


Researchers study which parts of the brain are engaged when a person evaluates a computer program

December 22 2022, by Steve Nadis



Credit: Pixabay/CC0 Public Domain

Functional magnetic resonance imaging (fMRI), which measures changes in blood flow throughout the brain, has been used over the past

couple of decades for a variety of applications, including "functional anatomy"—a way of determining which brain areas are switched on when a person carries out a particular task. fMRI has been used to look at people's brains while they're doing all sorts of things—working out math problems, learning foreign languages, playing chess, improvising on the piano, doing crossword puzzles, and even watching TV shows like "Curb Your Enthusiasm."

One pursuit that's received little attention is [computer programming](#)—both the chore of writing code and the equally confounding task of trying to understand a piece of already-written code. "Given the importance that computer programs have assumed in our everyday lives," says Shashank Srikant, a Ph.D. student in MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL), "that's surely worth looking into. So many people are dealing with code these days—reading, writing, designing, debugging—but no one really knows what's going on in their heads when that happens."

Fortunately, he has made some "headway" in that direction in a paper—written with MIT colleagues Benjamin Lipkin (the paper's other lead author, along with Srikant), Anna Ivanova, Evelina Fedorenko, and Una-May O'Reilly—that was [presented earlier this month](#) at the Neural Information Processing Systems Conference held in New Orleans.

The new paper built on [a 2020 study](#), written by many of the same authors, which used fMRI to monitor the brains of programmers as they "comprehended" small pieces, or snippets, of code. (Comprehension, in this case, means looking at a snippet and correctly determining the result of the computation performed by the snippet.)

The 2020 work showed that code comprehension did not consistently activate the language system, [brain regions](#) that handle language processing, explains Fedorenko, a [brain](#) and cognitive sciences (BCS)

professor and a coauthor of the earlier study. "Instead, the multiple demand network—a brain system that is linked to general reasoning and supports domains like mathematical and logical thinking—was strongly active." The current work, which also utilizes MRI scans of programmers, takes "a deeper dive," she says, seeking to obtain more fine-grained information.

Whereas the previous study looked at 20 to 30 people to determine which brain systems, on average, are relied upon to comprehend code, the new research looks at the brain activity of individual programmers as they process specific elements of a computer program. Suppose, for instance, that there's a one-line piece of code that involves word manipulation and a separate piece of code that entails a mathematical operation.

"Can I go from the activity we see in the brains, the actual brain signals, to try to reverse-engineer and figure out what, specifically, the programmer was looking at?" Srikant asks. "This would reveal what information pertaining to programs is uniquely encoded in our brains." To neuroscientists, he notes, a physical property is considered "encoded" if they can infer that property by looking at someone's brain signals.

Take, for instance, a loop—an instruction within a program to repeat a specific operation until the desired result is achieved—or a branch, a different type of programming instruction that can cause the computer to switch from one operation to another. Based on the patterns of brain activity that were observed, the group could tell whether someone was evaluating a piece of code involving a loop or a branch. The researchers could also tell whether the code related to words or mathematical symbols, and whether someone was reading actual code or merely a written description of that code.

That addressed a first question that an investigator might ask as to

whether something is, in fact, encoded. If the answer is yes, the next question might be: where is it encoded? In the above-cited cases—loops or branches, words or math, code or a description thereof—brain activation levels were found to be comparable in both the language system and the multiple demand network.

A noticeable difference was observed, however, when it came to code properties related to what's called dynamic analysis.

Programs can have "static" properties—such as the number of numerals in a sequence—that do not change over time. "But programs can also have a dynamic aspect, such as the number of times a loop runs," Srikant says. "I can't always read a piece of code and know, in advance, what the run time of that program will be." The MIT researchers found that for dynamic analysis, information is encoded much better in the multiple demand network than it is in the language processing center. That finding was one clue in their quest to see how code comprehension is distributed throughout the brain—which parts are involved and which ones assume a bigger role in certain aspects of that task.

The team carried out a second set of experiments, which incorporated machine learning models called neural networks that were specifically trained on computer programs. These models have been successful, in recent years, in helping programmers complete pieces of code. What the group wanted to find out was whether the brain signals seen in their study when participants were examining pieces of code resembled the patterns of activation observed when neural networks analyzed the same piece of code. And the answer they arrived at was a qualified yes.

"If you put a piece of code into the neural network, it produces a list of numbers that tells you, in some way, what the program is all about," Srikant says. Brain scans of people studying computer programs similarly produce a list of numbers. When a program is dominated by

branching, for example, "you see a distinct pattern of brain activity," he adds, "and you see a similar pattern when the machine learning model tries to understand that same snippet."

Mariya Toneva of the Max Planck Institute for Software Systems considers findings like this "particularly exciting. They raise the possibility of using computational models of code to better understand what happens in our brains as we read programs," she says.

The MIT scientists are definitely intrigued by the connections they've uncovered, which shed light on how discrete pieces of [computer programs](#) are encoded in the brain. But they don't yet know what these recently-gleaned insights can tell us about how people carry out more elaborate plans in the real world.

Completing tasks of this sort—such as going to the movies, which requires checking showtimes, arranging for transportation, purchasing tickets, and so forth—could not be handled by a single unit of code and just a single algorithm. Successful execution of such a plan would instead require "composition"—stringing together various snippets and algorithms into a sensible sequence that leads to something new, just like assembling individual bars of music in order to make a song or even a symphony. Creating models of [code](#) composition, says O'Reilly, a principal research scientist at CSAIL, "is beyond our grasp at the moment."

Lipkin, a BCS Ph.D. student, considers this the next logical step—figuring out how to "combine simple operations to build complex programs and use those strategies to effectively address general reasoning tasks." He further believes that some of the progress toward that goal achieved by the team so far owes to its interdisciplinary makeup.

"We were able to draw from individual experiences with program analysis and neural signal processing, as well as combined work on machine learning and natural language processing," Lipkin says. "These types of collaborations are becoming increasingly common as neuro- and computer scientists join forces on the quest towards understanding and building general intelligence."

More information: Paper: [Convergent Representations of Computer Programs in Human and Artificial Neural Networks](#)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Researchers study which parts of the brain are engaged when a person evaluates a computer program (2022, December 22) retrieved 20 April 2024 from <https://techxplore.com/news/2022-12-brain-engaged-person.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--