

Building a smarter interpolator with AI

INTERVIEW BY JESSICA MCKAY

EDITOR'S PICK FOR NOVEMBER

This month MRM met Akshay Chaudhari, Zhongnan Fang and Brian Hargreaves, whose recent paper “Super-resolution musculoskeletal MRI using deep learning” asked “how can we make a thick slice thinner?” Co-first author Zhongnan explains, “In many applications people prioritize the in-plane resolution, and therefore many sequences acquire thick slices. This can complicate matters for radiologists assessing oblique images, or surgeons planning an operation and needing to look at the whole structure, for example. Our project set out to establish whether the latest AI technology can reduce slice thickness without more scanning.”

6
Where most AI
in MRI is used
for detection,
we are using
it for image
acquisition.

– Zhongnan Fang

9

Co-authors Zhongnan and Akshay after winning the best poster award at the GPU Technology Conference in San Jose, California, for work featuring MRI super-resolution.

MRMH: Tell me about yourselves. How did you end up developing a deep learning technique for musculoskeletal (MSK) imaging?

Akshay: I am currently a first year postdoc at Stanford, where I also got my PhD in Bioengineering, focusing on fast, quantitative MSK imaging. In the past, we've tried many different methods to obtain a fast, diagnostic, quantitative MR sequence, such as parallel imaging and compressed sensing. All have tended to have limitations. Having already tried some traditional approaches, I was excited by new literature emerging in the field of computer vision and figured, why not give it a shot? They say, “There's no such thing as a free lunch”, but this might be the nearest we can get! So, we evaluated it, and luckily it works relatively well!

Zhongnan: I am a senior research scientist at LVIS Corporation. We are a startup out of Palo Alto working on medical image visualization. I also got my PhD at Stanford in Electrical Engineering, working on compressed sensing fMRI with Dr Jin Lee. My goal is to see images better reconstructed with less sampling. This project started at the 2017 ISMRM conference where I met Akshay. We had an idea, why not try this new AI technology with knee imaging and see how it works?

MRMH: Brian, I don't think I've ever heard the story of how you ended up doing MRI.

Brian: I came to Stanford in 1995. I was interested in coming back to graduate school to do something more satisfying than working in my previous consumer prod-



Chaudhari, A.S., Fang, Z., Kogan, F., Wood, J., Stevens, K.J., Gibbons, E.K., Lee, H.J., Gold, G.E., Hargreaves, B.A. Super-resolution musculoskeletal MRI using deep learning. *Magn Reson Med*. 2018;80: 2139-2154. DOI:10.1002/mrm.27178
<https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27178>

uct-oriented field. I took Prof. Dwight Nishimura's course and liked the way he taught. He very quickly got me interested in MRI and I've never left! Now we have this project involving high-resolution, quantitative knee imaging. We are looking at it from many angles,

from the acquisition right through to the reconstruction phase.

MRMH: How do you benefit in MSK imaging from going faster and getting quantitative data?

Akshay: There is plenty I could say in reply to that, but essentially MSK imaging has a clinical side and a research side. From the research perspective, we are interested in diseases, like osteoarthritis, which affect many collagen-rich soft-tissues. A lot of promising research suggests that quantitative MR parameters such as T_2 , $T_{1\rho}$, gagCEST, etc., are correlated to the collagen structure and can tell you about the biochemical status of these tissues over time.

Why do we want to faster? Because we can! Clinically, there's this big push towards value-based imaging since the high cost of MRI is a large burden on the healthcare system, and from a research perspective, it allows us to implement many promising new sequences in studies.

Brian: In the context of the osteoarthritis initiative, of course, we are studying a slow degenerative disease that involves numerous different factors. Arguably, the best approach would be to study changing patterns in different groups within a broader population. This is a huge motivation to bring down the cost of research scans. Shorter exams would cost less and offer a widespread test suitable for research purposes. Fighting the clinical cost is harder, for many reasons, so we should leave that discussion for now.

MRMH: Why did you choose to draw your data from this osteoarthritis initiative?

Akshay: Because it has a lot of data sets. With deep learning, you need a lot data, but it isn't easy to know how much is enough, so we started where there are thousands of datasets available. We only worked with ~170 to start with, knowing we could easily access more if needed.

MRMH: "Deep learning" is a big buzzword right now. How does your work compare with other great uses of AI in MRI?

Zhongnan: The similarity is that we are all trying to reduce the cost of MRI. However, where most AI in MRI is used for detection, we are using it for image acquisition, to get better image quality with less sampling.

Brian: I like to think of this work as a smart interpolator that uses multiple dimensions.

MRMH: What about the network? How did you decide what components and parameters to use?

Zhongnan: Our network consists of 20 layers of cascading convolutional layers and ReLU layers. Instead of learning the high-resolution image directly, we actually try to learn the residual between the high-resolution and interpolated images because the network converges faster when learning residuals than the high-resolution image itself.

MRMH: Why do so many neural networks use ReLU layers?

Zhongnan: The ReLU is a rectified linear unit that introduces nonlinearity into the function. ■



Akshay and Brian enjoying a well-earned beer at Akshay's Bioengineering Ph.D. graduation.

Akshay: It took me a long time to understand the significance of having nonlinearity in the neural network. Let's say the convolutional neural network has 20 layers and we look at the output of each individual one. The first convolution generates features like the edges in the x-direction, etc. In the next layer, the convolution gathers information about how these edges sit relative to one another. As you cascade through different convolutions, you pick up features of different levels. The ReLU is important because it can introduce nonlinearity into your system, allowing the network to learn high-level features. Without the nonlinearity, 20 convolutional layers could be represented using a single convolution.

MRMH: Looking back on this work, what have been your biggest successes and your biggest frustrations?

Zhongnan: I enjoyed working on this project, which was a great collaboration between LVIS and Stanford. We got feedback from many people from different fields and with different expertise. The outcome was positive and we published the paper! Training the network could be frustrating though! Sometimes, even after 3-4 days of training, your image was no better than the linear interpolation. I would have to go back to the source code and debug everything from the beginning. But that was also a fun part: seeing how AI technology can help us get a smarter interpolator.

Akshay: I loved the fact that this project had a very technical component but we also worked closely with clinicians. It was fascinating to be constantly getting feedback helping us to establish what was "good enough". After all, you can keep optimizing these networks forever, but it's important to learn to recognize when to stop.

Brian: I find it really exciting that AI is being led by people who are younger and newer to the field. It's great to have this fresh energy that is pushing the limits and asking some difficult questions. ■

“
You can keep
optimizing
these networks
forever, but
it's important
to learn to
recognize when
to stop.”

– Akshay Chaudhari