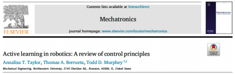
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

ELSEVIER

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

- **The entropy of an event measures the information that is provided by the event**
- The entropy of a probability distribution is defined as:

Information Gain

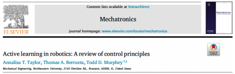
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- **The entropy of an event measures the information that is provided by the event**
- The entropy of a probability distribution is defined as:

For discrete distributions

$$H(X) = -\sum_{x} P(x) \log \left(P(x)\right)$$

Information Gain

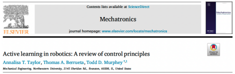
Contents lists available at ScienceDirect
Mechatronics
journal homepage: www.elsevier.com/locate/mechatronics
ELSEVIER
Active learning in robotics: A review of control principles
Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey [*,1]
Mechanical Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, 60208, IL, United States

- **The entropy of an event measures the information that is provided by the event**
- The entropy of a probability distribution is defined as:

For discrete distributions

$$H(X) = -\sum_{x} P(x) \log \left(P(x)\right)$$

For continuous distributions

$$H(X) = -\int_{-\infty}^{\infty} p(x) \log \left(p(x)\right) dx$$

Information Gain

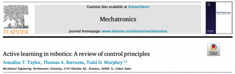
Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

Active learning in robotics: A review of control principles

Annalisa T. Taylor, Thomas A. Berrueta, Todd D. Murphey

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

- **The entropy of an event measures the information that is provided by the event**
- The entropy of a probability distribution is defined as:

For discrete distributions

$$H(X) = -\sum_{x} P(x) \log\left(P(x)\right)$$

For continuous distributions

$$H(X) = -\int_{-\infty}^{\infty} p(x) \log\left(p(x)\right) dx$$

- **An event x that maximises the information gain is a good candidate to be explored**:

$$\text{gain}(X, x) = \arg\max_{x} H(X) - H(X|x)$$

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure
- The KL divergence is defined as:

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure
- The KL divergence is defined as:

For discrete distributions

$$D_{KL}(P, Q) = \sum_{x} P(x) \ln \left(\frac{P(x)}{Q(x)} \right)$$

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure
- The KL divergence is defined as:

For discrete distributions

$$D_{KL}(P, Q) = \sum_x P(x) \ln \left(\frac{P(x)}{Q(x)} \right)$$

For continuous distributions

$$D_{KL}(P, Q) = \int_{-\infty}^{\infty} p(x) \ln \left(\frac{p(x)}{q(x)} \right) dx$$

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure

- The KL divergence is defined as:

For discrete distributions

$$D_{KL}(P,Q) = \sum_x P(x) \ln \left(\frac{P(x)}{Q(x)} \right)$$

For continuous distributions

$$D_{KL}(P,Q) = \int_{-\infty}^{\infty} p(x) \ln \left(\frac{p(x)}{q(x)} \right) dx$$

- The larger the divergence between the current belief distribution and a distribution after collecting a certain datum, the greater the benefit of that datum

Kullback-Leibler Divergence (Relative Entropy)

- **The Kullback-Leibler (KL) divergence is a measure of discrepancy between two different probability distributions** and can be used as an information gain measure
- The KL divergence is defined as:

For discrete distributions

$$D_{KL}(P,Q) = \sum_x P(x) \ln\left(\frac{P(x)}{Q(x)}\right)$$

For continuous distributions

$$D_{KL}(P,Q) = \int_{-\infty}^{\infty} p(x) \ln\left(\frac{p(x)}{q(x)}\right) dx$$

- The larger the divergence between the current belief distribution and a distribution after collecting a certain datum, the greater the benefit of that datum
- Note: The KL divergence is an asymmetric measure

Active Learning and Motivation

- The objective of active learning — that of finding experiences that maximise an information measure — can be seen as equivalent, or at least **compatible, with that of acting by maximising a robot's motivation**

Active Learning and Motivation

- The objective of active learning — that of finding experiences that maximise an information measure — can be seen as equivalent, or at least **compatible, with that of acting by maximising a robot's motivation**

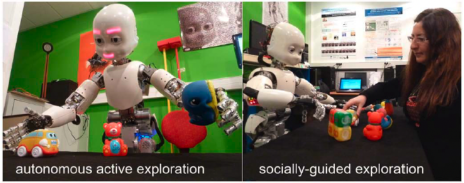

S. Ivaldi et al., "Object Learning Through Active Exploration," in *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 1, pp. 56–72, Mar. 2014. Available: https://doi.org/10.1109/TAMD.2013.2280614

- Some implementations of active learning — in the context of developmental learning — make use of this connection by enabling motivation-driven experience selection for both autonomous and interactive learning

Active Learning and Motivation

- The objective of active learning — that of finding experiences that maximise an information measure — can be seen as equivalent, or at least **compatible, with that of acting by maximising a robot's motivation**

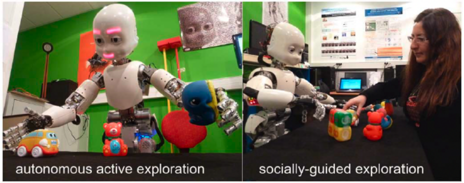

S. Ivaldi et al., "Object Learning Through Active Exploration," in *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 1, pp. 56–72, Mar. 2014. Available: https://doi.org/10.1109/TAMD.2013.2280614

- Some implementations of active learning — in the context of developmental learning — make use of this connection by enabling motivation-driven experience selection for both autonomous and interactive learning

- In fact, **information measures can be thought of computational models of curiousness** — and satisfying curiosity is a motivating factor for many individuals

Summary: Active Learning

- Active learning is a paradigm based on which a robot takes an active role during the learning process, by choosing experiences that are maximally informative and thus likely to provide a useful update to its model
- Active learning can be used for a variety of purposes, such as in conjunction with active perception, to achieve sim-to-real policy transfer, or for task learning from human feedback
- The main benefit of active learning compared to other learning paradigms is that it can be used in scenarios with limited data availability
- The selection of examples to consider during active learning is generally guided by an information measure, such as the information gain or the Kullback-Leibler divergence