Pamela: Integrating Modelling and Machine Learning for Autonomous Robotics

Dr Paul Robertson

Scientia

PAMELA: INTEGRATING MODELLING AND MACHINE LEARNING FOR AUTONOMOUS ROBOTICS

Machine learning is rapidly advancing the decision-making capabilities of today's computers, yet without an in-depth knowledge of the programming it involves, many engineers and researchers find current technology inaccessible. **Dr Paul Robertson** at Dynamic Object Language Labs (DOLL) in Massachusetts believes that a solution to the issue may have been hidden in plain sight: machine learning itself. His ideas have now culminated in 'Pamela': a universal, open-source language capable of modelling real-world systems, and building plans to overcome challenges. The language and its related tools could soon open up significant opportunities in the emerging field of artificial intelligence.

Machine Learning

Until very recently, the human mind has remained the only decision-making system smart enough to operate the many systems we use to keep our civilisation ticking. To deal with the intricate webs of interconnected variables that link systems together, and to their surrounding environments, our intelligence has so far been critical to making the complex, rapid decisions needed to keep them functioning.

Yet this picture is now changing quickly with the emergence of artificial intelligence (AI), which uses algorithms to first learn about systems and their surrounding environments, and then to make useful decisions about their operation.

'Systems ranging from photocopiers to spaceships are designed by human engineers,' explains Dr Paul Robertson of Dynamic Object Language Labs (DOLL), a Massachusetts-based company. 'Unlike the messy natural world, they can be described accurately and fairly completely. By turning knowledge of their engineering into models and having those models as part of the runtime system, systems can use AI reasoning to intelligently diagnose and operate such a system.'

Known as 'model-based autonomy', this area of research has exploded over the past 15 years. However, Dr Robertson acknowledges that a significant barrier lies in the way of the technology's widespread adoption.

Bridging the Gap

To apply machine learning in the real world today, researchers must first build models of the systems it will be applied to. These essentially consist of sets of variables, each associated with their own probability distributions, which approximate the characteristic behaviours of systems in different conditions. Subsequently, data gathered

by a system's sensors are linked to a separate 'solving' program, which uses models to make smart decisions about the system's operation.

The lack of an appropriately expressive modelling language has forced researchers to make do with simple models, seriously limiting the advances

in autonomous systems and machine learning. So far, this problem has delayed the many potential benefits of machine learning, especially for real-world systems, but Dr Robertson believes that a solution has been hiding in plain sight all along: machine learning itself. Instead of manually building models and writing the tedious programs required to link them with solving languages, he proposes a language and toolchain that together are capable of building models themselves, using only sensor data.

Understanding the Open World

To lay the groundwork for this new language, Dr Robertson and his colleagues first needed to consider how an AI system should deal with the complex realities it experiences through its sensors. Instead of executing actions based on fixed prerequisites, as most pre-programmed robots do today, more intelligent robots must consider a range of possibilities when making decisions about their future actions, based on the patterns they observe. 'As we relax the constraints on our models, we move towards systems that can be observed to exhibit certain regularities, and thus, their models can be learnt by AI,' explains Dr Robertson.

This behaviour is comparable to how wild animals learn how best to survive given the conditions of their habitats and the capabilities of their bodies.

'It allows us to change our models in the face of new evidence,' Dr Robertson continues. 'This finally leads to exploration and learning rather than engineering design. As modern AI systems venture out into the world, as robots, they need as much of the reactive learning found in primitive animals as they do for high-level reasoning. Furthermore, they need to learn by themselves.'

Introducing Pamela

This line of reasoning has now culminated in a new modelling language, named 'Pamela', developed by Dr Robertson and his colleagues at DOLL. In contrast to previous modelling languages, Pamela was designed to work with state-of-the-art machine learning and planning tools designed to operate in the real world. In this case, the probability distributions of each variable in a model themselves can be modelled by Pamela. 'These "probabilistic variables" allow us to reason over uncertainty in models, and an autonomous system's beliefs about them,' Dr Robertson explains.

For the first time, this brings the two worlds of modelling and probabilistic programming together into a single, unified language. 'Pamela was developed to add probabilistic variables and probabilistic reasoning to model-based systems,' Dr Robertson continues. 'Finally, the models themselves can be learned, and the learned models can change as learning makes them more reliable.'

Applying a Model

In their latest research, the DOLL team has developed a toolkit for applying their language to real-world situations, named the 'Pamela Autonomy Toolchain' (PAT). To explain how it works, the researchers describe a simple problem in which a robot is low on battery and must find a charging point to plug itself in to. The charging point is locked, however, and the robot must first find the key to unlock it somewhere in the room. To achieve this goal, PAT first builds a five-stage plan: go to the key; take the key; go to the charger; insert the key; and start the charge.

To enact these steps, several parameters must first be set out in a model of the situation. For example, if the robot is close to an object, then it can pick it up; and in order to move the key, it must first be holding it. With machine

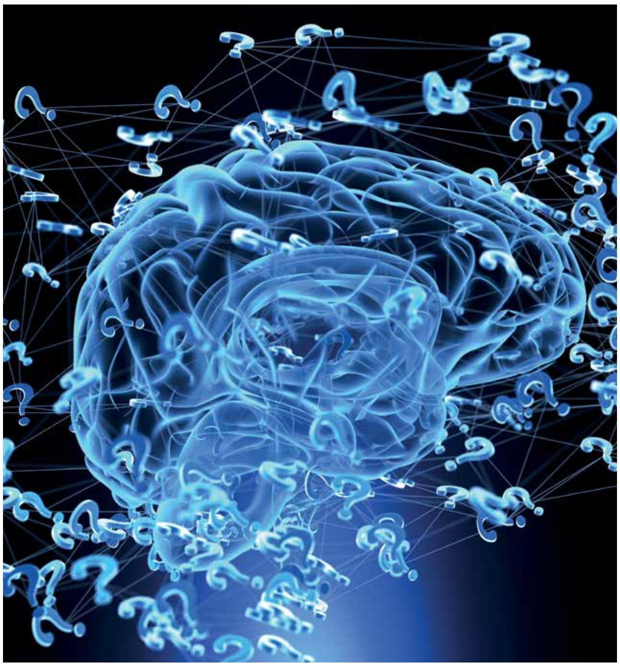

learning, the physical actions associated with each of these parameters can be learnt through experience. However, through previous approaches, researchers needed to establish the instructions for this themselves. With PAT, on the other hand, the probabilities of the variables measured by the robot's sensors can be used as the variables required to build a realistic model of the situation, with no human input required.

Updating Belief States

Dr Robertson first illustrates how this would be done in current approaches to machine learning. 'The robot believes that there is a key on one of two tables, based on prior history, but the probability of the key being on any one of the tables is low,' he says. 'Therefore, it plans to get close enough to one of the tables in order to increase the certainty in a key being there. The closest table has the least cost, so it will probably be picked.' When it is close to the table, the robot will then either confirm that the key is there or change the probability of it being there to zero. In this case, it will move to the second table and repeat the process.

More realistically, however, the key may not be on either of the tables, but could be placed anywhere on the floor, making the situation far more complex. 'If you have to search for the key before proceeding, it is clear that the whole plan cannot be produced in advance,' continues Dr Robertson. 'A highly contingent plan that covers every possibility of where the key might be would be a very complex plan. In this kind of situation, the robot needs to take actions to improve its knowledge about the world, and it also needs to take actions to complete its mission.'

This can be done through PAT, which establishes a set of 'beliefs' about the actions a system must undertake to realise the parameters of its plan, by assigning higher or lower probabilities to them. Therefore, the robot learns to look around the room in order to improve its belief state, allowing it to assign higher probabilities to the key being in particular areas. When its task is complete and the robot finally plugs in, a 'reward' system adds value to each step that successfully contributed to the overall goal. This means that in future runs, the robot will believe more strongly that the key will be in a similar area to where it was before, reducing the need for planning. Ultimately, this equates to a more reliable model of the situation than a human researcher could ever build.

Engaging with Research Groups

Having established a rigid framework for Pamela, Dr Robertson and his colleagues at DOLL are now starting to encourage members of the wider programming community to make their own contributions to the language. Since Pamela is now available online as open-source software, the team strongly hopes that its capabilities will quickly grow and evolve through the inputs of specialists around the world. If realised, this could bring about new machine learning algorithms, better planning systems, and new systems for diagnosing problems when plans are derailed.

With the diversity of expertise afforded by this approach, Pamela could soon equip researchers and engineers across a wide range of fields – including those with little knowledge of the language used in machine learning – to apply their models in the real world. With future research, Dr Robertson's team hopes to integrate their 'Continuous Effective Robot Learning' (CARL) platform, which they first developed in 2019. If added to PAT, the team believes that CARL would help make autonomous systems more independent. The level of human involvement in AI systems is one of today's greatest challenges.

Meet the researcher

Dr Paul Robertson
Dynamic Object Language Labs Inc
Lexington, MA
USA

Dr Paul Robertson holds a DPhil from the University of Oxford, for a research project entitled 'A self-adaptive architecture for image understanding'. Since then, he has worked at numerous prestigious institutions and companies, including the University of Massachusetts at Boston, the University of Texas at Dallas, Artelligence Inc, Symbolics Inc, Massachusetts Institute of Technology, and BBN Technologies. He has had a lifelong interest in reflection in intelligent systems, and the role that reflection can play in adapting responses to changes in the environment. A key interest has been advanced planning, especially in the context of self-adaptation. Today, Dr Robertson's research at Dynamic Object Language Labs (DOLL) Inc centres around the role of emotions and consciousness in robotic systems as a natural extension from his earlier work on reflection and self-organisation.

CONTACT

E: paulr@dollabs.com
W: https://www.dollabs.com/people/dr-paul-robertson-chief-scientist-president
W: http://paulrobertson-anglais.bcdoll.com/

KEY COLLABORATORS

Professor Patrick Winston and Adam Kraft, MIT
Dr Andreas Hofmann, Dan Cerys and Prakash Manghwani, DOLL Inc

FUNDING

This work was supported by Contract FA8650-11-C-7191 with the US Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory.