Title
Using Deep Reinforcement Learning to Peer Into the Unconquerable Mind: How Do Animals Learn to Track Odor Trails?

Journal
Berkeley Scientific Journal, 26(2)

ISSN
1097-0967

Author
Iyer, Anisha

Publication Date
2022

DOI
10.5070/BS326258275

Undergraduate

Using Deep Reinforcement Learning to Peer Into the Unconquerable Mind: How Do Animals Learn to Track Odor Trails?

BY ANISHA IYER

Nearly all pet owners and park frequenters have witnessed an animal searching furiously along a scent trail. Dogs and other olfaction-dependent animals are expert trail trackers, often relying on odor trails to perform life-sustaining tasks like foraging for food or navigating complex environments. However, odor trails in nature are laden with gaps, intersecting routes, and sporadic or incomplete odor cues which interfere with an animal's ability to develop a clear picture of a trail. Such constraints provide much to overcome for olfaction-dependent animals who rely almost solely on olfaction for vital behaviors. With such a computationally intensive task, one must wonder: What takes place inside an animal's brain to evaluate and optimize such a complex set of dynamic variables in real-time?

Olfaction-based trail tracking is a complex and precise behavior, which makes it a challenging investigative endeavor for theoretical neuroscientists and researchers in related fields. Scientists can simulate odor trail tracking in laboratory settings and have described experimentally observed tracking strategies using statistics and geometry. However, questions regarding the animal's cognition still remain unanswered, largely due to the difficulty of representing such far-reaching questions in a testable manner.

As a result, there is a broader goal in asking how scientists might reach conclusions about the interconnected systems of the infinitely complex brain. The quest to find ways to represent cognitive processes spans further than this specific research question and is complicated by a scientist's need for control in scientific experiments. If scientists require an observable version of the system in question and a controlled way to manipulate its variables, how might they find an equivalent for the vast and unconquerable brain?

QUESTIONS IN SYSTEMS NEUROSCIENCE

Despite centuries of neuroscience research, the precise workings of the brain's olfactory and navigational systems remain relatively unclear. Neuroscience is studied at a cellular and anatomical level, but processes like odor trail tracking require a robust theoretical framework for scientists to develop a system-wide understanding. While scientists conceptualize odor trail tracking as searching consecutive trail sectors using complex mathematical techniques, it seems unreasonable to credit an animal's impressive tracking abilities to conscious and deliberate statistical approximations. Rather, an animal following a scent must develop some kind of navigational intuition to deduce probabilities without calculating them.

Figure 1: Dog tracking odor trail

To understand these biological systems, scientists combine experimental and computational research to acquire a system-wide understanding, which is the basis of a field called systems biology. By simulating a subset of the system's components, scientists use models and theoretical exploration to glean insights into the behavior of widely expansive and intricately complex biological systems. Systems biology relies on data engineering and machine learning techniques to obtain

FEATURES

Figure 2: Art depicting neuronal migration. *The heterogeneous makeup of this image, in both artistic and symbolic cellular contexts, represents the complex task of trying to capture or recreate elements of neuroscience that we do not fully understand.*

a system-wide understanding of biological problems. For odor trail tracking, scientists have used machine learning techniques to recreate an animal's learning process as it tracks a trail by modeling the way an animal learns through trial and error in a simulated environment.[1]

MACHINE LEARNING

One of today's most rapidly growing technical fields, machine learning (ML), focuses broadly on constructing self-improving computer systems and understanding the statistical, computational, and theoretical laws that govern all learning systems on a fundamental level. Across a number of disciplines, techniques, and applications, ML models make inferences based on trends in existing data. Furthermore, these inferences, which operate by using machine-based computation, actually reflect physical neuroscientific learning processes that occur in animal life.

Throughout computer science and a wide range of other industries, ML is revolutionizing technology: In commercial sectors, ML optimizes consumer services and logistic chains, and in academia, ML serves laboratory curiosity. Accordingly, various empirical sciences have recruited ML methods to analyze high-throughput data in new ways.[2]

ML can be distilled into three branches: supervised learning, unsupervised learning, and reinforcement learning. Supervised learning (SL) is the most straightforward branch of ML, in which models make inferences by training on labeled and categorized data. SL requires a knowledgeable, external supervisor to label the training data and subsequently simplify the task for the machine. For instance, if an SL model were trained on a dataset of labeled images of cats and dogs, the model could learn to predict whether a new image is a cat or a dog, performing as a classifier.

Unsupervised learning (UL) is similar to SL in that it also requires previously obtained data, but it uses chaotic and unfiltered data that has not been labeled or categorized. ML models for unsupervised learning would recognize trends and patterns in chaotic data and use these observations to make inferences and predictions about new data. Beyond SL and UL, reinforcement learning employs supervised, unsupervised, and novel learning techniques to learn optimal strategies for success in a goal-oriented environment.

"Furthermore, these inferences, which operate by using machine-based computation, actually reflect physical neuroscientific learning processes that occur in animal life."

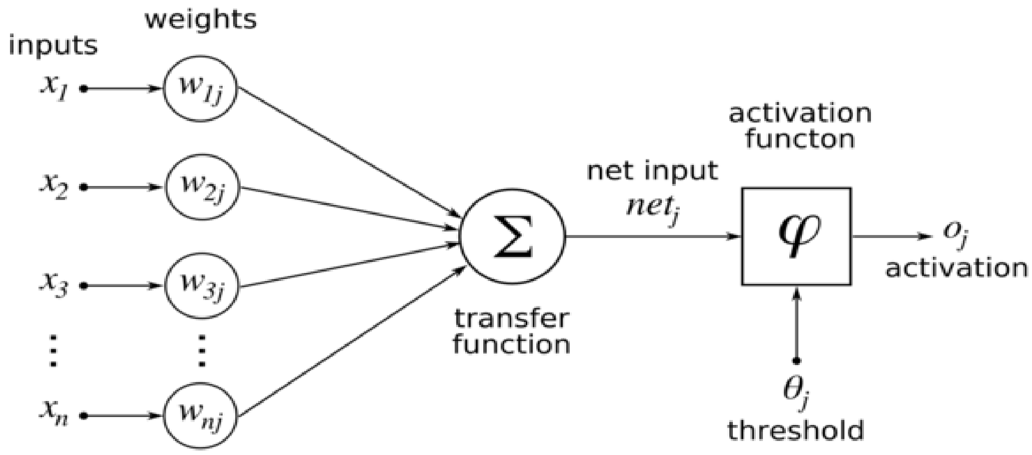

Figure 3: A schematic of an artificial neural network (ANN), modeled after a biological neuron with input-receiving dendrites, a cell body that takes in electrochemical information from dendrites, an axon hillock which conditionally triggers an electrical action potential if the voltage surpasses a threshold, an axon down which the action potential propagates, and a synaptic terminal where the output is sent to the next neuron in the circuit. *Here, the ANN takes in several numerical inputs, which undergo a linear transformation by the transfer function, or net input function, to integrate biases and tunable weights representing the mathematically-identified importance of data from that node. Transformed inputs are sent into an activation function which conditionally activates the neuron depending on whether the output surpasses a threshold. Deep neural networks have more complicated architectures than ANNs, with several hidden layers and more sophisticated mathematical processing.*

REINFORCEMENT LEARNING

Reinforcement learning models concern an agent as it fine-tunes its strategy to seek reward in an environment, solely through trial and error. Rather than identifying whether an image's features match more to a dog or a cat, RL problems typically center around an active agent trying to learn an optimal strategy, or policy, in a dynamic environment, such as a winning strategy for a game of chess. Broadly, RL agents learn to interact with their unknown and dynamic environment with no prior knowledge, where the in-progress policy is responsible for selecting present and potential future actions that affect each progressive state of the environment. While there are certain elements that must be controlled or simplified by scientists when setting up an RL model, these models are arguably the most accurate representations of policy-based human learning developed thus far, where policy-based learning refers to the fine-tuning of strategy that someone like a chess player would use to win a chess game.[2]

In every state of the environment, there are pathways through which the agent can achieve reward, which are closely dependent on the agent's own actions. With an RL agent's capacity to affect the state of the environment, the agent's actions play a large role in opening or closing pathways to reward. In the short term, an agent takes actions that change small aspects of the state and alter the potential for expected reward. However, if the agent takes a suboptimal action, it could limit future avenues for reward, metaphorically closing a door to a particular outcome. Building upon its short-term goal to optimize statistical descriptions of reward in the environment, the agent pursues its ultimate goal of traversing the maximal reward pathway through complicated and highly unintuitive "black box" algorithms. Such algorithms are termed "black box" algorithms because it is too complex to try to obtain a comprehensive understanding of their inner-workings. Resultantly, scientists ignore the question of how "black box" algorithms work, when working with them, focusing only on which inputs to

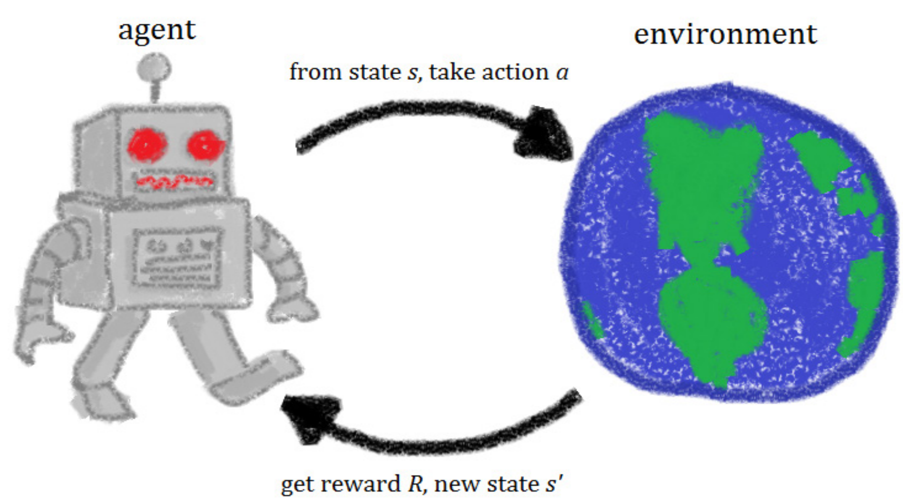

Figure 4: Reinforcement learning; depicts agent-environment interaction and relationship.

"The foundational principle of using reward to reinforce behaviors connects directly to neuroscience, rendering RL algorithms intentional microcosms of the adaptable brain."

feed in and what outputs to expect. As a result of this architectural composition, RL has great potential to aptly represent situations where a conscious being crafts an optimal policy through trial and error.

RL AS IT RELATES TO NEUROSCIENCE

RL algorithms allow us to recreate and optimize models of complex tasks, like playing a game of chess or tracking a surface-born odor trail. For the latter, an RL model can recapitulate the behavior of an animal searching for the source of an odor, a task that requires complex computation and statistical optimization that exceeds any animal's conscious computational capacity. Through this simulation, scientists have access to a model whose variables can be manipulated to understand the extent to which conditions like environmental, spatial, and physical constraints affect or limit the agent's behavior. By testing the agent's ability to overcome simulated constraints, RL algorithms can lend insight into the extents of an animal's tracking ability.

Foundationally, RL theory is based on psychological and neuroscientific principles of learning and reward. Much like how an infant is born with no prior knowledge and only a sensorimotor connection to its environment, RL agents begin with no background and only a means to take actions to affect the state of the environment. From feedback, an infant modifies its behaviors to develop a more powerful understanding of how to optimally behave in its environment. In the same way, a reinforcement learning agent, which would begin naive, builds upon a nonexistent understanding of its environment through corrective feedback to reach an optimal strategy. The foundational principle of using reward to reinforce behaviors connects directly to neuroscience, rendering RL algorithms intentional microcosms of the adaptable brain.

Methodologically, the broader goal of RL algorithms to achieve maximal reward is emblematic of the brain's limbic system and neurochemical reinforcement during the formation of neuronal connections. Animals learn from neurochemical reinforcement signals that are naturally embedded in the process of trial and error.[3] As a result, the parallel between the neurochemical learning process of an animal and the RL training of an agent is a curious one that enables multidirectional biomimicry. Roughly speaking, RL learning mimics the formation of synaptic connections from neurochemical reinforcement by upticking the model's quantification of reward to reinforce smart actions. Deepening the connection to neuroscientific learning, RL algorithms function in this manner to reinforce actions with simulated neurochemical reinforcement, once again maintaining a strong, but general, connection to the growth and reinforcement of neural synapses on a cellular and molecular level.

Moreover, using RL to represent a system of animalian learning is particularly significant because the method of RL learning is based on principles of neuroscientific learning. An agent's encounters with positive reinforcement, for the computational purposes of training, are indicative of the positive neurochemical reinforcement an animal receives while it learns a task in nature. This proposes a conceptual translation of modeled neurochemical reinforcement, via numeric upticking of reward, into actual neurochemical reinforcement, via excitatory neurotransmitters and other biophysical potentiation mechanisms during the training of the real animal.

As a result, RL is highly applicable to cognitive neuroscience, with strong quantitative components. Through the direct application of RL for simulation in neuroscience, as well as foundationally and methodologically, RL maintains intentional, direct, and symbolic connections to neuroscience.

CONCLUSION

Questions in cognitive neuroscience span the vast inner-workings of the brain and its interconnected systems, augmenting the need to develop scientific approaches which can answer them. A common barrier in systems biology is the task of creating an accurate representation of the system in question, which would offer insight into the complex process of interest, such as odor trail tracking. RL enables scientists to establish a model of the neurological systems needed to create this level of precision and optimization in an animal incapable of consciously calculating high-level statistical computations.

For odor trail tracking, RL algorithms allow scientists to obtain further insight into the requirements for cognitive neuroscience's abilities and the extents of neuroscientific systems. With RL's ability to recreate cognitive neuroscience problems and to represent an agent's development through models, the field of RL opens up possibilities to further understand cognitive neuroscience systems and peer into the unconquerable mind. As science continues to deepen its interest in the use of RL for neuroscience, the field holds a promising future for how scientists can use RL to shed light on cognitive neuroscience's complex processes and systems.

"An agent's encounters with positive reinforcement, for the computational purposes of training, are indicative of the positive neurochemical reinforcement an animal receives while it learns a task in nature."

REFERENCES

1. Kitano, H. (2002). Computational Systems Biology. Nature, 420(6912), 206–210.
2. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, Perspectives, and prospects. Science, 349(6245), 255–260. https://doi.org/10.1126/science.aaa8415
3. Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. Science, 275(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593
4. Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. Nature, 445(7126), 406–409. https://doi.org/10.1038/nature05464
5. Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2019). Continuous control with deep reinforcement learning (arXiv:1509.02971). arXiv. http://arxiv.org/abs/1509.02971
6. Jinn, J., Connor, E. G., & Jacobs, L. F. (2020). How Ambient Environment Influences Olfactory Orientation in Search and Rescue Dogs. Chemical Senses, 45(8), 625–634. https://doi.org/10.1093/chemse/bjaa060
7. Khan, A. G., Sarangi, M., & Bhalla, U. S. (2012). Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. Nature Communications, 3(1), 703. https://doi.org/10.1038/ncomms1712
8. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533. https://doi.org/10.1038/nature14236
9. Reddy, G., Murthy, V. N., & Vergassola, M. (2022). Olfactory Sensing and Navigation in Turbulent Environments. Annual Review of Condensed Matter Physics, 13(1), 191–213. https://doi.org/10.1146/annurev-conmatphys-031720-032754
10. Reddy, G., Shraiman, B., & Vergassola, M. (n.d.). Sector search strategies for odor trail tracking. 28.
11. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484–489. https://doi.org/10.1038/nature16961
12. Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning an introduction. A Bradford Book.

IMAGE REFERENCES

1. Greg Bradshaw. (2011, April). New Guinea Singing Dog sniffing the ground [Photograph].Wikimedia Commons. https://commons.wikimedia.org/wiki/File:New_Guinea_Singing_Dog_sniffing_the_ground.jpg
2. Harris, P. B. (n.d.). Neuronal migration is an artwork depicting many very young neurons that have been produced in the neuroepithelium migrating to their appropriate destinations in the brain. This image highlights the future of neuroscience showing different classes of cells colour coded. There is no available technique to do this now, but it is not far off considering the advances that have been made with brainbow mice. The brainbow technique allows for different cell types to be tagged with fluorescent proteins to track their development and connections with other cells. Wellcome Collection. Retrieved from https://wellcomecollection.org/works/u2mrc7w5.
3. Saini, G. (2017). Artificial Neural Network.png. Wikimedia Commons. Retrieved from https://commons.wikimedia.org/wiki/File:Artificial_neural_network.png.
4. Notfruit. (2017). Rl agent.png. Wikimedia Commons. Retrieved from https://commons.wikimedia.org/wiki/File:Rl_agent.png.