

This Computer's Always in School

Thanks to the supercomputer that is our brain, we can make lightning-fast inferences and associations between images and situations. For a real computer, though, the same task is a bit harder. That kind of advanced visual processing requires significant artificial intelligence (AI) — the ability to perform humanlike cognitive tasks such as reasoning, generalizing and learning from past experience.

Yet, since summer 2013, NEIL — the Never Ending Image Learner — has been hard at work at Carnegie Mellon University analyzing and forming relationships between images from all over the Internet. The better the system gets, the closer we are to truly powerful AI and a new era of smart technology. Made up of two computer clusters housing a total of 200 processing cores, NEIL is programmed to organize its database into three categories: objects (such as computer or Corolla), scenes (alley or church) and attributes (blue or modern).

Researchers left NEIL to itself to analyze online images, using an algorithm that allows it to build connections — the heart of its AI. Those connections include: object-object relationships (“eye is part of baby”); scene-object relationships (“bus is found in bus depot”); object-attribute relationships (“pizza has round shape”); and scene-attribute relationships (“alleys are narrow”). NEIL then adds these relationships to its database, giving it more data so it can become even better at finding new associations.

“Gathering visual common sense is an extremely difficult problem,” says Abhinav Gupta, principal investigator on the NEIL project. “The problem is considered to be among the hardest in all of AI because the breadth and richness of common sense is enormous.”

Gupta adds that it's important to develop strategies, like NEIL's learning algorithms, which allow computers to recognize, categorize and respond to images, as the machines become more incorporated into our lives. “Over the past decade, AI researchers have made tremendous advances in the field of computer vision — for example, object and scene recognition. NEIL is a small step toward the long-term dream of making truly intelligent machines.”

While NEIL may one day learn to make new kinds of connections — and Gupta's team hopes to develop novel applications of the software — there's no real endpoint to the project. “In a manner similar to humans,” Gupta says, “we expect NEIL to keep learning in a never-ending fashion.” So far, NEIL has analyzed more than 10 million images and created 5,000 likely relationships between them. As some of the examples at right show, sometimes NEIL does a great job linking the concepts behind images, and sometimes ... not so much. —MICHAEL FRANCO

DOES AN IMAGE-LEARNING A.I. MAKE THE GRADE?

Some of NEIL's greatest hits (and misses)

SPOT ON

Parasol can look similar to flower.



Gondola can be found in Venice.



Candle can be a kind of lantern.



Supernova can look similar to blizzard.



Book can be part of an accordion.




Pac-man can look similar to comet.



Cathedral can be a kind of boomerang.



Turkey can have a part of gorilla.



Axe can be a kind of hairbrush.



TOTAL MISSES