

# CAN MACHINE LEARNING COPE WITH THE ERRATIC AND UNCERTAIN NATURE OF THE REAL WORLD?

UNTIL RECENTLY, SELF-DRIVING CARS AND ROBOTIC HOUSEKEEPERS WERE THE STUFF OF SCI-FI MOVIES. HOWEVER, IN THE LAST DECADE, THE CAPABILITIES OF ARTIFICIAL INTELLIGENCE HAVE BEEN EXPANDING AT BREAK-NECK SPEED. AT THE **UNIVERSITY OF SOUTHERN CALIFORNIA, USA, DR HAIPENG LUO** HAS BEEN CONDUCTING INNOVATIVE NEW RESEARCH IN THE FIELD OF MACHINE LEARNING TO BRING THESE EXCITING TECHNOLOGIES WITHIN REACH.

## TALK LIKE A MACHINE LEARNING SCIENTIST

**ALGORITHM** – a sequence of well-defined rules that allows computers to make calculations and solve problems

**MACHINE LEARNING** – a type of artificial intelligence in which computers learn from past data in order to predict outcomes more accurately

**NON-STATIONARITY** – the quality of being subject to unexpected changes. A non-stationary environment has rules that may change over time

**REINFORCEMENT LEARNING (RL)** – an area of machine learning that investigates how algorithms can learn about their environments and make good decisions through reward signals

**REWARD SIGNAL** – feedback from an environment that informs an algorithm whether an action it has taken was good or bad

Do you ever get the feeling that you are being watched? It can feel slightly disconcerting when an advert for a band you have just been listening to or a TV show you have just been watching pops up on your social media feed. Although

you do not need to worry that a government spy is observing you through your webcam, this targeted advertising is no coincidence.

Companies such as Google, Netflix and Facebook use a type of artificial intelligence known as machine learning to help them better understand their users. Each time you search for something on the internet, machine learning algorithms track your activity to learn about your interests. They then use this information to tailor your adverts and film recommendations based on what they think you will want to buy or see.

Reinforcement learning (RL) is a branch of machine learning in which an algorithm learns about its environment through a reward system that provides feedback on its actions. Each time the algorithm performs an action, the environment will change state and issue a reward signal to the algorithm. This signal will be positive when the algorithm makes a good decision and negative when it makes a bad one. At first, the algorithm may make lots of mistakes. However, over time, the algorithm will learn how its actions affect the environment and which of these actions are desirable. Thanks to the positive and negative reinforcement provided by the reward signals, the algorithm will learn to make fewer mistakes and so will become better adapted to its environment.

While RL algorithms are currently used in robotics and telecommunications, they have certain characteristics that limit their ability to function in wider real-world settings. At

the University of Southern California, Dr Haipeng Luo has developed a new approach to reinforcement learning that may solve these problems.

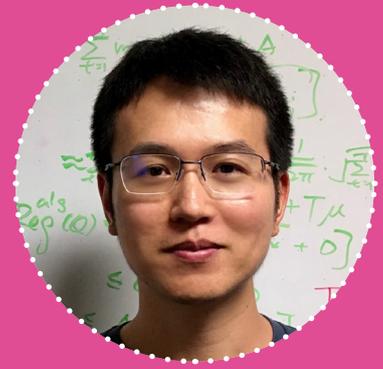
## LIMITATIONS OF REINFORCEMENT LEARNING

Traditional RL algorithms only work in stationary environments. For example, RL has been used to train computers to play Atari arcade games, such as Space Invaders. Since these games are designed according to fixed rules, the environment in which the RL algorithms work is stationary.

In stationary environments, the consequences of the algorithm's actions will always follow the same rules. Similarly, the reward signals that the algorithm receives from the environment will remain consistent. In other words, "if the same action is taken in the same state, the distribution over the outcome or 'next-state' will be the same, and the reward signal will be the same" says Haipeng.

"This is a huge limitation," he continues, "because in practice, environments are constantly changing." For example, if there are multiple RL algorithms in an environment, their behaviour is likely to change as they learn and improve. Such an environment would be highly non-stationary, and a traditional RL algorithm would struggle to perform effectively.

## THE COMPLEXITY OF NON-STATIONARY ENVIRONMENTS



## DR HAIPENG LUO

University of Southern California, USA

.....

### FIELD OF RESEARCH

Machine Learning

.....

### RESEARCH PROJECT

Developing new algorithms to help machine learning systems perform in real-world environments

.....

### FUNDERS

This work is supported by the NSF (grant number IIS-1943607) and the Google Faculty Research Award. The contents are solely the responsibility of the author and do not necessarily represent the official views of the funders.

paper award at the Annual Conference on Learning Theory, the world's most prestigious conference on machine learning theory. "This is a huge honour for us," says Haipeng.

Being selected as the best paper from hundreds of other submissions is a sign of how widely regarded Haipeng's research is. "While I have won similar awards previously, this award is especially meaningful for me as a young assistant professor, because the work was done by my student and myself only," he says.

### NEXT STEPS

Haipeng's current work considers the perspective of a single RL algorithm within a non-stationary environment. One key next step is to consider how several RL algorithms would work simultaneously within a non-stationary environment. There are many more questions to consider in a multi-agent system. For example, how would each algorithm's learning process be affected by the other algorithms? And how would other causes of non-stationarity in the environment affect this? While there is still a lot of work to be done, Haipeng and Chen-Yu's solution to the complex problem of non-stationarity is a major step forward.

Most real-world environments are non-stationary, with conditions that are constantly changing, and so it is hard for traditional RL algorithms to function. It is difficult to predict how non-stationary an environment is going to be before the algorithm starts interacting with it. Even once the algorithm has started to explore its environment, the feedback which it receives is limited. "The agent only receives the reward for an action that it has taken," says Haipeng. "Therefore, it cannot figure out what would have happened if it took a different action."

This means these RL algorithms cannot be used for many applications. For example, self-driving cars must navigate the particularly non-stationary environments of busy roads. "On the road, conditions can change quickly and unexpectedly, and the behaviour of other cars and pedestrians can be unpredictable and random." For such complex environments, a new approach to designing RL algorithms is needed. This is the focus of Haipeng's research.

### HOW HAS HAIPENG TACKLED THIS CHALLENGE?

There are two key ideas in Haipeng's new approach, the first of which involves using a pre-existing algorithm to detect the non-stationarity of the environment. This pre-existing algorithm is known as a base algorithm. Haipeng selects a base algorithm that performs optimally in a stationary environment. If the base algorithm does not perform as well as expected in the environment under investigation, then that environment must be non-stationary.

Haipeng's second idea was to use multiple copies of the base algorithm and run them for different lengths of time. The algorithms that run for short durations can detect large changes in the environment, while the longer-running algorithms detect smaller changes. Once an algorithm has finished, Haipeng conducts statistical tests to determine how non-stationary the environment is.

### WHAT ARE THE BENEFITS OF HAIPENG'S APPROACH?

One of the main benefits of Haipeng's new approach is that it is applicable to a wide range of situations, as long as a base algorithm exists. Through his approach, any RL algorithm that works in a stationary environment can be transformed so that it works in a non-stationary environment. Another benefit is that this transformation is easy to implement. The base algorithm can remain unchanged and would simply need to be 'wrapped' in Haipeng's new algorithm to greatly improve its performance.

Haipeng has also proven that his method produces the best statistically possible performance for any given environment. This is particularly impressive when one considers that this is achieved without even knowing how non-stationary the environment is.

### HOW GROUND-BREAKING IS THIS RESEARCH?

Haipeng's research has had a significant impact in the field of machine learning. In 2021, Haipeng and his student, Chen-Yu Wei, won the best

# ABOUT MACHINE LEARNING

Machine learning is a subfield of artificial intelligence. Researchers in this field, sometimes known as machine learners, develop algorithms that can figure out how to improve their own performance. There are many applications for machine learning algorithms, from social media and online streaming platforms to agriculture and medicine.

Due to its huge potential, research into machine learning techniques has surged in recent years. As a result, the global machine learning market is expected to be worth over \$30 billion a year by 2024. This makes it an exciting time to be a machine learner.

Machine learners use a mix of mathematics and computer science to develop and test their algorithms. While the maths can be complicated and coding is a skill that requires a lot of hard work to master, the rewards are worth it. Machine learners address all kinds of real-world problems, and their algorithms may help society

solve some of its biggest challenges.

## WHERE IS MACHINE LEARNING FOUND IN OUR DAY-TO-DAY LIVES?

Machine learning is already fully integrated in our daily lives. “Examples include automatic translations, fraud and spam detectors, and face and voice recognition software,” says Haipeng. Streaming services such as Netflix and Spotify use machine learning to recommend new movies, TV programmes and songs they think you will want to watch or listen to, based on your previous viewings. In the field of medicine, machine learning is used to help doctors make correct diagnoses and prescribe appropriate medicines to their patients.

## HOW WILL MACHINE LEARNING ADVANCE IN THE FUTURE?

Currently, machine learning relies on training gigantic statistical models with a huge amount of data. “To make machine learning more general and versatile, I think being able to learn from a

small dataset, either via transferring knowledge from other similar tasks or via some kind of logistical reasoning, would be an important next step,” says Haipeng.

A key challenge when developing machine learning systems is the consideration of societal and ethical impacts. “The next generation of computer scientists will have to ensure fairness, accountability, transparency and ethics (FATE) in artificial intelligence systems,” says Haipeng.

## WHAT ARE THE REWARDS OF WORKING AS A COMPUTER SCIENTIST?

Haipeng enjoys being able to do maths that has applications in the real world. “I enjoy developing algorithms, deriving theory for them, and, finally, implementing them as a computer program to solve problems,” he says. “I like that machine learning combines two of my favourite subjects, math and computer science, to solve real-life problems in an unprecedented way.”

## EXPLORE CAREERS IN MACHINE LEARNING

- There is currently a huge demand for machine learners and there are a variety of career paths to follow. For example, you could be a software engineer at a tech start-up, an analyst in the finance industry, a research scientist in an industrial lab or a professor at a university.
- Haipeng recommends searching online to find useful and interesting resources about machine learning, from forums and blog posts to courses and YouTube videos. “Learning to find the materials you need online is one of the most useful skills in the age of the internet!”
- The School of Engineering at the University of Southern California provides a wealth of resources for students and teachers, including many to help you learn programming languages and how to code: [viterbik12.usc.edu/resources](http://viterbik12.usc.edu/resources)
- The School of Engineering also offers high school students the chance to participate in an 8-week research project through the SHINE programme: [viterbik12.usc.edu/shine](http://viterbik12.usc.edu/shine)
- The following article discusses how to become a machine learning engineer, what skills you will need to develop, and the work you can expect to do: [www.csuglobal.edu/blog/how-become-machine-learning-engineer](http://www.csuglobal.edu/blog/how-become-machine-learning-engineer)

## PATHWAY FROM SCHOOL TO MACHINE LEARNING

- Haipeng recommends studying computer science courses to learn about programming, data structures and algorithms. He also recommends studying maths subjects such as algebra, calculus, probability and statistics.
- Most universities offer degrees in computer science, and some may offer specialised degrees in machine learning.



## HOW DID HAIPENG BECOME A MACHINE LEARNING SCIENTIST?

### WHAT WERE YOUR INTERESTS WHEN YOU WERE YOUNGER?

When I was young, my interest was mainly in mathematics. Growing up, I had no access to computers, however this changed when I started studying at Peking University in Beijing, China. I was able to take all kinds of computer science courses and learnt to program computers, and I quickly fell in love with this subject.

### WHAT INSPIRED YOU TO BECOME A COMPUTER SCIENTIST?

During my first year of undergraduate study, I heard about a programming competition called the International Collegiate Programming Contest. I was also made aware of some legendary Chinese contestants. I thought it was such a fascinating event, so I started training myself to take part in it. Eventually, I became part of Peking University's team and started competing in the contest. Although I never made it to the top level, I still think this is the most important thing that inspired me to continue studying in this field and to become a computer scientist.

### DURING YOUR PHD STUDIES, YOU CONDUCTED RESEARCH INTERNSHIPS WITH AT&T, YAHOO AND MICROSOFT. WHAT DID YOU GAIN FROM THESE EXPERIENCES?

These experiences greatly expanded my research capabilities, especially from a practical side. For example, I learnt what it is like to work in an industrial lab. This helped me understand what kind of practical problems scientists in these labs work on and how to test the performance of proposed algorithms on a real system.

### WHAT PERSONAL QUALITIES HAVE MADE YOU SUCCESSFUL AS A SCIENTIST?

The most important one is probably always being curious. I always ask myself why things work the way they do, how to improve them and whether they can be generalised to other problems or be united with other approaches. I find that these are often the most important questions to ask as a scientist.

### WHAT DO YOU ENJOY DOING OUTSIDE OF WORK?

I have many hobbies, including playing ping pong, badminton and video games. I also enjoy playing the guitar and the piano, as well as hiking, travelling and watching movies. Unfortunately, I rarely have enough time to enjoy all these things regularly!

### HAIPENG'S TOP TIPS

1. Find something you feel passionate about. I think that working on your true passion is the key to success.
2. Be curious and ask questions about how things work.
3. Explore all your options. Even within the field of computer science there are so many different directions you could pursue.