FEBRUARY 2022

Computer Vision News

The Magazine of the Algorithm Community

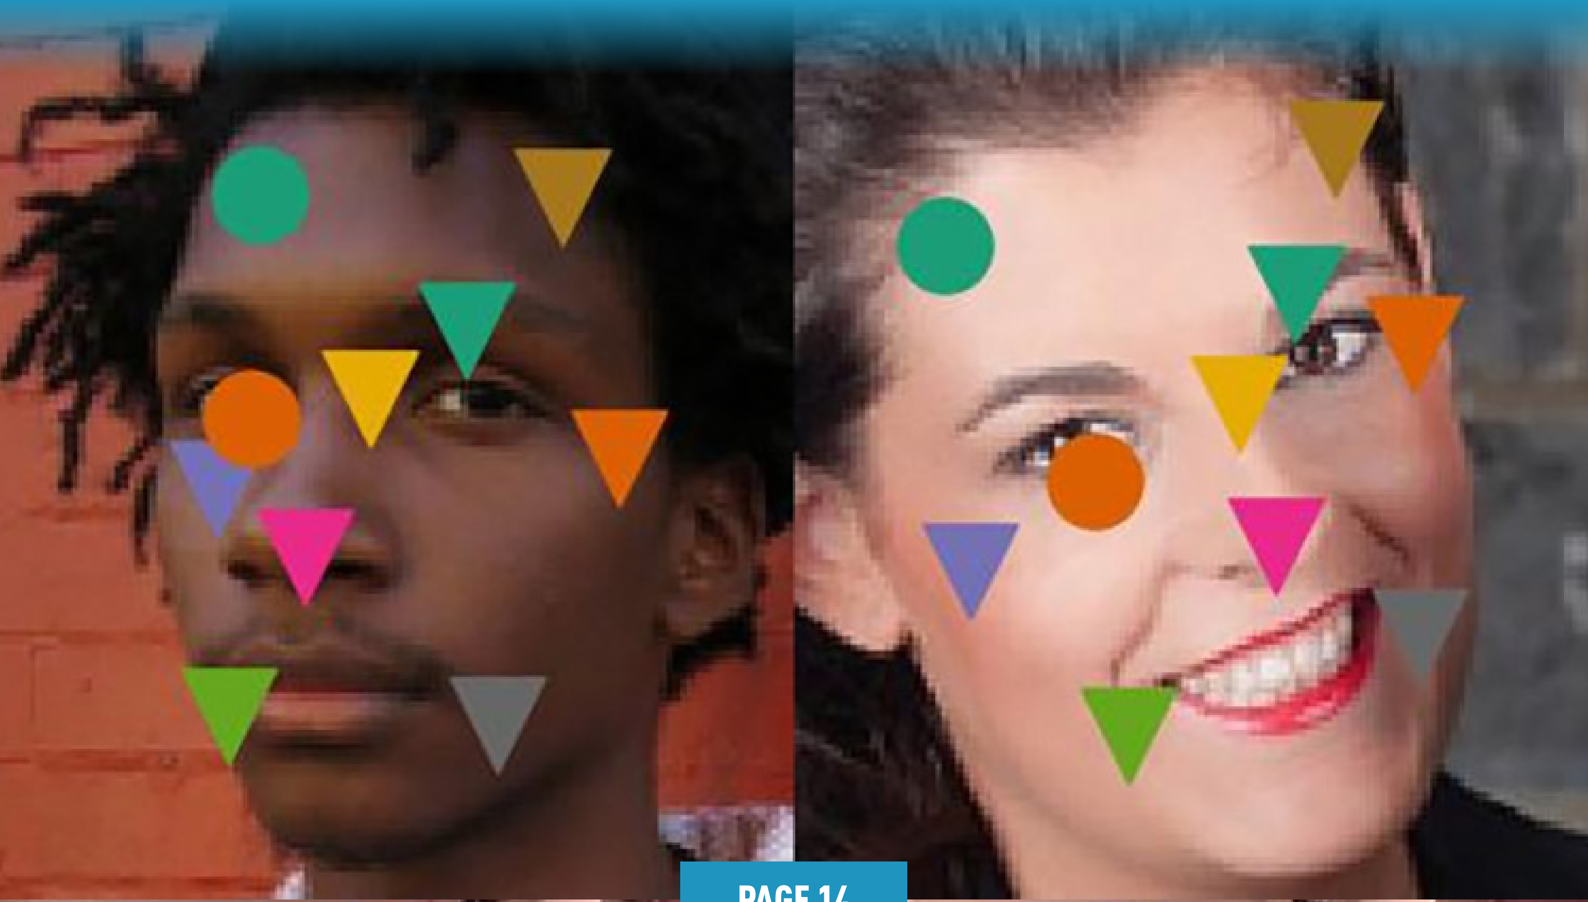

PAGE 14

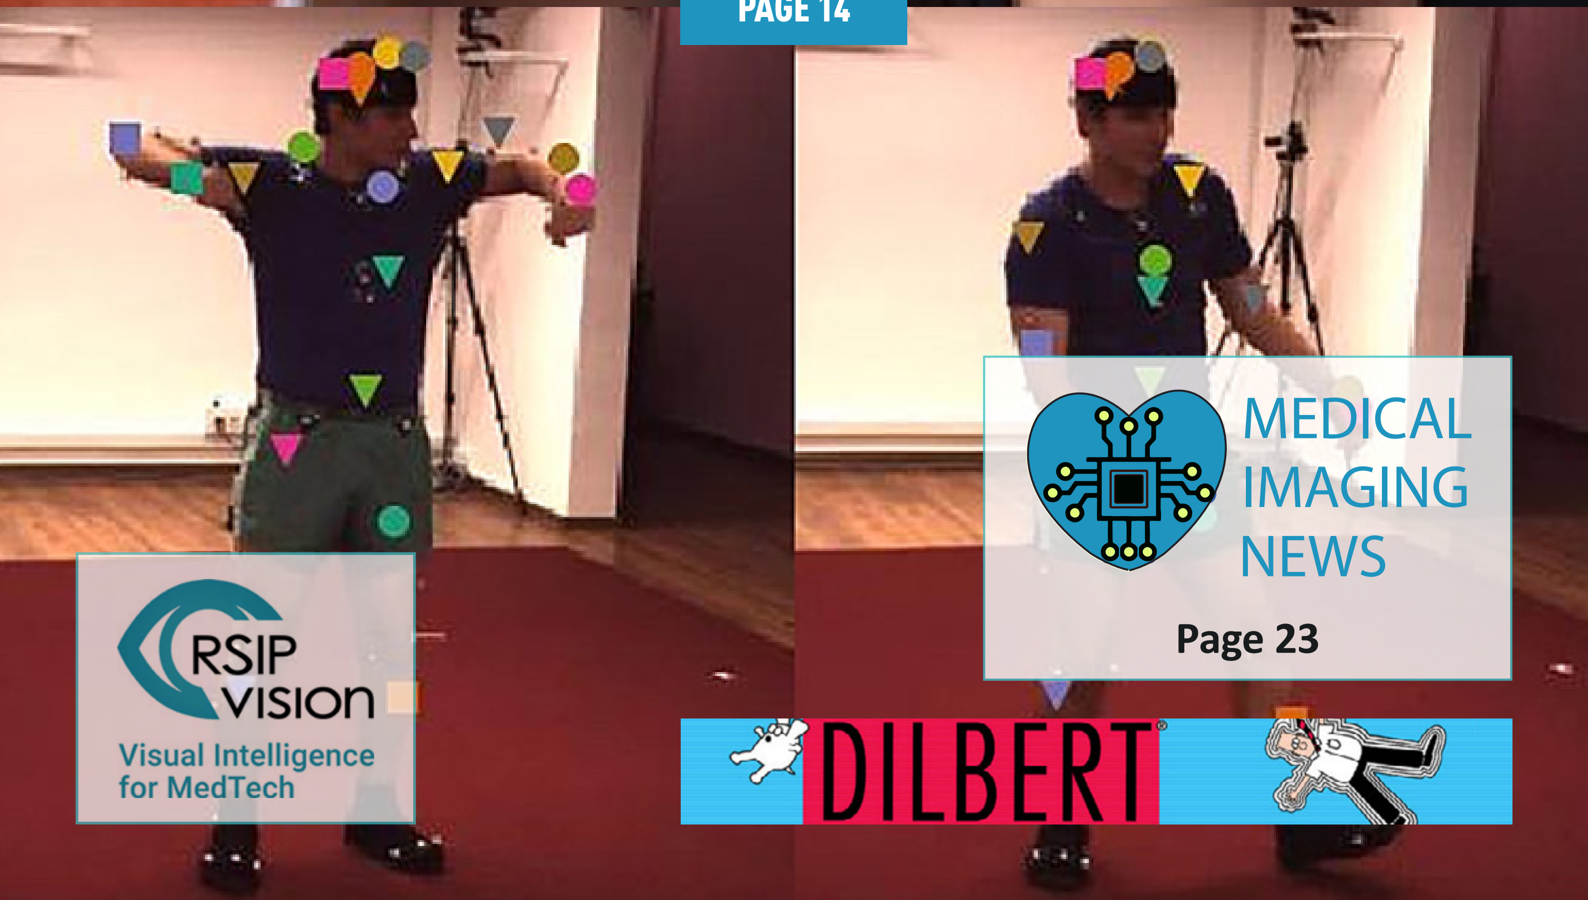

MEDICAL IMAGING NEWS

Page 23

RSIP vision
Visual Intelligence for MedTech

DILBERT

Computer Vision News

Editor:
Ralph Anzarouth

Engineering Editors:
Marica Muffoletto
Ioannis Valasakis

Publisher:
RSIP Vision

Contact us

Free subscription

Read Previous Magazines

Dear reader,

Last month saw **WACV 2022**, the **Winter Conference on Applications of Computer Vision**, take place in Hawaii. In this February issue of **Computer Vision News**, we report on three great papers from the event. Starting over the page with the **Best Paper Award** winner, **Agree to Disagree - When Deep Learning Models with Identical Architectures Produce Distinct Explanations**, authored by a top team from **Durham University**. For two more impressive papers, head to page 8 in the main magazine and page 48 in our supplement, Medical Imaging News.

Our research of the month is a fascinating paper in the field of AI for cardiology: **Automated Echocardiographic Detection of Severe Coronary Artery Disease Using Artificial Intelligence**. Can an AI system be developed to automate stress echocardiography analysis and support clinical interpretation? Apparently, yes, and providing automated classifications to clinicians when reading stress echocardiograms could improve accuracy, inter-reader agreement, and more. We have a full report by **Marica Muffoletto** on page 30.

A great friend of our magazine, well known in both the **CVPR and MICCAI communities**, is both a Professor at **Harvard Medical School** and a Research Scientist at **MIT**. We are honored to host **Adrian Dalca**, who speaks to us about his work on machine learning models for medical image analysis. With such a busy professional life, he reveals he has started a repository for all the ideas he doesn't have time for, so if you're short of research inspiration, you don't want to miss his exclusive interview on page 24!

Enjoy reading all these articles and many more in this exciting February issue of Computer Vision News. Tell your friends and colleagues about us and **bring us along for your next MedTech project!**

Ralph Anzarouth
Editor, **Computer Vision News**
Marketing Manager, **RSIP Vision**

Follow Us

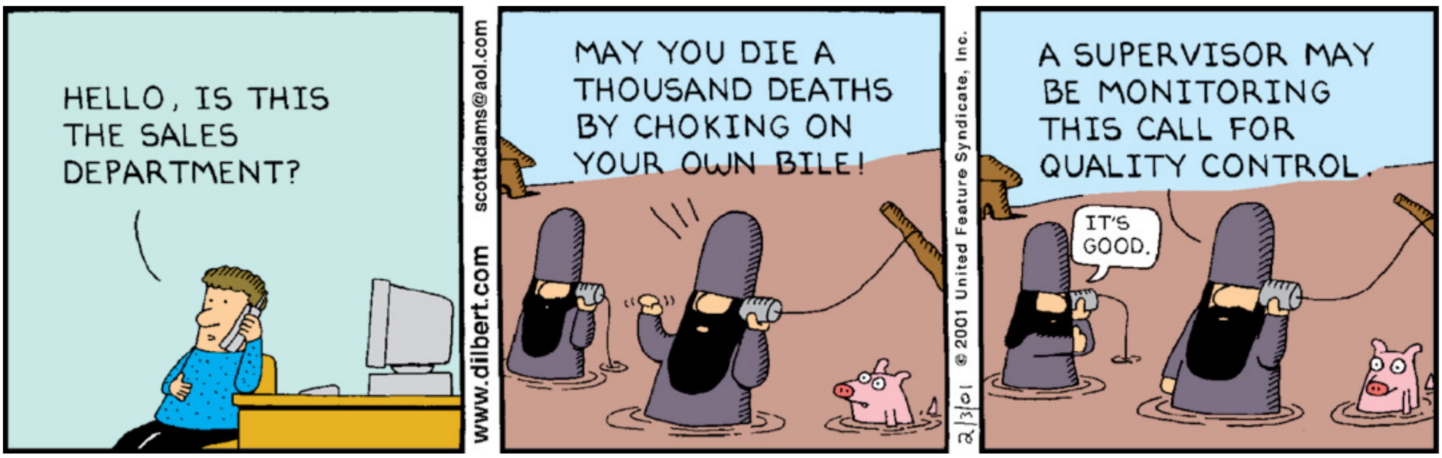

Computer Vision News

04


08
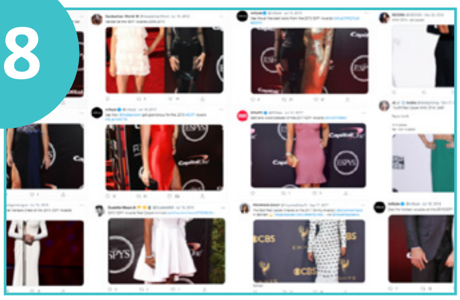

12


14


Medical Imaging News

24


30
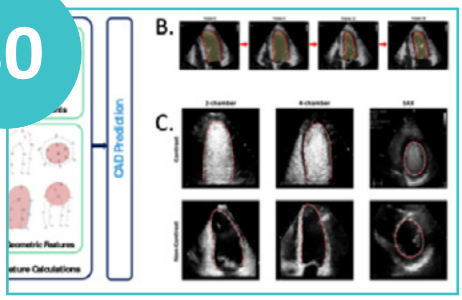

36
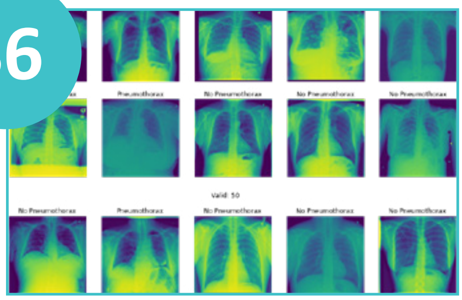

44
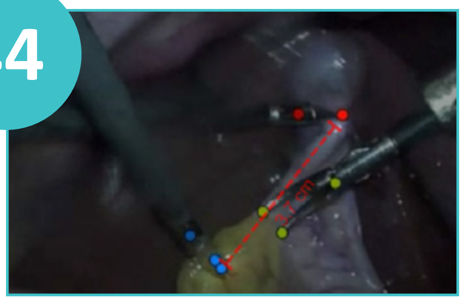

04 **Agree to Disagree ...**
Best Paper Award at WACV 2022

08 **Auditing Saliency Cropping Algorithms**
Paper Presentation from WACV 2022

12 **AI Spotlight News**

14 **Marianna Meo**
Women in Computer Vision

24 **Adrian Dalca - MIT / Harvard MS**
Exclusive Interview

30 **Automated Echocardiographic Detection...**
AI Research in Cardiology by Marica Muffoletto

36 **FMI and Deep Learning in Medical Imaging**
Medical Imaging Tools by Ioannis Valasakis

44 **Intra-op Virtual Measurements**
Laparoscopic and Robotic-Assisted Surgeries

AGREE TO DISAGREE:

WHEN DEEP LEARNING MODELS WITH IDENTICAL ARCHITECTURES PRODUCE DISTINCT EXPLANATIONS

Matthew Watson is a third-year PhD student under the supervision of Assistant Professor Noura Al Moubayed at Durham University. Bashar Awwad Shiekh Hasan is a visiting researcher at Durham and a speech scientist at Amazon Alexa. Their paper highlighting the problems with current explainability techniques for deep learning models, for which Matthew was first author, won the Best Paper prize at WACV 2022 last month in Hawaii. All three are here to tell us more about their award-winning work.

Deep learning methods have enjoyed much success in recent years, but some argue models can be inconsistent, unreliable, and untrustworthy, particularly in sensitive applications. It can be difficult to understand their inner workings, so **explainability techniques are now being used to open up the black box of a machine learning model** and explore what lies beneath.

However, this paper shows on multiple datasets and models that explainability methods themselves have major disadvantages and are inconsistent with each other. It demonstrates that a very minor modification of the models, such as changing dropout, adding noise, or shuffling the data, can produce different explanations.

"Current explainability methods probably aren't suitable to be applied in the real

world just yet," Matthew tells us. "I think it will be quite a while before we see suitable models that can be used in very sensitive scenarios, such as healthcare, where the decisions need to be perfect."

The team used the **MNIST handwritten digit dataset** for baseline results before extending their experiments to more sensitive applications using **MIMIC-CXR-JPG**, a large publicly available chest X-ray dataset.

"With the big tech companies, their applications are not very sensitive," Bashar points out. *"For them, it doesn't matter as much if you train the model a couple of times and there are small variations in the output. If you're classifying cats versus dogs, for example, you can afford to get it slightly wrong. But with chest X-ray, it's very important we get it right. The same problem occurs with both models, but clearly the impact is higher for medical applications."*

Noura adds: *"It will take a long time to see deep learning with explainability applied in industry in the real world. One major problem is industry itself is leading most of the research, financially at least, and the easy and successful applications always get more funding. Obviously, where the work is trying to break those models and show how vulnerable they are,*

Bashar Awwad Shiekh Hasan

Matthew Watson

Noura Al Moubayed

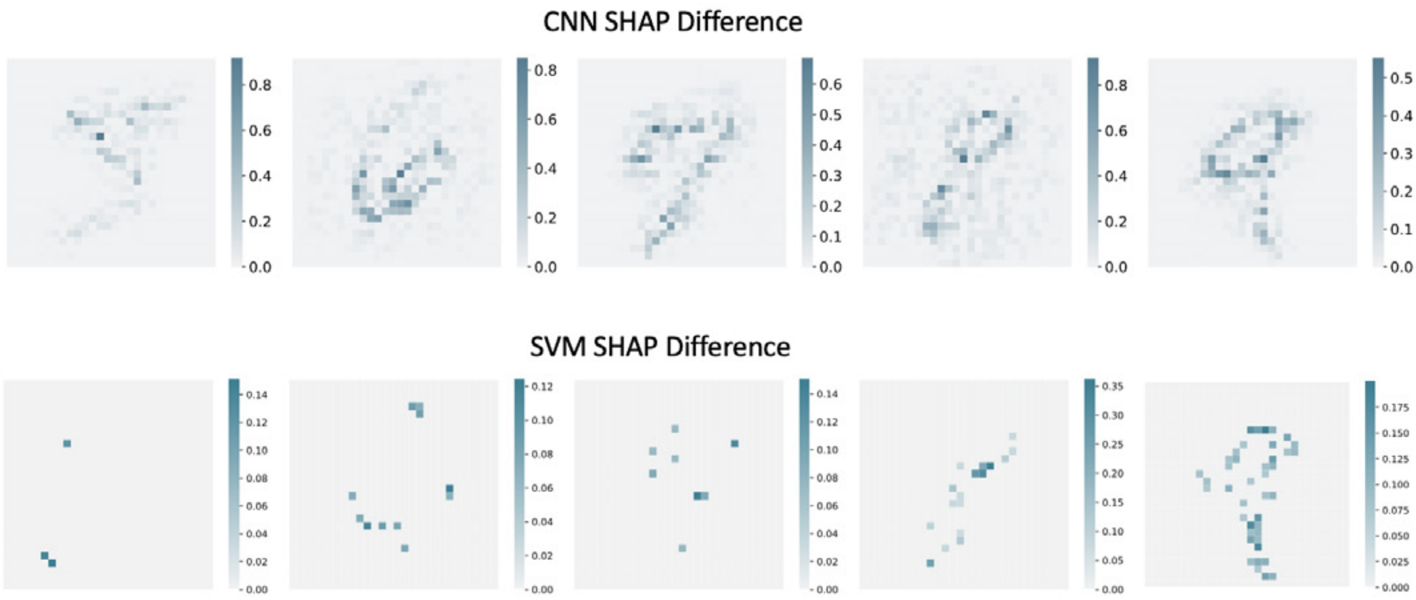

that gets less money!"

This work builds on previous work where the team used **explainability methods to detect adversarial attacks**. They were surprised to find by using explanations they could detect adversarial attacks that look very similar to the genuine data. Digging deeper and studying explanations and how consistent they were, they found they did not just detect adversary attacks, but also detected **minor changes in models**.

With the work taking home the coveted **Best Paper award at WACV 2022**, what do they all think convinced the jury?

"We're targeting a problem people don't think much about," Bashar tells us. *"It has always been there, but for some reason, the success of deep learning methods made it less obvious for people. We highlighted that problem and with the application, domain and approach we took, we made it really clear and concise."*

Matthew agrees: *"The area of explainability is very popular right now, so being able to expose some of the problems in this area is extremely important. We did that with very thorough experiments and achieved clear results."*

Matthew's interest in the explainability of machine learning grew out of being intrigued by understanding exactly what was going on inside a model.

"It's fascinating to get more of an idea of how these machine learning models work, but why they don't work is perhaps even more interesting," he observes. *"Once you figure out it doesn't work in some specific way it gives you more of an idea of how to make it better."*

Bashar describes how he came to the field:

"Previously, I worked in FinTech and experienced some of these problems in real life. We had very limited data to train on and explainability was very

important. When you're making financial decisions, you must be able to explain to your customers why you've made them. I've been wanting to solve these problems ever since!"

Outside of this paper, Noura works on a number of sensitive applications, including healthcare more generally, as well as being the **Head of Machine Learning at Evergreen**, a UK healthcare organization. She also works with a local FinTech company on financial data and has a project with her PhD students looking at detecting and reducing bias in facial recognition systems.

"All three are very sensitive applications, which is why we're doing more research on explainability and exposing their vulnerabilities to try to make them more robust," she tells us.

Noura and Bashar hail from Damascus in Syria originally. Bashar's day job is as a speech scientist at **Amazon**. He is working on building the next releases of **automatic speech recognition (ASR) models for Alexa**.

Matthew, who is from the UK, says his work is focused on the explainability of machine learning models.

"How explainability can be used to make better models, how these models can be more robust, and how this can improve the application of machine learning to healthcare," he reveals. He is currently working on another paper applying this work specifically to healthcare data, which the team are hoping to submit to **MICCAI** later this year. They are also working on a solution for producing more consistent explanations to take them a step further towards finding more **reliable explainable methods for deep learning in real-world applications**.

And we hope to see them all in the real world at MICCAI in Singapore this September!

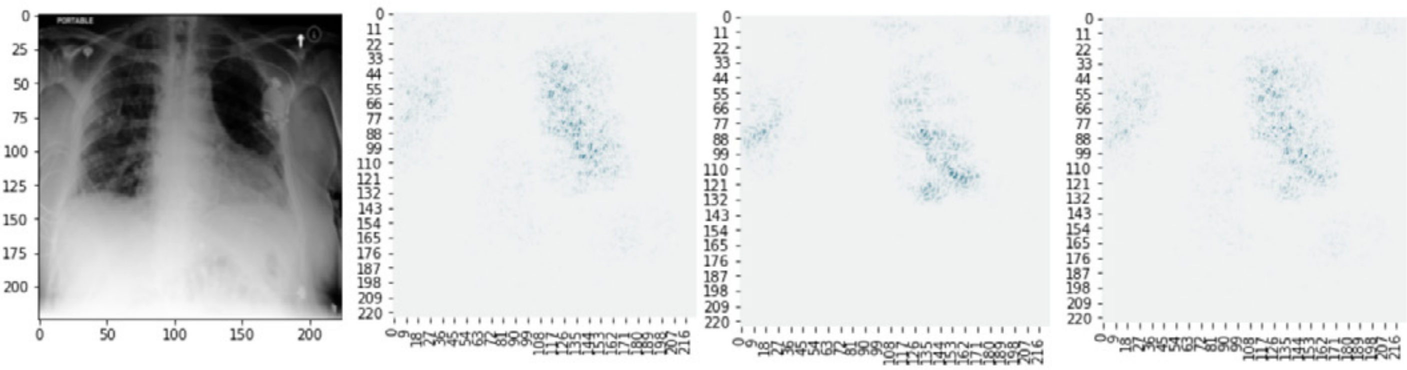

AUDITING SALIENCY CROPPING ALGORITHMS

Abeba Birhane just completed her PhD in Cognitive Science at the School of Computer Science, University College Dublin & Lero. Last month, she presented a paper at WACV 2022 in Hawaii, marking the first audit of saliency cropping algorithms carried out on three major tech platforms: Twitter, Apple, and Google. She speaks to us about this important and timely work.

In this work, Abeba collaborated with **computer vision researchers** to carry out a number of experiments to test **male-gaze-like artifacts and racial bias** in cropped images on **Twitter**, **Apple** and **Google**.

Ethiopian-born Abeba is a **cognitive scientist by training** but says she has been drawn more and more to the computer vision side of AI. Her associates on this work, **Vinay Uday Prabhu** and **John Whaley** from UnifyID Labs, are two experts in computer vision.

"My work sits at the AI end of cognitive science," she tells us. *"You can find all kinds of fascinating intersections between computer vision and cognitive science. Within this research, I'm particularly interested in how vision models present intelligence, emotion, or interestingness."*

The first the team knew that platforms such as Twitter use saliency cropping algorithms was when an image of US Senator Mitch McConnell and former President Obama went viral for all the wrong reasons. A Twitter user put photos of the pair at either end of a blank image to see which element the algorithm chose to present when cropped. It favoured McConnell almost every time.

"Many people don't realise these kinds of algorithms are created and deployed almost everywhere in these big ***social media and technology companies****,"* Abeba warns. *"Even within* ***vision research****, people don't know about it, because it's kept quiet, and the code and data aren't open access. It's difficult to carry out this kind of work because most of the datasets are protected by proprietary rights."*

The team carried out two experiments. The first was on the male-gaze-like artifacts.

"We created a Twitter account and put images through the platform to see what came out," Abeba explains. *"We kept noticing when images of women on the red carpet at the Emmy Awards, for*

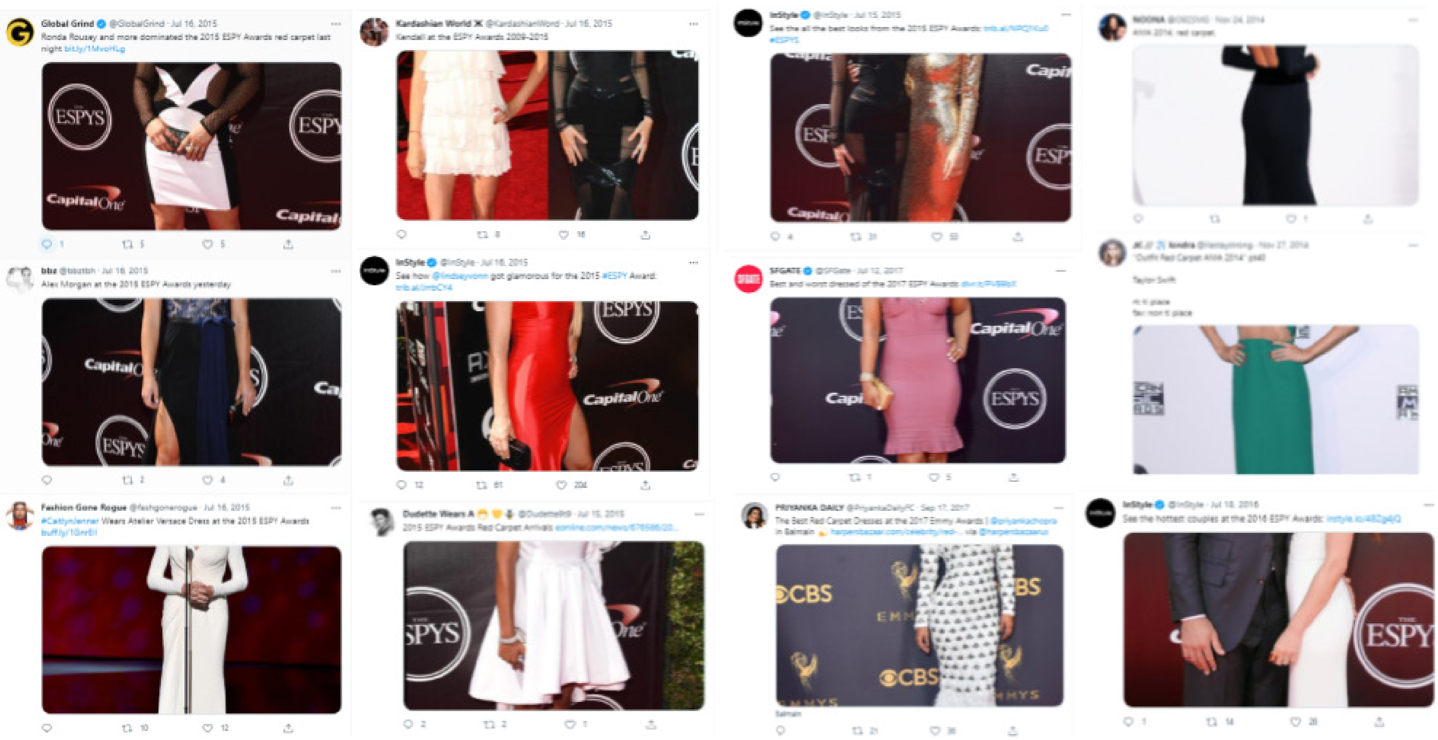

A collage of real-world user-uploaded images on Twitter that exhibited male-gaze-like artifacts

example, were passed through ***Twitter's saliency cropping algorithms****, it was cropping them in a way that was focused below the chest and above the knee. This is what experts term the male gaze. It's that part of a woman's body which tends to be objectified."*

The team curated its own dataset of 336 images of women at red carpet events and passed each image through the three cropping algorithms of Twitter, Apple, and Google. For Twitter in particular, as many as 79% of the images came out with a male-gaze-like crop.

The team also observed racial bias in the results, similar to the viral Obama-McConnell results. Based on that observation, they ran a second experiment.

"We used the ***Chicago Face Database****, with pre-labeled images for race and gender,"* Abeba tells us. *"We created the same 3 x 1 grid images and passed those through the three platforms again. Google's result was inconclusive in that 20% of the time it was just selecting the white space. Other times it selected the black face and other times the white face. But with Apple and Twitter, we found the algorithm preferred white faces over black faces."*

Part of the team's motivation for carrying out this work is that **these cropping algorithms exist and impact us every day**. It is very likely images we upload, or images we encounter when using our phones and computers have been through them.

Abeba looked at the reasons why companies and vision researchers use

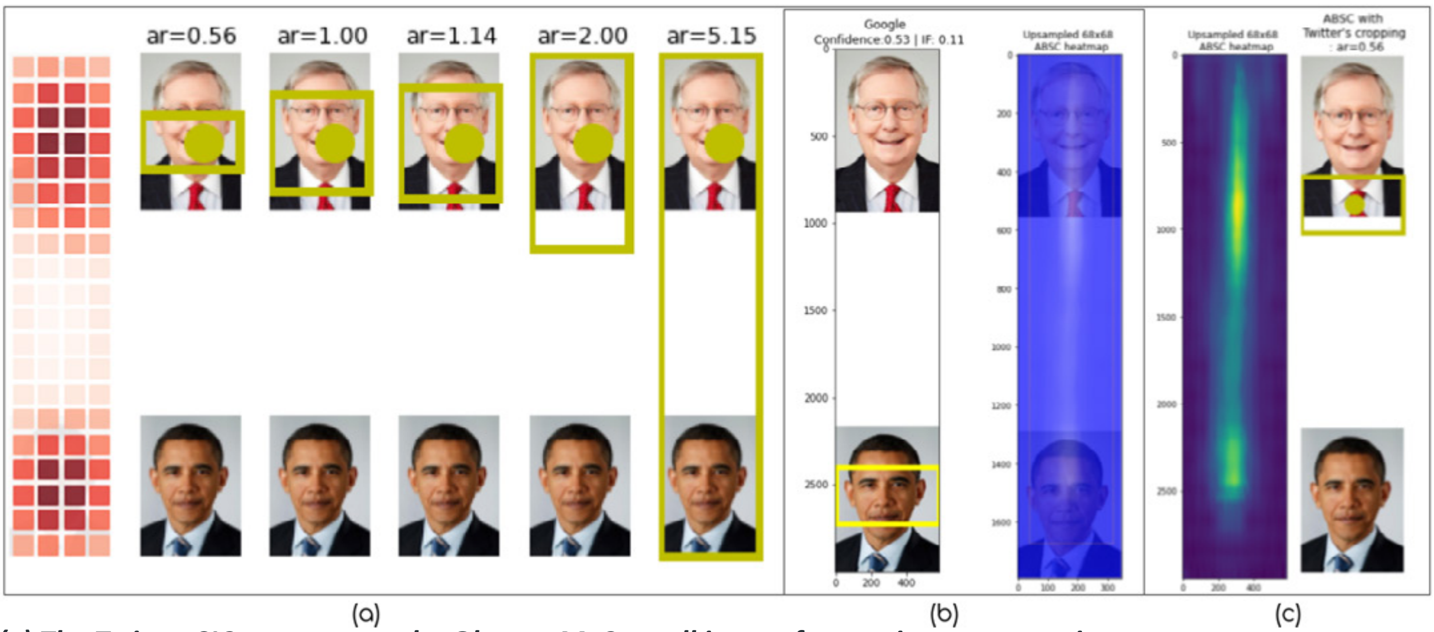

(a) The Twitter SIC response to the Obama-McConnell image for varying aspect ratios;
(b) The Google CROP_HINTS response to the Obama-McConnell image;
(c) The Apple ABSC response to the Obama-McConnell image.

cropping algorithms in the first place and found no consistency in the responses

"Some of reasons were even contradictory," she noted. *"There's no scientifically grounded reason for creating and deploying cropping algorithms on major platforms because they're all over the place and the science is very shaky."*

Twitter carried out its own audit and published a paper just before the team released a version of this work last year which was accepted to the **BeyondFairCV workshop at CVPR 2021**. It claims to have stopped using the algorithm on its platform. Abeba says that remains to be tested but recognizes Twitter is an exemplar overall because they are doing the required work, while companies like Apple, Google, and Facebook are still a closed book.

"It's encouraging Twitter have opened up the data, but they haven't answered many questions about the data," she argues. *"It's great they're carrying out their own audits, but it's important to let external auditors look at their code and algorithms. We need greater transparency. Maybe regulation is needed too to open source these algorithms. Our work marks the first audit work, but there is so much more left to do."*

We asked co-author (and old friend of the magazine) **Vinay Prabhu** to say a word about the computer vision work done for this research and he told us that the team has made the MGL-336 test-set accessible **here**.

"We are excited to see how the researchers working on ***saliency estimation and saliency cropping*** *in the computer vision community will utilize this test-set in their academic explorations"*, he added. *"Besides this, during the course of* ***WACV 2022****, we received plentiful feedback from computer vision researchers hailing from the two worlds of saliency research and we have summarized the resultant ideas and future directions of research* **on my github.** *We are hopeful that the computer vision community at large will take some interest into these exciting leads."*

Abeba says her research is interdisciplinary, ranging from the study of cognition to constructing cognitive models. More recently, she has been exploring audit work, auditing large image, text, and multimodal datasets.

"I'm getting more and more into ***auditing datasets****, but this is not by choice, it's more borne out of frustration!"* she reveals. *"With AI, because there's so much focus on producing the flashiest state-of-the-art models, people tend to ignore the datasets underlying these models. But they're crucial in* ***how accurate and how well your model performs****."*

Computer Vision News has found great new stories, written somewhere else by somebody else. We share them with you, adding a short comment. **Enjoy!**

University of Amsterdam Spinoff Launches Elderly Monitoring Tech

We have already talked in these pages about AI applications aimed at **monitoring elderly activities** to improve the protection of this vulnerable population. What is special at **Kepler Vision Technologies** is that **privacy of the patient** is protected in that the nurse does not see the video itself: sensor and software solution are designed to transform a monitored video feed into text to alert the care staff when they are needed. For example, a video feed showing an elderly resident fallen on the floor is indicated by the message *"In room 13 a client has fallen who needs help"*. They claim to eliminate 99% of false alarms. **Read More**

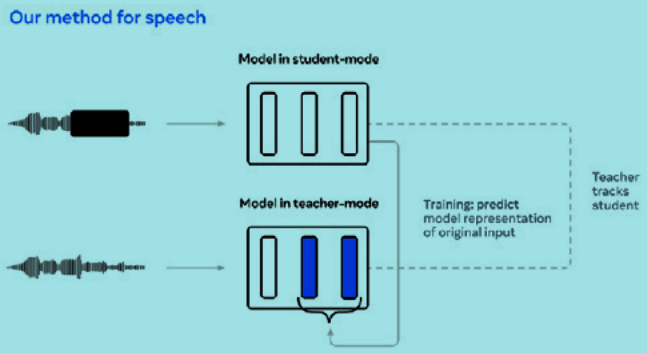

This 'Seasam' Drone Can Autonomously Follow Divers and Performs Underwater Tasks

As you know, this magazine is very fond of **scuba diving**. We proved it with this **pic from Hawaii!** Today's news is in this underwater drone - **iBubble** - which autonomously follows and records the actions of scuba divers, while being capable of wirelessly following them . Manufactured by a French marine tech company - **Notilo Plus** - it was even featured in a French horror flick that we don't care to mention. It works on acoustic signals emitted by a control unit carried by the diver and on an onboard 1080p/20fpss camera and visual recognition system able to detect each single diver. **Watch the Video**

Introducing the First Self-Supervised Algorithm for Speech, Vision and Text

We all know that convergence of modalities **speech, vision and text** is a thing. Facebook - oops **Meta** - is announcing **data2vec**, the first **high-performance self-supervised algorithm** that learns the same way in multiple modalities. This is fully consistent with what Yann LeCun told us during his famous interview with Computer Vision News and his interest in self-supervised learning for the coming years. With data2vec, Meta is closer to building machines that learn about different aspects of the world around them **without having to rely on labeled data**. Read it again here.

Read More

England Has an £850 Million Problem with Litter. This Startup Is Using AI to Fix It

British tech company **Littercam** has developed a technology that uses **AI and computer vision** to detect **littering from cars** and match it with the offender's license plate. The software is already being trialed and it is able to **detect very small items being littered**, like cigarette butts; BTW, did you know that it takes a cigarette butt about 10 years to decompose? But that's another story. What this story tells is that authorities are now using AI in all kinds of weather conditions to have a penalty notice issued to the car's registered owner, in the hope of preventing littering in the future. **Watch the Video**

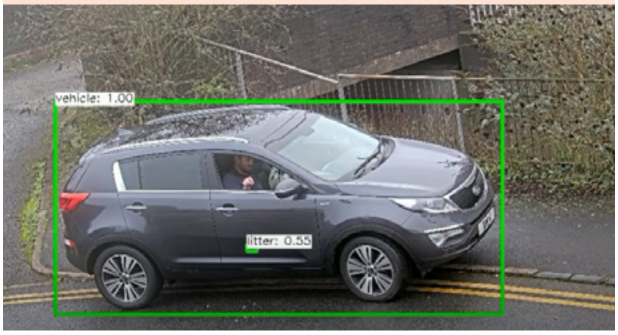

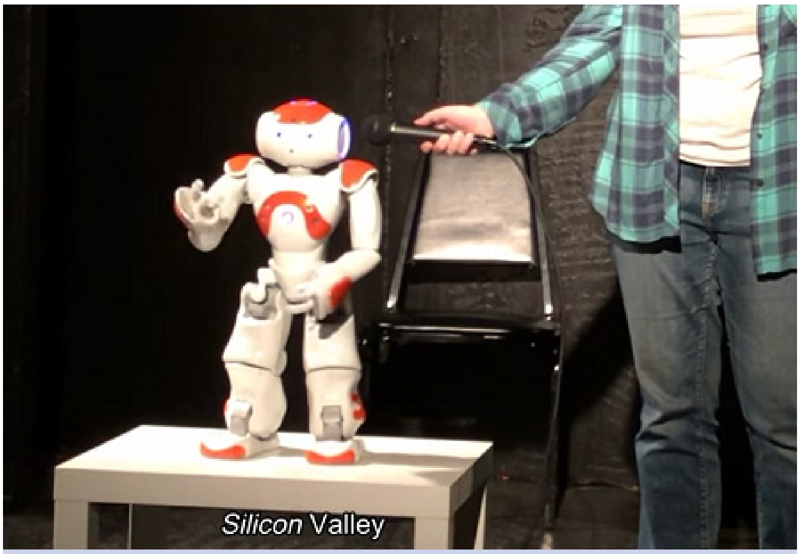

Artificial Intelligence Can Now Make Jokes - Jon the Robot

Robots are apparently able now to tell jokes and the novelty is that they adapt their lines to the feedback they receive from their audience. Created by **Naomi Fitter**, an assistant professor at Oregon State University, **Jon the Robot** performs live while at the same time it learns how to respond to its audience. Jon does not understand why we laugh at its jokes, and it is certainly not ready yet to craft its own repertoire of funnies. Still, you can say that it "feels" the listeners and adapt to their reactions (or lack of) in almost real time. *"You can probably tell from looking at me that I am from the Valley!"* **Watch the Video**

Synthetic Data: A Solution for Improving Driver Safety

Beginning in 2022, all new cars entering the EU must be equipped with advanced safety systems: distraction recognition and alert systems on trucks and buses to warn about vulnerable pedestrians or cyclists in close proximity. Meeting the new safety regulations will not be an easy task for car manufacturers, who will be faced with dedicating vast amounts of resources to build and deploy cars to collect diverse datasets to train AI models. One of their allies is synthetic datasets – computer-generated simulations that ensure an ample supply of diverse and anonymous training data. **Read More**

Tomas Jakab has recently completed his PhD at the Visual Geometry Group (VGG), University of Oxford. His research is focused on self-supervised learning of structural representations of objects. He developed self-supervised learning methods for object keypoints detectors, 3D keypoints for shape control, and 3D object category reconstruction from a single image. During his PhD, he also worked as a research intern at Google Research New York. Currently, he continues his work at VGG improving self-supervised 3D object reconstruction. Congrats, Doctor Tomas!

Artificial intelligence agents operating in the real world need to create internal representations of objects to reason about them. For instance, self-driving cars need to reason about pedestrians, cyclists, and other cars and anticipate their actions so they can prevent collisions with them. They can represent these objects using structural representations like object keypoints, parts, or full 3D shape models such as meshes.

Deep learning has demonstrated great success in learning these structural representations from images. This success heavily relies on large datasets of annotated images that are expensive to collect. This complicates the deployment of AI into novel environments. Self-supervised learning aims to reduce this annotation bottleneck by using the actual input images as the source of supervision.

A simple example of this is an autoencoder that is tasked to reconstruct the input image that was passed through a tight information bottleneck. We extend this autoencoding framework by engineering bottlenecks that, given a collection of images or videos of an object category, disentangle the object structure from other factors, such as the object pose from the appearance. Moreover, the bottlenecks and regularizers are designed to enforce the structural representation to be in the form of 2D [1, 2], 3D object landmarks [3], or 3D mesh [4].

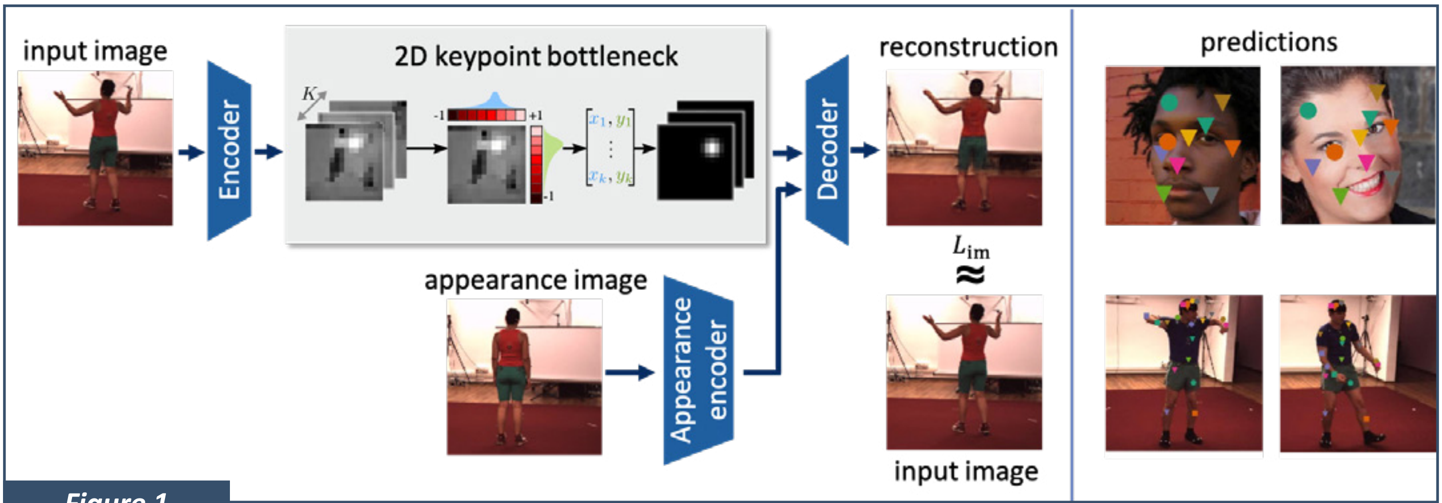

Figure 1

Figure 1 illustrates our conditional autoencoder for self-supervised object keypoint discovery [1]. The decoder is tasked to reconstruct the input image but it can only see 2D points that were extracted from the input image using a 2D keypoint bottleneck. A second image containing the same object but a different pose is supplied to provide appearance information. The network is trained using only image reconstruction loss and learns to use the 2D points as object keypoints that encode the pose by dedicating each of the points to semantically consistent locations on the objects. Our follow-up work [2] allows the use of a prior over the keypoints to learn human-interpretable keypoints.

The idea of conditional generation was later also leveraged in our method for self-supervised shape deformation controlled by 3D keypoints [3].

The autoencoding framework can be also adapted for more complex structural representations like 3D meshes. Figure 2 shows our photo-geometric autoencoder with a bottleneck that disentangles the input image into a rigid pose, 3D canonical shape, articulation, and texture [4]. A differentiable renderer working as a decoder reconstructs the input image. The model is trained end-to-end with reconstruction losses, without any explicit 3D supervision. For more information see our **project website**.

Figure 2

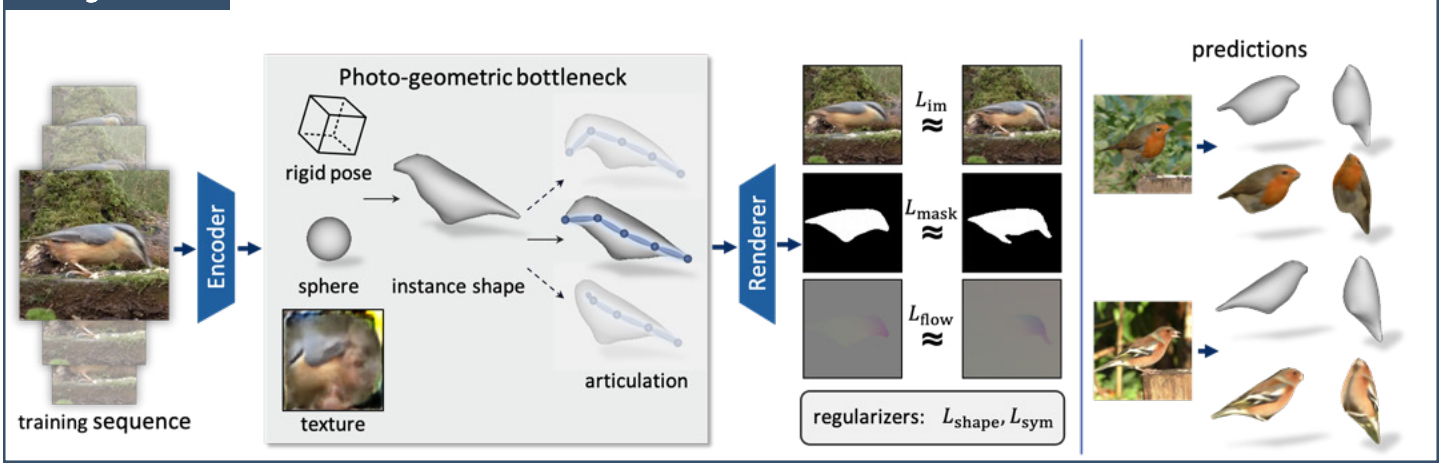

Marianna Meo is a research scientist in the advanced development team at EPD Solutions - a Philips company in the Netherlands

More than 100 inspiring stories of Women in Science here!

Marianna, can you tell us about your work?

Sure, I work in the advanced development team, which is something that comes a little bit before Research & Development, as we know it. What we try to do is propose concepts to be integrated into our product, which I will describe in a minute, in order to prove the feasibility before delivering them to Research & Development, that is going to integrate the concepts in the product.

What is our product? EPD solutions produces systems for cardiac mapping for electroanatomical reconstruction of the heart chambers in the treatment of cardiac arrhythmia. There are a variety of tools that are integrated into these kinds of systems. We have imaging. We have catheter navigation. These surgeries are not performed open heart, but you have to introduce a catheter through a vein, so it's minimally invasive. You also have a lot of signals that are recorded through these catheters, and this is my specialty. I work with my colleagues, my team, to develop signal processing tools to analyze the information that is brought by these signals that are recorded at the surface of the heart chambers.

Did you know in advance that one day you would master such deep scientific and medical concepts?

I follow the path full of opportunity. In the very beginning, there was the need, on my side, to do something related to healthcare and try to provide a service to patients. I was always driven by this need. This is why my university studies were focused on biomedical engineering. I wanted to apply technology and put them in the service of patients and doctors. Then, when I moved to France for my first job opportunity, that's when I found another need that brought me to the academia world. So, I started a PhD thesis on cardiac signal processing. I started with that, and I started building knowledge on electrophysiology, the medicine of the heart. I started building engineering knowledge. I learned how to use programming tools and how to implement algorithms. Then, little by little, other opportunities came across my path. I had a postdoc in Boston at Harvard Medical School at Brigham and Women's Hospital. Then another one in Bordeaux in cardiac signal processing. Through all these opportunities, I built my knowledge: it was not something that I expected in the beginning. That's also the beauty of research: new fields open up. We get interested in them. We gain knowledge. Then we propose and spread this knowledge through our applications. Now, I'm very happy because I'm also working for a company, which is also a different perspective.

A very respected company.

Yes, and actually, it's something new also. Philips was not really doing cardiac mapping. It was more focused on other imaging technologies. It's a nice way to merge signal processing aspects. It's a field with a lot of patients, since many patients in the world suffer from cardiac arrhythmia. I'm glad that, even with my smaller perspective, I can bring something useful to them.

I imagine that the little girl from Napoli, when she sees all these impressive places like Harvard and Philips, would open her eyes and say, ***"Wow! Is that really me working here?"***

Absolutely! I still do that actually. [*laughs*]

Can you explain how you feel?

I feel honored because, of course, it's a big honor to play a role at all of these institutes, at Harvard and in Bordeaux, which is also very well known for cardiac arrhythmia management. It's also a responsibility because I try to give 100 percent of myself in my work. So, I have to give 1,000 percent probably. I want to be at a good level for this role. So yes, I feel the responsibility of my role, but I'm also glad for having these opportunities, for exploring my knowledge at these kinds of institutes where you have a lot of resources, which is very important for our research.

I also learned to work with this vibe.

I am here now because I built my path.

Have you ever been present for these procedures in the operating room? Either catheter or open-heart operations?

Yes, in particular in Bordeaux, I was working on clinical projects. My lab is closely collaborating with the Bordeaux University Hospital, which has all of these cardiac arrhythmia treatment procedures. Yes, I was there to record my data for my research projects. Of course, there are specialists handling the devices for treatment because they are there for that purpose. I was there to get data from the procedure. In these kinds of procedures, the patient is in the same room but usually separated by glass at a certain distance. Also, in this case, I feel the responsibility because this data comes from a patient who wants to be treated. All these people have hopes to be treated and to successfully manage their disease. It gives me motivation. This is why I like working in the hospitals with the doctors.

Were you ever scared by watching these procedures, by all the blood?

Let's say, there is not so much blood outside because, as I mentioned, these procedures are minimally invasive. They just insert the catheter into the femoral vein, so let's say at the level of the upper leg. The blood is really minimal. There may be complications, not so often fortunately, but there are. What I've seen in Bordeaux is that people master these kinds of side complications. The nurse staff is there, ready to come up with solutions. So nothing very bad happens. I know there are more invasive procedures, but I'm not 100 percent sure I would be able to stand it! During my studies in biomedical engineering, I have watched stomach surgeries. They were open organ surgeries... and I survived! [*laughs*] I didn't faint!

Did the patient survive?

Always. I've been working a lot on atrial fibrillation. It is known to be dangerous for the patient in the long term, in the sense that it can lead to stroke. There are very severe long term consequences, but it's less life threatening than ventricular fibrillation for example, when the patient has to be defibrillated immediately. For these patients, the treatment, the catheterization must be done with attention. I have not seen lethal risks in these kinds of interventions.

We are like men, we can do everything! We can study! We can learn! We can explore!

By now, our readers know that you are Italian like me. Would you like to describe the girl that grew up in Naples some time ago?

I left my hometown of Naples really early because I wanted to study in Rome. I studied at the biomedical engineering campus in Rome. Yes, I was very motivated at that time to explore, to learn, to do something more outside of my neighborhood. I always liked, and I still love, my hometown. But when I finished my high school studies, there were not enough opportunities in technology for engineers, and especially for women engineers. I think that now, more than ten years later... It's difficult to say, but it's a long time [*laughs*]... I think now

Yes, I was very motivated at that time to explore, to learn, to do something more outside of my neighborhood.

there are a little bit more opportunities for women, but we are still fighting a lot. I really wanted to do something to demonstrate to myself and to other women around me, that… [*changes voice*] we are like men, we can do everything! We can study! We can learn! We can explore! I still feel very little sometimes, but I know that I have something to offer. I have tried to do this from the very beginning. This has brought me relatively far. I am here now because I built my path.

And the story is not over yet!

Absolutely! When I was in France, I thought that was the end, let's say. I found my achievement. I reached a stable position. I had a good team of collaborators. Then I found this opportunity at EPD Solutions. Even if I am almost 40 years old, it is not over yet!

Tell me something about Naples that we don't know.

It has been changing a lot, and there are people that know more than me actually. Naples is a colorful city. That means that there is the black and the white, and a lot of shades in between. There are a lot of problems, especially in some areas. But there are also a lot of great people, great people who want to bring something positive to their life. They don't want to mix with the wrong crowd.

You certainly know the saying about Naples that is in almost every language.

Which one?

"See Naples and die."

This is true! I recommend it! There are a lot of beautiful places in the world. But your hometown is probably the dearest place to you. Also, I get feedback from my friends. I have a lot of foreign friends who have visited Naples that were astonished, because there is a particular vibe.

That's also the beauty of research: new fields open up. We get interested in them. We gain knowledge.

Tell us about the vibes at Philips. Do you think they also have special vibes?

I think so, although it has been a great challenge for everyone, and also for me starting at a new company. I started my new role in the middle of the lockdown in the Netherlands. It has been very challenging to meet people because there was no physical office. Also, just to get the workstation to run my studies… That was also a strength at the same time because I learned to work in a new environment with a wide network of people geographically distanced. I found that amazing. Philips was ready for that, much earlier than COVID. I also learned to work with this vibe. I'm pretty lucky also because I have a very international team, in the sense that I'm based in the Netherlands, but my team is in Cambridge, Massachusetts. We work with EPD research in Israel. It's great! There is a lot of cultural cross interaction. Everybody brings a different competency and culture. I hope to enjoy it even more after COVID lockdowns. We just ended a lockdown here.

What is the first thing that you will do after COVID?

I don't know! Even despite COVID, I try to travel a little bit. I like traveling. I would really like to go back to conferences onsite. I went once in September, and it was great because I finally managed to see a lot of people. I was a bit scared though! After two years, talking in front of a lot of people was a bit weird. After the first ten seconds, it was fine!

What's your final message for the community?

It's about your **"Women in Computer Vision"** interviews. I think that many women paved the way before me. But it's still important to send the message. Women should probably believe a little bit more in themselves, in what they can achieve. Try to strive for higher roles. I see that there are still a lot of postdocs and PhD students who do not try to go farther. It's a pity because it's important. Everybody can bring a big value to science, including women. There is no reason why it should be different. Always try to change. It's important to catch the opportunities. We can improve, and we can do something better.

So go women! And go Italy!

Absolutely! Forza Napoli! [*both laugh*]

Through all these opportunities, I built my knowledge.

More than 100 inspiring stories of Women in Science here!

COMPUTER VISION EVENTS

AAAI Conf. on AI Virtual 2-9 Feb	AI Ethics Summit S.Francisco, CA 17-18 February	Deep Learning Summit S.Francisco, CA 17-18 February
ICARA Aut., Robotics and Apps Prague, Czech Rep. 18-20 February	SPIE Medical Imaging S.Diego, CA 20-24 February	Ai4 Healthcare 2022 Fully Digital 23-24 February
CTO Plus (cardiology) New York, NY 24-25 February	IVRHA 2022 Nashville, Tennessee March 3-4	
LSI Emerging MedTech Summit Dana Point, CA March 15-18		
SAGES 2022 Denver, CO March 16-19		

SUBSCRIBE!

Join thousands of AI professionals who receive Computer Vision News as soon as we publish it. You can also visit our archive to find new and old issues as well.

FREE SUBSCRIPTION

(click here, its free)

Did you enjoy reading Computer Vision News?
Would you like to receive it every month?

Fill the Subscription Form

it takes less than 1 minute!

We hate SPAM and promise to keep your email address safe, always!

Due to the pandemic situation, most shows are considering going virtual or to be held at another date. Please check the latest information on their website before making any plans!

MEDICAL IMAGING NEWS

FEBRUARY 2022

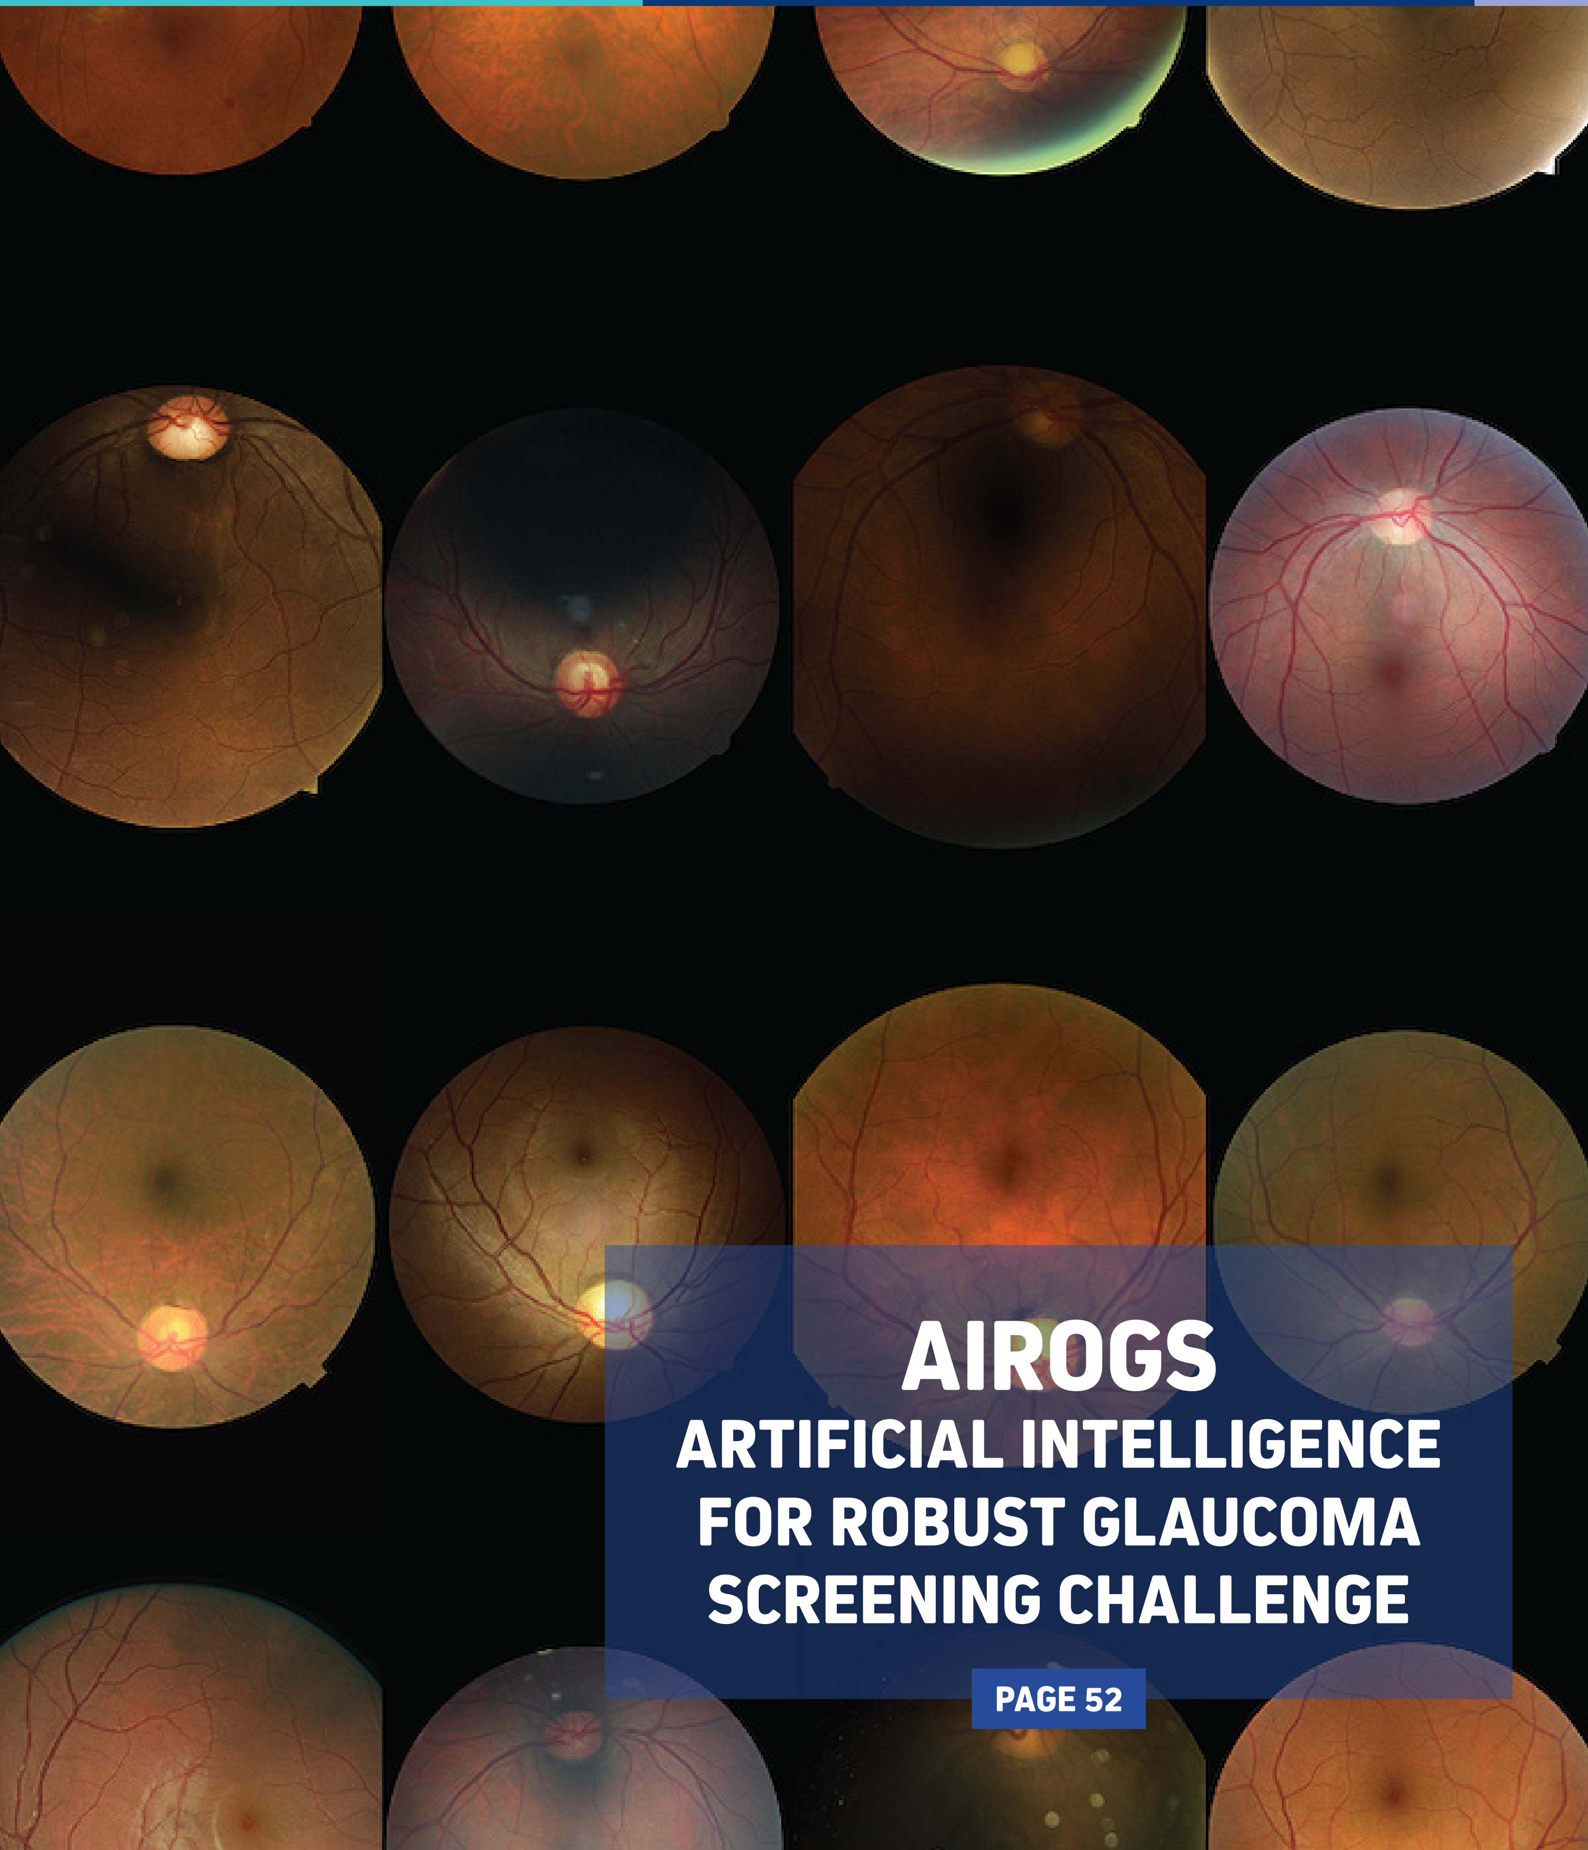

AIROGS

ARTIFICIAL INTELLIGENCE FOR ROBUST GLAUCOMA SCREENING CHALLENGE

PAGE 52

Adrian Dalca is a Assistant Professor at Harvard Medical School and a Research Scientist at MIT, where he is part of Professor John Guttag's lab working on machine learning for healthcare.

His main position is at Harvard, but his projects often cross over between the two.

He speaks to us about this unique position and reveals all about his latest work.

Welcome, Adrian! It must be so interesting to work at both Harvard Medical School and MIT. Can you tell us more?

At MIT, we have strong technical students who create a very technical environment. We talk about the latest machine learning trends. At Harvard, we have some people who are very technical, but the work I do is more around healthcare and neuroscience – things the clinicians care about. My projects are heavily inspired by both institutions. I get the technical aspects more from MIT and the applied aspects more from Harvard. If you look at any of my papers from the last couple of years you will see people from both places on there.

That's really nice – combining their talents.

Yes, it's a very nice set-up. I'm lucky because there are always good people on either side. If you only work on one, it's very easy to miss the other, or not be realistic about it.

What is your work about?

I work on machine learning models for medical image analysis. I like to try to solve problems or think about problems that were not easily solvable before.

Can you give us an example?

One of the common problems we have is to make our algorithms work, we need to tune hyperparameters. Armies of students end up repeatedly tuning and retraining. It's a horrible, lengthy process. That doesn't show up in our papers. We never say, *"We tuned this for six months!"* We just give the best result. But the truth is it's really challenging and frustrating: it makes it hard to deploy these algorithms because you give them to clinicians and then the clinicians have to tune them too.

Which isn't their job.

Exactly. A new line of our projects is looking at how we can substantially alleviate this. Rather than giving you the best hyperparameter, we're acknowledging there isn't a best hyperparameter, it varies, but we've created a framework so that tuning the best hyperparameter is not much effort. We have an algorithm that says if you change this hyperparameter by this much, here's the result, and you get it instantly. There's no more retraining. A clinician can tune it in real time, and they'll see the result in real time. Our goal is to eliminate the nuisance processes. I don't want clinicians to worry about how the software works. I don't want graduate students to worry about retraining. Just train one model and be done with it. Then at test time you can tune it a little bit, but it's substantially less mundane work. That's just one example of a task I feel we've swept under the rug.

"You're putting me on the spot here!

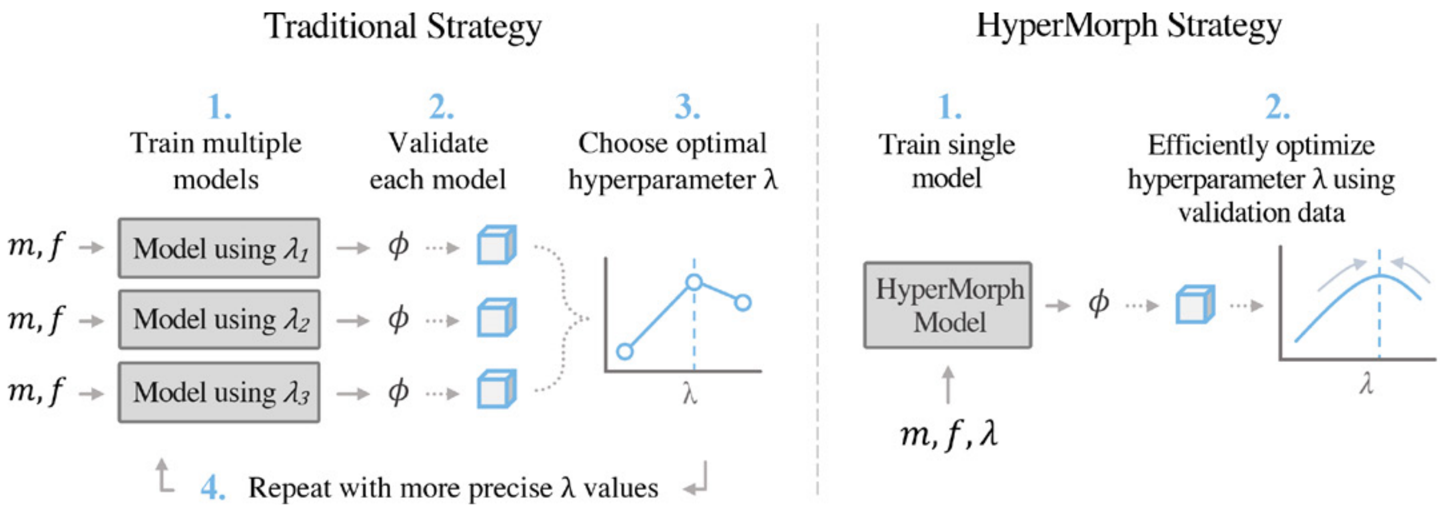

Are there any other tasks you're working on at the moment?

We have a series of projects on how we can get more and more general networks within medical imaging. So many of our models are enormously narrow. Segmenting the hippocampus in the brain, for example, is a popular task, but it's narrow. You're only segmenting the hippocampus, and you're only segmenting it on a particular type of MRI, and it's only for the brain. It would be nice to have a system that's independent of input modality or task. You can tell it what you want to segment and it just segments. You don't have to retrain it every time and pretend it's a new problem and restart the project. How can we generalize from these things in a way that is completely robust to image data type, modality, or even task?

Have you had any eureka moments recently?

The project that had the most impact on me was VoxelMorph with Guha Balakrishnan, Amy Zhao, Mert Sabuncu and John Guttag. It was a really weird moment when Guha and I came up with it. Usually, you come up with an idea, you try it out, that takes a few days, it doesn't work, you iterate. But with Guha, we had this idea, we tried it, and in two days, he came back to me and said, *"Okay, it works. Now what?"*

[*we both laugh*] I was like, *"No, no, no, it can't work. It's the first attempt."* He was like, *"No, no, it works. It just works! Look, here's the result."*

Yeah, we were really happy, but at the same time, I felt he wasn't getting the real student experience! What I'm really excited about with VoxelMorph is not just it's this registration algorithm

"A lot of the time I would love to try out an idea with just one experiment and then put it online as like a mini paper and that's it. We don't have a mechanism for that in the community yet…"

that's faster and more accurate, but it's enabled us to look at problems we couldn't look at before. We have a follow-up paper about building atlases - representative images of the brain - conditioned on attributes we care about like age. We couldn't have done that without VoxelMorph. And that project with the hyperparameters, without VoxelMorph, we couldn't build a network on top of it that estimated the effect of the hyperparameter.

Have other people built on VoxelMorph?

VoxelMorph is pretty well cited. Some people took it and added all kinds of properties to it. There are a few papers out there right now that use the ideas but replace the core architecture.. People applied it to domains I'd never heard of, but it was so interesting to get all these emails from people trying to use it!

Have you thought about basing a challenge or a competition around it?

That's an interesting idea! I like that. We do have a general medical image registration challenge, Learn2Reg, organized by Alessa Hering, Lasse Hansen, Mattias Heinrich and me.

What work would you like to be doing in the future?

In medical imaging, a lot of the work starts from the problem, and we come up with a solution, but I like part of my work to be about evolving technology and how we can use it in our field. Students get really excited about that. Things like attention, or implicit representations. Then the idea is, can we use this technology in our work to start something new? That might sound a bit silly, like a hammer without a nail, but to me, it's exciting!

Adrian with Guha Balakrishnan (now Assistant Professor at Rice University) and Amy Zhao (now a researcher working on Oculus at Facebook Reality Labs)

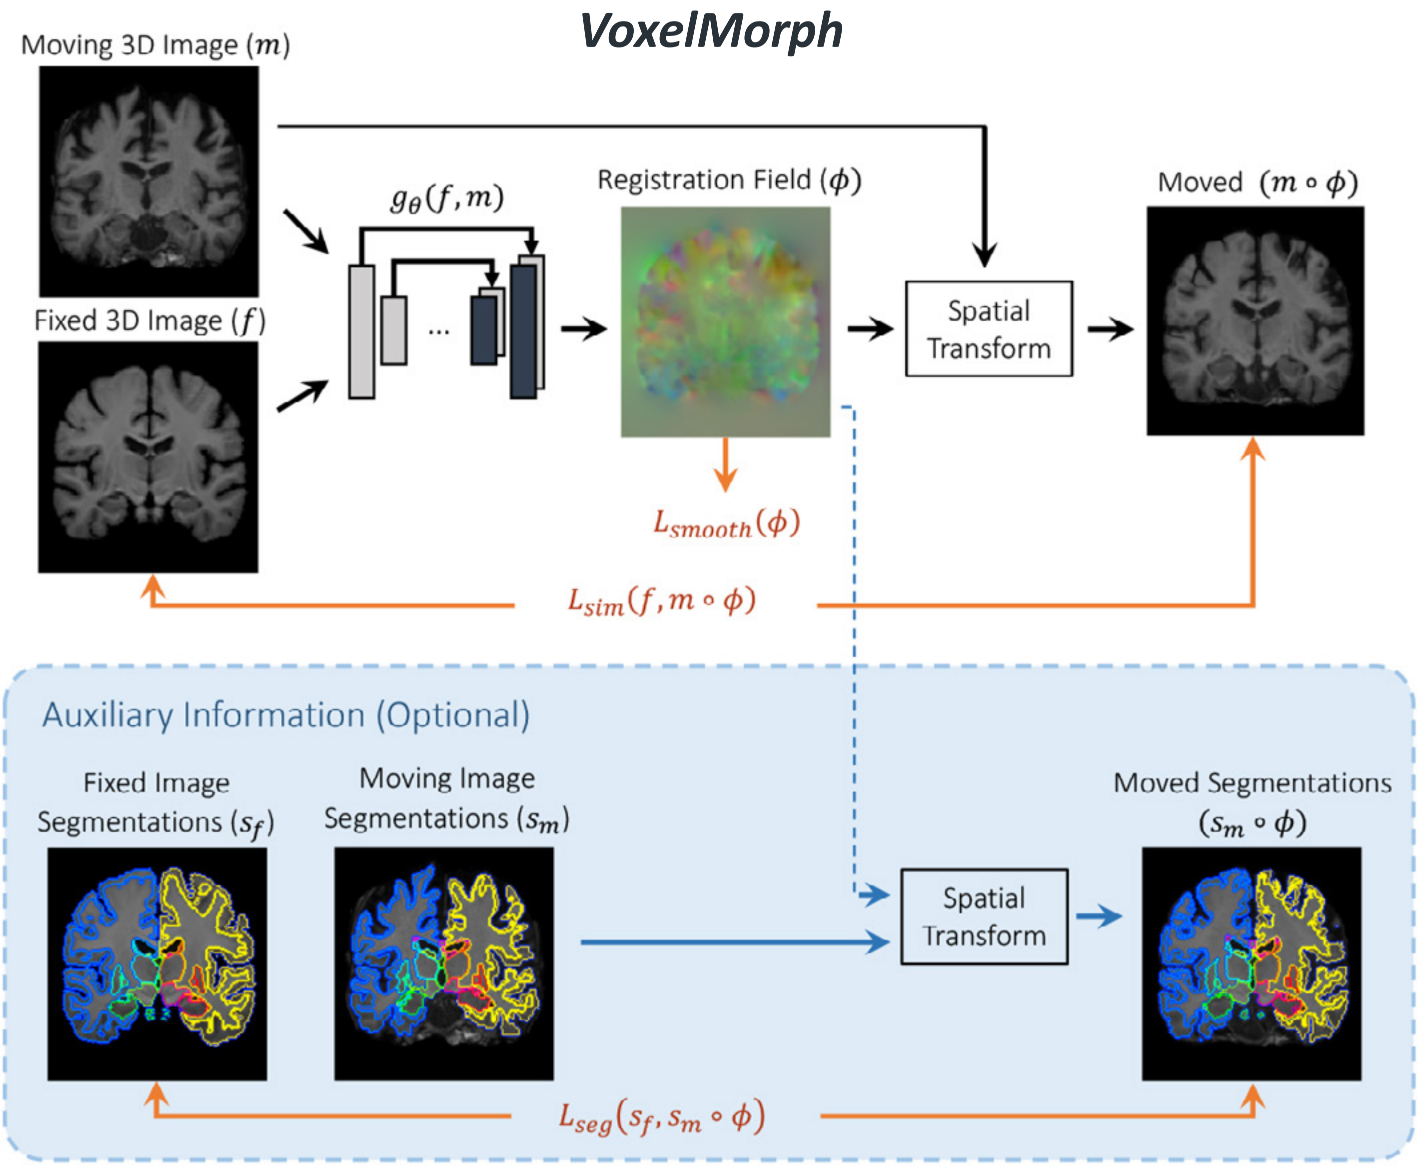

Are you working on anything like that at the moment?

One of the directions we're looking at is hypernetworks, which are basically neural networks that output other neural networks. The great thing about that is you can start looking at larger systems. You might have a hypernetwork for registration that gives you different networks for registering different things. A bit like meta learning. That's a direction I'd really like to look at. How do we use systems that are not just about solving one problem, but solving several problems at once, with interesting technology?

Your mind must be full of ideas and presumably you cannot go off in every direction!

You're right. Like everyone, I have lots of ideas and it's difficult to choose, partly because you don't know which ones will work! In fact, I've thought about starting a GitHub repository of ideas I don't have time for.

Is there an idea you would love to take forward, but don't have the time, so would be happy for somebody else to have a go?

[*he pauses*] You're putting me on the spot here! A lot of the time I would love to try out an idea with just one experiment and hen put it online as like a mini paper and that's it. We don't have a mechanism for that in the community yet. Sometimes I play around or somebody in my lab plays around, and we say, okay, that's cute, but we don't have time to follow it up. Maybe just putting stuff up on GitHub is a good start. You've given me the motivation!

I can only give you motivation, not time!

[We both laugh. But Adrian is serious: indeed, here is the new GitHub repository!]

Most of our readers have never been to Harvard or MIT. Could you give us all a taste of what they are like?

I know this sounds cliche, but it's all about the people. With both, the name and the prestige mean people really want to be there, and those who end up there are just really great people. What makes them as good as they are is they're all driven by something different. When you get a bunch of these people together and you put them in a room, they create a great environment.

Is there something that works well for the people at Harvard that could be of benefit to those at MIT, or vice versa?

[*he laughs*] That's a good question! I don't know if either should learn something the other does, but I will say the most important thing by far that both sides can learn is the ability to communicate with someone who's different to you. When you're in your own bubble, it's easy to dismiss the other side and assume things about them, but it's incredibly difficult for a clinician to talk to an AI expert and vice versa. Not only talk but talk in a way that leads to productive projects.

Your final message for the community.

We've all felt isolated over the last couple of years, but I have faith that we're going to come out of this soon. That's because - despite all the problems we have in society - science has come through for us. Within two years, we've got vaccines, treatments, better filtration. It's remarkable. I regret students missed out on the opportunity to go to conferences. What I know drives me and I think drives most progress is interaction. When we come out of this, I would urge everybody to go to as many conferences as you can, communicate as much as you can, make friends as much as you can! We need to get students into those places because that's what will lead to the next generation of great research.

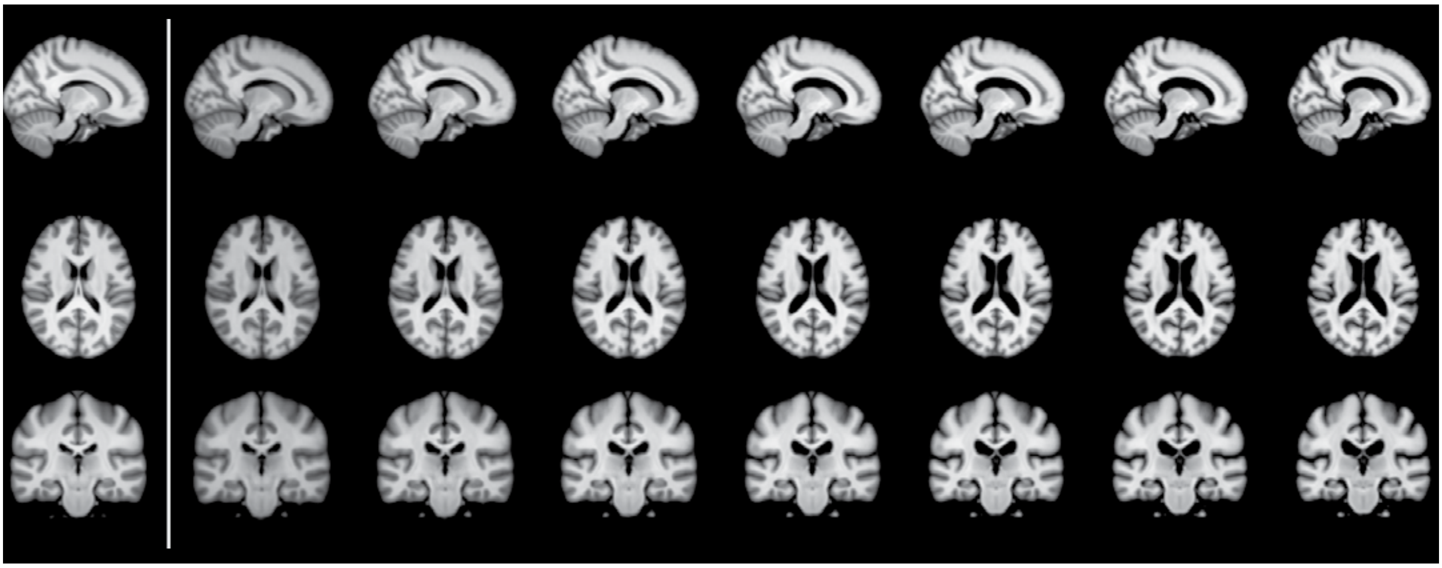

AUTOMATED ECHOCARDIOGRAPHIC DETECTION OF SEVERE CORONARY ARTERY DISEASE USING ARTIFICIAL INTELLIGENCE

by Marica Muffoletto @maricaS8

Hi everyone and welcome to another great issue of Computer Vision News! This month we review a paper called **Automated Echocardiographic Detection of Severe Coronary Artery Disease using Artificial Intelligence** whose authors (*Ross Upton, Angela Mumith, Arian Beqiri, Andrew Parker, William Hawkes & all others*) we deeply thank for allowing us to use their images. The paper is born from a collaboration between academics (at the University of Oxford, the Canterbury Christ Church University, the Knight Cardiovascular Institute in Oregon and the Department of Imaging Sciences and Biomedical Engineering in King's College London) and the company Ultromics, a start-up leader in proposing fully automated solutions for echo imaging. We also thank Sarah Jackson, Marketing Manager at Ultromics, for her kind support.

ULTROMICS™

This paper is a real gem which bridges clinical relevance with novel AI solutions, hence we are eager to share its details with all our readers.

The focus of the study, very much in line with the vision of the company itself, is to automatically detect Coronary Artery Disease (CAD), one of the major causes of mortality and whose current diagnostic methods are time-consuming, can suffer from inter- and intra-observer variability and need to rely on adequate training from clinicians. Hence, the paper aims to specifically show ***how AI can be a reliable source*** of disease detection and a viable option for future classification of CAD, to reduce the

variability among doctors, facilitate their job, and significantly impact the clinical workflow and patient care, providing an earlier and improved detection. The authors had previously worked on an AI pipeline that can automatically process and extract novel imaging features from stress echocardiograms. In this work, they incorporate the features into a machine learning model trained to identify patients with significant coronary artery disease.

In the image on the left, we can have a peak through the vision of Ultromics: a cardiologist can provide a patient's diagnosis based on both her expertise and a software powered by AI.

This is truthfully a vision shared with most of the medical AI research out there. But for Ultromics and their team, the clinical translation of the developed AI models is not just a dream anymore: it is a successfully achieved goal.

Data

This study leverages on a dataset of stress echocardiography (SE) images which were provided by a multicenter study (EVAREST) and include images from different operators and US vendors. This modality is a valid **non-invasive, non-ionizing and low-cost alternative imaging test** to diagnose CAD, and it relies on the manual observation by clinicians who identify regional wall motion abnormalities on the echocardiogram "by eye", comparing pre- and post-induced stress images. In fact, in areas affected by flow-limiting CAD, the contractility of the myocardium reduces as heart rate rises with stress (induced by exercise or medications). A set of images with apical 4 chamber (A4C), apical 2 chamber (A2C), parasternal short-axis mid-ventricular (SAX) views are included in the study, with sufficient frames between end-systole and end-diastole, from patients with no previous history of coronary artery bypass or other surgery.

Ground truth answers for severity of disease were based on decisions made by a committee with at least 1 cardiologist, blinded to the SE, who inspected the results of invasive coronary angiography (ICA).

Methodology

The methodology proposed in this paper is a full end-to-end pipeline which takes as input videos of SE images from a DICOM file. These are first classified into one of 8 view classes (A2C, A3C, A4C & SAX, contrast & non-contrast) by a bespoke CNN built using 10 convolutional layers for view classification. The **view classification** determines which of the three **auto-contouring models** must be applied to the respective video clips. Three U-net based CNN segmentation frameworks are developed to contour the left ventricle (LV) endocardium in A2C, A4C and SAX views using contrast enhanced and non-contrast studies, and their performance is assessed using Sørensen-Dice Coefficient (DC). The LV contours are then passed to a **cardiac cycle selection classifier** to algorithmically select end-diastolic and end-systolic contours within a respective cardiac cycle, and R-wave triggers. Here, various geometric features and clinical measurements are computed (**feature calculations**) and input as features for the machine learning ensemble model used for the prediction of CAD. Among these, both routine clinical measures and novel specifically engineered features were developed.

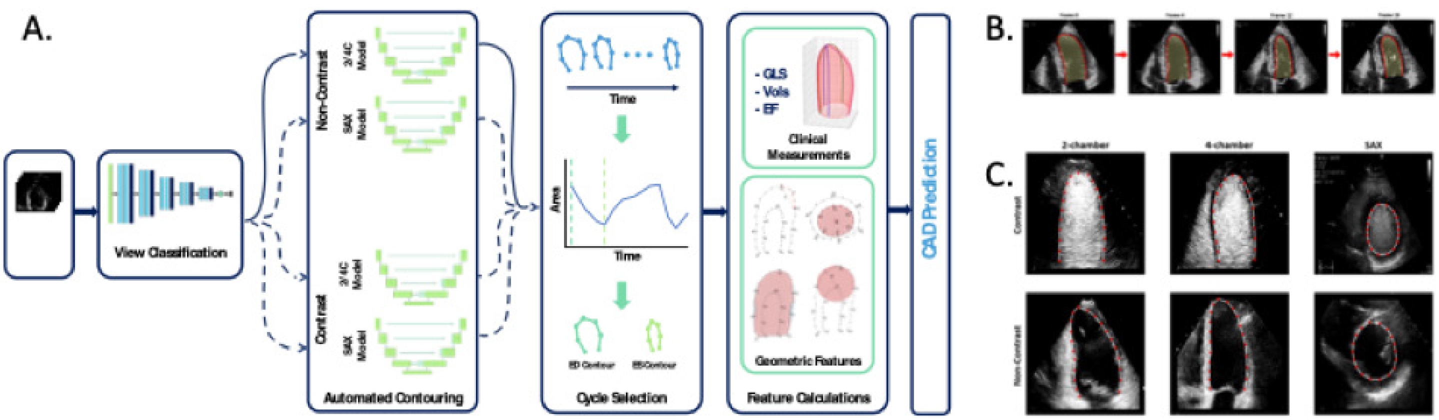

The CNNs produced contours that can track the endocardial walls smoothly through time, as it can be observed in parts B and C of the figure above. As shown there, this method well generalizes to the different types of images, including contrast-enhanced images.

The whole pipeline, which includes the LV contouring and the extraction of features, is a fundamental step which leads to the novel contributions of this paper. The first one answers the question:

How to automatically process SEs and classify whether the patient is likely to have severe coronary disease?

To pursue this, different ML classifiers are employed: support-vector machine, random forest and logistic regression models are all trained and compiled in a soft voting strategy. Stratified and repeated 2-fold nested cross validation is performed for hyperparameters optimization and feature selection. Model predictions are finally compared to the ground truth outcomes at a range of class prediction probability thresholds.

Moreover, the models are also independently tested using a retrospective study of patients who had undergone SE and had follow-up to at least 6 months through clinical record review. All images are processed using the same automated AI pipeline and then the model is used to identify patients with severe CAD.

The second contribution of the paper is instead a direct clinical translation:

Do clinicians benefit if provided an automatic classification?

This question is answered by the authors through a multiple reader multiple case (MRMC) randomized crossover design study. This uses medical records of 148 SE studies which were selected to be suitable from the test set of the previous study. Four experienced physicians/echocardiographers who were independent of any other part of the investigation were selected to classify CAD. For each reader, 50% of the studies were randomly provided with a report stating if the AI-based classifier deemed patients as having severe coronary artery disease. After a 1-month wash out period, the reader was presented with the same images but on this occasion AI-based classification was provided for the alternative 50% of patients.

Results

Let's now look at how the authors managed to answer these issues.

First, it's worth showing part of the pre-processing pipeline which does the auto-contouring (red) accuracy against manual contours (teal) from BSE accredited echocardiographers in contrast enhanced and non-contrast images in the apical 2-chamber, apical 4-chamber and parasternal short axis at the level of the mitral valve views. The accuracy is indicated at the bottom of each image and provides confidence that the future steps of feature extraction and disease classification can be properly performed.

Now, getting into the novel findings from the paper, the feature selection process discovered 31 features, whose individual evaluation demonstrated efficacy in identifying patients with severe disease. On the left of the image below, you can observe two plots which show the capability of individual model features to differentiate outcome where the vertical and horizontal lines indicate example cut-off values for disease classification, and the red and green dots identify CAD negative and positive patients.

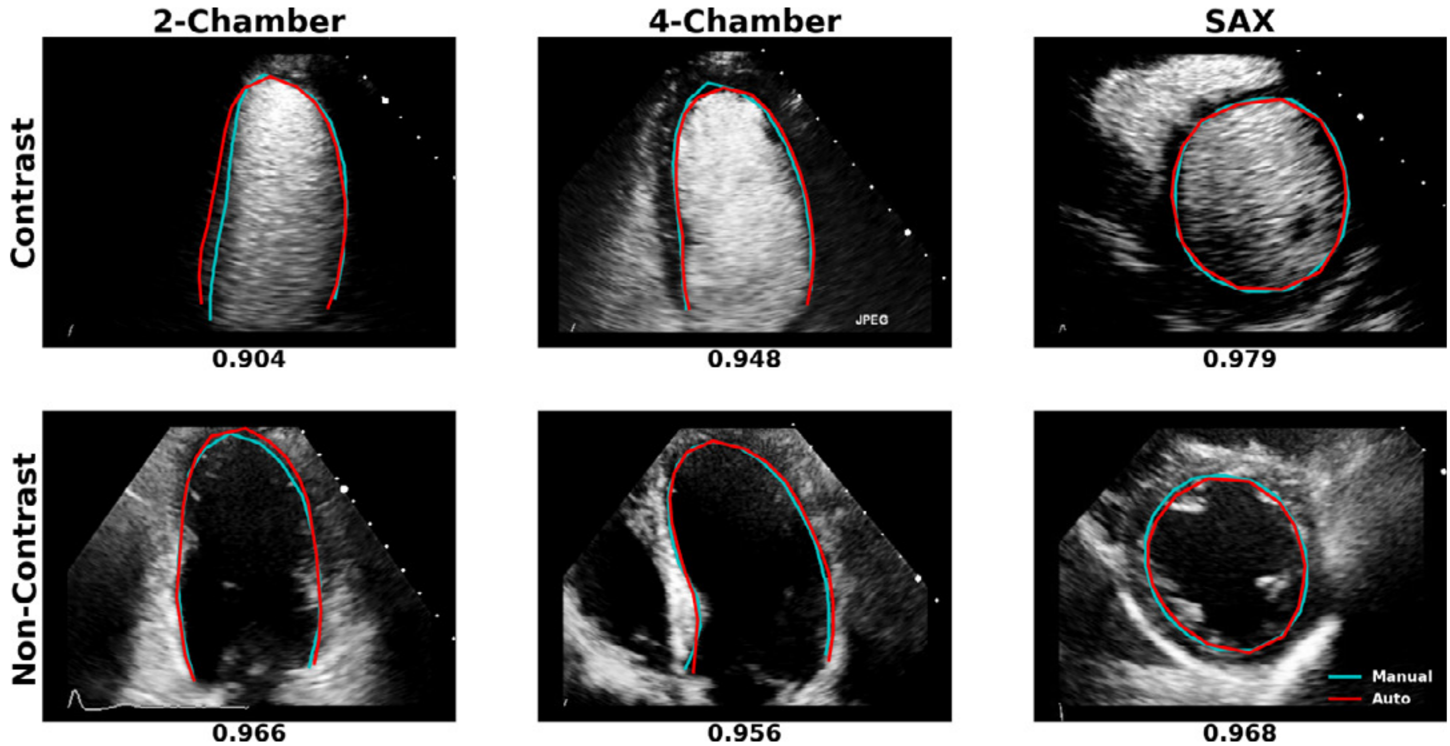

Performance of the final ML classifier is based on an ensemble model built from the 31 novel features, and it is shown on the top left of the figure. The model has an AUROC of 0.934 with an optimal specificity of 85.7% and sensitivity of 86.7% for the training set, and an AUROC of 0.927 with a specificity of 92.7% and a sensitivity of 84.4% for the independent testing set.

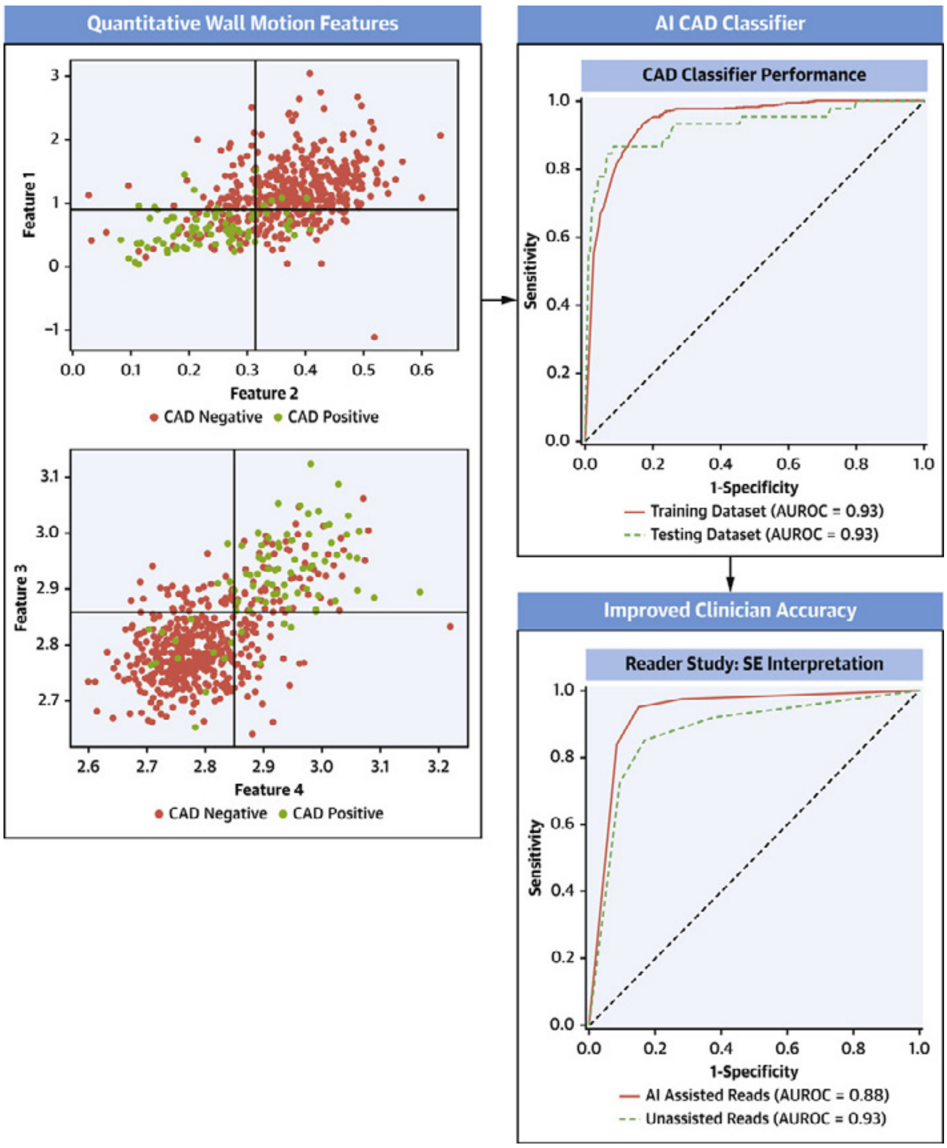

About the second study, the findings show that the readers exhibit significant increases in the mean ± SE AUROC from 0.877 ± 0.019 without to 0.931 ± 0.028 with the assistance of the classifier, resulting in a 10% increase in number of confident reads (sensitivity).

When they look at individual reader performance, the authors detect that, with the assistance of the AI-based classifier, the AUROC of two readers exceeded the performance of the AI classifier on its own (AUROC ¼ 0.927), with one of the readers achieving 100% sensitivity.

Conclusions

This paper strongly suggests that AI alone can achieve good diagnostic accuracy, and, more importantly, in a time when we are all wondering how the integration of AI models with humans would work, it shows that such a system can only improve the performance and confidence of a clinician.

The study perfectly fits with the mission of **Ultromics**, optimize the performance of echo exams with the power of AI, **providing Machine Learning decision support for Coronary Artery Disease**. As such, it contributes to one of their two software packages, Echo Go Pro, designed to provides reports to support physicians in the treatment of patients with suspected CAD.

EchoGo PRO™

Their other developed software, EchoGo Core, is instead built for the automatic analysis of the Left Ventricle, which interprets echocardiogram scans and provides accurate measurements of metrics such as global longitudinal strain (GLS), ejection fraction (EF), left ventricle end-diastolic volume (LV EDV) and so on from multi-views echo images.

Due to their accuracy and cloud-based service, the solutions of Ultromics are really at the cutting-edge of medical imaging and technology. They are something we always look forward to, hoping that they will represent the future of medicine.

Application of fmi Library to fastai in Medical Imaging and Deep Learning

IOANNIS VALASAKIS, KING'S COLLEGE LONDON

@WIZOFE

I hope everyone in those parts of the world that changes the year, had a great celebration for the new year! I also hope that everyone had a wonderful time in this quite hard winter, hoping for many new, bright days to come...

This month we will investigate the **application of the fmi library to the existing, great fastai approach to medical imaging and deep learning**. As you would probably remember from the earliest articles, fastai has created a user-friendly deep learning library, that can be used with almost zero overhead, and has great documentation.

We are going to improve this experience using the fmi library! Keep reading to find out and maybe get a few innovative ideas for your future projects! Let's start with basic image processing, imports for the libraries needed and check the system information. **Important to have a GPU** to run those examples! You can always find a free GPU instance using **Google's Collab** or other similar platforms ☺

```
import pydicom,kornia,skimage
 from fastai.vision.all import *
 from fastai.medical.imaging import *
 from torchvision.utils import save_image
 import seaborn as sns

from fmi.pipeline import *
 from fmi.explore import *
 from fmi.preprocessing import *
 from fmi.examine import *
 from fmi.train import *
 from sklearn.model_selection import train_test_split

import timm
 matplotlib.rcParams['image.cmap'] = 'viridis'

import warnings
 warnings.filterwarnings("ignore", category=DeprecationWarning)
```

system_info conveniently lists the current *fastai, fastcore* versions as well as *cuda, pydicom* and *kornia* versions

```
system_info()
```

```
fastai Version: 2.4
```

```
 fastcore Version: 1.3.20
 python Version: 3.8.10 (default, May 19, 2021, 13:12:57) [MSC v.1916 64 bits (AMD64)]
 torchvision: 0.8.2+cu110
 torch version: 1.7.1+cu110

Cuda: True
 Cuda Version: 11.0
 GPU: NVIDIA GeForce RTX 3070 Laptop GPU

pydicom Version: 2.1.2
 kornia Version: 0.2.0
```

Looking at the data

Getting Data

The data consists of 250 dicom images from the [SIIM-ACR Pneumothorax Segmentation dataset] which consists of chest X-rays with and without pneumothorax. If you remember, we did use a similar dataset for our **Kaggle expert tutorial**, a few months ago. Of course, you can use your own dataset, or anything that you feel comfortable working with!

Let's unzip (oh well, untar!) our data and have a look at the csv file coming with the dataset.

```
pneu = untar_data(URLs.SIIM_SMALL)
 p_items = get_dicom_files(f'{pneu}/train')
 p_items
```

```
(#250) [Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000000.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000002.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000005.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000006.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000007.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000008.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000009.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000011.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000012.dcm'),Path('C:/Users/avird/.fastai/data/siim_small/train/No
Pneumothorax/000014.dcm')...]
```

```
df = pd.read_csv(f'{pneu}/labels.csv')
 df[:5]
```

	file	label
0	train/No Pneumothorax/000000.dcm	No Pneumothorax
1	train/Pneumothorax/000001.dcm	Pneumothorax
2	train/No Pneumothorax/000002.dcm	No Pneumothorax
3	train/Pneumothorax/000003.dcm	Pneumothorax
4	train/Pneumothorax/000004.dcm	Pneumothorax

Metadata from DICOM files

Dicom metadata has a wealth of information like the example below.

```
p_items[0].dcmread()
```

```
Dataset.file_meta -------------------------------
 (0002, 0000) File Meta Information Group Length  UL: 200
 (0002, 0001) File Meta Information Version       OB: b'\x00\x01'
 (0002, 0002) Media Storage SOP Class UID         UI: Secondary Capture Image Storage
 (0002, 0003) Media Storage SOP Instance UID      UI: 1.2.276.0.7230010.3.1.4.8323329.6904.1517875201.850819
 (0002, 0010) Transfer Syntax UID                 UI: JPEG Baseline (Process 1)
 (0002, 0012) Implementation Class UID            UI: 1.2.276.0.7230010.3.0.3.6.0
 (0002, 0013) Implementation Version Name         SH: 'OFFIS_DCMTK_360'
 -------------------------------------------------
 (0008, 0005) Specific Character Set              CS: 'ISO_IR 100'
 (0008, 0016) SOP Class UID                       UI: Secondary Capture Image Storage
 (0008, 0018) SOP Instance UID                    UI: 1.2.276.0.7230010.3.1.4.8323329.6904.1517875201.850819
 (0008, 0020) Study Date                          DA: '19010101'
 (0008, 0030) Study Time                          TM: '000000.00'
 (0008, 0050) Accession Number                    SH: ''
 (0008, 0060) Modality                            CS: 'CR'
 (0008, 0064) Conversion Type                     CS: 'WSD'
 (0008, 0090) Referring Physician's Name          PN: ''
 (0008, 103e) Series Description                  LO: 'view: PA'
 (0010, 0010) Patient's Name                      PN: '16d7f894-55d7-4d95-8957-d18987f0e981'
 (0010, 0020) Patient ID                          LO: '16d7f894-55d7-4d95-8957-d18987f0e981'
 (0010, 0030) Patient's Birth Date                DA: ''
 (0010, 0040) Patient's Sex                       CS: 'M'
 (0010, 1010) Patient's Age                       AS: '62'
 (0018, 0015) Body Part Examined                  CS: 'CHEST'
 (0018, 5101) View Position                       CS: 'PA'
 (0020, 000d) Study Instance UID                  UI: 1.2.276.0.7230010.3.1.2.8323329.6904.1517875201.850818
 (0020, 000e) Series Instance UID                 UI: 1.2.276.0.7230010.3.1.3.8323329.6904.1517875201.850817
 (0020, 0010) Study ID                            SH: ''
 (0020, 0011) Series Number                       IS: "1"
 (0020, 0013) Instance Number                     IS: "1"
 (0020, 0020) Patient Orientation                 CS: ''
 (0028, 0002) Samples per Pixel                   US: 1
 (0028, 0004) Photometric Interpretation          CS: 'MONOCHROME2'
 (0028, 0010) Rows                                US: 1024
 (0028, 0011) Columns                             US: 1024
 (0028, 0030) Pixel Spacing                       DS: [0.168, 0.168]
 (0028, 0100) Bits Allocated                      US: 8
 (0028, 0101) Bits Stored                         US: 8
 (0028, 0102) High Bit                            US: 7
 (0028, 0103) Pixel Representation                US: 0
 (0028, 2110) Lossy Image Compression             CS: '01'
 (0028, 2114) Lossy Image Compression Method      CS: 'ISO_10918_1'
 (7fe0, 0010) Pixel Data                          OB: Array of 126284 elements
```

Dicom image info

However sometimes it is better to get more targeted information from the metadata.

get_image_info supplies image specific information such as the Modality, Photometric Interpretation, Pixel Spacing, SliceThickness, Bits Allocated, RescaleIntercept and RescaleSlope. **This information is important when modelling the data**. Note that in this dataset the dicom images do not have any *Rescale Intercept* or *Rescale Slope*

```
get_image_info(p_items[0])
```

```
(0008, 0060) Modality                              CS: 'CR'
 (0028, 0004) Photometric Interpretation           CS: 'MONOCHROME2'
 (0028, 0030) Pixel Spacing                        DS: [0.168, 0.168]
 No SliceThickness
 (0028, 0100) Bits Allocated                       US: 8
 (0028, 0101) Bits Stored                          US: 8
 (0028, 0103) Pixel Representation                 US: 0
 No Pixel Padding Value
 No Rescale Intercept
 No Rescale Slope
```

What about PII?

PII or **Personally Identifiable Information** refers to information that can be used to distinguish or trace an individual's identity, either alone or when combined with other personal or naming information that is linked or linkable to a specific individual. **Understanding what identifiable information is contained within the metadata is important in creating a fair non-biased model**.

get_pii conveniently lists this information as well as looking for the *Image Comments* tag within the metadata that may hold useful information.

It is also sometimes useful to know if there is any information about with de-identification information is available in the dicom meta-data.

```
get_pii(p_items[0])
```

```
(0010, 0010) Patient's Name                        PN: '16d7f894-55d7-4d95-8957-d18987f0e981'
 (0010, 0030) Patient's Birth Date                 DA: ''
 (0010, 0040) Patient's Sex                        CS: 'M'
 (0010, 1010) Patient's Age                        AS: '62'
 No Image Comments
 No De-Identification method
 No De-identification Method Code Sequence
```

DICOMS method from fastai

fastai has a **handy method** *from_dicoms* that can convert dicom metadata into a dataframe. By default, from_dicoms returns more stats such as minimal pixel value, largest pixel value, mean pixel value, image standard deviation and pct_in_window value depending on the specified window value. These stats are returned at the end of the dataframe.

I successfully submitted a *PR* that allows you to toggle the creation of more stats on or off. Note that if you have a big dataset having this possibility one could significantly increase the compute time.

```
%%time
patient_df = pd.DataFrame.from_dicoms(p_items, window=dicom_windows.lungs, px_summ=True)
patient_df[:5].T.tail(5)
```

```
Wall time: 7.45 s
```

	0	1	2	3	4
img_min	0	0	0	0	0
img_max	254	250	246	255	250
img_mean	160.398039	114.524713	132.218334	153.405355	166.198407
img_std	53.854885	70.752315	73.023531	59.543063	50.008985
img_pct_window	0.358613	0.62032	0.561823	0.359524	0.345166

```
%%time
patient_df = pd.DataFrame.from_dicoms(p_items, window=dicom_windows.lungs, px_summ=False)
patient_df[:5].T.tail(5)
```

```
Wall time: 1.52 s

                                                                                     0  \
 LossyImageCompression                                                              01
LossyImageCompressionMethod                                                ISO_10918_1
fname                        C:\Users\avird\.fastai\data\siim_small\train\No Pneumothorax\000000.dcm
MultiPixelSpacing                                                                    1
PixelSpacing1                                                                    0.168


                                                                                     1  \
 LossyImageCompression                                                              01
LossyImageCompressionMethod                                                ISO_10918_1
fname                        C:\Users\avird\.fastai\data\siim_small\train\No Pneumothorax\000002.dcm
MultiPixelSpacing                                                                    1
PixelSpacing1                                                                    0.143


                                                                                     2  \
 LossyImageCompression                                                              01
LossyImageCompressionMethod                                                ISO_10918_1
fname                        C:\Users\avird\.fastai\data\siim_small\train\No Pneumothorax\000005.dcm
MultiPixelSpacing                                                                    1
PixelSpacing1                                                                    0.143


                                                                                     3  \
 LossyImageCompression                                                              01
LossyImageCompressionMethod                                                ISO_10918_1
fname                        C:\Users\avird\.fastai\data\siim_small\train\No Pneumothorax\000006.dcm
MultiPixelSpacing                                                                    1
PixelSpacing1                                                                    0.171


                                                                                     4
LossyImageCompression                                                               01
LossyImageCompressionMethod                                                ISO_10918_1
```

```
fname                        C:\Users\avird\.fastai\data\siim_small\train\No Pneumothorax\000007.dcm
MultiPixelSpacing                                                                                   1
PixelSpacing1                                                                                   0.171
```

Let's explore some important demographics and basic statistics:

Male/Female distribution

```
patient_df['PatientSex'].value_counts()
```

```
M    125
 F    125
 Name: PatientSex, dtype: int64
```

```
sns.displot(data=patient_df, x="PatientAge", kde=True, col="PatientSex", hue='PatientSex', height=7,
aspect=2);
```

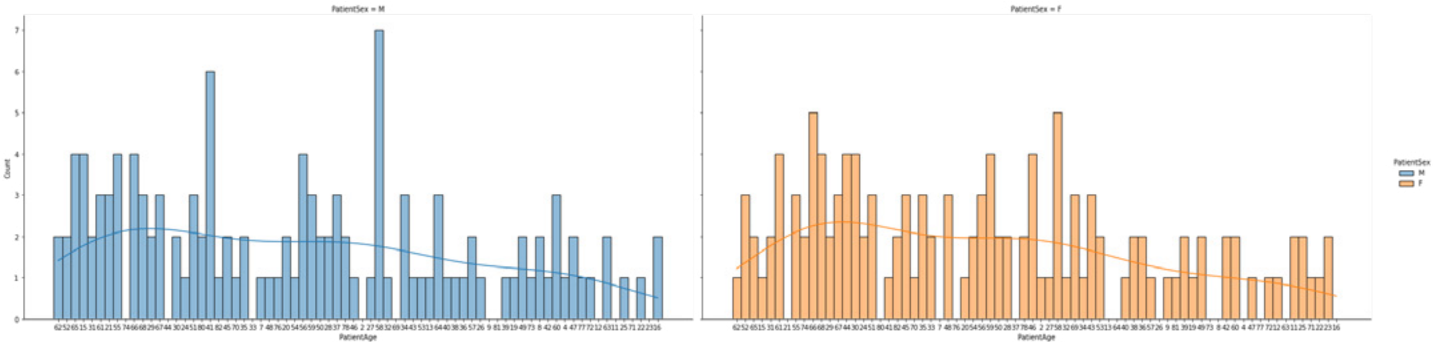

PreProcessing Splitting Data

It is customary practice to split the dataset into *train* and *valid*, for example *RandomSplitter* splits the data with an 80:20 split. However, it is also important to ensure that **the same patient is not present in both the train and valid splits**.

```
trn,val = RandomSplitter(valid_pct=0.2, seed=7)(p_items)
```

You can now see the indices of each image in the *train* and *valid* sets

```
trn, val
```

```
((#200) [33,65,231,167,74,127,184,89,122,79...],
  (#50) [115,233,139,163,161,177,57,21,34,99...])
```

The *SIIM_SMALL* only holds 250 images and you can easily check if duplicates will exist when splitting the data by using *check_duplicates* and specifying a *seed* value.

```
check_duplicate(p_items, valid_pct=0.2, seed=7)
```

```
Train: 200
 Original Validation: 50
 Updated Validation: 50
```

check_duplicates display the number of train and valid images and if there are duplicates it will also display the updated valid count with the duplicates removed from the valid set. It also displays images from the train and valid sets

The dataset as is does not have any duplicate images and hence the reason the updated valid count is the same as the original valid count of 50.

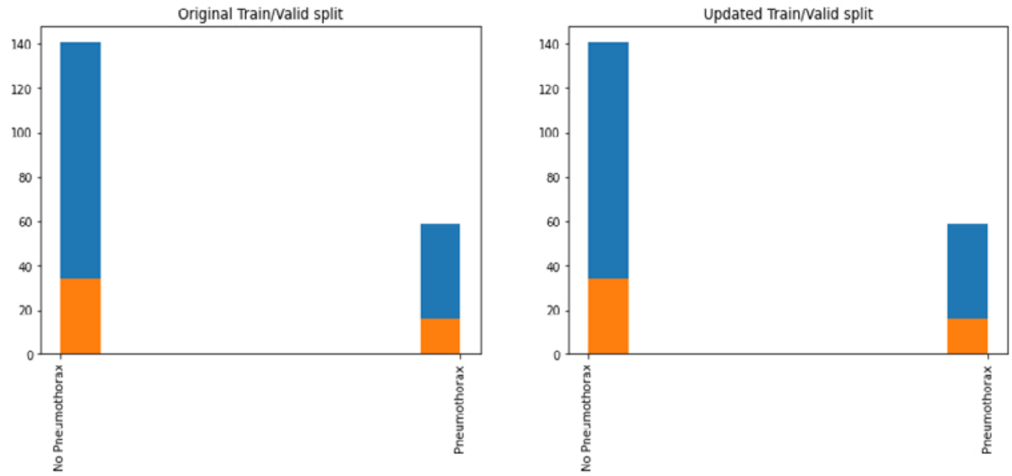

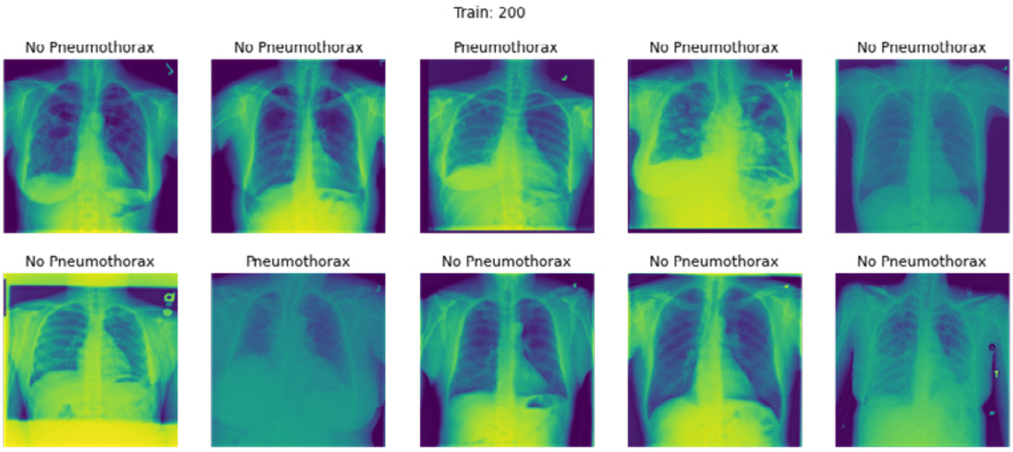

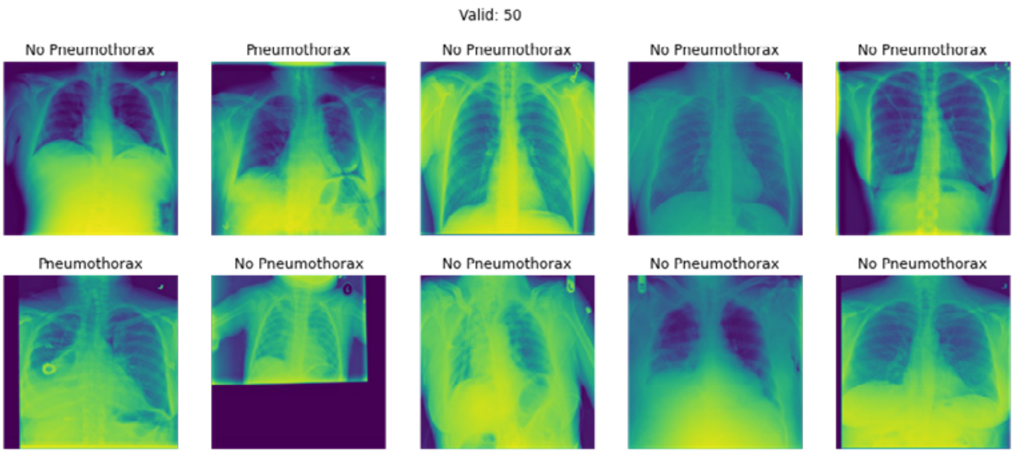

What happens when there is/are duplicates in the train and valid sets?

To see what happens I have another dataset that has duplicate images

```
duplicate_ds = get_dicom_files(f'{pneu}/sm')
```

```
check_duplicate(duplicate_ds, valid_pct=0.2, seed=7)
```

```
Train: 211
 Original Validation: 52
 Updated Validation: 48
```

In this case the original split had 211 images in the train set and 52 images in the validation set. *check_duplicates* was able to find 4 duplicates in the validation set and remove

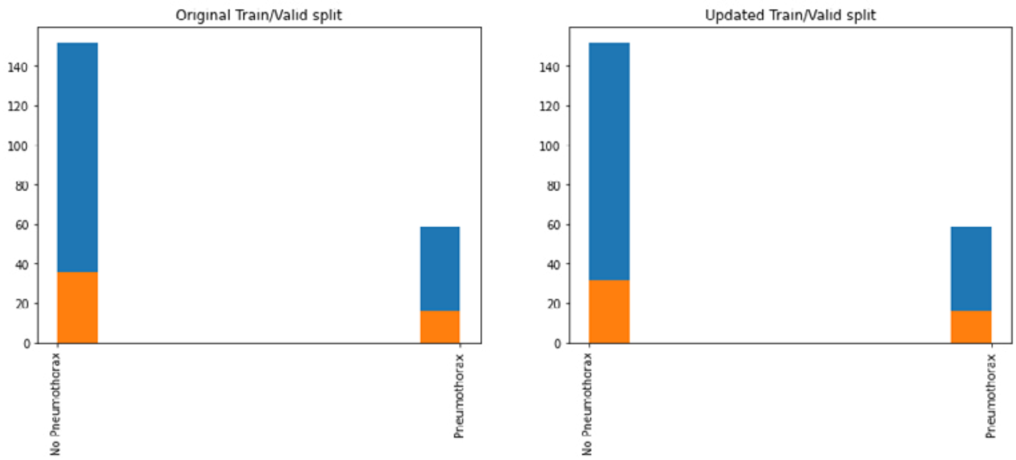

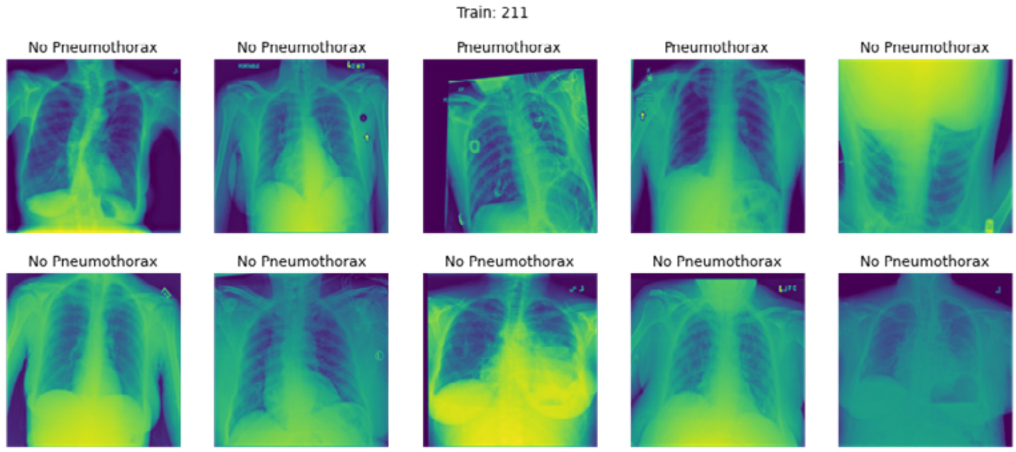

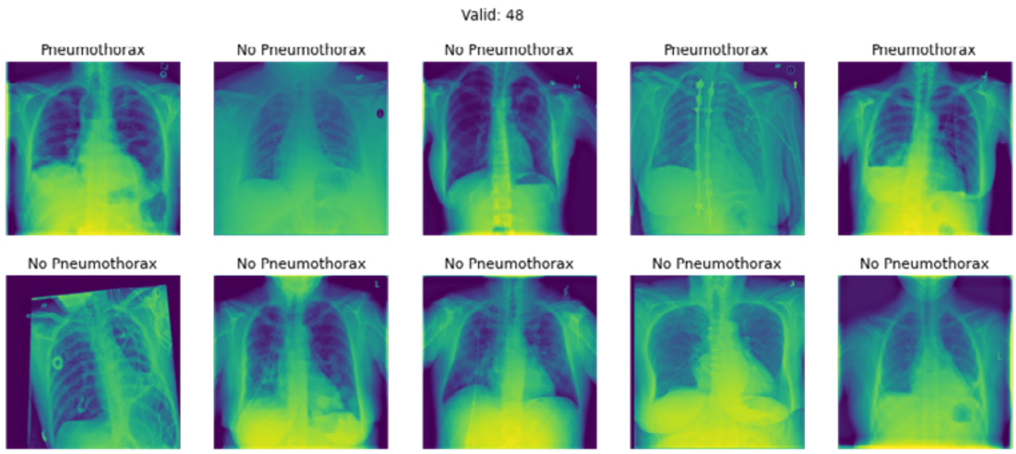

them from there so the updated validation count is now 48. However, **we now know that there are no duplicate patients in the train or validation sets**.

Epilogue

There's much more to come, as you may have guessed! After the next month's research review, we are going to have a second part, exploring **the sensitivity and specificity of the model**, how to expand its layers and many more!

As always, please send any ideas, corrections and recommendations my way! Thanks for keeping in contact with social media. Enjoy :-)

INTRA-OP VIRTUAL MEASUREMENTS IN LAPAROSCOPIC AND ROBOTIC-ASSISTED SURGERIES

Laparoscopic and Robotic-Assisted Surgeries (RAS) are taking up a larger portion of the surgical activity world-wide. This change is a result of the accuracy of such procedures which leads to superior clinical outcomes. Both are performed using video cameras for visualization - a good solution for visibility in the close confinement of the peritoneal cavity. However, when the need for true dimensions arises, this view does not suffice.

There are several minimally invasive procedures which require **real-time measurements**:

1. **Sleeve Gastrectomy**: during sleeve gastrectomy, the stomach is resected in proximity to the pylorus. It is essential that this resection is 2cm below the pylorus. Too close may lead to a tear, and too far may lower the success rate of the procedure.
2. **Intestine Resection**: procedures like Roux-en-Y gastric bypass, bilio-pancreatic diversion, or gastrectomy for oncological purposes require resection of the intestine. Here the surgeon wants to avoid bile reflux on the one hand and malabsorption on the other hand, so the length of the resected portion must

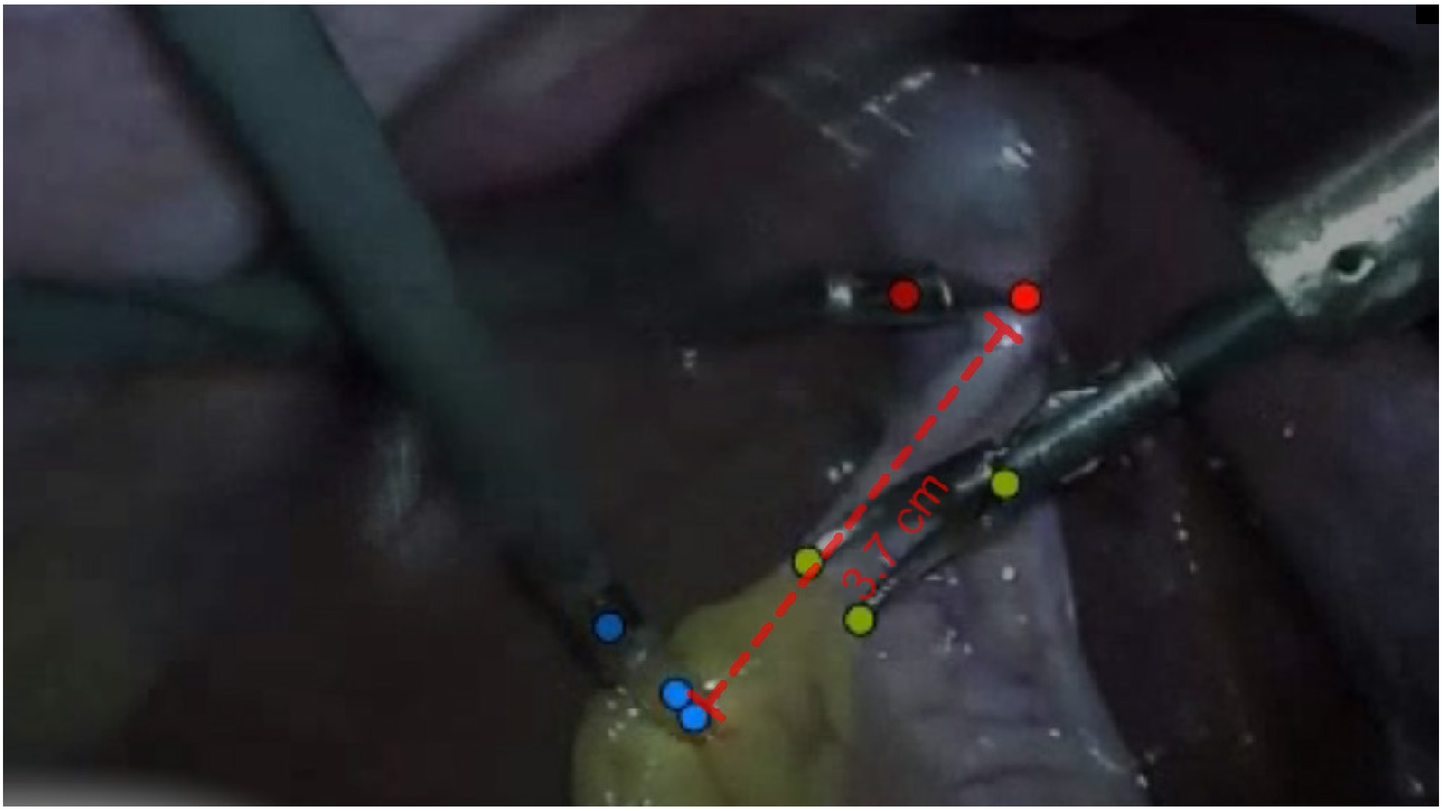

be accurately measured.

Currently, to measure a distance in laparoscopic or Robotic-Assisted Surgeries, the surgeon can perform an estimation based on their surgical experience. Needless to say, this method is **highly subjective and lacks accuracy**. Another method is to introduce an actual ruler into the image and inserting a foreign object into the body is never ideal.

RSIP Vision established an innovative method to **allow measurements in real time without additional hardware**. To perform measurements within an image, the following are required:

1. Camera calibration parameters
2. Known dimensions in the image

The camera parameters can be obtained from the system in use. Surgical tools are standardized, therefore their dimensions are known as well. Dedicated deep learning based tracking algorithms can track the tools throughout the procedure, and at any given moment the image can be calibrated using the camera parameters and the tool dimensions. The surgical tools serve as a **virtual ruler** instead of the actual one. **Now the surgeon can measure organ dimensions within the field-of-view, with absolute precision**. This implementation even allows measurements in 3D.

Using this technology enables surgeons to measure lengths which were only estimated before, with certainty and accuracy. Harnessing Artificial Intelligence to assist in Robotic-Assisted Surgeries improves surgical accuracy and clinical outcome. RSIP Vision has the experience and expertise to implement this technology quickly and efficiently. **Contact us today for more information!**

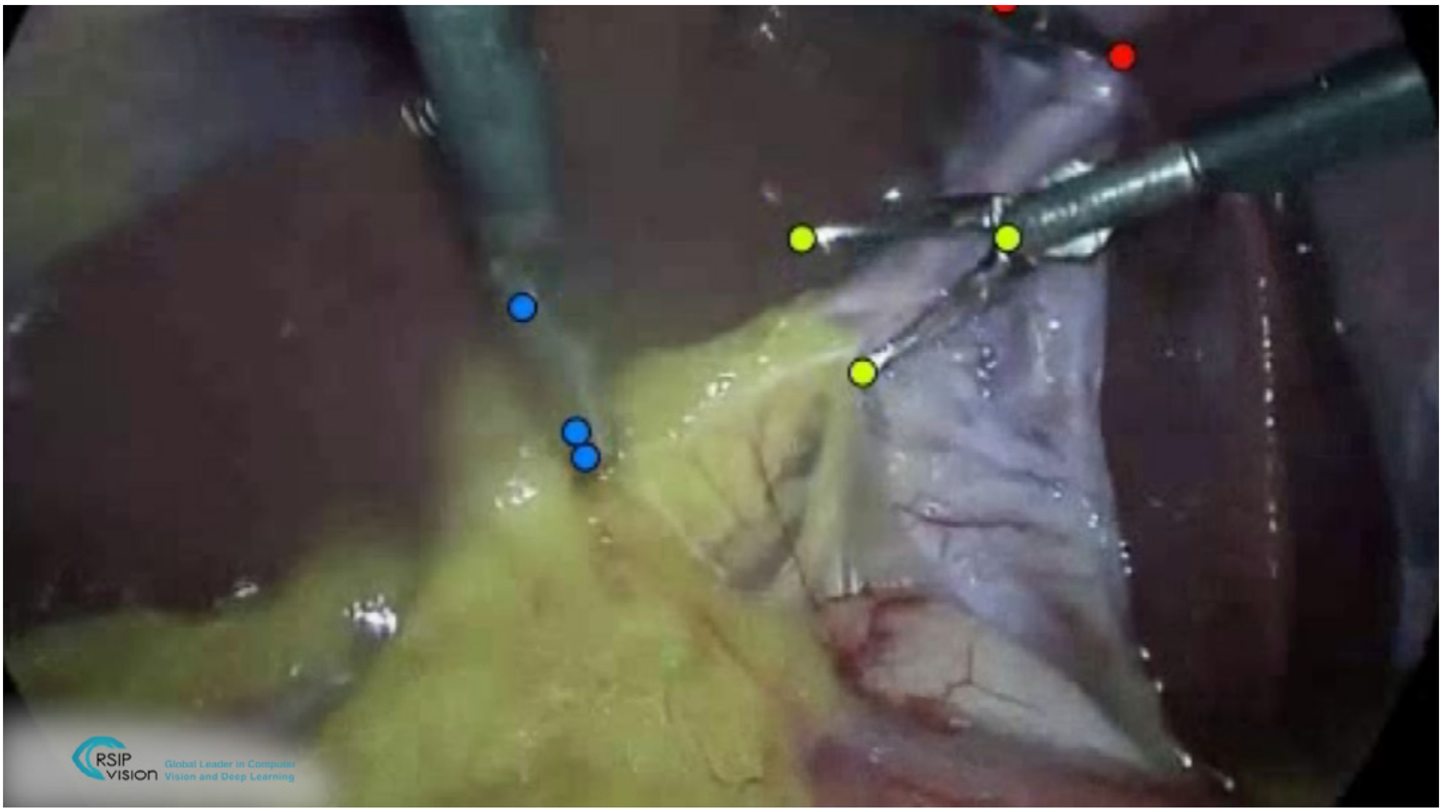

Ylenia Giarratano has recently completed her PhD in Precision Medicine at the University of Edinburgh. Her research interests focus on retinal imaging, computational modeling, and machine learning with the aim of revealing retinal biomarkers of microvascular diseases. She is currently pursuing her work at the University of Edinburgh as a research fellow, investigating retinal microvascular predictors of compromised brain haemodynamics in small vessel disease. Congrats, Doctor Ylenia!

Studies over recent years have reported that changes in the retinal vasculature due to disease may happen years before symptoms manifest. There has thereby been an increased interest in using retinal imaging for the discovery of biomarkers of both ocular and systemic disease. Among retinal imaging modalities, optical coherence tomography angiography (OCT-A), a recent technology, has emerged as a key modality for this investigation. OCT-A is a fast and dyeless technique that allows the visualisation of the microvasculature in *vivo* and to the deepest level of detail. Nonetheless, the potential of this technology in clinical diagnostics is still under investigation. In this research, we develop a fully automated framework for the analysis of OCT-A scans, we propose novel retinal measurements to characterise the microvasculature in the eye, and we show the application of our framework to a wide spectrum of conditions with the aim of demonstrating the full potential of this technology.

OCT-A computational framework

The first step of the OCT-A computational framework is image processing. Since OCT-A is a fairly recent technology, standard blood vessel segmentation has not yet been established. Hence, we first explored the best image processing pipeline by creating an original **OCT-A dataset** with manually segmented scans (publicly available) by investigating a large set of blood vessel segmentation methodologies. Handcrafted filters and neural network architectures were surveyed and evaluated using standard performance metrics and newly customised OCT-A vascular measures. Results showed the superiority of deep learning approaches over handcrafted filters and the susceptibility of clinical measurements to each of the **segmentation approaches**, suggesting the need to pay special attention to the image processing step when performing meta-analyses.

The second step of the OCT-A computational framework is conducting **network analysis for the discovery of retinal biomarkers** in the segmented OCT-A images. Common

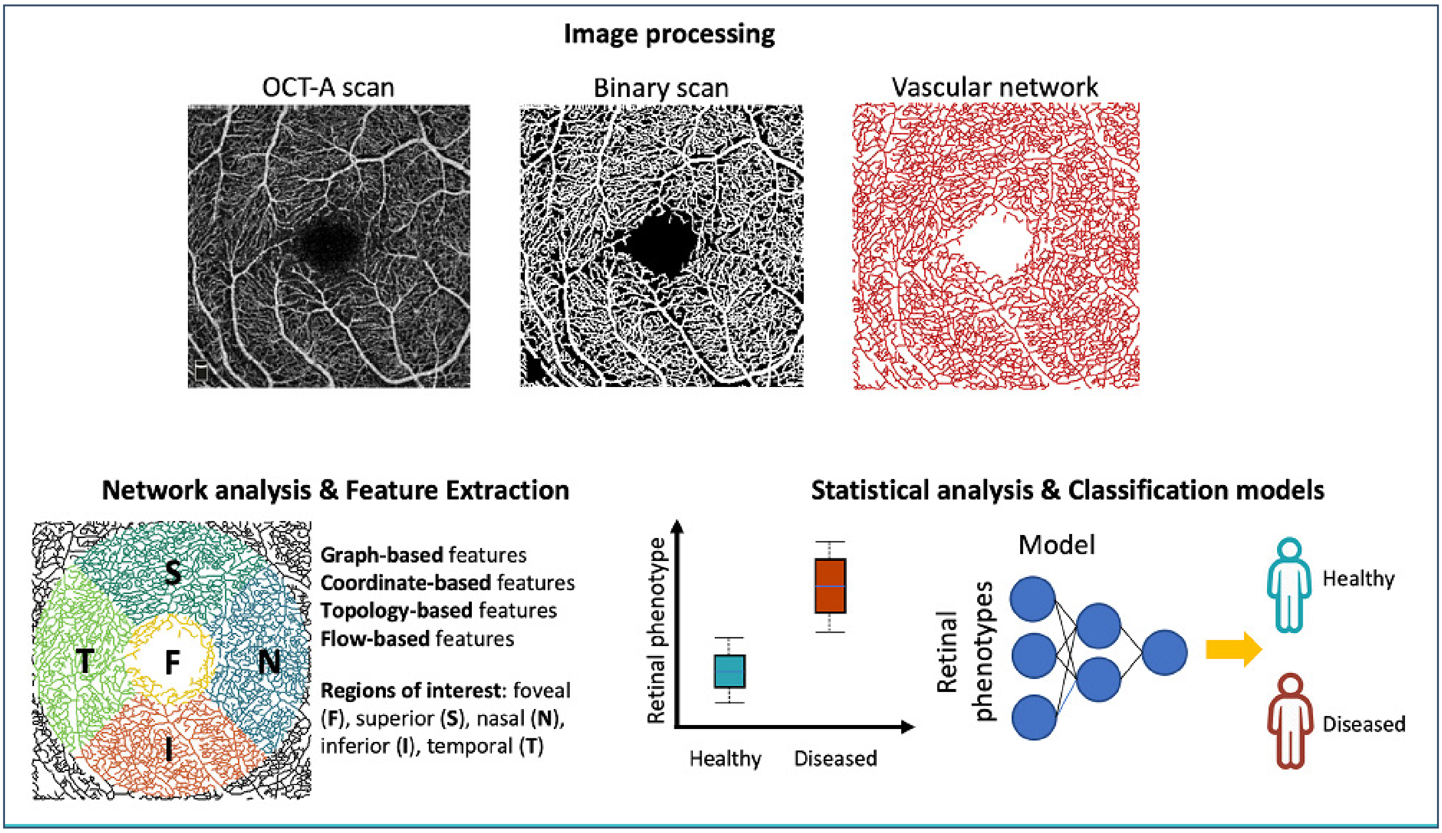

retinal characteristics such as vessel density, vessel radius, and vessel tortuosity were computed based on current clinical knowledge. In addition, we developed novel microvascular metrics based on geometrical, topological, and functional properties of the vascular network to cover the full spectrum of possible retinal measurements, enabling the hypothesis-free discovery of new clinically relevant biomarkers of diseases (see Figure with OCT-A computational framework).

Applications

The proposed candidate retinal biomarkers computed by our OCT-A framework can be used to explore associations in ocular and systemic dysfunctions and to build machine learning classifiers for the prediction of patient status. We have demonstrated the application of the OCT-A computational framework in three case studies with a retinal vascular footprint: diabetic retinopathy (DR), chronic kidney disease (CKD), and living kidney donation. In the DR study, we were able to reproduce previously reported results on changes at the microvascular level and achieve good performances in classifying patients with DR. Whereas changes in the eye of patients with CKD were more subtle and challenging to detect. Finally, despite living kidney donors being considered near-healthy patients, they remain at higher risk of developing cardiovascular disease and CKD. We investigated, for the first time, the possibility of OCT-A retinal imaging as a tool to improve the targeting of patients at risk. Our results suggested that OCT-A microvascular phenotypes may provide further insights into the long-term risk assessment of kidney donation.

METGAN: GENERATIVE TUMOUR INPAINTING AND MODALITY SYNTHESIS IN LIGHT SHEET MICROSCOPY

Izabela Horvath

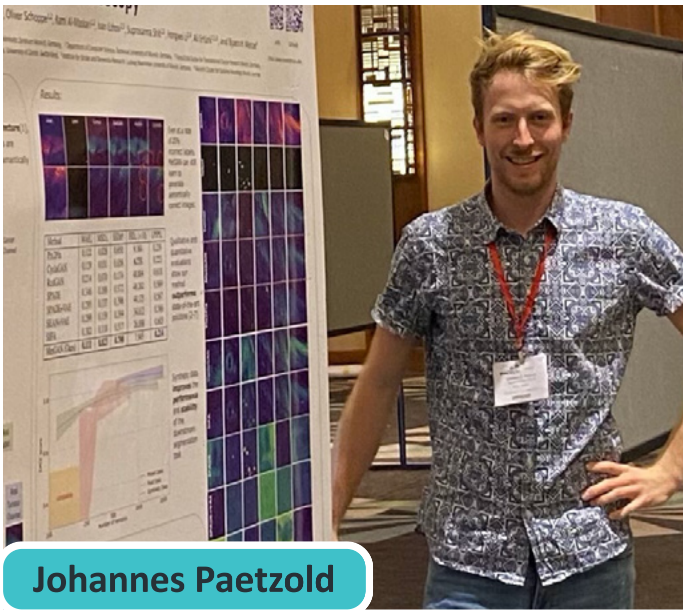

Johannes Paetzold

Johannes Paetzold @jocpae is coming to the end of his PhD at the Technical University of Munich (TUM), under the supervision of Professor Bjoern Menze. Izabela Horvath is pursuing a PhD at TUM under the supervision of Professor Menze, as well as being a PhD candidate at the Erturk-Lab, Helmholtz Munich, under Ali Ertürk, where Johannes leads the AI Team. Their joint paper proposing a novel generative method for light sheet microscopy was just presented by Johannes at WACV 2022 in Hawaii. Sadly, Izabela was unable to travel due to administrative reasons. However, we are pleased to say they are both here now to tell us more about their work.

This paper uses a **generative adversarial network approach** to solve the problem of **needing high-quality labeled ground truth data to develop robust segmentation methods**. The team found state-of-the-art methods did not prove useful or sufficient for generating the labeled data required. They sought to develop a new method to generate more diverse data and improve the segmentation for novel specimens.

"The imaging method we use in our lab is called ***light sheet microscopy****,"* Izabela tells us. *"With this method, we can image the same specimen in multiple channels to highlight different aspects of the organism. We have an autofluorescence channel at the base to better image the anatomy. We impose a label that we use as a ground truth label to obtain a contrast channel – what we call in our paper the* ***cancer channel*** *– where we highlight objects of interest. In our case, it was* ***metastases****. These are small tumors that are spreading throughout the body."*

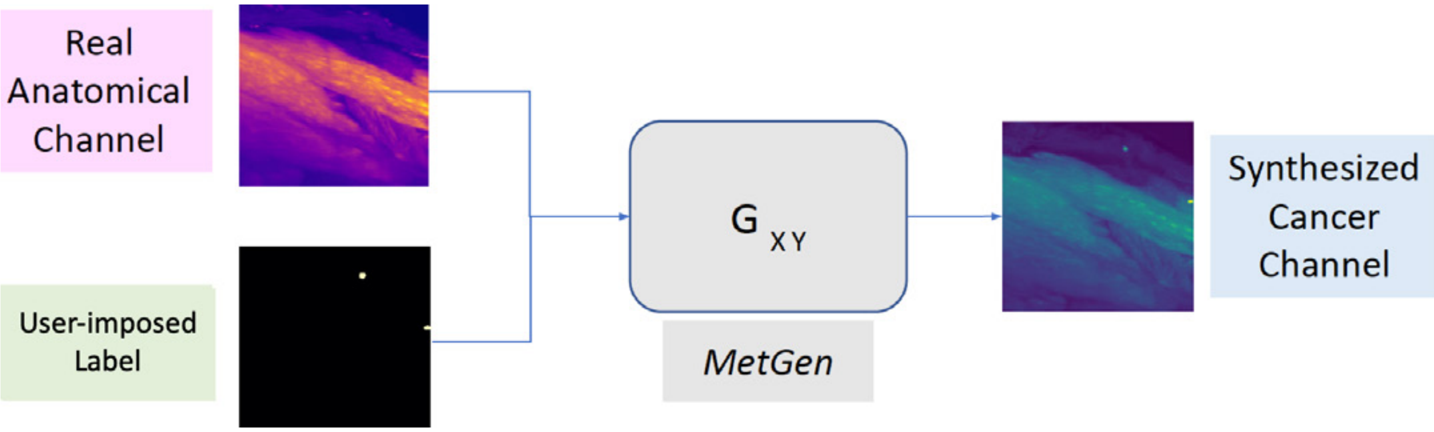

Inference. In order to generate realistic annotated data, we can use a combination of a real background and a pre-determined, user-specified label.

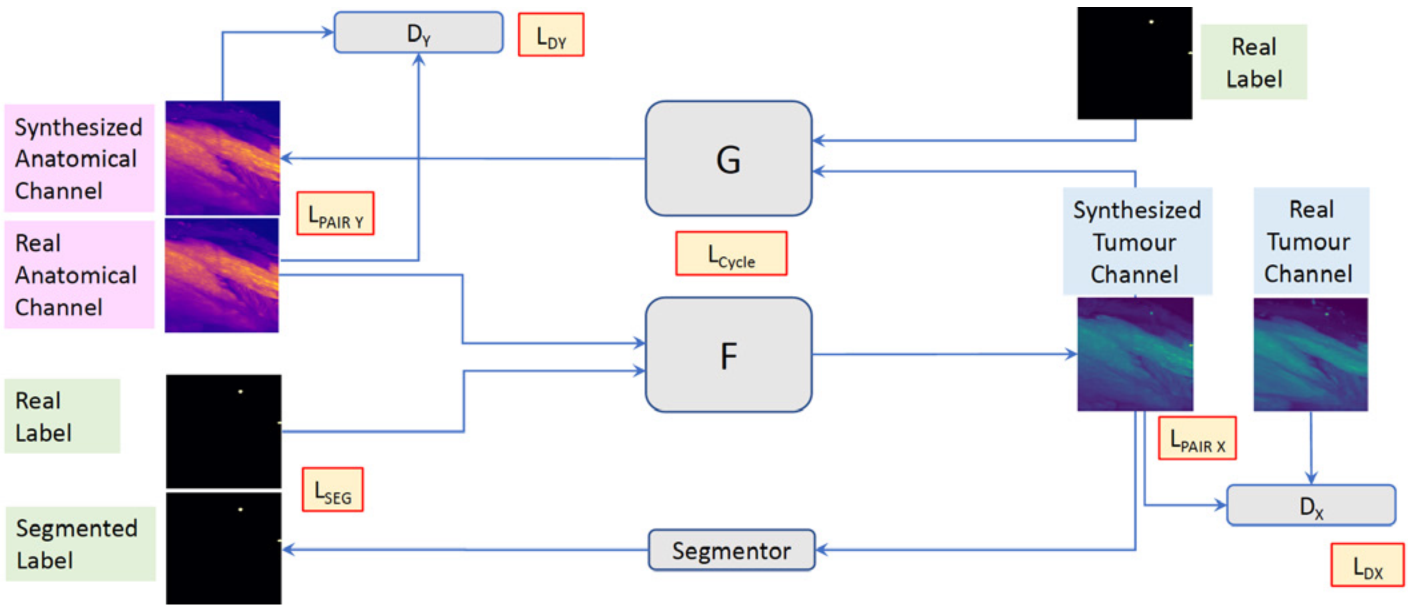

Within our pipeline, we train two generators and two discriminators in a cycle-consistent manner. The generator that is creating the synthetic tumour domain images is further constrained by a pre-trained segmentor.

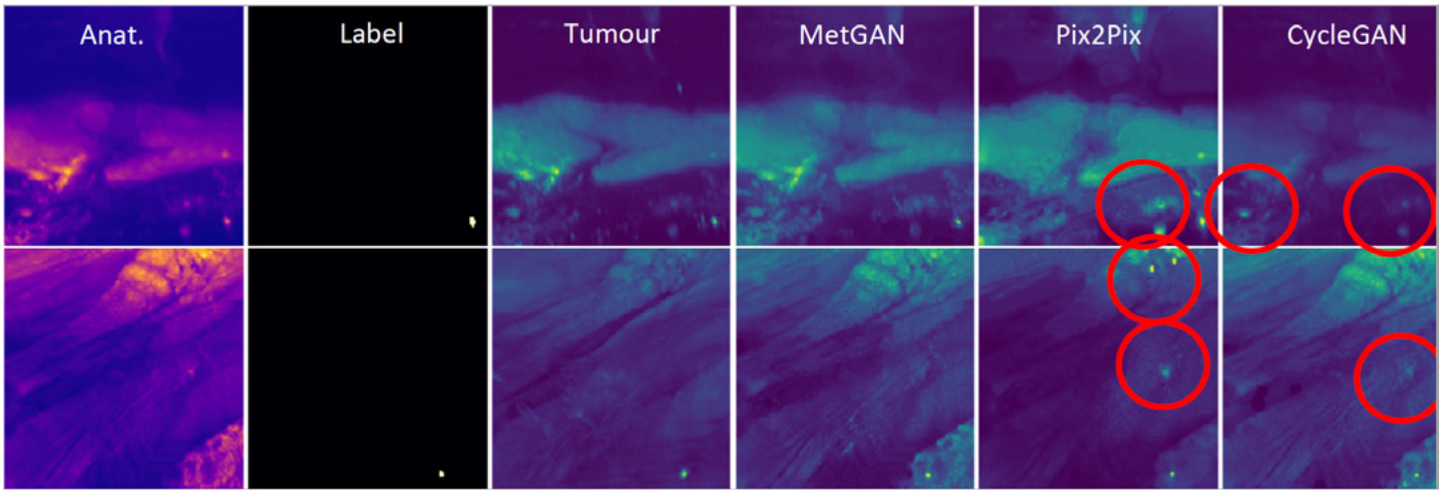

State of the art methods are not designed to handle multi-channel data with often inconsistent annotations, thus justifying the design of a novel architecture and training pipeline tailored to our problem.

Previously, biologists would be tasked with manually analyzing this kind of data and counting tumor cells. They did not have the capabilities to look at the whole specimen at the same time, so they would only look at small parts. **Deep learning** has been an enabler for unbiased analysis, making it possible to explore multiple whole animals or organs in rapid time.

"The data is so large," Johannes points out. *"We're talking about single images of 10,000 by 10,000 by 10,000 pixels. Machine learning is the only way to analyze this in an efficient manner."*

It is clear when they speak that Izabela and Johannes are both highly passionate about this work.

"What excites me is that we are working with truly experimental imaging data," Johannes tells us.

Izabela adds: *"I'm passionate because I can see the problems we work on have real applications. We can drive discoveries,*

***Vascular structure of a whole mouse brain imaged with light sheet microscopy. Source: [Nature Methods](). ***

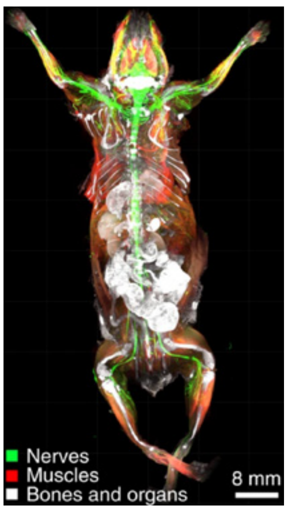

3D reconstruction of Light Sheet Microscopy-imaged mouse. Source: Nature Neuroscience.

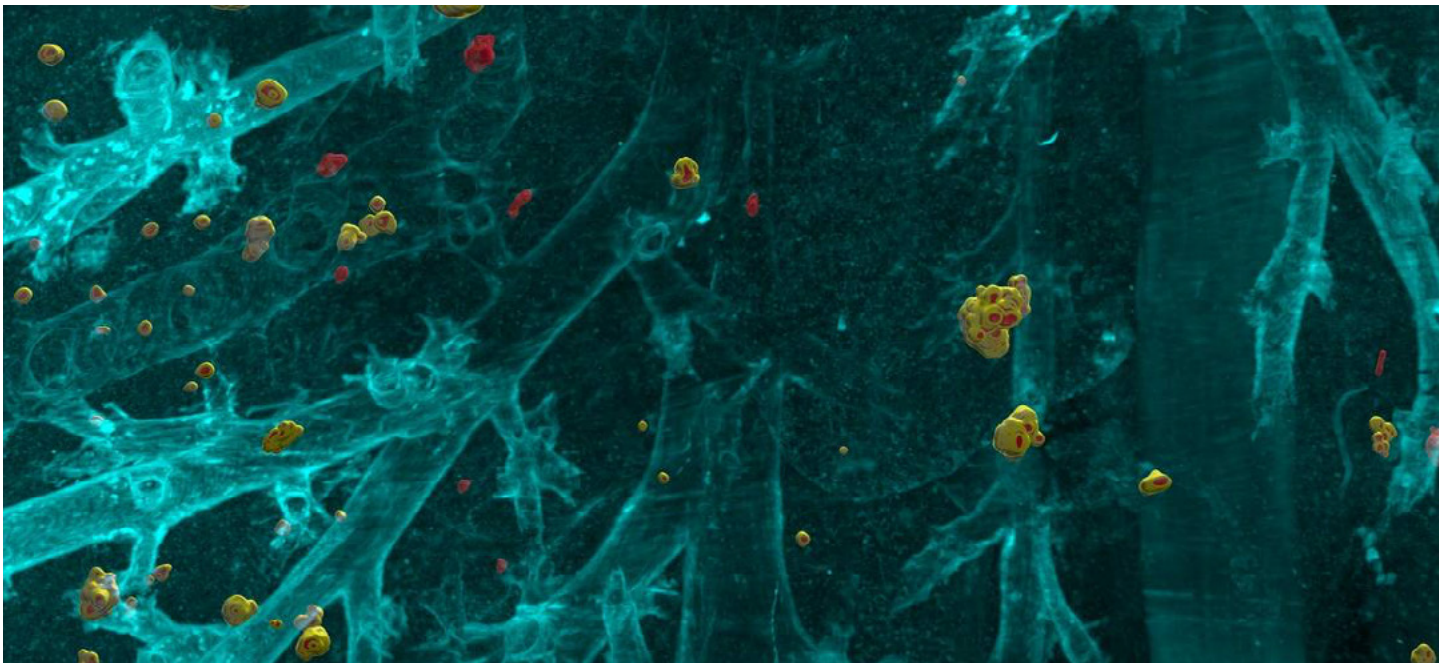

Detection of cancer micrometastases in mouse lungs. **Source: Cell.**

not only in computer science, but also in biology. I think that's a very cool thing!"

Outside of this paper, Izabela's usual work involves research on machine learning for light sheet microscopy applications. Johannes's main interest is in understanding large biological network structures, such as vessels and neurons, using machine learning applications.

"My favorite tools for this used to be CNNs and segmentation, but I have shifted my research focus more towards graph learning or geometric deep learning," he reveals.

We tell Johannes we have heard on the grapevine he has an exciting new position in the pipeline.

"Yes, in fact I was really interested to read your interview last month with my future boss, **Daniel Rueckert***! I will be starting a postdoc at* ***Imperial College London*** *soon while retaining my remote position as team lead for AI in the Erturk-Lab."*

Thinking about next steps for this work, the team say their approach could be extended to other state of the art medical imaging domains.

"The multimodal data we see in microscopic images are also present in other imaging domains, like MRI, where you have different kinds of sequences," Johannes explains.

Izabela agrees: *"Yes, it would be interesting to try contrast MRI. Also, in our lab, we already have novel experiments in the pipeline, and we want to apply this method to improve our segmentations on those – not only for cancer imaging, but for other pathologies as well."*

AIROGS: ARTIFICIAL INTELLIGENCE FOR ROBUST GLAUCOMA SCREENING CHALLENGE

Coen de Vente is a third-year PhD student at the University of Amsterdam, under the supervision of Full Professor of AI and Health Clarisa Sanchez. They are two of the co-organizers of AIROGS, the Artificial Intelligence for Robust Glaucoma Screening Challenge, which is taking place as we speak and will present its results at ISBI 2022 in India next month. Coen and Clarisa tell us more about the challenge as it enters its final test phase.

The **AIROGS Challenge** is about the automated detection of an eye disease called **glaucoma**, which affects the optic disc in the eye and can cause blindness. Usually, patients are identified through **color fundus photography graded by human experts**, but this can be a slow and inefficient process with so many of these images to analyze.

To solve this, AI approaches have been proposed to automate the screening process and make it more cost-efficient. These have proved promising in the lab; however, **they do not perform as well in real-world scenarios**. This challenge aims to **develop automated methods with a focus on screening performance and robustness of results**.

Participants have been given a training set with gradable images for developing their models. There is also a closed test set including ungradable images to encourage the **development of methods with robustness mechanisms to detect out-of-distribution samples**. Methods cannot be trained with the ungradable images.

"We have a large dataset, which we want to be as close to a real-world scenario as possible," Coen explains. *"It includes data from many different screening sites, optometrists, and hospitals. All images have been manually graded with a label of either referable glaucoma, non-referable glaucoma, or ungradable. We wanted a truly diverse dataset that will result in much more robust solutions than already exist."*

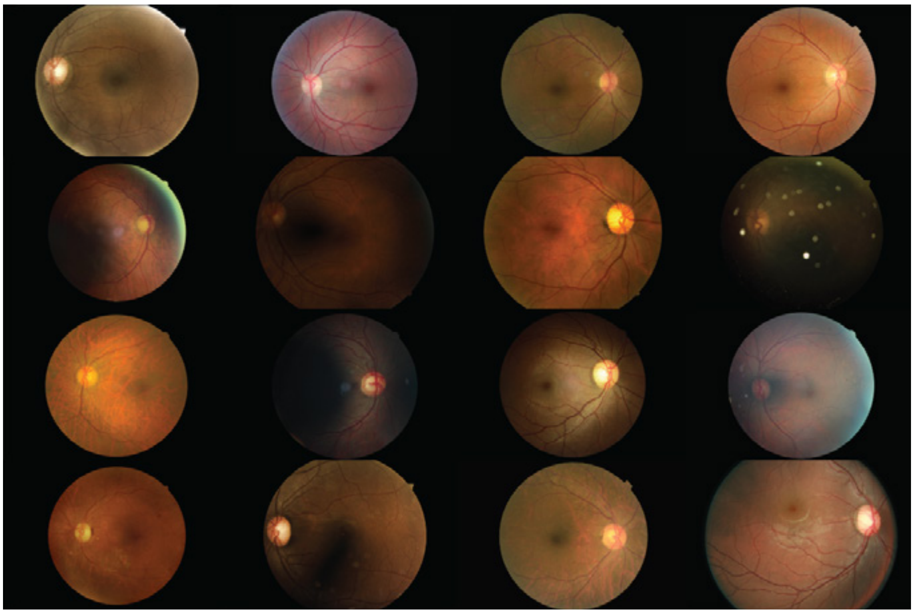

Automated methods can fail in the real world due to unexpected out-of-distribution data and poor-quality images. These images are often kept out of curated evaluation datasets, but part of this challenge is evaluating how well solutions can distinguish between **gradable and ungradable images**.

"We realized that when AI solutions are applied in real clinical practice, they encounter situations they can't cope with, and performance drops," Clarisa tells us. *"We wanted to focus on a challenge that represents real-world problems, so we're trying to reproduce some of these problems to assess how the algorithms perform."*

Clarisa says you should **keep in mind any issues** you've encountered when developing algorithms in the past if you're taking part in the challenge.

"Think back to times you've processed datasets and faced problems due to bad quality or other errors," she advises. *"That happens in real practice, and maybe it appears here too."*

Coen adds:

*"**Be careful not to focus on only one of the two tasks** but consider there are two sides to the challenge. The training set and the test set are different. We have deliberately set aside this ungradable data to simulate what would happen in the real world in a different setting."*

Looking ahead, Clarisa tells us there is **a hot topic** they will need to consider for future challenges.

*"Everybody is talking about **trustworthy AI**. There are already European guidelines, and they're also coming in the US, so we're all looking at what's needed. We'll need to think about how we incorporate that into new challenges."*

Clarisa and Coen have enjoyed organizing this challenge, but with a mix of technical people and clinicians on board, it has not been without its obstacles.

"We've had to learn how to communicate with each other!" Clarisa laughs. *"I remember once we had a long discussion, and we were all talking about the same thing but in different ways. After that, we said, okay, we need to find a common language! That really helped going forward."*

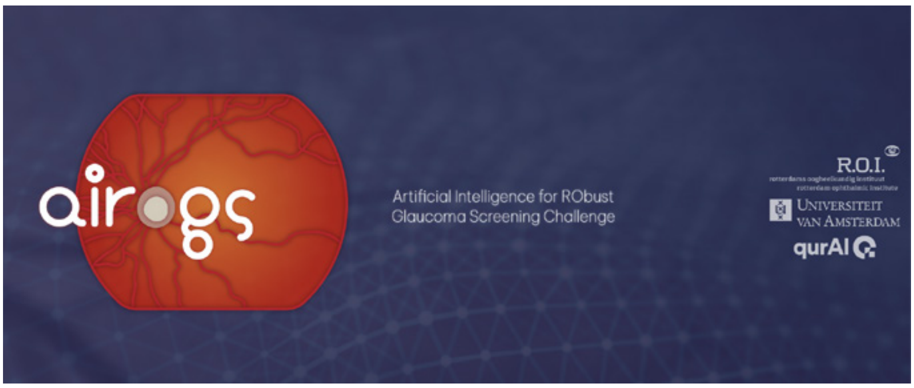

Computer Vision News has found great new stories, written somewhere else by somebody else. We share them with you, adding a short comment. **Enjoy!**

Kernel Flow Headset: A Wearable Device for Noninvasive Optical Brain Imaging

Believe it or not, this wearable helmet-shaped device monitors brain activity. In fact, functional near-infrared spectroscopy (fNIRS) is a widely used noninvasive imaging technique that employs near-infrared light to determine the relative concentration of hemoglobin in the brain. Current scanning systems being expensive and complex, researchers from a neurotechnology company called Kernel are now able to overcome these challenges: they have developed this wearable headset which offers accurate results and possibly making neuro measurements mainstream. Read More

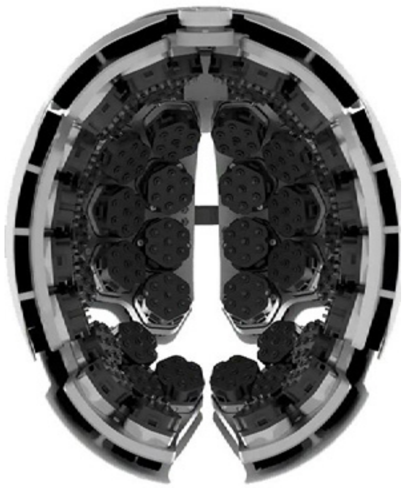

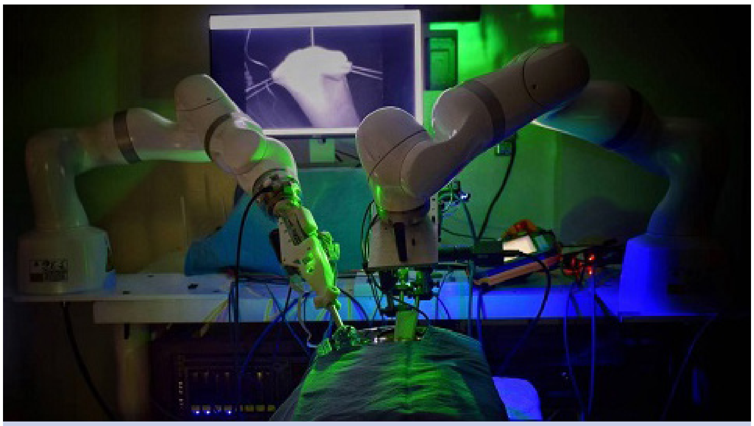

Smart Tissue Autonomous Robot Performs Laparoscopic Surgery on Pig Soft Tissue Without Human Help

Scientific breakthroughs in Robotic-Assisted Surgeries are fast and very significant, as readers of Medical Imaging News and followers of RSIP Vision know well. The one described in this story by Johns Hopkins University and Science Robotics tells of a robot that has performed laparoscopic surgery on the soft tissue of a pig without the guiding hand of a human. This could lead in time to robotics-enabled fully automated surgery on humans. Results show that the Smart Tissue Autonomous Robot (STAR) was able to successfully reconnect two ends of an intestine. Read More

Protecting Healthcare from The Growing Threat of Cybercrime

While Covid-19 is presenting a very real threat to healthcare systems all over the world, a more ominous and human-initiated threat is growing in the background. There's been a massive jump in the number of criminals targeting the sector with cyberattacks. What makes it even more heinous is the fact that these crimes come at a moment when health systems are vulnerable and they need clear access to all their data and medical images. The answers are: education to cyber threats, keeping updated software and maintain secure backups of critical data and images. Read More

RSIP VISION
Visual Intelligence for MedTech

TUM

Imperial College London

BAY VISION VIRTUAL
MEETUP

AI and the future of radiology

How AI is elevating medical imaging while preserving privacy & providing trustworthy, explainable clinical solutions

Missed the Meetup?

Don't miss the video!

Be sure not to miss next time :-)

GUEST SPEAKER

Prof. **Daniel Rueckert**,
PhD, FREng, FMedSci, FIEEE
Alexander von Humboldt Professor for AI in Medicine and Healthcare at the Technical University of Munich. Professor of Visual Information Processing in the Department of Computing at Imperial College London.

HOSTED BY

Moshe Safran
CEO of RSIP Vision USA
Defining & developing innovative Medical Visual Intelligence solutions, in partnership with MedTech industry leaders.

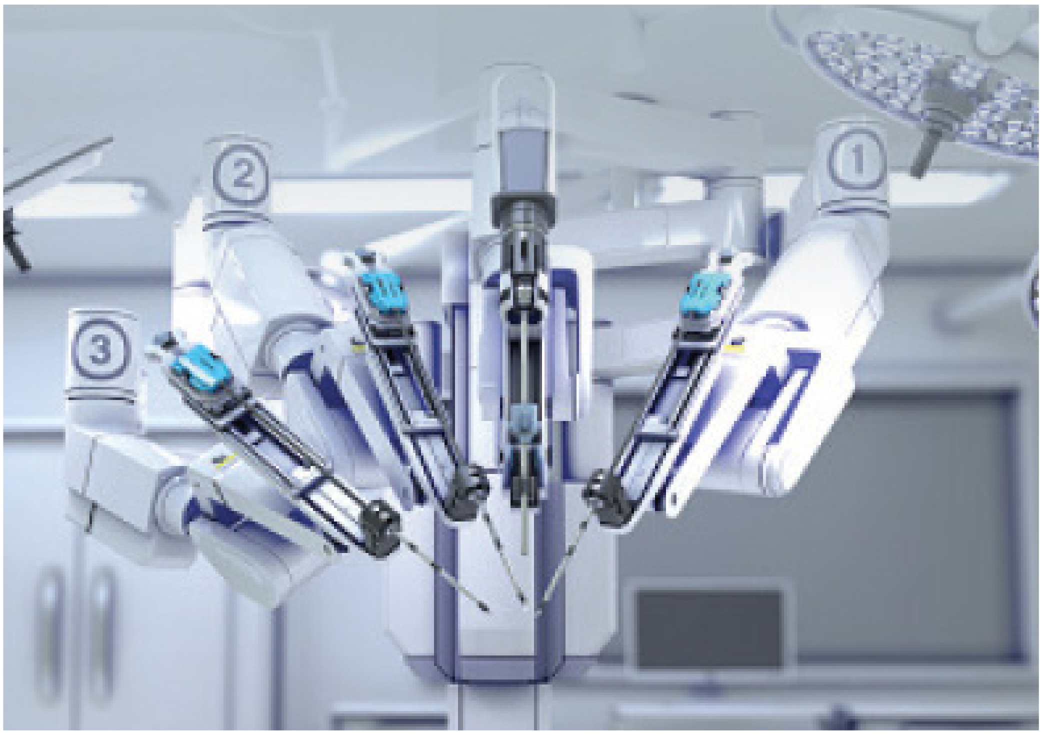

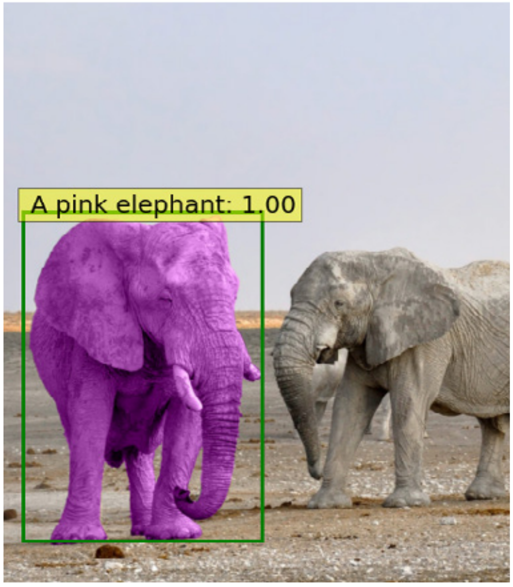
A pink elephant: 1.00