

# Modelling minds: can computers mimic human intelligence?

The human brain is often referred to as the most complex device known in the Universe. It gives rise to the amazing abilities of the human mind, including our capacity for intelligence. For many years, psychologists and neuroscientists have sought to uncover how the mind functions, while computer scientists have attempted to reproduce human-like capabilities in machines. Collaborations between these disciplines have resulted in computational models that are helping scientists to uncover the secrets of the mind and impart this intelligence to computers.

**A**re computers smarter than you? Or are you smarter than computers? Well, it depends on how you define what it means to be 'smart'. In recent years, we have developed computers that greatly exceed our own abilities in many tasks. For example, computers can perform thousands of calculations per second, meaning they can solve some mathematical problems far more quickly and accurately than humans. And computers have beaten human world champions at games of chess and go. But does this mean they are smarter than us? While computers can outperform humans in many domains, including arithmetic, object recognition and some aspects of language processing, no machine yet exists that can match human performance in *all* these domains. We are still unique in our ability to use our intelligence to adapt to a remarkably wide range of tasks.


## What is human intelligence?

As humans, we have the unique ability to make connections between things we have previously learnt and use those to solve challenges in new situations. For example, your language processing abilities mean that if you hear an unfamiliar word, you can often infer its meaning from the context of the conversation, or know to look it up in a dictionary. While many animals display specialised forms of intelligence, humans are flexible and creative when dealing with unfamiliar circumstances as our brains have the capacity for adaptive learning. Your brain is also incredibly efficient, running on only about 20 Watts of power to control your whole body and mind – similar to a standard lightbulb!

Understanding how the human brain functions so efficiently and flexibly remains one of the greatest mysteries of science. Psychologists, neuroscientists and computer scientists hope to uncover this mystery with the help of computational models.

## What are computational models?

A mathematical expression, or **function**, transforms an input to an output. For example, the function  $+10$  transforms the inputs  $[1, 2, 90]$  to the outputs  $[11, 12, 100]$ , by adding 10 to each input. An **algorithm** links multiple functions together in a rule-based sequence. For example, the algorithm  $+10, \times 2$  is a different algorithm than  $\times 2, +10$  (because  $1 + 10 = 11, 11 \times 2 = 22$ ; while  $1 \times 2 = 2, 2 + 10 = 12$ ). A **variable** is an element of a function that can be assigned different **parametric values**. For example, if the function is  $+b$ , then the variable can be set to  $b = 1$  (the parametric value is 1, transforming  $[1, 2, 90]$  to  $[2, 3, 91]$ ) or  $b = 314$  (the parametric

 Talk like a ...

## computational cognitive neuroscientist

**Cognition** — the mental process of acquiring, evaluating and/or using knowledge

**Cognitive task** — an activity that involves processing and/or recalling information

**Computational model** — the expression of a set of mathematical functions or algorithms in the form of a computer program designed to emulate a natural process, such as cognition

**Intelligence** — the ability to acquire and flexibly apply knowledge and skills to solve new problems

**Language processing** — the ability to formulate and understand language

**Mind** — the expression of cognition implemented by a physical system, such as the brain

**Neuron** — a nerve cell that transmits messages to and from other neurons and other cells in the body and brain

**Neuroscience** — the study of the nervous system, including the brain


**Stimulus** — any input that can be processed to generate an internal or external response by the nervous system

**Synapse** — the part of a neuron that communicates its signals to other neurons

value is 314, transforming  $[1, 2, 90]$  to  $[315, 316, 404]$ ). **Computational models** are built from mathematical functions linked together into algorithms with specific parametric variables. A model can then be used to explore how the assignment of different inputs to its variables affect its behaviour. This is helpful, as models can quickly become too complex to solve mathematically; in this case, a computer program can be run to see how the model behaves with different inputs.

## How do scientists use computational models?

Scientists aim to create computational models that reproduce the patterns of behaviour exhibited by what is being modelled, such as human performance in a task. This involves exploring which functions to use, how to arrange these functions into algorithms, and testing different parametric values for each variable until the behaviour of the model matches the observed data (for example, how



© H\_Ko/Shutterstock.com

someone behaves when performing a task). Once scientists have created a model that replicates observed data, it can be used to generate predictions of future data.

Computational models allow scientists to investigate how different components in complex systems interact by providing a way to mathematically express their theories of how they think the system may behave. When creating a computational model, scientists are forced to acknowledge all the assumptions on which their theory relies, as each assumption must be translated into an element of the model. Then, scientists can use the model to test the behaviour of the system, and any assumptions about it, by observing the system 'in action' in the form of a computer simulation. One common objection to computational models is that they make too many assumptions. However, the advantage of a model is that it makes those assumptions clear and concrete. One way to think about them is as hypotheses made precise by the model, that can then be tested in future experiments.

### **Weather forecasting: an example of a computational model**

Some natural phenomena, such as the orbit of the Earth around the Sun, can be well described by a single mathematical equation. This means that scientists can accurately calculate the position of our planet for any given day in the future. However, other phenomena, such as weather systems, involve complex dynamic interactions among a vast number of components (such as the molecules in the atmosphere) that cannot be simplified into a single equation to describe the behaviour of the full system. This means it is hard to predict future weather. However, scientists can tackle this complexity using computational models to forecast what the future weather might be like.

Weather is influenced by a variety of variables, including land temperature, air temperature, humidity, air pressure, wind speed and direction, and topographic features (such as mountains and

water bodies), all of which influence the interactions among the molecules that make up the weather. Meteorologists create computational models of weather systems to examine how these different variables lead to different weather outcomes. For example, it is likely to rain if the temperature and humidity are both high but then the temperature drops, or if the wind carries air over a water body and then over a mountain. These statements can be converted into algorithms and combined to create computational weather models.

These models are tested on previously observed weather data. For example, if yesterday's conditions of temperature, pressure, humidity, etc., are entered into the model as inputs, and the model produces an output that matches today's weather, then this indicates that the functions, algorithms and parametric values used in the model do a good job of predicting future weather. If they do not accurately predict the weather, they can be modified in an effort to improve the model, which can then be used to better predict future weather.

Predictions of complex phenomena are always challenging, in part because there are so many interacting variables, and because uncertainties in the model can multiply and overwhelm initial assumptions. This means that computational models of the weather can make predictions with high accuracy for only a short distance into the future. For example, while the forecast of the weather in two hours' time or tomorrow will be fairly accurate, a prediction of the weather in ten days' time is likely to be less accurate.

### **Modelling minds: a computational model of the human brain**

Like the weather, the functioning of the mind involves highly complex interactions over a vast number of parts - in this case, the neurons and synapses of the brain. To better understand these, scientists attempt to recreate them using computational models. At their core, computational models of the human brain are algorithms, many

of which are designed and trained with the goal of reproducing intelligent human behaviour. For example, a language processing model could use the previous words in this sentence to predict what word is going to appear \_\_\_\_\_. If you guessed 'next' or 'here', then the language processing mechanisms in your brain work in a similar way to most people's when making language predictions. These mechanisms are accurately predicted by many existing computational models of language processing.

Computational cognitive neuroscience is an emerging field that seeks to understand human cognition by building computational models that express theories of how the brain gives rise to cognition. To test whether these models operate in a similar way to the human mind, scientists compare their performance with findings from work in cognitive psychology studying human performance.

For over a century, psychologists have conducted experiments that use cognitive tasks to characterise the mental functions people use when solving problems, generating a wealth of data on human perception, memory, decision making, language learning, emotions, etc. These experiments have also revealed how the mind functions when the brain is disrupted, for example through injury or psychiatric conditions. In more recent years, neuroscientists have uncovered the structural organisation of the brain at different scales, from individual neurons to large brain regions. They have characterised the functions that different brain structures perform in tasks such as perception, memory and movement.

These data are combined to build computational models of the human brain, enabling scientists to test how the operation of different brain mechanisms gives rise to different psychological functions. Understanding this relationship will not only provide insights into how disruptions to brain functions can be corrected but will also help scientists understand how to better implement intelligence in machines.

# Examples of computational cognitive neuroscience models

## Reward learning and prediction error

Many of the ways in which we exhibit intelligence are not exclusive to humans. Animals are born into complex environments where they learn about potential rewards (e.g., food) and dangers (e.g., predators). While millions of years of evolution have programmed certain instincts into animals (such as 'seek food' and 'run away from mouths with lots of sharp teeth'), they also exhibit the capacity for adaptive learning.

The ability of (human and non-human) animals to adaptively learn how to behave can be simplified into a computational model. For example, a simple model might simulate how an animal holds a recent sensory input (such as a scent) in its memory, then uses that information as a predictor for the outcome of what happens next. For example, if the animal smells a berry then finds food, or smells a fox then almost gets eaten, then the next time the animal encounters that same sensory input, it can predict that a similar outcome will occur and either approach the reward or avoid the danger. When the prediction is accurate (the animal smells another berry then finds more food), this strengthens the association between the input stimulus and the response, while if the prediction is wrong, the association is weakened. Over time, the animal adaptively learns what stimuli will lead to different rewards and adversities.

However, while this sort of stimulus-response learning can lead to efficient behaviour, it can be inflexible. What happens if the circumstances change (e.g., the bush contains no more berries)?

## Cognitive control and flexibility

While such association formation is a powerful and ubiquitous form of learning in animals, it can also be inflexible. Humans exhibit a remarkable ability to flexibly overcome such associations — and even basic reflexes — when needed. For example, if a person is instructed *not* to scratch an itch, they are able to do so (though it may be uncomfortable!). Simple as this seems, it shows that with a simple instruction, a person can instantaneously overcome the millions of years of evolution that produced that reflex (even reptiles scratch itches!). This ability is not matched by any other species.

This flexibility can also be demonstrated using a simple cognitive experiment known as the Stroop task. In this task, a person is shown the written name of a colour displayed in a different colour than the word (e.g., **GREEN**). If they are told to respond to this stimulus out loud, they almost universally do so by reading the word (i.e., saying 'green') and almost never report the colour. This is because the association between the written word and verbal response is useful in many settings (e.g., reading) and so is strongly learnt, much like the association of the location of a berry bush with reward.

However, if the person is instructed to report the colour (i.e., say 'red'), they can do so (although they may find it challenging). This relies on the capacity for cognitive control — the ability to take an instruction (or internal intention) and use it to override a strongly learnt association or instinctive reflex to achieve a different goal (e.g., naming the colour or not worsening the inflammation causing the itch).

## Cognitive control and task evaluation

When performing tasks, there is a trade-off between flexibility and efficiency. The human brain has a remarkable ability to switch between these two modes of operation, from efficient behaviour based on past learning to more flexible forms of behaviour that, although not quite as efficient, allow us to adapt quickly and effectively to new situations. One example of this is typing: almost anyone can quickly learn to use a keyboard, by typing one figure at a time. This works for any keyboard, whether it is on a computer or a piano, but it is not very efficient. With practice, people can learn to use a particular keyboard with tremendous efficiency, but they might not transfer very well to a different keyboard (e.g., from a computer to a musical instrument). One of the primary functions of cognitive control is to decide when each type of behaviour is best, and guide behaviour accordingly.

## Human memory

Memories are characterised by their duration (**long-term** memories last days to years; **short-term** memories last minutes to hours; **working** memories are actively maintained in your current awareness) and function (**episodic** memories are recollections of specific events, such as your tenth birthday party; **semantic** memories are ones about regularities and structure in the world, such the knowledge that birds fly while mammals (for the most part) do not).

Different kinds of memory serve different functions. For example, episodic memories will help you make predictions in specific contexts, while semantic memories will help you make predictions across a wide variety of contexts. Long-term memories contain information that is occasionally useful and can be retrieved from storage when that information is relevant, while short-term memories contain information that was recently accessed and still needed for the current task (either newly encountered information or a long-term memory that was recently retrieved from storage). Another important role of cognitive control is to decide which type of memory is most useful for a particular problem, and to use that memory system for storage and/or retrieval.

Computational modelling has been useful in gaining a better understanding of how the brain carries out all these functions, from reward learning to the role of cognitive control in the use of memory for planning and problem solving, all of which contribute to the capacity for intelligent behaviour.



# Tools for computational modelling

The hardware and software used for computational models of intelligence are continually evolving as computer scientists, psychologists and neuroscientists work to solve problems and drive progress in the field. As computational practices change over time, with computational tools and programming languages continually evolving, computational modelling faces many challenges regarding the development of models in each field.

Ideally, there would be a universal programming language that could be used by all researchers who were interested in building models of the human brain and intelligence, which would make it easy for them to exchange and make use of each other's models. In reality, there are many different types of models in use, written in many different programming languages, and each research project must consider trade-offs and make decisions that will influence the investigation and resulting computational models.

To ensure continued progress in the field of computational cognitive neuroscience, it would be helpful if these many different computational

models of human brain function and intelligence could be brought into the same framework so they could be evaluated, deconstructed and recombined into new models. PsyNeuLink is one experimental effort to do this.

## What is PsyNeuLink?

PsyNeuLink is a platform developed by researchers in the Neuroscience of Cognitive Control Lab at Princeton University in the US. It is a modelling environment for building the computational models used by neuroscientists, psychologists and computer scientists to explore brain functions and psychological behaviours, and their relationship with artificial intelligence. In PsyNeuLink, scientists can construct different components of their models and integrate them into the modelling environment, where they can run the models as simulations and analyse their interactions. The aim is to develop a common repository for model-sharing, allowing scientists working on different computational cognitive neuroscience projects to access the best modelling tools for uncovering how the brain gives rise to human intelligence, and how this can be implemented in machines.

## Pathway from school to computational cognitive neuroscience

- As computational cognitive neuroscience is an interdisciplinary field combining psychology, neuroscience and computer science, a degree in any of these subjects could lead to a career in the field.
- A wide range of other subjects is also relevant for uncovering the mechanisms behind intelligence. For example, members of the Neuroscience of Cognitive Control Lab come from many different backgrounds, including machine learning, data science, electrical engineering, medicine, physics and philosophy.
- Coding and computer programming are key skills for developing computational models, so look for free online tutorials where you can learn, such as Code Academy: [www.codecademy.com](http://www.codecademy.com)



**Neuroscience of Cognitive Control Lab**  
Princeton Neuroscience Institute,  
Princeton University, USA

**Fields of research**  
Neuroscience, Psychology,  
Computational Modelling

**Research project**  
Investigating the neuroscientific and psychological mechanisms behind cognition and intelligence using computational models

**Funders**  
US National Institutes of Health (NIH),  
National Science Foundation (NSF),  
Office of Naval Research

**Websites**  
[ncclab.princeton.edu](http://ncclab.princeton.edu)  
[www.psyneulink.org](http://www.psyneulink.org)

# Meet members of the Neuroscience of Cognitive Control Lab



## Dr Javier Masís

**Role:** Postdoctoral researcher

**Fields of research:** Neuroscience, Cognitive Science

**I have always been very curious about the world.** I am not only interested in science, but also in history, politics, literature, art and music. I chose to become a researcher in the hope of contributing a little bit to humanity's body of knowledge.

**As a cognitive scientist, I study learning** and treat it as an investment, and I investigate when it is worthwhile to make that investment and when it isn't. For example, if your teacher suddenly announces there will be a test tomorrow, you could lose sleep trying to review all your notes, but you might not be able to retain much information. However, if your teacher announces there will be a test next week, then you could dedicate some

time every day to revise and you are likely to learn and remember more, so your learning investment will be more worthwhile.

**I use a combination of simulations and behavioural measurements during games** to study how we decide to invest in learning. I've found that this decision depends on several factors, such as how much time you have to learn, how fast you think you'll learn, what your reward for learning will be and how far in the future it is. I've also discovered that there's a key trade-off between satisfying immediate needs (i.e., instant gratification) and how quickly you'll learn a new task.

**Science is about communication and collaboration,** plus some well-intended conflict, to resolve which theories are correct and which are incorrect. I love working with my colleagues as there's nothing more satisfying than encountering curious, energetic people who want to work together to figure out cool stuff!

**You don't have to actually study brains to 'study the brain.'** Research is so vast that there

are many different ways to contribute to the field of neuroscience. For example, while I am technically a brain scientist, I hardly ever deal with specific brain regions. Instead, I study the problems the brain solves, which generates important findings for other scientists who study the brain itself.

**When I'm not studying neuroscience, I play drums in an alternative rock band.** We are called Burne Holiday ([www.burneholiday.com](http://www.burneholiday.com)), we have been playing together since college and we now have two albums and several cool cover songs. I love playing with the group!

## Javier's top tip

Cultivate your curiosity! The world becomes a much more interesting place when you start wondering how things work and why things are the way they are.



## Tyler Giallanza

**Role:** Graduate student

**Field of research:** Cognitive Psychology

**I was interested in many topics when I was younger,** including reading fiction, economics, computer science, robotics and debating. I didn't know what I wanted to do with my career, so I explored all these different interests through high school. It wasn't until my second year of college that I decided I wanted to be a cognitive scientist.

**I work on understanding how memory can be used in abstract problem solving.** I am particularly

interested in how humans can learn so quickly from such little data – something that is currently very challenging for AI. I've discovered that we can take advantage of the relationships between concepts we've already learnt, which helps us learn new skills quickly, but current AI is less able to make connections between its past experiences.


**I enjoy uncovering the 'algorithms' humans use to learn and make decisions.** It was originally thought that humans struggle to multitask because our mental capacity is limited, but recent research hypothesises that it is the flexible nature of cognition that limits our ability to do multiple things at the same time. The human brain is good at reusing algorithms for different purposes. However, this causes issues if we try to simultaneously do different activities that use the same algorithms. I think this theory is interesting as it connects common experiences (such as the

inability to drive safely while texting) to underlying principles.

**In my free time, I enjoy rock climbing and cooking.** As a graduate student, my mind is very busy when I'm working, so these activities help me clear my mind and relax.

## Tyler's top tip

If you want to become a scientist, don't just study science. Studying non-scientific subjects, such as English and history, has been unexpectedly useful in my scientific career. Scientific discoveries require creative and critical thinking, rather than fact memorisation, and at school, these skills were most exercised in my humanities classes.



### Dr Harrison Ritz

**Role:** Postdoctoral researcher

**Field of research:** Computational Cognitive Neuroscience

**I've always been interested in music.** I played in indie bands during high school and college, and I'm an avid music collector. Sports were also a big part of my life in high school – I ran cross country and played soccer. Academically, I've wanted to study psychology for as long as I can remember. The more I learnt about psychology, the more interested I was.

**I study how we switch between thinking in different ways.** For example, if you switch between conversations in different languages, what's going on in your brain during that transition? To investigate this, my team has developed a new cognitive task for measuring changes in attention that involves someone reading text while listening to someone else speaking, forcing them to switch

between different tasks. By analysing people's behaviour as they completed this cognitive task, we discovered that they follow two distinct steps when switching between reading and listening. First, they decide what to think about, then they gradually switch their attention to that task. The ability to measure these steps separately is important for developing a theory of how people switch their attention.

**I enjoy the excitement of finding patterns in human behaviour that help us uncover the structure of the human mind.** With the right experiments, you can understand then predict how people will behave in different situations. It feels like you are doing physics or chemistry experiments, but with the brain!

**I think it's incredible that we can now decode the contents of the mind.** By measuring someone's brain activity, we can now reconstruct simple aspects of their thoughts. For example, imagine you are shown an X or an O, then asked to think about which shape you were looking at. By analysing images of your brain activity, scientists can tell (with high accuracy) whether you are

thinking of an X or an O. It's like mind-reading!

**In my free time, I enjoy jogging and hiking.** Running really helps me clear my head and improves my sleep. This year I ran my first half marathon. My wife and I often go hiking on weekends; there are lots of great trails in New Jersey. I also love to watch movies. Sometimes my friends and I sneak into the Princeton lecture halls in the evening and watch movies on the big projector screens!

---

### Harrison's top tips

1. There are many ways to succeed, but the best direction towards progress is to push yourself to try things that don't come naturally to you.
2. Psychology and neuroscience increasingly depend on computational modelling, so develop your quantitative skills and take classes in programming and statistics.
3. Get involved in research labs to learn what research involves and whether this is the right career path for you.