EchoNous Brings AI to Ultrasound

Niko Pagoulatos, PhD is the COO of EchoNous, a pioneer in applying the emerging field of artificial intelligence with the extreme miniaturization of ultrasound to solve common every day problems in healthcare.

The Catalyst Driving Ultrasound Usage

EchoNous, which is the Greek term for "intelligent ultrasound", was formed in mid 2015 on the basis that AI will be the catalyst for driving ultrasound usage to wider adoption making it a ubiquitous and accessible tool in medicine, one that can be used daily to solve common everyday problems in healthcare.

AI, if implemented properly and focused on the tasks associated with using an ultrasound-based device for specific clinical applications targeted, can literally revolutionize ultrasound usage by significantly reducing its learning and confidence curves thereby making it simple for providers of all kinds to learn to use it effectively and efficiently in their healthcare delivery.

EchoNous' Expertise in AI

EchoNous was founded on the premise of having arguably the best team of veterans in the extreme miniaturization of high-grade ultrasound technology.

We started our AI journey in early 2016 setting up an exclusive partnership with the University of Adelaide, more specifically the Australian Center of Visual Technologies (ACVT), a worldwide leader in AI technologies, particularly in deep learning which is one of the most promising areas in the field of AI.

ACVT has been participating in the Google and Stanford sponsored "Image Large Scale Visual Recognition" (ILSVR) challenge, a prestigious international competition with participation from leading academic institutions and companies including Qualcomm, Microsoft, Google and many other leading technology organizations.

The objective of the ILSVR challenge is the automatic object classification and detection on hundreds of object categories (e.g., piano, laptop, person) and millions of images. ACVT has been developing deep learning algorithms for the ILSVR competition and has been able to rank among the top 5 positions. This was what attracted us to them as a partner.

In the course of the last 2 years EchoNous has been working closely with ACVT; we have developed a robust framework and platform for developing deep learning algorithms to solve focused clinical problems. In parallel, EchoNous (Seattle) has been at work building its own local team of AI & deep learning experts, based at its Redmond headquarters.

This has been a strong success with the recent addition of two high-profile hires, a veteran machine learning scientist and 10-year veteran of Microsoft AI, and a Yale PhD with research focus on deep learning and ultrasound (https://globenewswire.com/news-release/2018/04/19/1481933/0/en/EchoNous-Hires-Two-Top-Machine-Learning-Scientists-To-Accelerate-Ongoing-Development-of-Intelligent-Medical-Tools.html).

Defining Deep Learning

As previously mentioned, deep learning is a highly successful approach in the field of AI and it has been applied to many applications and industries with high success. As an example, deep learning can enable computers to understand the content of images (what is called computer vision) and make rapid decisions based on this understanding; such technologies are used in many applications, the most notable of which are fully-autonomous cars and computer-based medical imaging interpretation. Likewise, deep learning can also be applied to natural language processing thus creating a fluid understanding of actual spoken language, with applications in voice assistants such as Alexa from Amazon.

Deep learning algorithms are based on the theory of neural networks. A neural network is a mathematical and computational framework inspired by the human cognition mechanisms powered by the network of neurons found in the human brain. A neural network is a network of nodes, with each node (also called perceptron - from the word perception) being analogous to a brain neuron. Deep learning leverages neural networks and it is based on an architecture of multiple layers of nodes/neurons totaling thousands or even millions of nodes/neurons, all of them working together and forming the basis for a computer to learn in similar ways as humans.

One of the most successful deep learning architectures used in applications involving imaging is the convolutional neural networks (CNN) architecture. CNNs consist of multiple layers of nodes/neurons, where the initial layers of the network recognize low-level image features such as edges, and the final layers of the network recognize high-level image features such as an organ like urinary bladder.

EchoNous' Development of a CNN

At EchoNous we have developed a CNN for the automatic detection of the urinary bladder using high quality ultrasound images acquired with Uscan, taking advantage of the "high spatial density fanning technique". Based on this CNN-based urinary bladder automatic detection we can compute the urinary bladder volume with much higher accuracy and repeatability.

The development of our CNN for automatic urinary bladder detection involved two steps:

- Training: In this step our CNN was trained to recognize the bladder based on a large amount of collected data. We ensured that the data we used for training was a proper representation of various parts of the population that would be encountered and scanned day to day in medicine by our users.

- Computational performance optimization of inference: Once the CNN is trained, it can be used to automatically identify bladders during real time imaging; this is what is known as the "inference step" in a CNN. We have optimized the computational performance of the "inference step" of our trained CNN to provide results sufficiently fast as dictated by the expected clinical workflow, and within the power consumption limits dictated by our hardware. To achieve this optimization we used cutting-edge processors, including current day GPUs, and cutting-edge deep learning computational frameworks.

It is important to understand that one of the key elements of a successful CNN for automatic bladder detection is the data used for training, both in terms of the amount of data used but also in terms of the data covering a broad patient population for different types of subjects that will be scanned with our device.

To ensure continuous training and improvement of our CNN bladder detection algorithm we have added an optional feature in Uscan that customers can choose, that enables automatic transmission of completely anonymized, deidentified and HIPAA compliant data to our servers. Uscan bladder data moved to our servers can then be used for improving the performance of our CNN.

We have also set up collaborations with medical institutions around the world that allow us to have access to Uscan bladder data targeted at specific categories of subjects, such as pregnant and post-partum women, where bladder identification is challenging. This approach allows us to improve our CNN's bladder detection performance in these challenging cases.

The EchoNous Mission

Automatic bladder detection is our first clinical application, in the process we have established a highly valuable and unique AI asset base, including an exclusive partnership with a leading AI academic institution, and a world-class Seattle-based AI team building a robust framework for the ongoing development of deep learning algorithms. We are leveraging all of these AI assets to develop additional AI algorithms for new (upcoming) ultrasound products and applications for vein access, as well as heart and lung functional assessment.

As history in medical innovation has repeatedly shown, the winning products are those that perform in a "frictionless" and focused manner precisely what customers are hiring them to do.

We at EchoNous are obsessed with empowering our customers to focus on their patients, rather than our products, with our ultimate mission aimed at developing highly-impactful clinical tools that solve common everyday problems in healthcare.

EchoNous

"AI Fraud" Detection Questionnaire

When inquiring as to the legitimacy of a medical device company's AI capability & truthfulness, we suggest asking these questions to which you should readily expect an answer. We are always happy to engage on these.

1. What field or discipline of AI are you using?

2. Are you using deep learning and if yes what type of deep learning? Is it CNN?

3. How did you train your deep learning algorithms?

4. What kind of data did you use to train the deep learning algorithms? How did you get access to the data used for training?

5. How are you able to run deep learning on your hardware during the bladder scanning procedure? Do you have capable enough processors on your device (GPUs, CPUs, DSPs) to run deep learning quick enough to provide timely results for the users?

6. How many machine learning scientists do you have on staff? For how long? Could we have a discussion with them? What experience do they have?