

Algorithm makes the process of comparing 3-D scans up to 1,000 times faster

June 18 2018, by Rob Matheson



MIT researchers describe a machine-learning algorithm that can register brain scans and other 3-D images more than 1,000 times more quickly using novel learning techniques. Credit: Massachusetts Institute of Technology

Medical image registration is a common technique that involves

overlaying two images, such as magnetic resonance imaging (MRI) scans, to compare and analyze anatomical differences in great detail. If a patient has a brain tumor, for instance, doctors can overlap a brain scan from several months ago onto a more recent scan to analyze small changes in the tumor's progress.

This process, however, can often take two hours or more, as traditional systems meticulously align each of potentially a million pixels in the combined scans. In a pair of upcoming conference papers, MIT researchers describe a machine-learning [algorithm](#) that can register [brain scans](#) and other 3-D images more than 1,000 times more quickly using novel learning techniques.

The algorithm works by "learning" while registering thousands of pairs of images. In doing so, it acquires information about how to align images and estimates some optimal alignment parameters. After training, it uses those parameters to map all pixels of one image to another, all at once. This reduces registration time to a minute or two using a normal computer, or less than a second using a GPU with comparable accuracy to state-of-the-art systems.

"The tasks of aligning a [brain](#) MRI shouldn't be that different when you're aligning one pair of brain MRIs or another," says co-author on both papers Guha Balakrishnan, a graduate student in MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and Department of Engineering and Computer Science (EECS). "There is information you should be able to carry over in how you do the alignment. If you're able to learn something from previous image registration, you can do a new task much faster and with the same accuracy."

The papers are being presented at the Conference on Computer Vision and Pattern Recognition (CVPR), held this week, and at the Medical Image Computing and Computer Assisted Interventions Conference

(MICCAI), held in September. Co-authors are: Adrian Dalca, a postdoc at Massachusetts General Hospital and CSAIL; Amy Zhao, a graduate student in CSAIL; Mert R. Sabuncu, a former CSAIL postdoc and now a professor at Cornell University; and John Guttag, the Dugald C. Jackson Professor in Electrical Engineering at MIT.

Retaining information

MRI scans are basically hundreds of stacked 2-D images that form massive 3-D images, called "volumes," containing a million or more 3-D pixels, called "voxels." Therefore, it's very time-consuming to align all voxels in the first volume with those in the second. Moreover, scans can come from different machines and have different spatial orientations, meaning matching voxels is even more computationally complex.

"You have two different images of two different brains, put them on top of each other, and you start wiggling one until one fits the other. Mathematically, this optimization procedure takes a long time," says Dalca, senior author on the CVPR paper and lead author on the MICCAI paper.

This process becomes particularly slow when analyzing scans from large populations. Neuroscientists analyzing variations in brain structures across hundreds of patients with a particular disease or condition, for instance, could potentially take hundreds of hours.

That's because those algorithms have one major flaw: They never learn. After each registration, they dismiss all data pertaining to voxel location. "Essentially, they start from scratch given a new pair of images," Balakrishnan says. "After 100 registrations, you should have learned something from the alignment. That's what we leverage."

The researchers' algorithm, called "VoxelMorph," is powered by a

convolutional neural network (CNN), a machine-learning approach commonly used for image processing. These networks consist of many nodes that process image and other information across several layers of computation.

In the CVPR paper, the researchers trained their algorithm on 7,000 publicly available MRI brain scans and then tested it on 250 additional scans.

During training, brain scans were fed into the algorithm in pairs. Using a CNN and modified computation layer called a spatial transformer, the method captures similarities of voxels in one MRI scan with voxels in the other scan. In doing so, the algorithm learns information about groups of voxels—such as anatomical shapes common to both scans—which it uses to calculate optimized parameters that can be applied to any scan pair.

When fed two new scans, a simple mathematical "function" uses those optimized parameters to rapidly calculate the exact alignment of every voxel in both scans. In short, the algorithm's CNN component gains all necessary information during training so that, during each new registration, the entire registration can be executed using one, easily computable function evaluation.

The researchers found their algorithm could accurately register all of their 250 test brain scans—those registered after the training set—within two minutes using a traditional central processing unit, and in under one second using a graphics processing unit.

Importantly, the algorithm is "unsupervised," meaning it doesn't require additional information beyond image data. Some registration algorithms incorporate CNN models but require a "ground truth," meaning another traditional algorithm is first run to compute accurate registrations. The

researchers' algorithm maintains its accuracy without that data.

The MICCAI paper develops a refined VoxelMorph algorithm that "says how sure we are about each registration," Balakrishnan says. It also guarantees the registration "smoothness," meaning it doesn't produce folds, holes, or general distortions in the composite image. The paper presents a mathematical model that validates the algorithm's accuracy using something called a Dice score, a standard metric to evaluate the accuracy of overlapped images. Across 17 brain regions, the refined VoxelMorph algorithm scored the same accuracy as a commonly used state-of-the-art registration algorithm, while providing runtime and methodological improvements.

Beyond brain scans

The speedy algorithm has a wide range of potential applications in addition to analyzing brain scans, the researchers say. MIT colleagues, for instance, are currently running the algorithm on lung images.

The algorithm could also pave the way for image [registration](#) during operations. Various scans of different qualities and speeds are currently used before or during some surgeries. But those images are not registered until after the operation. When resecting a brain tumor, for instance, surgeons sometimes scan a patient's brain before and after surgery to see if they've removed all the tumor. If any bit remains, they're back in the operating room.

With the [new algorithm](#), Dalca says, surgeons could potentially register scans in near real-time, getting a much clearer picture on their progress. "Today, they can't really overlap the [images](#) during surgery, because it will take two hours, and the surgery is ongoing" he says. "However, if it only takes a second, you can imagine that it could be feasible."

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: Algorithm makes the process of comparing 3-D scans up to 1,000 times faster (2018, June 18) retrieved 10 May 2024 from <https://techxplore.com/news/2018-06-algorithm-d-scans-faster.html>

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.