IT matters

# Crowdsourced feedback to help train robots

To teach an AI agent a new task, like how to open a kitchen cabinet, researchers often use reinforcement learning - a trial-and-error process where the agent is rewarded for taking actions that get it closer to the goal.

In many instances, a human expert must carefully design a reward function, which is an incentive mechanism that gives the agent motivation to explore. The human expert must iteratively update that reward function as the agent explores and tries different actions. This can be time-consuming, inefficient, and difficult to scale up, especially when the task is complex and involves many steps.

Researchers from MIT, Harvard University, and the University of Washington have developed a new reinforcement learning approach that doesn't rely on an expertly designed reward function. Instead, it leverages crowdsourced feedback, gathered from many nonexpert users, to guide the agent as it learns to reach its goal.

While some other methods also attempt to utilize nonexpert feedback, this new approach enables the AI agent to learn more quickly, despite the fact that data crowdsourced from users are often full of errors. These noisy data might cause other methods to fail.

In addition, this new approach allows feedback to be gathered

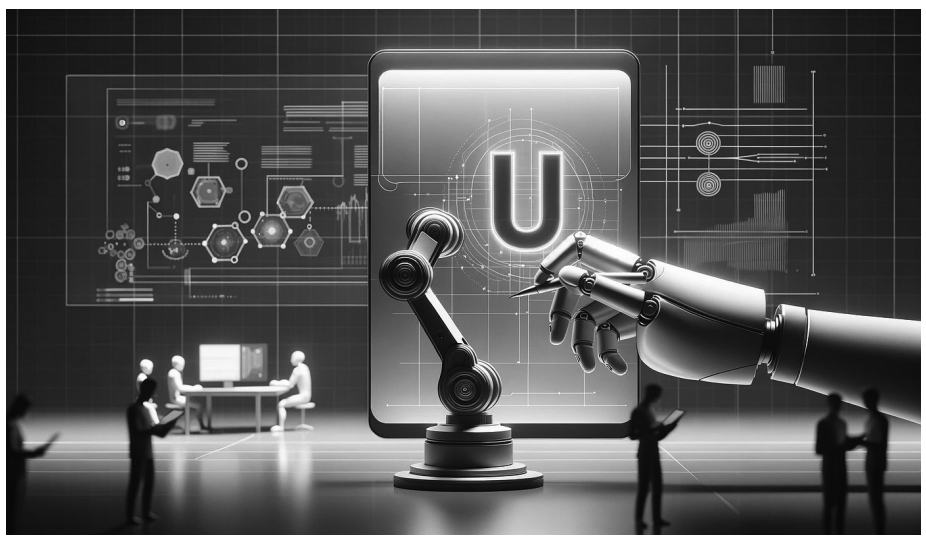

asynchronously, so nonexpert users around the world can contribute to teaching the agent.

"One of the most time-consuming and challenging parts in designing a robotic agent today is engineering the reward function. Today reward functions are designed by expert researchers - a paradigm that is not scalable if we want to teach our robots many different tasks. Our work proposes a way to scale robot learning by crowdsourcing the design of reward function and by making it possible for nonexperts to provide useful feedback," says Pulkit Agrawal, an assistant professor in the MIT Department of Electrical Engineering and Computer Science (EECS) who leads the Improbable AI Lab in the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL).

In the future, this method could help a robot learn to perform specific tasks in a user's home quickly, without the owner needing to show the robot physical examples of each task. The robot could explore on its own, with crowdsourced nonexpert feedback guiding its exploration.

"In our method, the reward function guides the agent to what it should explore, instead of telling it exactly what it should do to complete the task. So, even if the human supervision is somewhat inaccurate and noisy, the agent is still able to explore, which helps it learn much better," explains lead author Marcel Torne '23, a research assistant in the Improbable AI Lab.

One way to gather user feedback for reinforcement learning is to show a user two photos of states achieved by the agent, and then ask that user which state is closer to a goal. For instance, perhaps a robot's goal is to open a kitchen cabinet. One image might show that the robot opened the cabinet, while the second might show that it opened the microwave. A user would pick the photo of the "better" state.

Some previous approaches try to use this crowdsourced, binary feedback to optimize a reward function that the agent would use to learn the task. However, because nonexperts are likely to make mistakes, the reward function can become very noisy, so the agent might get stuck and never reach its goal.

"Basically, the agent would take the reward function too seriously. It would try to match the reward function perfectly. So, instead of directly optimizing over the reward function, we just use it to tell the robot which areas it should be exploring," Torne says.

# Closer to designing artificial bio machines

Biological materials are made of individual components, including tiny motors that convert fuel into motion. This creates patterns of movement, and the material shapes itself with coherent flows by constant consumption of energy. Such continuously driven materials are called "active matter." The mechanics of cells and tissues can be described by active matter theory, a scientific framework to understand shape, flows, and form of living materials. The active matter theory consists of many challenging mathematical equations.

Scientists from the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG) in Dresden, the Center for Systems Biology Dresden (CSBD), and the TU Dresden have now developed an algorithm, implemented in an open-source supercomputer code, that can for the first time solve the equations of active matter theory in realistic scenarios. These solutions bring us a big step closer to solving the century-old riddle of how cells and tissues attain their shape and to designing artificial biological machines.

Biological processes and behaviors are often very complex. Physical theories provide a precise and quantitative framework for understanding them. The active matter theory offers a framework to understand and describe the behavior of active matter - materials composed of individual components capable of converting a chemical fuel ("food") into mechanical forces. Several scientists from Dresden were key in developing this theory, among others Frank Jülicher, director at the Max Planck Institute for the Physics of Complex Systems, and Stephan Grill, director at the MPI-CBG.

With these principles of physics, the dynamics of active living matter can be described and predicted by mathematical equations. However, these equations are extremely complex and hard to solve. Therefore, scientists require the power of supercomputers to comprehend and analyse living materials. There are different ways to predict the behavior of active matter, with some focusing on the tiny individual particles, others studying active matter at the molecular level, and yet others studying active fluids on a large scale. These studies help scientists see how active matter behaves at different scales in space and over time.

Scientists from the research group of Ivo Sbalzarini, TU Dresden Professor at the Center for Systems Biology Dresden (CSBD), research group leader at the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG), and Dean of the Faculty of Computer Science at TU Dresden, have now developed a computer algorithm to solve the equations of active matter. Their work was published in the journal "Physics of Fluids" and was featured on the cover. They present an algorithm that can solve the complex equations of active matter in three dimensions and in complex-shaped spaces.

"Our approach can handle different shapes in three dimensions over time," says one of the first authors of the study, Abhinav Singh, a studied mathematician. He continues, "Even when the data points are not regularly distributed, our algorithm employs a novel numerical approach that works seamlessly for complex biologically realistic scenarios to accurately solve the theory's equations. Using our approach, we can finally understand the long-term behavior of active materials in both moving and non-moving scenarios for predicting their dynamics. Further, the theory and simulations could be used to program biological materials or create engines at the nano-scale to extract useful work."

The other first author, Philipp Suhrcke, a graduate of TU Dresden's Computational Modelling and Simulation MSc program,me adds, "thanks to our work, scientists can now, for example,

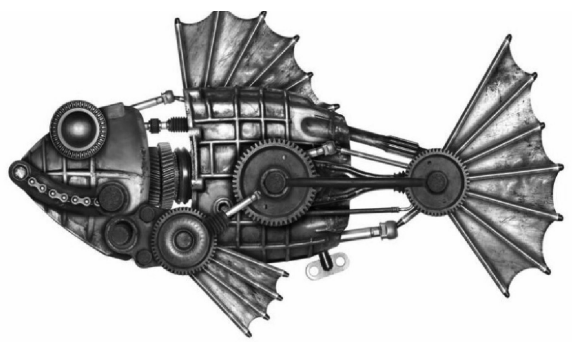

predict the shape of a tissue or when a biological material is going to become unstable or dysregulated, with far-reaching implications in understanding the mechanisms of growth and disease."

The scientists implemented their software using the open-source library OpenFPM, meaning that it is freely available for others to use.

# Quantum randon numbers as cloud encryption keys

Distributed cloud storage is a hot topic for security researchers around the globe pursuing secure data storage, and a team in China is now merging quantum physics with mature cryptography and storage techniques to achieve a cost-effective cloud storage solution.

Shamir's secret sharing, a known method, is a key distribution algorithm. It involves distributing private information to a group so that "the

secret" can be revealed only when a majority pools their knowledge.

It's common to combine quantum key distribution (QKD) and Shamir's secret sharing algorithm for secure storage - at an utmost security level.

But utmost security solutions tend to bring substantial cost baggage, including significant cloud storage space requirements.

The team members recently presented method that uses quantum random numbers as encryption keys, disperses the keys via Sharmir's secret sharing algorithm, applies erasure coding within ciphertext, and securely transmits the data through QKD-protected networks to distributed clouds.

Their method not only provides quantum security to the entire system but also offers fault tolerance and efficient storage -- and this may help speed the adoption of quantum technologies.

"In essence, our solution is quantum-secure and serves as a practical application of the fusion between quantum and cryptography technologies," said corresponding author Yong Zhao, vice president of QuantumCTek Co. Ltd., a quantum information technology company.

"QKD-generated keys secure both user data uploads to servers and data transmissions to dispersed cloud storage nodes."

The team explored whether quantum security services could expand beyond secure data transmission to offer a richer spectrum of quantum security applications such as data storage and processing. They came up with a more secure and cost-effective fault-tolerant cloud storage solution.

"It not only achieves quantum security but also saves storage space when compared to traditional mirroring methods or ones based on Shamir's secret sharing, which is commonly used for distributed management of sensitive data," said Zhao.

When the team ran the solution through experimental tests ranging from encryption/decryption, key preservation, and data storage, it proved to be effective.

The solution is currently feasible from both technological and engineering perspectives: It meets the requirement for relevant quantum and cryptographic standards to ensure a secure storage solution capable of withstanding the challenges posed by quantum computing.

"In the future, we plan to drive the commercial implementation of this technology to offer practical services," said Zhao. "We'll explore various usage models in multiuser scenarios, and we're also considering integrating more quantum technologies, such as quantum secret sharing, into cloud storage."

# Gaming new tool to foster cooperation

Intense focus pervades the EEG laboratory at the University of Konstanz on this day of experimentation. In separate labs, two participants, connected by screens, engage in the computer game Pacman. The burning question: Can strangers, unable to communicate directly, synchronize their efforts to conquer the digital realm together?

Doctoral candidate Karl-Philipp Flösch is leading the experiment. He states: "Our research revolves around cooperative behaviour and the adoption of social roles." However, understanding brain processes underlying cooperative behaviour is still in its infancy, presenting a central challenge for cognitive neuroscience.

How can cooperative behaviour be brought into a highly structured EEG laboratory environment without making it feel artificial or boring for study participants?

The research team, led by Harald Schupp, Professor of Biological Psychology at the University of Konstanz, envisioned using the well-known computer game Pacman as a natural medium to study cooperative behaviour in the EEG laboratory.

"Pacman is a cultural icon. Many have navigated the voracious Pacman through mazes in their youth, aiming to devour fruits and outsmart hostile ghosts," reminisces Karl-Philipp Flösch.

Collaborating with colleagues, co-author

Tobias Flaisch adapted the game.

In the EEG version, two players instead of one must collaboratively guide Pacman to the goal. Flaisch explains: "Success hinges on cooperative behaviour, as players must seamlessly work together."

However, the researchers have built in a special hurdle: the labyrinth's path is concealed.

Only one of the two players can see where Pacman is going next.

Flösch elaborates: "The active player can communicate the direction to the partner, but only indirectly using pre-agreed symbols, communicated solely through the computer screen." If you do not remember quickly enough that a crescent moon on the screen means that Pacman should move right, and that only the banana on the keyboard can make Pacman move to the right, you're making a mistake.

"From the perspective of classical psychological research, the game combines various skills inherent in natural social situations," notes Harald Schupp.During each game, the players' brain reactions were measured using EEG. Calculating event-related potentials provides a detailed view of the effects elicited by different game roles with millisecond-level temporal precision.The team hypothesized that the game role significantly influences brain reactions.

Therefore, they examined the P3 component, a well-studied brain reaction exhibiting a stronger deflection in the presence of significant and task-relevant stimuli.

The results confirmed their assumption: "The P3 was increased not only when the symbol indicated the next move's direction but also when observing whether the game partner selected the correct symbol," says Flösch.The team concludes that the role we take on during cooperation determines the informational value of environmental stimuli situationally.EEG measurements allow the brain processes involved to be dynamically mapped. "Cooperative role adoption structures our entire society," summarises Schupp, providing context for the study. "An individual achieves little alone, but collectively, humanity even reaches the moon. Our technological society hinges on cooperative behavior," says Flösch, adding that children early on take individual roles, thereby learning the art of complex cooperation. Consequently, this role adoption occurs nearly effortlessly and automatically for us every day. "Our brains are practically 'built' for it, as evidenced by the results of our study."

# Heat used as energy source in computing

Physicists at Martin Luther University Halle-Wittenberg (MLU) and Central South University in China have demonstrated that, combining specific materials, heat in technical devices can be used in computing. Their discovery is based on extensive calculations and simulations. The new approach demonstrates how heat signals can be steered and amplified for use in energy-efficient data processing.Electric current flow heats up electronic device. The generated heat is dissipated and energy is lost. "For decades, people have been looking for ways to re-use this lost energy in electronics," explains Dr Jamal Berakdar, a professor of physics at MLU.This is extremely challenging, he says, due to the difficulty in directing and controlling accurately heat signals. However, both are necessary if heat signals are to be used to reliably process data.

Berakdar carried out extensive calculations together with two colleagues from Central South University in China.

The idea: instead of conventional electronic circuits, non-conductive magnetic strips are used in conjunction with normal metal spacers.