

Searching Is Intelligence

Image retrieval – the next revolution in pathology

By Hamid Tizhoosh

The human brain is the result of millions of years of evolution – and, as such, it's an extremely capable recognition machine. Every time we see somebody we know, we effortlessly recognize their face, an astonishing ability that we perceive as trivial thanks to our visual cortex (responsible for processing images). For machines, however, this has – until recently – been an impossible task.

Almost eighty billion neurons (each one connected to approximately ten thousand others, on average) serve our innate thinking and recognition abilities, so mimicking it is far from easy. Many details of image recognition in the central nervous system are still unknown, yet we may justifiably deduce that at least some, if not most, of our impressive cognitive

At a Glance

- *Artificial intelligence is increasingly advancing on pathology – but has yet to be implemented in the most practical ways*
- *One useful application of AI is image search and retrieval – a task that computers can perform much faster than humans*
- *New approaches using artificial neural networks can help overcome challenges with computer-based image recognition*
- *Content-based image retrieval may rely on AI, but it's a pathologist-centric application that cannot function without a human element*



capabilities are literally based on “recognition.” We re-identify an image that we have previously seen and, depending on the depth of the memory in which that image is stored, recognize it instantly (or after a short while, with some mental effort – for instance, when encountering someone we don't know well or have not seen for many years). Image information, in whatever format it may be stored in our brain, is certainly subject to sophisticated comparisons and inferences for the purpose of identification. Neuroscience will continue to amaze us with more discoveries and conclusions that we can hopefully translate into more capable algorithms for computer vision.

“Seeing” pathology

In medical image analysis, we have a large collection of computer algorithms that perform different operations on digital images: quality enhancement, filtering, registration, and segmentation, to mention just a few. The latter has been the focus of extensive research to quantify cell nucleus morphology and distribution. As important as these

measurements may be, they have not been able to bring about a disruptive change in diagnostic imaging. Why? Chiefly because conventional quantification is often fed into a “smart” algorithm to output a “classification” – a category of some sort, generally either a yes/no decision or some type of disease grading. As valuable as these quantifications may be, they have not fundamentally altered the diagnostic process, perhaps because such computer algorithms do not reduce uncertainty to increase pathologists' confidence in a diagnosis. More importantly, classification-oriented computer algorithms have not been able to truly assist pathologists because they provide no clues for writing the pathology report. And so the pathology community has instead turned to well-organized second opinions through telepathology to reduce inter-observer variability (an apparent manifestation of diagnostic error).

Image search, as an alternative approach to medical image analysis, offers the historical chance to perform “virtual telepathology,” consulting other pathologists by accessing their knowledge

without requiring their physical presence to examine specimens. It also allows us to consult not just one pathologist but as many as we would like within a given healthcare institution or network. Image search lets us access the expertise of multiple pathologists in a very short time and at much lower costs than doing so in person, or even via real-time telepathology. And it can establish a reliable framework to move toward quality control through computational consensus-building.

But why do we assign such immense expectations to image search? Although synaptic connections (with their binary states of excitatory and inhibitory) are the building blocks of the human brain, the actual inference is granular, fuzzy, implicit, and qualitative – as opposed to specific, certain, explicit, and quantitative – characteristics that seem to enable us to process highly complex, ambiguous information like variable tissue patterns and the intricacies of polymorphism. The diagnostic process commonly ends in writing a report, an activity we can describe as “computing with words.” The contradiction is that we – both the computer vision community and the artificial intelligence (AI) community – understand “computing” to mean merely crunching and producing numbers. We may ignore what algorithms do internally, but what they output could be decisive if it helps pathologists write better reports or have more confidence in their conclusions.

Given a large archive of diagnosed patients with corresponding data (images and reports on treatment and monitoring), we should be able to identify and retrieve images that are anatomically or pathologically similar to the biopsy sample of the patient being examined – as well as the annotated data for each case. The reports contain the medical knowledge of many other pathologists for similar cases, making them a treasure trove of high-quality diagnostic information. Next generation computer software may make

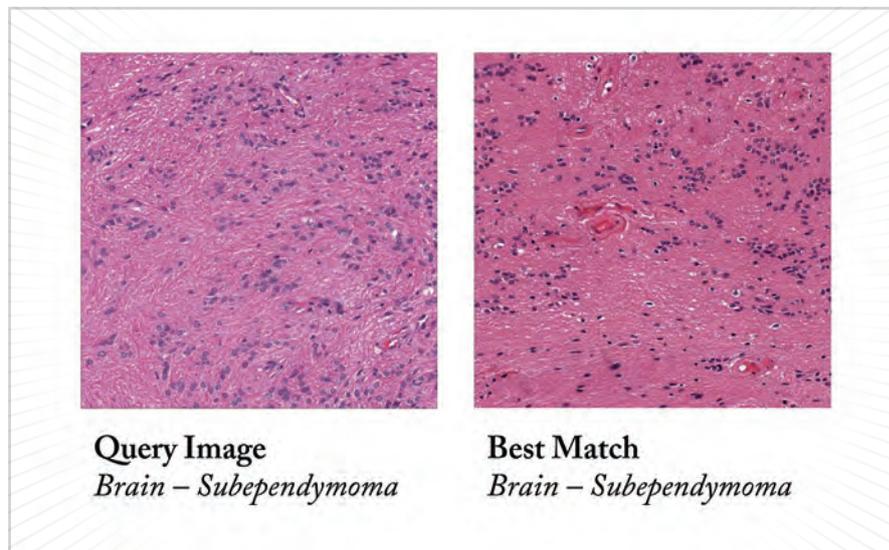


Figure 1. When given a query image (left), image search can find similar ones (right). This search was conducted among 2,000,000 patches extracted from scans of 300 patients with more than 85 conditions.

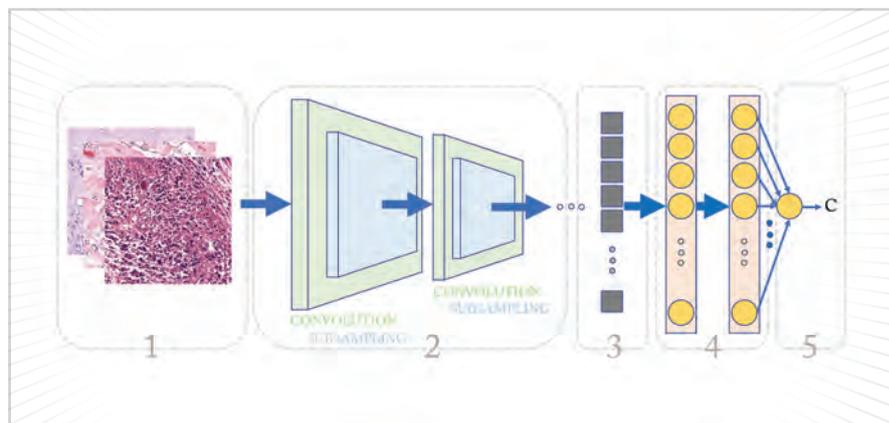


Figure 2. Schematic illustration of a deep convolutional neural network. 1) Many labeled – often meaning benign vs. malignant – images are used for training; 2) images go through many series of convolutions (image filtering) and subsamplings (image downsizing); 3) the end result of many convolutional layers is a large number of small image sections that capture significant information such as edges and corners; 4) all small image sections now go through “traditional” layers of artificial neurons; 5) one or more classification categories are assigned to each image.

the raw information directly available to the pathologist (showing retrieved images along with corresponding reports), or it may fuse the key information in retrieved reports to provide “auto-captioning” of whole slide images. The latter would even allow triaging and prioritization in real-time as glass slides go through digital scanners. The world of AI-based

image search opens up a vast range of options for advancing and optimizing the laboratory workflow.

Content-based image retrieval
Research into content-based image retrieval (CBIR) has been happening for almost three decades. So if our above expectations are justified, then why hasn't

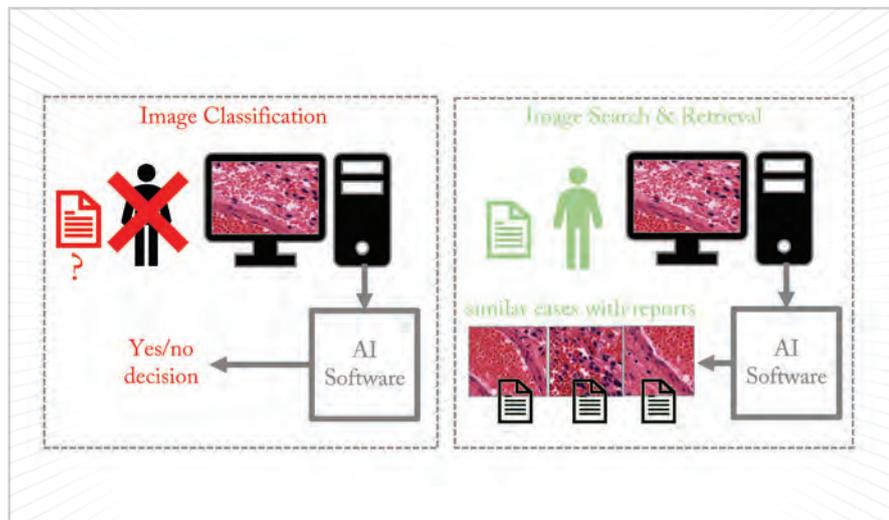


Figure 3. Left: AI-driven image classification makes decisions on behalf of the pathologist; it is not clear who should write the pathology report if an AI entity is in charge of diagnosis. Right: Image search strengthens the pathologist by providing similar images and their corresponding reports from archives.

CBIR delivered on these promises?

The most important reasons, from an engineering perspective, are computational and accuracy challenges. The former refers to the difficulty of performing image matching in large archives in real time; the latter is about matching images properly so that the identified images are actually similar to the query image (see Figure 1). But from a digital pathology perspective, the obstacles are slightly different. To us, the main reason CBIR systems haven't made it to the daily laboratory workflow is most likely the so-called "semantic gap." Image representations in computer vision are numerical and objective, whereas human pathologists use verbal and subjective representations that often can't be modeled or analyzed. The resulting gap between computers and human experts does not permit an unambiguous definition of similarity. Indeed, the semantic gap is arguably the paramount challenge in adopting CBIR into the laboratory workflow; the results of CBIR have not thus far been acceptable to pathologists. The path to the retrieved images is irrelevant if the pathologist doesn't agree that the matched images

are truly similar to the query image – a wrong answer is wrong, no matter how it was reached. But, in recent years, this has started to change; CBIR is going through a renaissance with the promise of a revolution.

AI is a general term used for a class of computer algorithms capable of instructional and sample-based learning. From its birth 70 years ago with some simple abstractions of the way a neuron operates in the human brain, AI has become an indispensable tool for computer vision applications. Most notably, artificial neural networks (ANNs) have gained great popularity due to their impressive recognition capability when implemented with many layers of artificial neurons (processing units that can perform simple aggregation of incoming synaptic values originating from other units). These "deep" ANNs recognize the content of a digital image by learning a compact representation of the image – an elegant encoding that we can assume to be a primal, but functioning, computational model for what happens to a retinal image when it

“Content-based image retrieval is going through a renaissance with the promise of a revolution.”

travels through the optic nerve to reach the visual cortex in the human brain.

Convolutional neural networks (CNNs) are among the most successful such solutions to extract relevant features from digital images (see Figure 2). A typical example is to learn 1,024 deep features to represent a face or an object depicted in a 240x240 image, reducing the information to less than 2 percent of its original size. To create such compact representations, deep networks usually adjust several hundred thousand artificial synapses to achieve their learning goal, a training process dominated by trial and error in the design phase and many hours or even days of actual training. Countless papers and articles report high recognition accuracies for face and object recognition using deep networks. Many papers have also begun to report similar findings for medical imaging in general, and for digital pathology in particular. Most, however, use deep features for the purpose of classification (that is, to tell us whether or not an image depicts a malignancy). Image search solutions in medical CBIR refrain from this approach.



Spotlight on the pathologist

Medical CBIR is fundamentally pathologist-centric, in contrast to classification-based AI, which essentially attempts to make decisions on behalf of the pathologist. You may be understandably opposed to the latter – but the former makes valuable use of AI solutions. Instead of letting CNNs and other deep ANNs use the extracted image representations (deep features) as a basis for a “yes/no” cancer classification (see Figure 3), we can use them to index and retrieve whole slide images, which draws upon several advantages. First, the image recognition capabilities of deep networks have empirically shown that the semantic gap between computer and human perceptions can be closed. Second, AI offers a multitude of versatile techniques for recognition, indexing and search. And third, advances in software and hardware have made it possible to perform millions of image comparisons in a fraction of a second. The fact that we are currently undergoing a transition from microscopy to digital pathology is just an amazing coincidence that further benefits computer vision adoption in pathology.

Despite the obvious opportunities, there are, of course, still many hurdles to overcome if we want to bring CBIR systems to pathology laboratories – not least the need for thorough and comprehensive validation of image search for different purposes in pathology. Unlike image classification, which can be validated in the engineering lab, image search cannot be validated without the presence and intensive involvement of pathologists. But there’s a silver lining to this cloud: the technology places the focus on human pathologists, rather than seeking to replace them. CBIR systems exist to help pathologists – and they cannot be designed and validated without our direct involvement. Moreover, once in use, they cannot continue to learn without pathologists at the heart of the process.

The design, validation, and regulatory clearance of image search solutions will certainly not happen overnight. In the meantime, we can identify practical use cases for image search that demonstrate how it can propel us toward computational consensus-building. With the recent success of AI in a multitude of computer vision applications and the rapid growth of digital pathology, we’re moving ever closer to the horizon of pathologist-computer partnerships.

Hamid Tizhoosh is the director of Kimia Lab (Laboratory for Knowledge Inference in Medical Image Analysis) in the Faculty of Engineering at University of Waterloo. He is also a member of the Waterloo AI Institute, and a faculty affiliate to the Vector Institute. As part of his commercial activities, he is presently the AI advisor of Huron Digital Pathology, St. Jacobs, Canada.