



# The value of abstraction

Mark K Ho<sup>1</sup>, David Abel<sup>2</sup>, Thomas L Griffiths<sup>1</sup> and Michael L Littman<sup>2</sup>

## Addresses

<sup>1</sup> Princeton University, Princeton, NJ, United States

<sup>2</sup> Brown University, Providence, RI, United States

Corresponding author: Littman, Michael L ([mlittman@cs.brown.edu](mailto:mlittman@cs.brown.edu))

Current Opinion in Behavioral Sciences 2019, 29:111–116

This review comes from a themed issue on **Artificial Intelligence**

Edited by **Matt Botvinick** and **Sam Gershman**

<https://doi.org/10.1016/j.cobeha.2019.05.001>

2352-1546/© 2019 Elsevier Ltd. All rights reserved.

## Introduction

In spite of bounds on space, time, and data, people are able to make good decisions in complex scenarios. What enables us to do so? And what might equip artificial systems to do the same? One essential ingredient for making complex problem solving tractable is a capacity for exploiting *abstract representation*.

To illustrate why abstractions are needed for adaptive decision making, consider hiking in a forest (**Figure 1**). A river runs from mountains in the north down south through the forest, and you have set up camp in a clearing just east of the river. As the sun begins to set, you find yourself west of the river and want to return to camp. How do you accomplish this task? You might reason, “Since I am on the *west side of the river*, I probably want to *cross the river* to get to the *east side of the river* and then *walk towards my campsite*.” Such a compact thought is possible because you have abstract representations for that support decision making: These include the general notion of being in a location relative to the axis of the river, or the rough ideas of crossing the river and moving towards the campsite.

Without abstractions, decision making would be confined to specific, low-level states and actions—e.g., “At the current state, rotate your left leg 25°, place it down beside the rock on the path, then swing your arm forward...”. In a huge and complex scenario like navigating through a forest, this approach is both cognitively unwieldy and, in a fundamental sense, computationally intractable [1,2]. The difficulty is due to decision-making costs that arise from two interrelated sources. First, agents need to learn about their environment, which costs time and experience (e.g., via exploration) [3].

Second, agents need to compute decision-theoretic quantities, which costs memory and thought (e.g., via planning) [4]. As problems grow, these costs grow exponentially. Abstractions are essential because they enable decision makers to manage the growth of these costs.

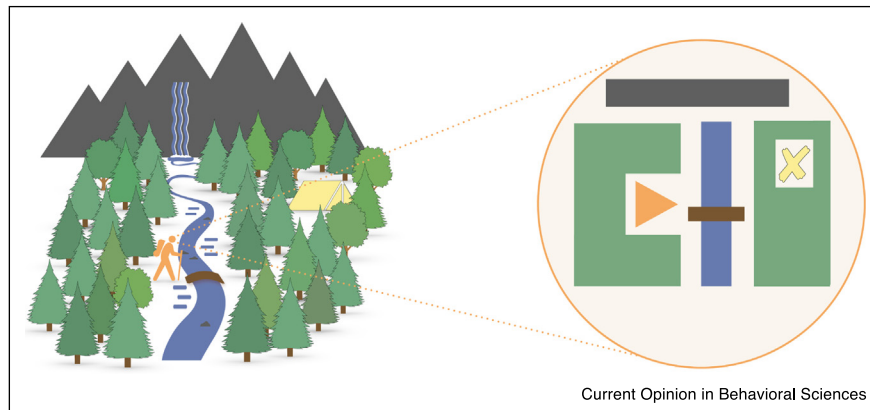
Within psychology and neuroscience, abstraction plays a key role in the hierarchical learning and organization of motor sequences [5–7], habits [8–10], and planning [11,12]. For this reason, researchers in artificial intelligence (AI) and reinforcement learning (RL) have long been interested in how to leverage abstraction to provide human-like solutions to human-level problems [13–18]. In this paper, we discuss work in AI and RL that helps elucidate how abstractions support efficient decision making in both people and machines.

Research in AI and RL focuses on two broad types of abstraction. First, *state abstractions* treat certain configurations of the environment as similar by aggregating them or assigning them shared features. For instance, in the hiking example, *being west of the river* is not a single concrete configuration of the environment, but rather an abstract representation that covers many concrete states (e.g., being different distances west of the river). Second, *temporal abstractions* (often called “options”) are temporally extended macro-actions that describe a general course of action. For example, the idea of *crossing the bridge* is not a single specific action, but rather captures many possible courses of action.

State and temporal abstractions enable compact representation of a domain. For instance, **Figure 2a** represents the hiking example as a graph where nodes are specific configurations of the environment (i.e., ground states) and edges are transitions to new configurations resulting from taking actions. The appropriate abstractions could induce the simpler problem representation in **Figure 2b**. This new model ignores irrelevant details like the color of specific trees while capturing useful distinctions such as your location relative to the river. As a result, it supports efficient learning and requires little thought to reason about compared to the original version.

The foregoing discussion sketches out how abstractions can support efficient decision making. But how could it work in practice? Here, we dive into recent work in AI and RL that provides some answers. Specifically, we focus on three ways in which abstractions have been shown to facilitate efficient and scalable decision making. First, abstractions *guide exploration and generalization* by systematically modifying the distribution of learning experiences. In particular, abstractions can guide how agents

Figure 1



An environment, an agent, and the agent's abstract representation of a domain.

explore and generalize based on environmental structure or representational simplicity. Second, abstractions *facilitate efficient tradeoffs in learning*. For example, they enable learning time and optimality to be exchanged as well as support optimal transfer between tasks. Finally, abstraction is essential for *simplifying computation*. For instance, abstractions influence the cost of computing a good plan by inducing simpler or more complex planning models. In the following sections, we discuss each of these benefits of abstraction.

### Abstractions guide exploration and generalization

Suppose during your camping trip you want to learn to fish. At first, you try a few spots along the river at random. But soon, you notice a pattern: Certain areas have more vegetation, others have less. This information provides a new and more efficient way for you to organize your fishing attempts. For example, you could fish areas with high and low vegetation to gain a range of experiences about how it affects your catch. Or, on your first trip you may learn that your best catches were in high vegetation areas, so that on your second trip to a different river you seek out similar areas. Here, the abstract concept of river vegetation guides your exploration in the current fishing task and tracks a generalizable feature relevant to future tasks, which both allow you to make better use of your limited time and experience.


In RL, abstractions similarly facilitate efficient learning by guiding exploration and generalization. But what is the basis of this guidance? Put another way, the concept of vegetation is useful, but what determines the identification of such a concept in general? Below, we discuss two ways in which abstractions can guide learning: domain structure and representational simplicity.

### Domain structure

Many domains have meaningful structure based on the topology of state connections. In the hiking example, *being west of the river* is an intuitive abstract state because all the locations west of the river are more connected to one another and less connected to locations east of the river. Early research in RL identified a number strategies for learning abstractions based on clustering or graph-theoretic notions like “bottleneck” states [19]. More recent methods generalize such approaches to learn abstractions based on the state connectivity induced by a domain’s structure. For instance, the *successor representation* [20–22,69–71] and *successor features* [23,24] are state abstractions based on how likely an agent is to visit other states or features. A similar idea underpins *eigen-options* [25,26,27]. Specifically, both successor representations and eigen-options involve decomposing a domain’s structure into a set of “principal components” that compactly represent how certain regions of a domain are connected or will be visited.

Thus, *being west of the river* and *being east of the river* could emerge as successor-based state abstractions due to the connectivity and visitation of their respective ground states. A corresponding eigen-option would express the temporally abstract action of *going from the west side to the east side* since action sequences falling in this category capture a principal source of variance in connectivity. Notably, abstractions based on state connectivity encode information about the dynamics of a domain separate from its reward structure. These reward-agnostic representations can be learned either offline or during task learning, but, in either case, they provide an especially powerful and efficient way to adapt when rewards change or when transferring to tasks with related dynamics.

Figure 2



(a) A domain in all of its complexity, and (b) the representation induced by an effective set of state and temporal abstractions for this domain. The representation in (b) facilitates efficient learning and computation, unlike the one in (a).

### Representational simplicity

Abstractions also guide exploration and generalization based on representational simplicity. Simplicity plays a key role in understanding representations in humans [28,29] and machines [30]. For finite decision makers, simplicity is important because it enables *compression*, the representation of information using less physical space (e.g., memory).

In RL, a number of methods have been developed that leverage a bias towards representational simplicity. For example, *curiosity-driven learning* [31–33,34,35,36] pits a motivation to learn a simple, compressed representation of the world against a desire for new experiences that do not easily fit within that representation. Placing these two processes in competition with one another prevents the drive for simplicity from compressing everything into a single undifferentiated representation that cannot support high quality decisions and instead pushes the system to adaptively explore a range of qualitatively different parts of a task. Doing so ultimately allows the architecture to both represent the world compactly and explore it efficiently.

Simplicity also plays a role in the Option-Critic framework [37,38,39], a neural network architecture that simultaneously learns abstract macro-actions and the value of low-level actions that constitute them. By concurrently learning at multiple levels of abstraction, the agent acquires a representation that leads to better transfer to new tasks. The process at play is closely related to the notion of a “blessing of abstraction” in causal learning [40] and regularization in supervised learning [41]: A bias

towards simple, unifying representations not only saves on representational space, but also prevents overfitting to a particular task. An important consequence of this bias is that the agent is able to transfer to new tasks more efficiently.

### Abstractions facilitate efficient tradeoffs

Decision making and learning involve tradeoffs when space, time, and data are limited. But, it is important that these tradeoffs be made *efficiently* (i.e., without unnecessary costs). Here, we discuss how abstractions can efficiently trade off learning time and optimality, as well as minimize negative transfer between tasks.

#### Learning time and optimality

Imagine that every morning you wake up at your campsite and need to navigate to the bridge crossing the river. What kind of abstract states would be useful? It depends on how much time you have to spend learning. One option is to learn every tree, rock, and shortcut through the forest in great detail. Doing so has the benefit of ensuring you learn a highly efficient route to the bridge. An alternative strategy is to learn a less detailed representation of the forest that takes you to an obvious landmark—e.g., the river—and then take a simple path to the bridge. This second approach may be more costly in terms of the physical actions taken or time navigating, but is simpler and faster to learn. Depending on how much time and data are available, it may be better for the agent to pursue the second strategy than the first.

Learning the simpler but suboptimal policy depends on the resolution of the state abstractions used—i.e., the

amount of detail one is willing to keep or ignore about the environment. In RL, a number of approaches use abstraction in this manner to efficiently balance optimality and learning time [42\*,43]. For instance, if abstractions are represented by aggregating ground states into clusters, adaptively modulating the size of the clusters and granularity of the state representation controls this tradeoff [44,33,45].

### Optimizing transfer between tasks

Abstractions also facilitate effective transfer from one task to another. Transfer can be positive [46,47,37\*\*], but it can also be negative [48]. Ideally, an agent would learn abstractions that avoid negative transfer as much as possible while increasing the possibility of positive transfer. One setting in which this approach is feasible is in *lifelong learning*, in which an agent is assumed to be continuously receiving tasks from a distribution of tasks. Given separate phases of discovering and using temporal abstractions, it is possible to ensure that any learned options only help later learning and decision making [49]. Used properly, temporal abstractions can also guarantee that the amount of data needed for learning is reduced even within a single task [50,51,72]. For agents with limited time, such abstractions can play a critical role in ensuring efficient and effective use of data.

### Abstractions simplify computation

Computation is costly for people [52] as well as artificial agents [53]. For example, if you were to plan a hike back to your campsite, you could simulate possible routes from your current location or “work backwards” from your final destination. These cognitive operations require time and mental effort. However, the total cost of computation depends directly on the model such operations occur over [16]. Planning can be easier given a good abstract representation (e.g., compare Figure 2a and Figure 2b). Recent work in RL leverages this intuition to efficiently scale planning to complex domains. Here, we focus on two broad approaches: *algorithm-specific abstractions*, in which an abstraction scheme is tailored to take advantage of particular features of a planning algorithm, and *end-to-end* approaches, in which learning abstractions for planning occurs in the context of a more general learning architecture.

### Algorithm-specific abstractions

Algorithms such as Monte Carlo Tree Search [54–56] are based on the notion of *forward simulation*: From an initial state, candidate plans are generated, evaluated, and the best is chosen. The major problem for these algorithms is branching: Each additional timestep into the future increases the number of plans to evaluate multiplicatively. To keep branching manageable, state-aggregation techniques simplify models by only storing relevant information. For instance, by treating all states with the same optimal action as identical, branching due to

stochasticity in action outcomes can be reduced [57–59]. Along similar lines, the model used for planning can be simplified by aggregating states that lead to the same possible outcomes within a few time steps [60].

An alternative to forward simulation is *backwards induction*, which involves starting from possible future states and propagating information backwards to possible precursor states. Abstractions can also assist computation in backward induction. For instance, value iteration [1] repeatedly “sweeps” over a set of states and backs up information to each precursor state. Since each iteration involves looking at every state in the planning model, it is extremely sensitive to the number of possible states. However, by incorporating temporal abstractions into the planning model, more information can be propagated further backwards on each sweep [61] and it becomes possible to more efficiently trade off error in the approximation of future rewards with computation [62].

Finally, relaxing the goal of planning provides opportunities to leverage powerful abstractions. For instance, *plan feasibility* seeks to determine whether a plan can be constructed that satisfies some set of constraints. Since the goal is no longer *optimality* but simply *satisfiability*, relevant aspects of the task can be translated into symbolic logic and efficiently tested [63–65].

### End-to-end efficient planning

Another approach is to learn planning abstractions while also learning everything else about a task in an *end-to-end* manner with neural networks. Using abstractions for planning is especially important because planning in a ground state space (e.g., the pixels of a video game) is only possible for short time horizons [66]. To address these limitations, architectures have been developed that learn and use abstract representations for planning. In particular, the abstract representation needs to be trained to predict future rewards, thus preserving the most relevant information for decision making in the planning model [67\*,68,73].

### Conclusion

Abstractions are key for efficiently scaling decision making. In particular, recent research in RL demonstrates how state and temporal abstractions guide exploration and generalization, facilitate efficient tradeoffs, and simplify computations. In the context of learning and decision making with limited resources, abstractions enable agents to strategically manage finite space, time, and data, which is necessary for scaling problem solving to complex tasks.

For psychologists, this analysis raises questions about why biological agents like humans have certain types of abstract representations. For instance, humans and other organisms may have certain abstractions not only because



they facilitate effective exploration and generalization but because they support fast planning or modulate critical tradeoffs during learning. Future work in both computer science and psychology will need to identify other pressures that can shape abstraction learning—such as the need to communicate and coordinate with others—to offer a clearer understanding of the value of abstraction.

## Conflict of interest statement

The authors declare no conflict of interest.

## References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. Bellman R: *Dynamic programming*. Princeton University Press; 1957.
  2. Bylander T: **The computational complexity of propositional STRIPS planning**. *Artif Intel* 1994, **69**:161-204.
  3. Kakade SM: 2003. On the sample complexity of reinforcement learning, (Unpublished doctoral dissertation), University of London, England.
  4. Littman ML, Dean TL, Kaelbling LP: **On the complexity of solving Markov decision problems..** *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* 1995:394-402.
  5. Lashley KS: *The problem of serial order in behavior. Cerebral mechanisms in behavior; the Hixon symposium*. Oxford, England: Wiley; 1951, 112-146.
  6. Rosenbaum DA, Kenny SB, Derr MA: *Hierarchical control of rapid movement sequences. (Vol. 9) (No. 1)*. US: American Psychological Association; 1983 <http://dx.doi.org/10.1037/0096-1523.9.1.86>.
  7. Rosenbaum DA, Inhoff AW, Gordon AM: **Choosing between movement sequences: A hierarchical editor model**. *J Exp Psychol: Gen* 1984, **113**(3):372-393 <http://dx.doi.org/10.1037/0096-3445.113.3.372>.
  8. Botvinick MM: **Hierarchical models of behavior and prefrontal function**. *Trends Cognit Sci* 2008, **12**(5):201-208 <http://dx.doi.org/10.1016/J.TICS.2008.02.009>.
  9. Dezfouli A, Balleine BW: **Actions, action sequences and habits: Evidence that goal-directed and habitual action control are hierarchically organized**. *PLoS Comput Biol* 2013, **9**(12): e1003364 <http://dx.doi.org/10.1371/journal.pcbi.1003364>.
  10. Ribas-Fernandes J, Solway A, Diuk C, McGuire J, Barto A, Niv Y, Botvinick M: **A neural signature of hierarchical reinforcement learning**. *Neuron* 2011, **71**(2):370-379 <http://dx.doi.org/10.1016/J.NEURON.2011.05.042>.
  11. Solway A, Diuk C, Córdova N, Yee D, Barto AG, Niv Y, Botvinick MM: **Optimal behavioral hierarchy**. *PLoS Comput Biol* 2014, **10**(8):e1003779 <http://dx.doi.org/10.1371/journal.pcbi.1003779>.
  12. Botvinick MM, Weinstein A: **Model-based hierarchical reinforcement learning and human action control**. *Philos. Trans. R. Soc. B: Biol. Sci.* 2014, **369**(1655):20130480 <http://dx.doi.org/10.1098/rstb.2013.0480>.
  13. Dayan P, Hinton GE: **Feudal reinforcement learning**. *Advances in Neural Information Processing Systems* 1993:271-278.
  14. Sacerdoti ED: **Planning in a hierarchy of abstraction spaces**. *Artif Intel* 1974, **5**(2):115-135 [http://dx.doi.org/10.1016/0004-3702\(74\)90026-5](http://dx.doi.org/10.1016/0004-3702(74)90026-5).
  15. Dietterich TG: **Hierarchical reinforcement learning with the MAXQ value function decomposition**. *J Artif. Intel. Res.* 2000, **13**:227-303.
  16. Givan R, Dean T, Greig M: **Equivalence notions and model minimization in Markov decision processes**. *Artif Intel* 2003, **147**(1-2):163-223.
  17. Andre D, Russell SJ: **State abstraction for programmable reinforcement learning agents**. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence; Menlo Park, CA, USA: American Association for Artificial Intelligence: 2002:119-125*.
  18. Barto AG, Mahadevan S: **Recent advances in hierarchical reinforcement learning**. *Discrete Event Dynamic Systems* 2003, **13**(1-2):41-77.
  19. Şimşek Ö, Barto AG: **Using relative novelty to identify useful temporal abstractions in reinforcement learning**. In *Proceedings of the Twenty-first International Conference on Machine Learning '04; New York, New York, USA: ACM Press. 95* doi:10.1145/1015330.1015353: 2004.
  20. Dayan P: **Improving generalization for temporal difference learning: The successor representation**. *Neural Comput* 1993, **5**(4):613-624.
  21. Gershman SJ, Moore CD, Todd MT, Norman KA, Sederberg PB: **The successor representation and temporal context**. *Neural Comput* 2012, **24**(6):1553-1568.
  22. Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ: **The successor representation in human reinforcement learning**. *Nat Hum Behav* 2017, **1**(9):680-692 <http://dx.doi.org/10.1038/s41562-017-0180-8>.
  23. Barreto A, Dabney W, Munos R, Hunt JJ, Schaul T, van Hasselt HP, Silver D: **Successor features for transfer in reinforcement learning**. *Advances in Neural Information Processing Systems* 2017:4055-4065.
- Introduces a generalization of successor representations [20] that enables efficient transfer across tasks with shared causal structure.
24. Kulkarni TD, Saeedi A, Gautam S, Gershman SJ: *Deep successor reinforcement learning*. 2016. arXiv preprint arXiv:1606.02396.
  25. Machado MC, Bellemare MG, Bowling M: *A Laplacian framework for option discovery in reinforcement learning*. 2017. arXiv preprint arXiv:1703.00956.
  26. Machado MC, Rosenbaum C, Guo X, Liu M, Tesauro G, Campbell M: *Eigenoption discovery through the deep successor representation*. 2017. arXiv preprint arXiv:1710.11089.
- Combines learning eigen-options with learning successor representations, which both leverage the transition structure of a domain for learning abstract representations.
27. Mahadevan S: **Proto-value functions: Developmental reinforcement learning**. *Proceedings of the 22nd International Conference on Machine Learning* 2005:553-560.
  28. Chater N, Vitányi P: **Simplicity: A unifying principle in cognitive science?** *Trends Cognit Sci* 2003, **7**(1):19-22 [http://dx.doi.org/10.1016/S1364-6613\(02\)00005-0](http://dx.doi.org/10.1016/S1364-6613(02)00005-0).
  29. Sanborn S, Bourgin DD, Chang M, Griffiths TL: **Representational efficiency outweighs action efficiency in human program induction**. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*. Edited by Rogers T, Rau M, Zhu X, Kalish C. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*Austin, TX: Cognitive Science Society: 2018:2400-2405.
  30. Li M, Vitányi P: *An Introduction to Kolmogorov Complexity and Its Applications*. New York, NY: Springer New York; 2008 <http://dx.doi.org/10.1007/978-0-387-49820-1>.
  31. Schmidhuber J: **A possibility for implementing curiosity and boredom in model-building neural controllers**. *Proc. of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats* 1991:222-227.
  32. Oudeyer P-Y, Kaplan F, Hafner VV: **Intrinsic motivation systems for autonomous mental development**. *IEEE Trans Evol Comput* 2007, **11**(2):265-286.
  33. Bellemare M, Srinivasan S, Ostrovski G, Schaul T, Saxton D, Munos R: **Unifying count-based exploration and intrinsic**

- motivation.** *Advances in Neural Information Processing Systems* 2016:1471-1479.
34. Pathak D, Agrawal P, Efros AA, Darrell T: **Curiosity-driven exploration by self-supervised prediction.** *Proceedings of the 34th International Conference on Machine Learning* 2017:2778-2787. JMLR. Implements curiosity as an intrinsic reward signal based on an agent's error when predicting the outcome of actions to enable effective representation learning even in the absence of external rewards.
35. Ostrovski G, Bellemare MG, Oord A, Munos R: **Count-based exploration with neural density models.** *Proceedings of the International Conference on Machine Learning* 2017:2721-2730.
36. Kulkarni TD, Narasimhan K, Saeedi A, Tenenbaum J: **Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation.** *Advances in Neural Information Processing Systems* 2016:3675-3683.
37. Bacon P-L, Harb J, Precup D: **The option-critic architecture.** *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence* 2017:1726-1734. Proposes *option-critic* architecture that demonstrates the effectiveness of jointly learning abstract and low-level policies.
38. Harb J, Bacon P-L, Klissarov M, Precup D: **When waiting is not an option: Learning options with a deliberation cost.** *Thirty-Second Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence* 2018.
39. Liu M, Machado MC, Tesauro G, Campbell M: *The eigenoption-critic framework.* 2017. arXiv preprint arXiv:1712.04065.
40. Goodman ND, Ullman TD, Tenenbaum JB: **Learning a theory of causality.** *Psychological Review* 2011, **118(1)**:110.
41. Bishop CM: *Pattern Recognition and Machine Learning.* Springer; 2006.
42. Jiang N, Kulesza A, Singh S: **Abstraction selection in model-based reinforcement learning.** In *Proceedings of the 32nd International Conference on Machine Learning.* Edited by Bach F, Blei D. *Proceedings of the 32nd International Conference on Machine Learning, Lille, France: PMLR: 2015:179-188.* Analyzes how state-aggregation-learning schemes can be used to efficiently trade off data and solution quality, which is key when data is costly or finite.
43. Ortner R: **Adaptive aggregation for reinforcement learning in average reward Markov decision processes.** *Ann Oper Res* 2013, **208(1)**:321-336.
44. Taiga AA, Courville A, Bellemare M: **Approximate exploration through state abstraction.** *International Conference on Machine Learning Workshop on Exploration in RL* 2018.
45. Mandel T, Liu Y-E, Brunskill E, Popovic Z: **Efficient Bayesian clustering for reinforcement learning.** *Proceedings of the International Joint Conference on Artificial Intelligence* 2016.
46. Sutton RS, Precup D, Singh S: **Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning.** *Artif Intel* 1999, **112(1)**:181-211.
47. Konidaris G, Barto AG: **Building portable options: Skill transfer in reinforcement learning.** *Proceedings of the International Joint Conference on Artificial Intelligence* 2007:895-900.
48. Jong NK, Hester T, Stone P: **The utility of temporal abstraction in reinforcement learning.** *Proceedings of the Conference on Autonomous Agents and Multiagent Systems* 2008:299-306.
49. Brunskill E, Li L: **PAC-inspired option discovery in lifelong reinforcement learning.** *Proceedings of the International Conference on Machine Learning* 2014:316-324.
50. Fruit R, Lazaric A: *Exploration-exploitation in MDPs with options.* 2017. arXiv preprint arXiv:1703.08667.
51. Fruit R, Pirota M, Lazaric A, Brunskill E: **Regret minimization in MDPs with options without prior knowledge.** *Advances in Neural Information Processing Systems* 2017:3169-3179.
52. Griffiths TL, Lieder F, Goodman ND: **Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic.** *Top Cognit Sci* 2015, **7(2)**:217-229 <http://dx.doi.org/10.1111/tops.12142>.
53. Russell SJ: **Rationality and intelligence.** *Artif Intel* 1997, **94(1-2)**:57-77 [http://dx.doi.org/10.1016/S0004-3702\(97\)00026-X](http://dx.doi.org/10.1016/S0004-3702(97)00026-X).
54. B.D. Abramson: 1987. The Expected-Outcome Model of Two-Player Games (Unpublished doctoral dissertation), New York, NY, USA. Columbia University.
55. Coulom R: **Efficient Selectivity and backup operators in Monte-Carlo Tree Search.** *Proceedings of the International Conference on Computers and Games* 2006:72-83.
56. Browne CB, Powley E, Whitehouse D, Lucas SM, Cowling PI, Rohlfshagen P, Colton S: **A survey of Monte Carlo tree search methods.** *IEEE Trans Comput Intel AI Games* 2012, **4(1)**:1-43.
57. Hostetler J, Fern A, Dietterich T: **State aggregation in Monte Carlo Tree Search.** *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence* 2014:7.
58. Anand A, Grover A, Mausam, Singla P: **ASAP-UCT: Abstraction of state-action pairs in UCT.** *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence* 2015.
59. Anand A, Noothigattu R, Mausam, Singla P: **OGA-UCT: On-the-Go abstractions in UCT.** *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling* 2016.
60. Jiang N, Singh S, Lewis R: **Improving UCT planning via approximate homomorphisms.** *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems* 2014:1289-1296.
61. Mann T, Mannor S: **Scaling up approximate value iteration with options: Better policies with fewer iterations.** *Proceedings of the International Conference on Machine Learning* 2014:127-135.
62. Mann T, Mannor S, Precup D: **Approximate value iteration with temporally extended actions.** *J Artif Intel Res* 2015, **53**:375-438.
63. G. Konidaris, L.P. Kaelbling, T. Lozano-Perez: 2014. Constructing symbolic representations for high-level planning, In: Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence, 1932-1938.
64. Konidaris G, Kaelbling LP, Lozano-Perez T: **From skills to symbols: Learning symbolic representations for abstract high-level planning.** *J Artif Intel Res* 2018, **61**:215-289.
65. McDermott D, Ghallab M, Howe A, Knoblock C, Ram A, Veloso M, Wilkins D: *PDDL - The planning domain definition language.* 1998.
66. Oh J, Guo X, Lee H, Lewis RL, Singh S: **Action-conditional video prediction using deep networks in Atari games.** *Advances in Neural Information Processing Systems* 2015:2863-2871.
67. Oh J, Singh S, Lee H: **Value prediction network.** *Advances in Neural Information Processing Systems* 2017:6118-6128. Proposes the *Value Prediction Network* architecture that learns an internal model composed of abstract states for efficient model-based planning.
68. Silver D, van Hasselt H, Hessel M, Schaul T, Guez A, Harley T, Degris T: *The predictor: End-to-end learning and planning.* Proceedings of Machine Learning Research; 2017.
69. L. Lehnert, M. L. Littman, Successor features support model-based and model-free reinforcement learning, 2019, arXiv preprint arXiv:1901.11437.
70. L. Lehnert, M. L. Littman, Transfer with model features in reinforcement learning, 2018, Lifelong Learning: A Reinforcement Learning Approach Workshop at FAIM, Stockholm, Sweden.
71. L. Lehnert, S. Tellex, and M. L. Littman, Advantages and limitations of using successor features for transfer in reinforcement learning, 2017, arXiv preprint arXiv:1708.00102.
72. Topin N, Haltmeyer N, Squire S, Winder J, desJardins M, MacGlashan J: **Portable option discovery for automated learning transfer in object-oriented Markov decision processes.** *Proceedings of the International Joint Conference on Artificial Intelligence* 2015:3856-3864.
73. Gelada C, Kumar S, Buckman J, Nachum O, Bellemare MG: **DeepMDP: Learning continuous latent space models for representation learning.** *Proceedings of the 36th International Conference on Machine Learning* 2019:2170-2179.