JUNE 2022

Computer Vision News

& Medical Imaging News

The Magazine of the Algorithm Community

This photo was taken in peaceful, lovely and brave Odessa, Ukraine.

Computer Vision News

Editor:
Ralph Anzarouth

Engineering Editors:
Marica Muffoletto
Ioannis Valasakis

Publisher:
RSIP Vision

Contact us

Free subscription

Read Previous Magazines

Follow Us

Dear reader,

It's a significant month for the computer vision community as **CVPR** returns with its first **in-person conference** in three years in the eclectic city of New Orleans. **RSIP Vision** and **Computer Vision News** will be there, and we're so excited to meet all of you! We'll be publishing four new **CVPR Dailies** across the week. Receive them in your inbox daily from 21 June by subscribing to our initiative: **Feel at CVPR as if you were at CVPR!**

With CVPR in mind, turn to page 4 for a preview of the **Computational Cameras and Displays workshop** with co-organizer **Emma Alexander**. Also, on page 10, **Marica Muffoletto** has reviewed a brilliant CVPR paper with Ukrainian co-authors: **Is Mapping Necessary for Realistic PointGoal Navigation**?

Computer Vision aficionados will adore our review of **Liberty Defense's HEXWAVE**, a revolutionary approach to security screening using **3D imaging and AI**. Find out more on page 18.

We love bringing you the BEST of the BEST, and that is why we review not one, not two, but three award-winning papers this month. On page 14, learn about **bootstrapped meta-learning** with the winner of an **Outstanding Paper Award at ICLR 2022**. We also speak to the winners of the **Best Paper Award** (page 36) and **Best Demo Award** (page 40) at **CRAS 2022** about innovative new platforms for **robot-assisted surgery**.

It's never too early to talk about **MICCAI**, and we've got the inside scoop on September's main event from **General Chair Shuo Li**. Read our interview on page 46. We also have a preview of the HECKTOR challenge on page 58.

Enjoy reading all this rich content and more in June's edition of Computer Vision News and our supplement Medical Imaging News. Next month, we'll have an extraordinary **BEST OF CVPR** section. You won't want to miss it!

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

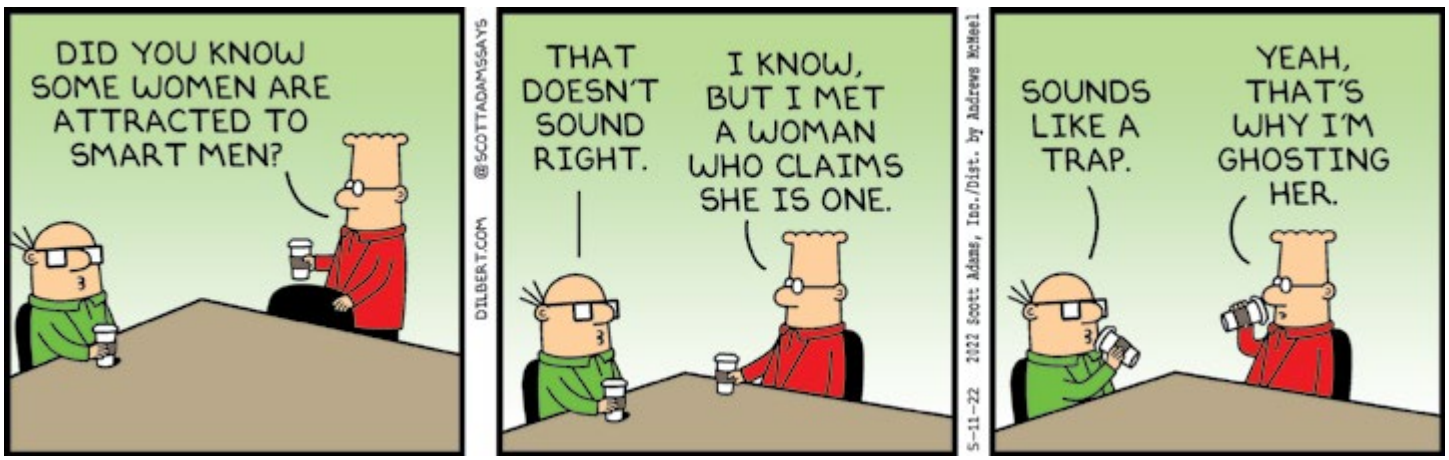

Computer Vision News

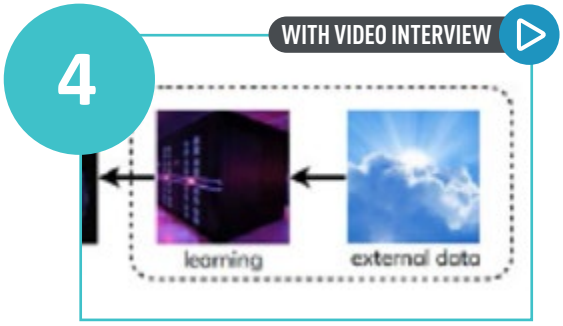

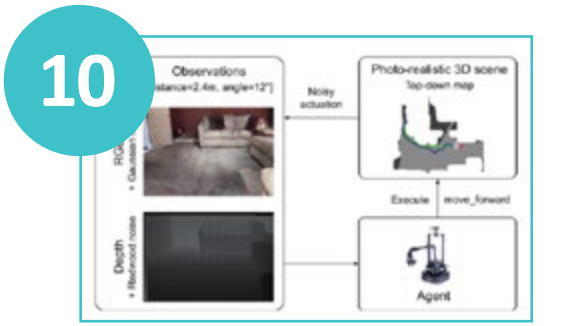

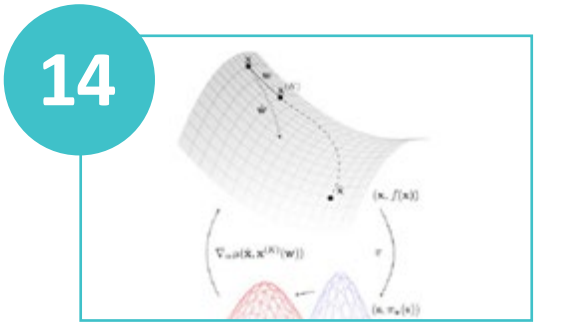

Medical Imaging News

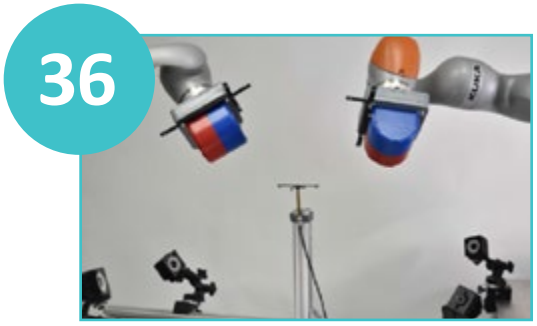

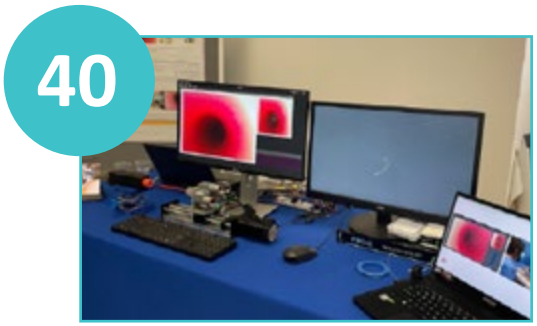

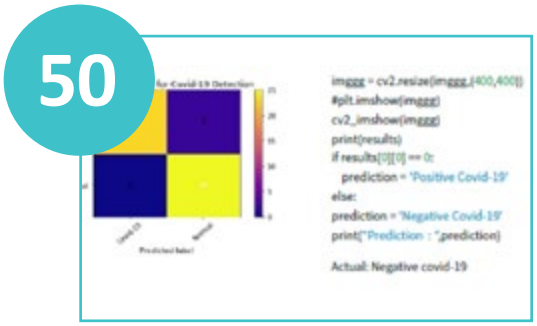

COMPUTATIONAL CAMERAS AND DISPLAYS WORKSHOP

Emma Alexander is an Assistant Professor at Northwestern University and a co-organizer of the upcoming Computational Cameras and Displays Workshop. She is here to tell us more about the event, which takes place at CVPR 2022 later this month.

Now in its 11th edition, the **Computational Cameras and Displays Workshop (CCD)** is an annual gathering of researchers and practitioners with an avid interest in the field.

A computational camera is

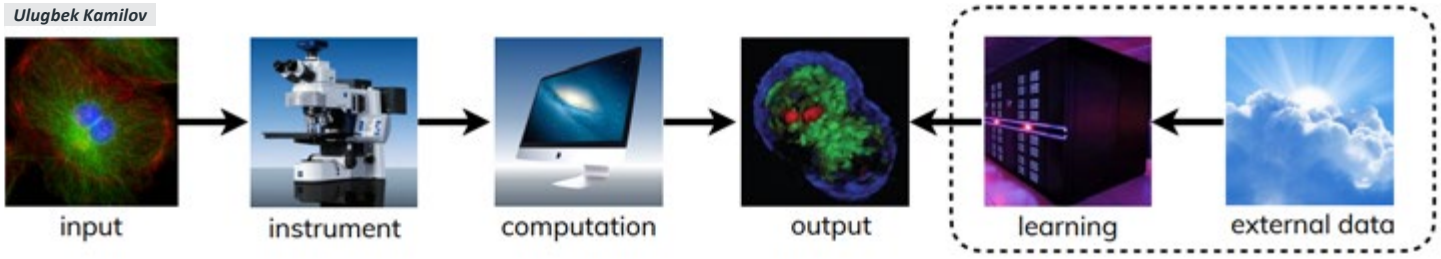

a camera where **the optics and the algorithm are co-designed to make scene information available**. That could be moving some of the computation into the optics to use the physics of light to pre-process data and save on computational efficiency, or it could be using lots of computation to exceed the native limits of the optical system. Doing this unlocks enormous potential in both cameras and displays, where instead of using standard datasets and standard photographic and video data, the optical pipeline can be altered to create new kinds of data to access and display information.

Ten years after its first outing, what keeps people coming back to CCD?

"*There's new work coming out every year and a great deal of space to explore in displays and cameras,*" Emma tells us.

"*Standard cameras are very simple. It's a lens and a sensor. But there are plenty of new designs and technologies even within the components that we know – for example, just arranging lenses and filters in new ways or pairing them with new algorithms. There's so much to explore. Every advance in mainstream computer vision gets multiplied and exponentially expanded in terms of the potential once you let every single component change. We're never going to run out of new computational cameras and displays.*"

One area where the workshop offers broad appeal to the mainstream CVPR audience is self-driving cars. It is widely acknowledged that regular cameras do not cut it in this domain. This year, the event will be expanding its optical horizons by looking at **LiDARs and SPADs**.

With the workshop schedule at CVPR typically full to capacity, what will make CCD stand out from the crowd?

"*It's going to be **a very creative workshop** – you're going to see and think about things in new ways,*" Emma reveals.

"*When you break out of the mold of what a camera or a display is, you can bring back **the magic of science and problem-solving**. I always learn, not just new facts or new techniques, but whole new perspectives on the field by finding out about these creative, innovative ways*

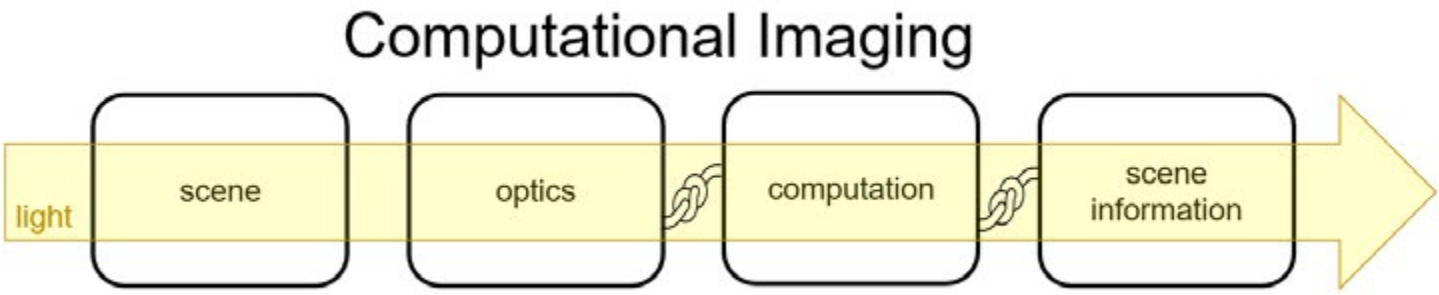

that people are thinking about things the rest of us take for granted."

The event aims to showcase the diversity of the field with a range of invited speakers at different points in their careers, covering an array of topics and geographies. With this year's guests all lined up, we ask if there is a dream booking Emma would love to have come along next year.

"*In terms of the big names, everyone we invited was able to make it, but I hope we get a poster submission from a grad student we've never heard of this year so that next year we can invite them,*" she responds.

"*I think the fear as an organizer is that you're going to miss new creative thinkers on the edge of things – people who aren't as well connected. This is a workshop that showcases new and unusual perspectives. It's about getting a wide view of the amazing work being done in different places in different ways by different people.*"

When CVPR returns this month, it will be its first in-person conference for three years. For many young scholars, their only experience with conferences and workshops will have been virtual. Does Emma have any advice for someone attending a physical event for the first time?

"*Well, they should come to my workshop – that's tip number one!*" she laughs.

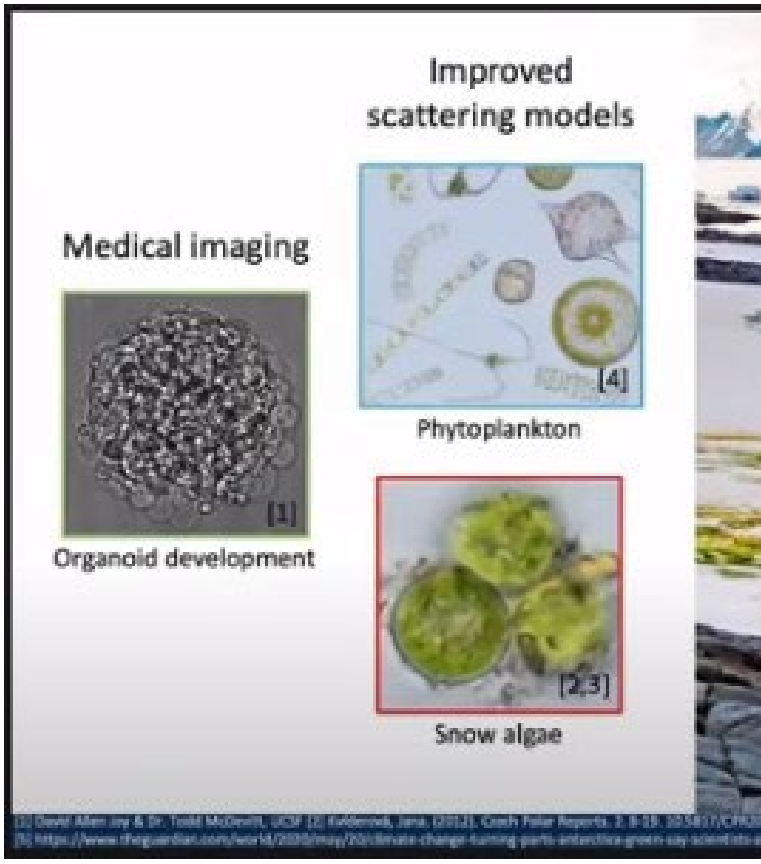

"CVPR is this amazing buffet table, but also this overwhelming circus, and what you think you're going to be interested in versus what catches your eye is unpredictable. ***Be open, explore widely, and don't be shy****. These folks have missed out on many years of getting to know each other.* ***They should make friends*** *and know that everyone who's a little more established has missed them and is looking out for them. Just talk to lots of people!"*

With so many exciting developments coming up in the field, CCD is sure to be a feature at CVPR for many years to come. There are **new kinds of optics and sensors being developed all the time**, with principled incorporation of the tools from machine learning and the fundamentals of physics and mathematical modeling.

"We're making much progress on both of those fronts, but the real missing piece that this workshop is actively addressing is how you put them together in a principled way to take advantage of both," Emma points out.

"For anyone who wants to combine the engineering tools highlighted at CVPR with a deeper scientific understanding of imaging, attend this workshop!"

Regina Eckert's talk at the CCD Workshop - CVPR 2021

Thank you to Antonino Furnari and Marco Del Coco for these photos from ICIAP, the International Conference on Image Analysis and Processing, held last week in Lecce, Italy. We recognize awesome Laura Leal-Taixé and Nicu Sebe, both great friends of our magazine!

Are you going to miss CVPR 2022? We can be in touch anyway!

Follow CVPR almost in real time:
click here and subscribe for free to CVPR Daily
(21-22-23-24 June)
with all the highlights from CVPR!

IS MAPPING NECESSARY FOR REALISTIC POINTGOAL NAVIGATION?

By Marica Muffoletto (twitter)

CVPR 2022 is again the inspiration for our monthly review, with a paper entitled: Is Mapping Necessary for Realistic PointGoal Navigation?

By choosing this paper, Ralph and I intend to show our continuous support to Ukraine and its people. We deeply thank all authors (Ruslan Partsey, Erik Wijmans, Naoki Yokoyama, Oles Dobosevych, Dhruv Batra, Oleksandr Maksymets) for allowing us to use their images.

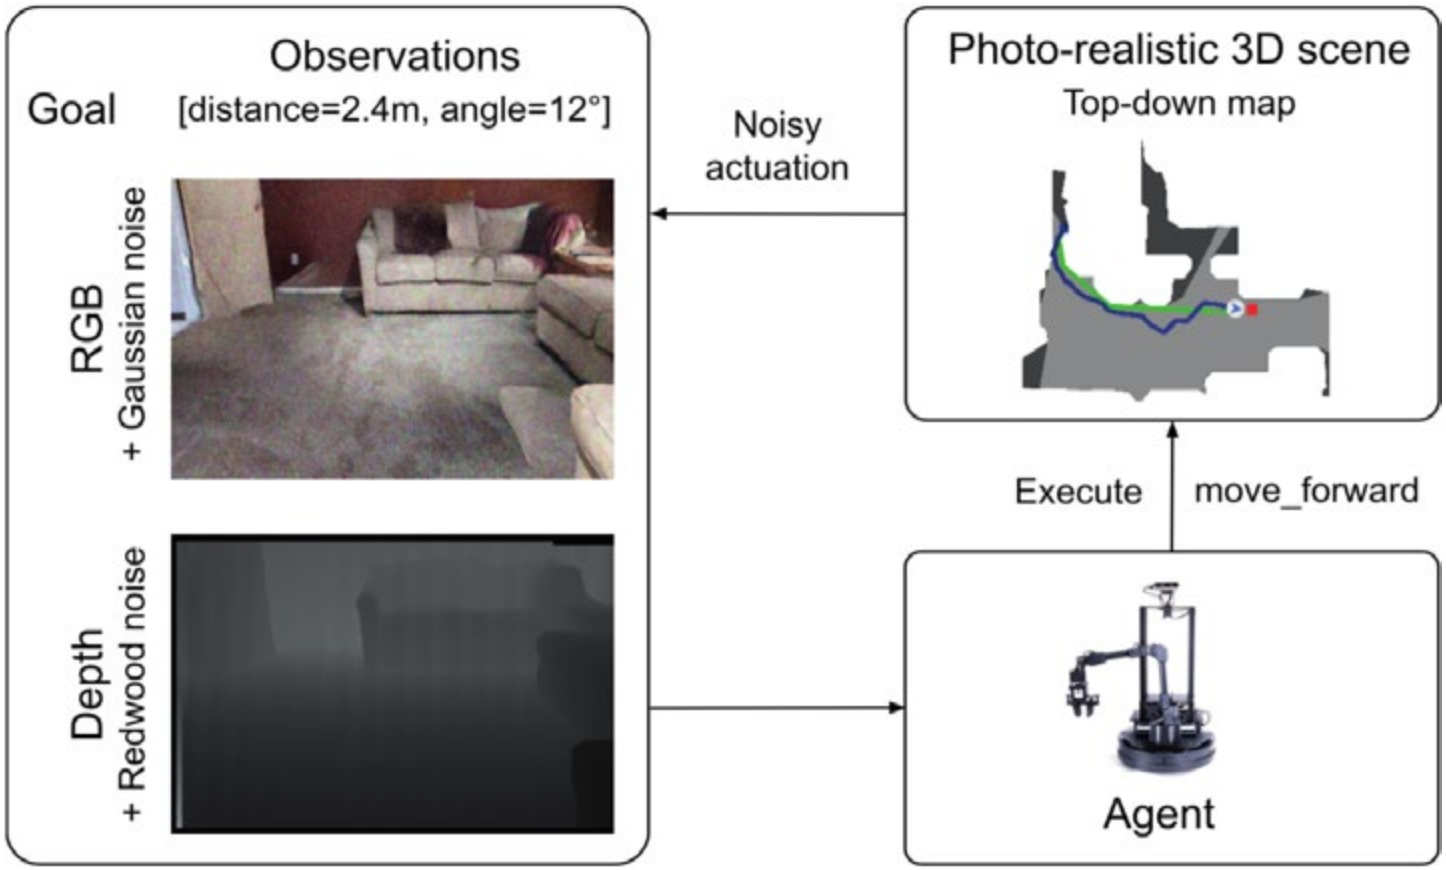

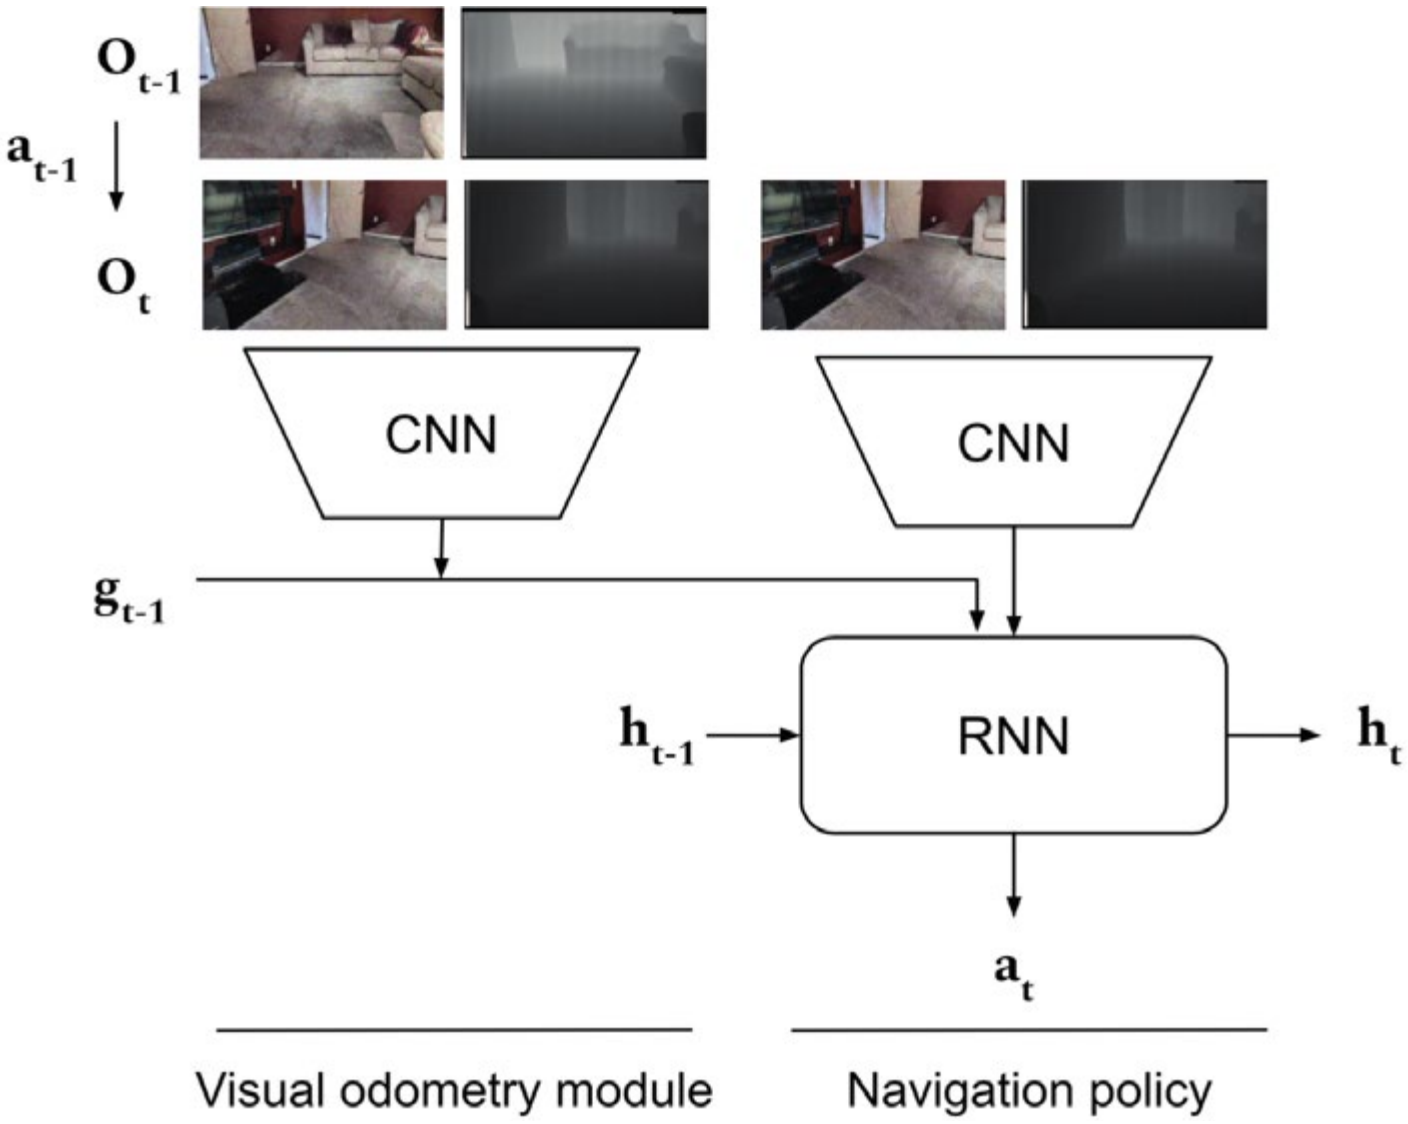

The problem of goal-navigation was the subject of a previous article in **Computer Vision News of December**. We are now again looking into this from a different perspective: **How is Realistic PointGoal Navigation different from idealized one, and how much does it rely on Mapping?** PointNav is a navigation task where an Agent is initialized in a previously unseen environment and is tasked to reach the goal specified relative to its starting location. The action space is discrete and consists of four types of actions: stop (to end the episode), move forward by 0.25m, turn left and turn right by a specified angle. In this work, the agent was equipped with an RGB-D camera mounted at a height of 0.88m and tilted -20°. Camera's resolution was 360x640 pixels with 70° horizontal field of view and base radius of 0.18m.

PointNav comes under two versions: v1) idealized setting: the agent is equipped with noise-free camera and access to ground-truth localization and movement is deterministic. In idealized setting, with no noise, map-less navigation models trained with large-scale reinforcement learning achieve 100% success. State-of-the-art approaches seem to have solved this problem; v2) realistic setting, where the agent must deal with actuation and sensing noise, and lack of high-precision localization in indoor environments. This is

considered yet an unsolved challenge.

In all experiments, the agent is evaluated via three primary metrics. 1) **Success**, where an episode is considered successful if the agent issues the stop command within 0.36m (2×agent radius) of the goal. 2) **Success** weight by (inverse normalized) **Path Length** (SPL), where success is weighted by the efficiency of the agent's path, calculated considering the geodesic distance (shortest path). 3) **SoftSPL**, where the binary success is replaced by progress towards the goal.

The authors' approach is built on the following components: 1) a **CNN-based visual odometry module**- referred to as VO- that, given two consecutive observations (O_t, O_{t-1}), predicts the change between t − 1 and t and then updates the goal wrt the current pose (g_{t-1} and 2) an **RNN-based RL navigation policy module**, which is given the estimated pose g_{t-1} and the current observation O_t, and predicts the next action a_t.

Below you can observe an example of the combined VO+Navigation approach on the validation dataset with performance: SPL = 0.63, Success = 1, SoftSPL = 0.62.

The Navigation policy consists of a two-layer Long Short-Term Memory (LSTM) and a half-width ResNet50 encoder. To evaluate the two components separately and understand the impact of localization on navigation, it was trained assuming perfect odometry (hence given ground-truth location) and, only later, the VO module was used to estimate the localization as a drop-in replacement without fine-tuning. With ground-truth localization, the agent achieves 99.8% Success and 80% SPL on Gibson-val PointNav-v2 dataset, showing that visual odometry is a limiting factor to a map-less approach to realistic point goal navigation, while noisy observations and actuations can be overcome easily.

The VO module is made of a ResNet encoder followed by a compression block and two fully connected (FC) layers, where BatchNorm is replaced with GroupNorm, and the compression block consists of 3×3 Conv2d+GroupNorm+ReLU. It is trained on a static dataset D = {(O_{t-1}, O_t, a_{t-1}, $\Delta poase$)} and decoupled from the navigation policy.

Ablation experiments to this module included several additions to the basic network and analyzed:

➔ The effect of action embedding, by incorporating knowledge of the action taken between two consecutive observations as an additional input. This is shown to improve performance, because the network received more context to learn more accurate egomotion for each action type.

- ➔ The effect of training with a larger dataset, which substantially improves performance.
- ➔ The effect of train-time and test-time augmentations (Flip and Swap), which are found to be more effective with a larger training set.
- ➔ The effect of a Deeper encoder (ResNet-50 vs ResNet-18)
- ➔ The effect of Dataset transfer, which unsurprisingly shows poor performance when the agent is trained with visual odometry and sets an open question for future need of a universal (cross-dataset) VO module.

The experiments lead to an attempt to deploy the learned agent on a real-world challenge. Across 9 episodes, the LoCoBot provided with this agent achieves 11% Success, 71% SoftSPL and makes it 92% of the way to the goal (SoftSuccess). Watch a demo!

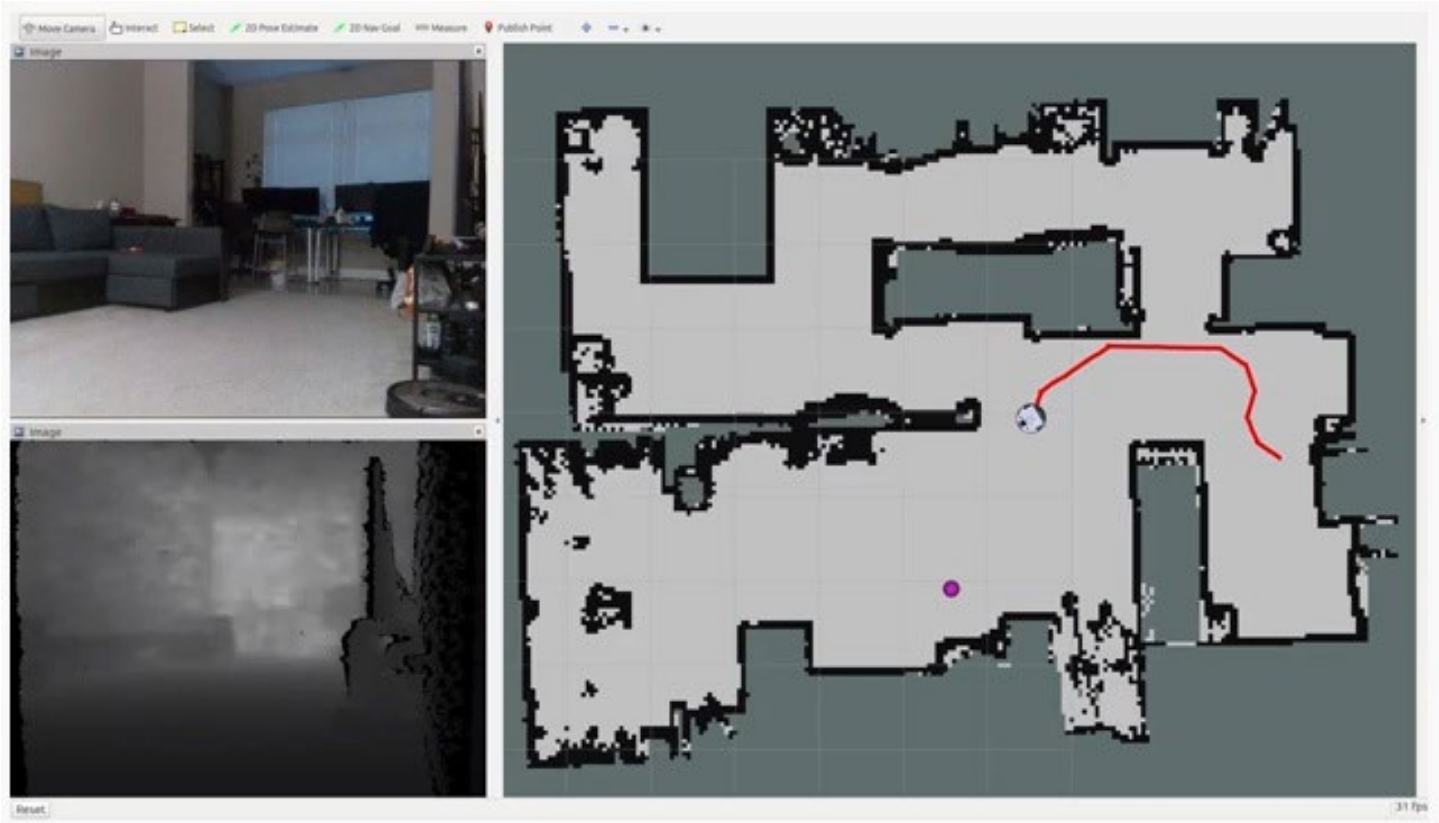

This work investigates the link between mapping and navigation, reaching the conclusion that this is a weak link, and proving that the only performance bottleneck in the PointNav task is the agent's ability to self-localize. The authors finally hint at some exciting future work on analysis of indirect links (from mapping to localization to navigation) and invariance of the approach to datasets and embodiment specificity. Looking forward already!

BOOTSTRAPPED META-LEARNING

Sebastian Flennerhag is a Research Scientist at DeepMind. He speaks to us after scooping an Outstanding Paper Award at ICLR 2022 in April for his work on bootstrapped meta-learning.

Learning to learn as an idea has been around for centuries, but early approaches in modern times date back to **Jürgen Schmidhuber** in the 1980s. It is usually applied by unrolling the update rule you are trying to meta-learn by some number of steps and then immediately evaluating performance.

"*In few-shot learning, we do this all the time,*" Sebastian explains.

"*We adapt for several steps and then ask: how did we do? We then optimize for that performance. The limitation is that* ***we have no idea what will come after****. If we'd trained for longer, we might have got even better.*"

This work aims to develop an algorithm that can aut**omatically tune another learning algorithm as it is being applied**. A simple example is online hyperparameter tuning. A more ambitious or advanced example would be discovering a learning algorithm directly from data. The core idea is to change the meta-objective – the way you're optimizing your learning rule.

"*I took inspiration from looking at how online optimization works and seeing if we could bring some of those ideas into the deep learning setting,*" Sebastian tells us.

"*The basic idea applies very broadly: you have something to optimize, you're going to look at how it behaves for a couple of gradient steps, and then optimize it such that if you were to train again, it would be much faster.*"

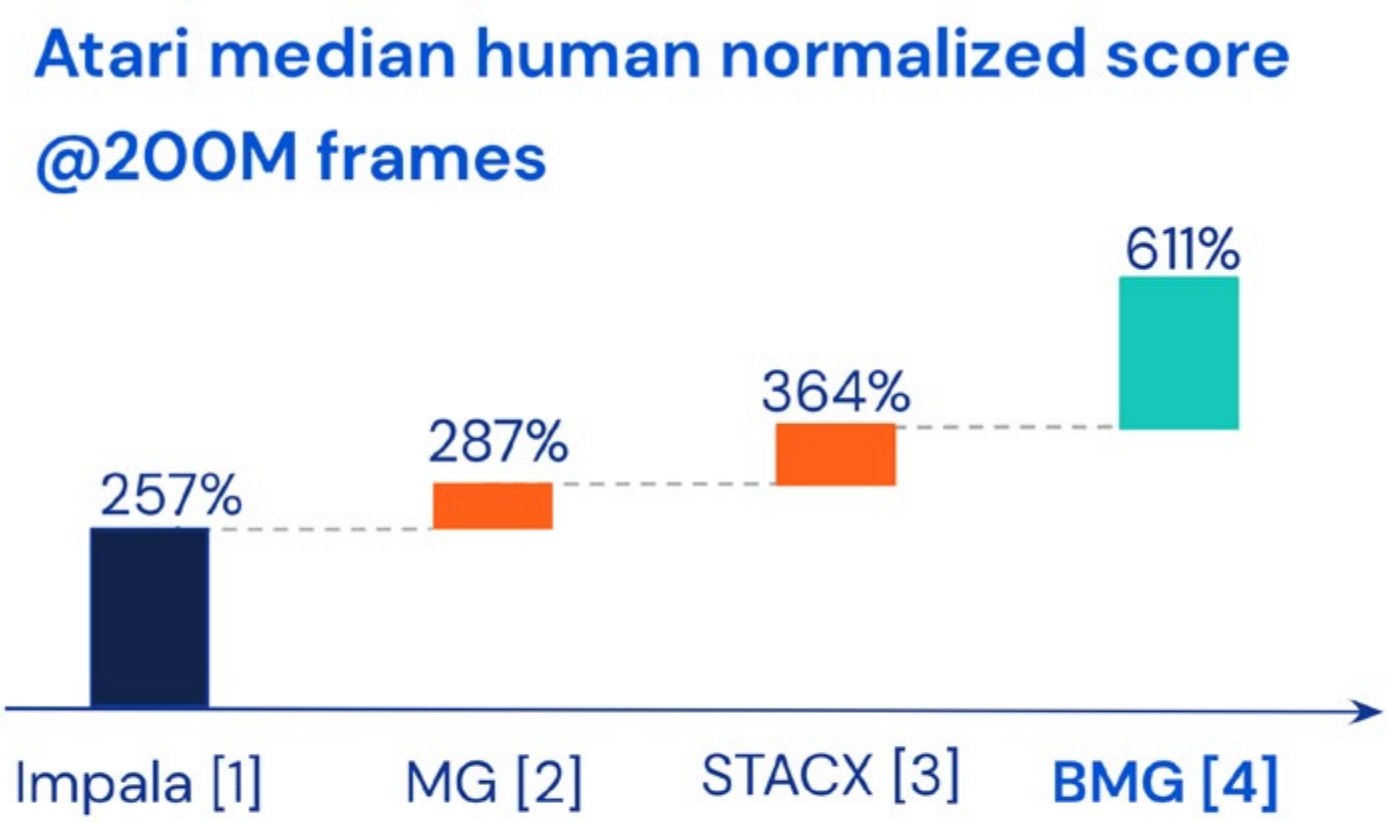

[1] Espeholt et. al. "IMPALA: Scalable ...." 2018.
[2] Xu et. al. "Metagradient Reinforcement Learning." 2018.
[3] Zahavy et. al. "A self-tuning actor-critic...." 2020.
[4] Flennerhag et. al. "Bootstrapped Meta-Learning." 2022.

The experiment people tend to default to in the gradient space is the Atari one, where **a vast neural network is trained using reinforcement learning to master a game like Pong or Breakout**. There are hundreds of thousands of parameters, and gradient descent optimizes the neural network given a reinforcement learning objective. These are typically called actor-critic objectives because two neural networks are interacting, with one trying to estimate how good the other is.

The mass of hyperparameters in these algorithms is critical to ensuring they are stable. If one algorithm is overestimating or underestimating, it can learn bad behavior. Researchers carefully set these parameters using **big hyperparameter sweeps or costly trial and error procedures**. The hyperparameters are automatically tuned in this work by running an extra optimization process – the **meta-learning** part. A **meta-gradient algorithm** predicts what the next best hyperparameters are going to be and then tunes that prediction again using **gradient descent**.

Sebastian tells us he found this an atypical project – it took him a couple of months of theoretical reading and toying around with ideas before he proposed a much simpler algorithm.

"I said there aren't any theoretical

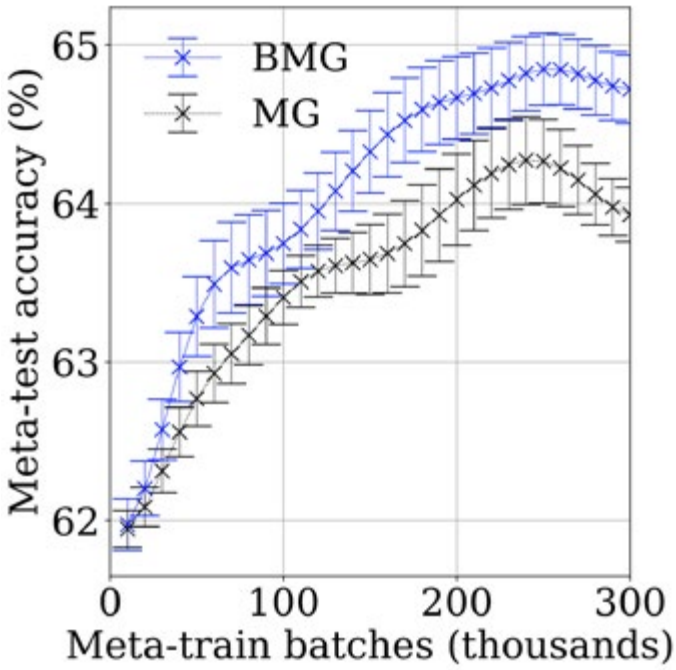

guarantees for this, but it looks nice on paper, and pretty quickly, we got very strong results," he recalls.

"*The empirical stuff is where you usually spend time making sure things work. But here, once we had the right algorithm and applied it, it worked straight out of the box. The hardest part is spending enough time on the abstract problem to have the insight to propose the right algorithm.*"

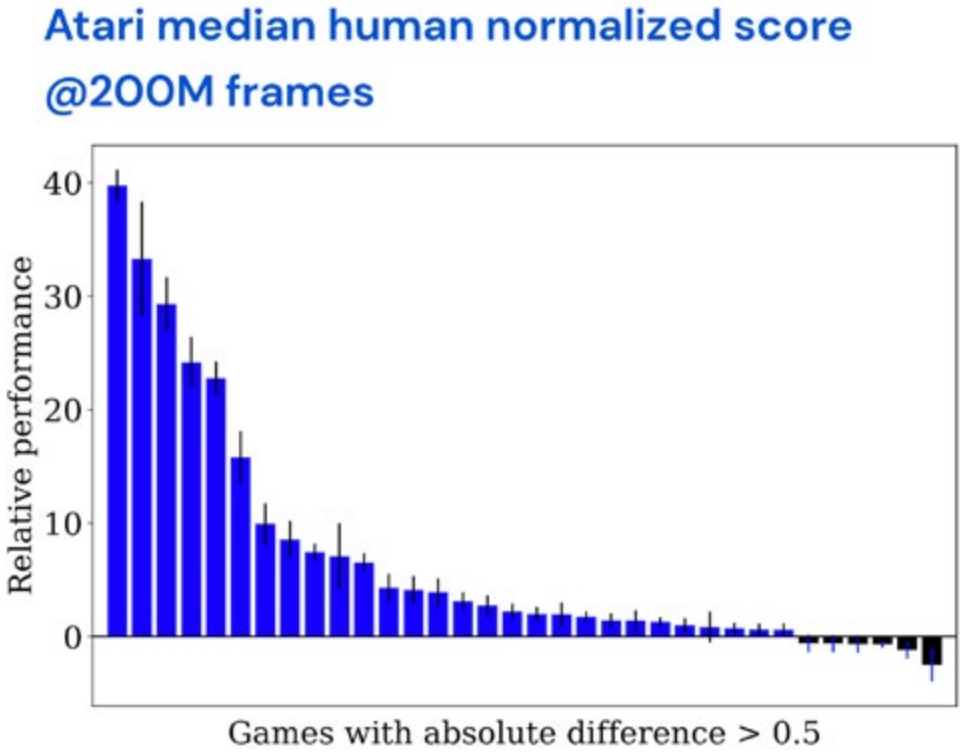

On behalf of the community, we have to ask, does he have any advice about what we should do when working on something and a solution is just not coming?

"*If I had a sure-fire answer to that, I would be halfway to the Nobel Prize!*" he laughs.

"*In my experience, when things aren't working out, simplify. Even if it doesn't turn out that the simplest possible algorithm is the right one, it's much easier to see what's going wrong. I apply a kind of* ***Occam's razor principle that the simpler the algorithm is, the more widely applicable it will be****. Rich Sutton's Bitter Lessons are all about this. Of course, you might have to increase complexity again at some point. Then it's important to know where you can afford to have that complexity so that it still works.*"

We ask Sebastian what he thinks convinced the jury that his paper deserved an award, and he points to three things. The first is that the idea is so simple and general, making it of interest to the community because it could be taken off in many different directions. Secondly, the striking results the team achieved on Atari. Thirdly, the way the paper was written seemed to appeal to people in terms of the mix of inspiration and theoretical and empirical insight.

"*I did everything I could to make sure the Atari results were a fair comparison and ablated all our choices,*" he adds.

"*It stands out that you can double your performance by tuning these hyperparameters online if you do it correctly.*

The key insight is how you deal with optimization and, in this case, meta-optimization. If you solve the dynamics of the optimization problem, you can get a long way. ***80% of the problem is probably the data, but the rest is just solving the optimization issues****."*

In terms of the next steps for this work, even though intuitively it makes sense, Sebastian says there is no theoretical reason for why and when this type of algorithm works, so a future task would be to gain more understanding of that. He would also like to use this work more directly for few-shot learning.

"We wrote this paper from a meta-learning point of view, but one fascinating experiment we ran shows that the algorithm can be used without meta-learning," he reveals.

"You can use the algorithm for zeroth-order optimization and even for optimizing non-differentiable parameters. I think that's a cool direction to go in as well."

Outside of this paper, Sebastian's work is focused on meta-learning research. Together with his peers at DeepMind, he is looking at ways to take extensive systems that train for a long time and give them more autonomy over their learning dynamics to make them less dependent on careful human engineering.

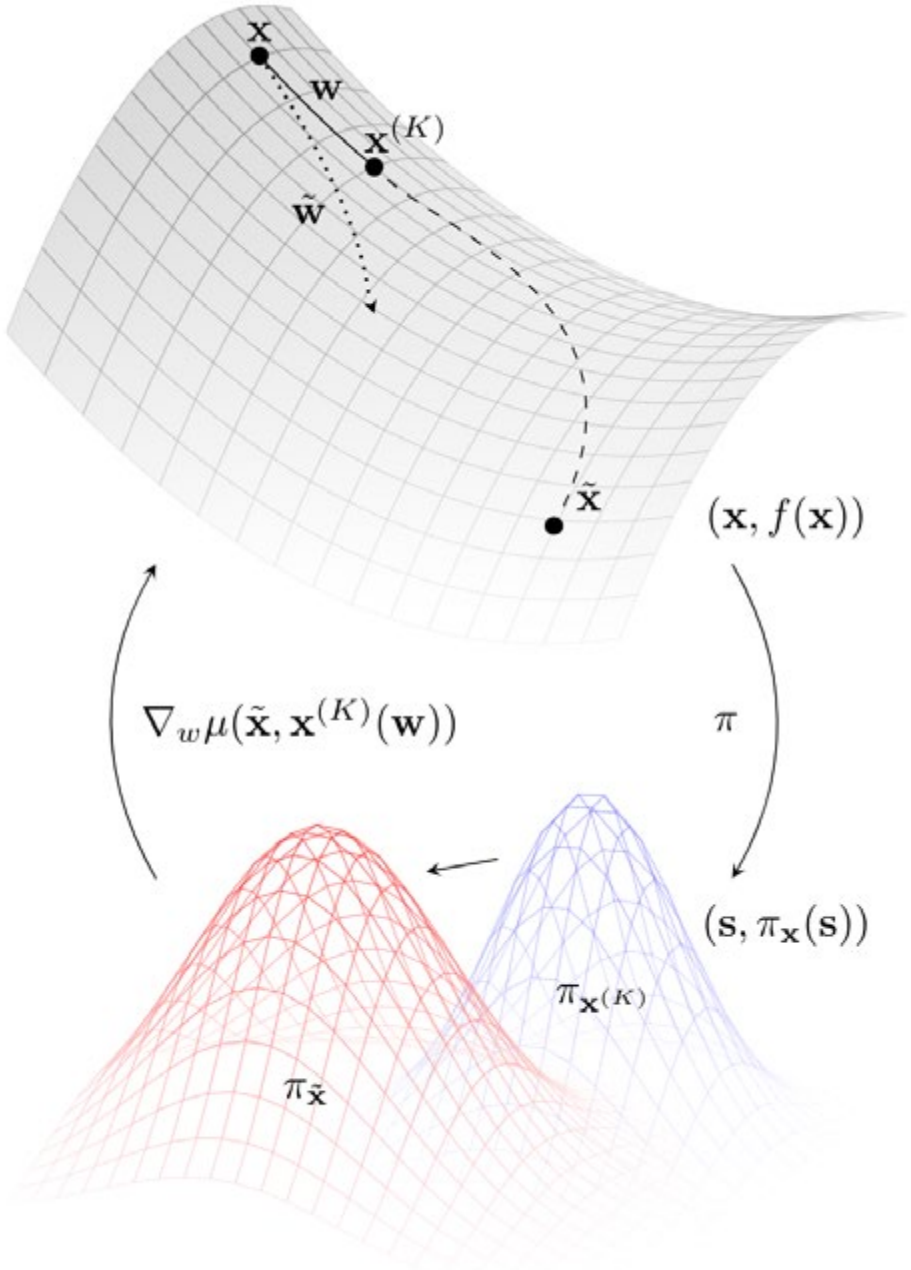

"I like to look at very simplified settings and try to understand the issue we're seeing in a toy environment and gain insights from that," he explains.

"Once you come up with a general algorithm, you can go back to the big systems to see if it scales."

Can he tell us something about **DeepMind** that we don't know?

He laughs before adding, finally: *"Well, if it's something you don't know, you're probably not supposed to know!"*

LIBERTY DEFENSE

Jeffrey Gordon is the Vice President of Engineering at Liberty Defense, overseeing the development of its breakthrough security detection products. Daniel Pineo is a Chief Engineer and AI Architect and has been working with Jeff for six months, focusing primarily on AI, but connecting with other areas of software and hardware. They are speaking to us about developing an innovative new personnel screening product.

Jeff and Dan are veterans of the **Homeland Security field**, involved in developing the first airport body scanner used across the US and throughout most of Europe. Their latest project is turning their attention outside of the airport to venues and events where screening infrastructure is still based around **decades-old metal security detectors**.

With the range of possible threats expanding, relying on a traditional metal-only approach is no longer sufficient. Joe Biden recently announced **a crackdown on ghost guns, which can be self-assembled, made of plastic, and even 3D printed. There are also ceramic knives and the proliferation of explosives, flares, and smoke bombs**.

"Detecting the spectrum of threats is one thing, but it also has to be done with very, ***very high efficiency and high throughput****,"* Jeff points out.

"Nuclear power plants, for example, have thousands of people going in on a shift within two or three hours, requiring something beyond a body scanner, where people stop and get screened, or a metal detector, which has a much lower throughput of people and objects. We've got to create ***a solution that operates much faster and has the intelligence to allow people***

***to pass through with permissible objects**, such as keys or a cell phone, without causing false alarms and without needing them to remove clothing items every time. That's where AI is going to be key."*

Security cannot be compromised to achieve this high throughput; it should be enhanced over current products. The process should be almost invisible to people, non-intrusive, and protect their privacy. To be successful and deployable, it must use a low-cost computational platform, as venues like schools and stadiums do not have the money to pay for costly solutions.

Jeff Gordon

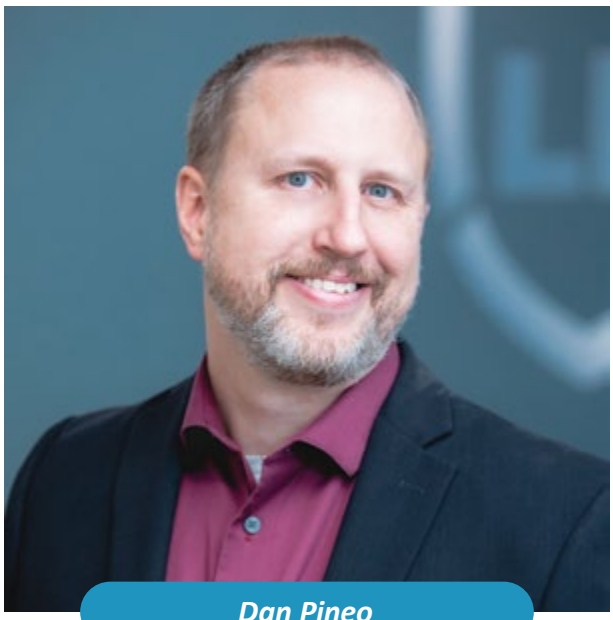

Dan Pineo

The team's new system uses **millimeter-wave imaging**, which can penetrate clothing but not skin. It supplies a complete data set at speed, with the AI making swift decisions.

*"We're creating whole datasets, with hundreds of thousands of data points available for reconstruction at 10 to 20 times a second – **eventually 100 times a second**,"* Jeff tells us.

"Think of a fast-frame camera. When someone walks by, it's as if they're frozen, but they're moving. You must go through image processing, reconstruction, and AI decision-making 20 times a second. It's a pipeline that's very demanding."

Dan adds:

*"To some degree, **it restricts the types of algorithms we can apply**. We need something that can finish – if we're doing 30 hertz, you've got about 30 milliseconds to be done, and that's the whole pipeline. Unlike a video camera where you get an image already formed, we have a processing step called holographic reconstruction. If you're familiar with CT, it's like tomographic reconstruction, but it's essentially reconstructing from the various wavefronts captured by the sensor."*

Traditionally, reconstruction comes first, then AI, but here, **combining the reconstruction process end to end with the AI process into one optimization** was the only way to achieve the speed and performance enhancements required.

The team looked to the autonomous driving community for guidance, where several **single-shot detector algorithms have been developed**, sacrificing some detection performance but being designed to be very fast.

"That ends up being a big win for us," Dan says.

"Unlike the autonomous driving problem where you're dealing with reconstructed 2D images, the result of our reconstruction is a volume. We use maximum intensity projection to get a view from that, and then we use several different viewpoints of the volume."

The holographic volume is collapsed into a 2D frame and put through a single-shot

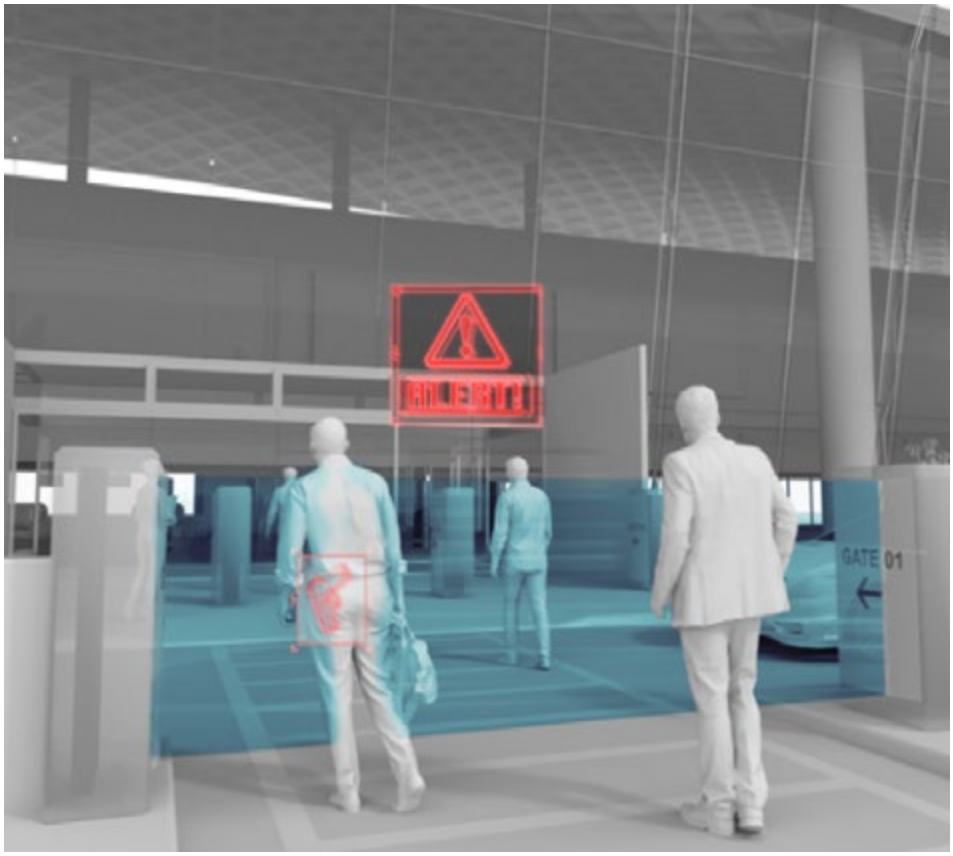

projector, but that alone is not enough, as threats can be occluded at many angles arbitrarily on the body. **3D convolutions** are one way of dealing with this, but the team chose a 2D approach with a specific number of views.

In autonomous driving, there is an endless range of unpredictable events, but the possibilities are more constrained in this situation. It has a specific **CONOPS, or concept of operations**, which means people have to use the system in a particular way and should not need to worry about unknowns. They know a person is going in, and they are looking for a threat on that person. For this, they need to know the space of threats – what weapons they might be carrying and where they may be hidden.

*"The biggest bottleneck we have with CONOPS is the **permutations of a training set**,"* Jeff explains.

"We must put in a representation of all the shapes and sizes of people going through and the range of threats to train the AI what normal is."

Every customer has a different concept of a threat and how different objects should be treated. The system uses **targeted or switchable algorithms** to solve this, but the challenge is creating permutations of scans and a training and scoring set. With **supervised learning**, every scan has to be marked and tagged, and thousands and thousands of scans require labeling to create a broad enough set to cover every eventuality.

*"Try finding **a labeling tool that can label volumetric videos** – I've looked around, and I couldn't,"* Dan reveals.

*"We effectively have a volume that is a video. We're collecting data at 30 frames per second when somebody walks through our machine. **That's terabytes every day to label!** If that's done naively, just on images, you're talking about labeling orders of magnitude slower than you're collecting.*

Part of the challenge here is creating a new workflow for efficient labeling that can overcome the problem of this firehose of data."

The AI community is used to visible-wave (RGB) photographs and the performance you get with algorithms on that type of data. However, **millimeter-wave reconstruction data behaves differently**. You run the image through similar algorithms, but there are some reconstruction artifacts because of the process.

"*Algorithms trained on this millimeter-wave data are very good,*" Dan explains.

"*They're way better than people! It was a big deal when visible-wave algorithms started outperforming people, but it took a lot of work to get to that point. I've noticed millimeter-wave algorithms are seeing things that I have a hard time seeing in the data, which causes another problem because labeling when you can't see something yourself, but you don't know for certain that the signal is not there, that's tricky! The ground truth is where the object was placed when the scan was performed. Whether or not a person can see that in the data is not as important as the fact that it is there.* ***The algorithm will see something that people don't see!***"

Jeff agrees:

"*These technologies have already surpassed humans as far as the ability to detect on a case-by-case basis, but we are now trying to get to levels that would be almost invisible in the image to a human.*"

Dan is currently working on several innovations, including **incorporating the time flow dimension into 4D labeling**, and

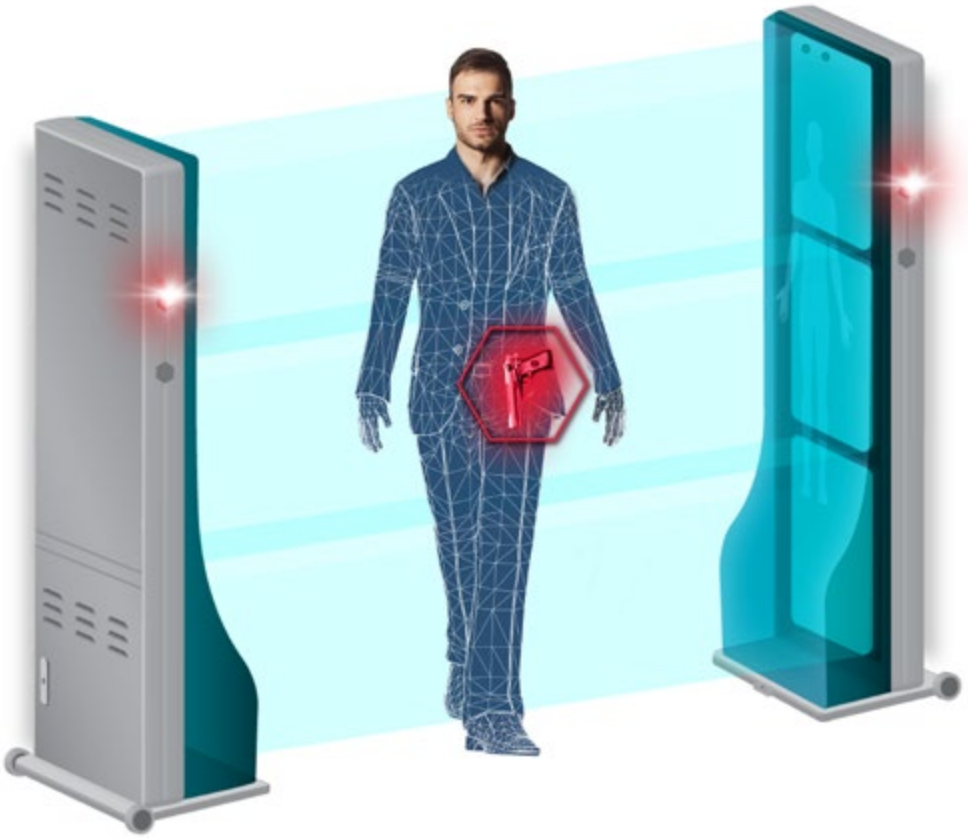

using spatial reinforcement to correlate images between different machines to combine as much data as possible and make the hard-to-see objects more apparent.

"*We're very excited about where we can go with detection and allowing, in an intelligent way, people through where they have permissible objects,*" Jeff adds, finally.

"*At the end of the day, we want to make this a frictionless process.*"

Congrats, Doctor!

Kaichun Mo has recently completed his Ph.D. at Stanford University, advised by Prof. Leonidas Guibas. Before that, he received his BS.E. degree from the ACM Honored Class at Shanghai Jiao Tong University. He has interned at Adobe Research, Autodesk Research (AI Lab), and Facebook AI Research. His research interests focus on learning visual representations of 3D data for various applications in 3D vision, graphics, and robotics. Congrats, Doctor Kaichun!

Humans accomplish everyday tasks by perceiving, understanding, and interacting with a wide range of 3D objects, with diverse geometry, rich semantics, and complicated structures. One fundamental goal of computational visual perception is to equip intelligent agents with similar capabilities. My Ph.D. research is motivated by exploring answers to one central research question -- what are good visual representations of 3D shapes for diverse downstream tasks and how do we develop general frameworks to learn them at scale? In my Ph.D., I explored along with two directions to tackle the huge complexity of 3D data and tasks. First, I worked on developing compositional approaches that smaller, simpler, and reusable subcomponents of 3D geometry, such as the parts of an object, are discovered and leveraged toward reducing the complexity of the 3D data. Next, I investigated learning visual actionable representations over 3D shapes for robotic manipulation applications, where large-scale self-supervised learning frameworks using simulated interaction are investigated and proposed for learning task-specific 3D shape semantics.

Figure 1

Learning Compositional Visual Representations for 3D shapes.

In PartNet [**1**] (Figure 1, left), we introduced a large-scale dataset providing fine-grained part annotations over ShapeNet [**2**] models and set up several 3D part segmentation benchmarks. We proposed 3D deep learning methods that are able to segment 3D shape inputs into semantic, instance-level, or hierarchical part instances. While the PartNet work is concerned with the tasks of 3D shape part segmentation, StructureNet [**3**] (Figure 1, right) focuses on investigating the inverse problem of synthesizing novel 3D shapes by composing 3D parts. We proposed a part-based and structure-aware 3D shape generative model that not only generates high-fidelity 3D part geometry assembled into a 3D shape but also captures the rich relationships and structural constraints among the 3D parts. We refer to the papers for the technical methods.

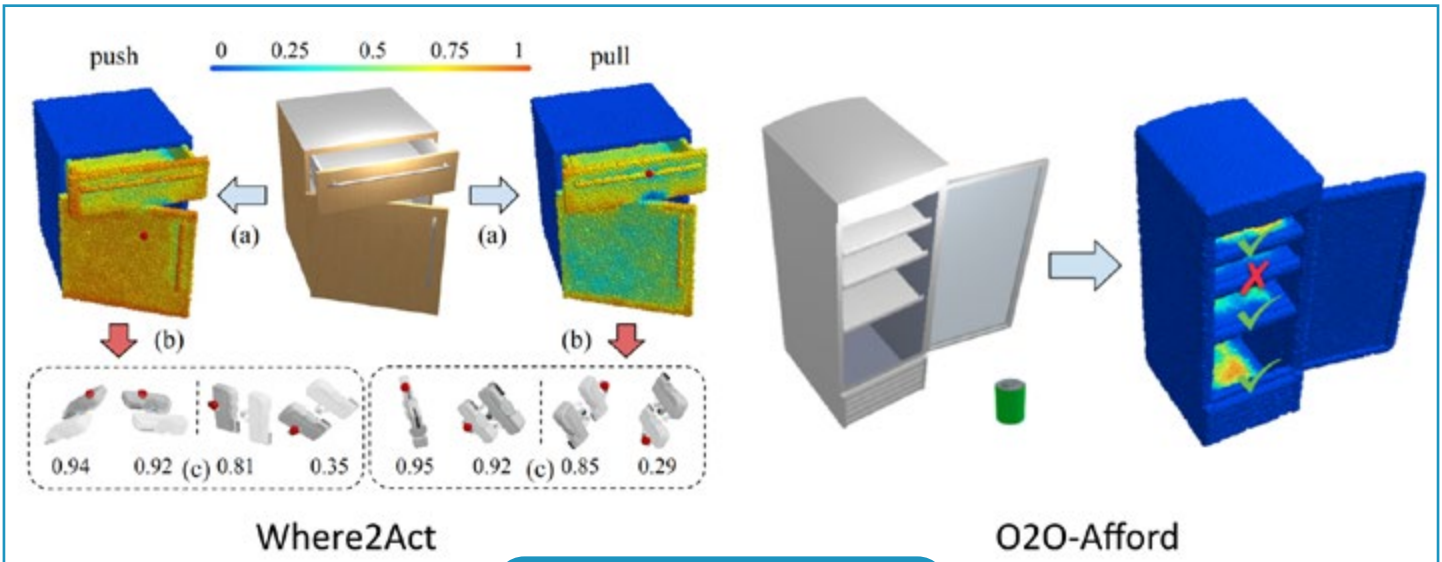

Figure 2

Learning Actionable Visual Representations for 3D shapes.

In Where2Act [**4**] (Figure 2, left), we proposed a general and self-supervised framework for learning actionable and task-specific visual representation for manipulating 3D articulated objects. Our method leverages scalable and inexpensive simulated interaction data collected from a physical simulator SAPIEN [**5**] to automate the robot-object affordance learning for various manipulation tasks, e.g., estimating where to push/pull a drawer/door on the cabinet. O2O-Afford [**6**] (Figure 2, right) extends the system to handle object-object interaction scenarios, such as fitting a bucket inside a cabinet. The critical challenge is to unify task specifications for a diverse set of downstream tasks, including fitting, placement, and stacking, so that the proposed method is generally applicable. Please check the papers for more details.

Ewa Nowara is a research scientist at Meta. She has a PhD from the Rice Computational Imaging Group, where she worked with Ashok Veeraraghavan.

More than 100 inspiring stories of Women in Computer Vision here!

Ewa, where are you from?

I am Polish. Many years ago, maybe 700 years ago, my family supposedly originated from Italy so the last name is Italian, but it is with a Polish spelling. We have many generations just from Poland.

It rings a bell with me because I am Italian. Let's start with your current work. Can you tell us what you do?

I just started at Meta about three months ago. I don't have a lot to share yet. I'm still learning as I go. I am working in an applied research team. We do research and come up with new ideas for enhancing the AR and VR experience. We also are interested in putting our ideas into products. It's an interesting, exciting, unique position. It's in between where you get to do innovative creative research and publish papers, but at the same time, you can see your work realized in products. My position builds a lot on my past research in computer vision about scene understanding, extracting information from images and from video. Prior to coming to Meta, I was a postdoc at Johns Hopkins University where I worked with Rama Chellappa on geolocalization. Given a single RGB image taken anywhere, maybe from your window in your house, we wanted to determine where that

image was taken and be as specific as what country, what city, or even what street the image was taken. Prior to that, I worked in a different area of computer vision for my PhD at Rice with Professor Ashok Veeraraghavan, where I worked on camera-based physiology. Given a video recording of someone's face, my interest was to extract very subtle color intensity variations in the skin related to their blood flow, from which we can measure heart rate, breathing rate, and other vital signs in clinical parameters. It's essentially how a smart watch works, like a Fitbit or an Apple Watch. They have this green light that illuminates the skin. Some light penetrates the skin. Some light returns to the sensor. The light that has interacted with the blood will have this pulsatile temporal component. It's a very tiny change in the intensity of the skin, but from that we can measure the pulsatile signal, from which we can extract heart rate and other parameters. I was working on doing it from a video recording without touching the skin, which is very challenging: lots of noise, motion, illumination variations. I have spent several years doing that and that has been the main focus of my research.

I am sure that we could find continuity between the different things that you just told me. One was innovation. The second was to put things into practice. It seems these two things matter the most to you. Is that right?

That's right.

You added a third thing. That it was challenging. Is the challenge also important to you?

It's definitely important. I've always been very inclined to work on things that are exciting but challenging. We don't know if they are possible or not because it's a very difficult problem. It hasn't been done before, but we have some basis to believe that it might be possible. That has been very

Yes, that was my dream job!

For a very long time, I wanted to be where I am now.

exciting. Of course, frustrating and stressful at times as well, when you have to graduate as a student. Even as a professional, you need to get some results, but it's very exciting that you can do something for the first time. Whatever you're working on is actually useful for a product or has some use hopefully for some positive application to improve people's quality of life or enjoyment of life. That has always been very important to me and has guided how I choose my project, the postdoc, and my full-time job. That has been the underlying theme, looking for work that is challenging, and innovative, where I can grow and also have an impact on products.

The more you choose challenging tasks, the higher the chance to fail. Didn't that scare you?

Definitely, yes, but I think I was very lucky that this was a part of my job description to do research. We know that when we work on something new, that is challenging. A part of research is that oftentimes it doesn't work. In many of the projects, either it did not work or the method I was pursuing did not work. We had to switch gears and make it work in a different way. I suppose this sort of fear or worry was always there. It was a part of the journey, and it was kind of expected that if things don't work, that is a part of recent research. I wasn't betting on a particular project.

In the beginning, did you imagine yourself where you are now?

Yes, that was my dream job! I have always

wanted to be a research scientist. For a very long time, I wanted to be where I am now. I wasn't sure if I would make it because these positions can be very competitive. Also, there aren't as many of them because there's only so much need for research in a company. You need to have strong publications. I'm very happy that it did work out.

You had doubts along the way. How did you deal with these doubts? What guided you through the moments of doubt?

I was reaching out to several mentors, people who I have interned with, my PhD advisor, several other professors, and also peers who have gone through a similar process. It has been really reassuring. They told me that sometimes it's more about the timing. It's a bad time right now to apply for these positions. They recommended a postdoc and trying to work on something new, get some more publications.

So you are living proof that mentoring is really fruitful. It helps a lot. What was the main lesson you learned from a mentor that was particularly successful?

My PhD advisor gave me many pieces of advice along the way. One stuck with me. He said that I should first go for my top positions, so interview at the top places. Because if I interview at my second and third choice places, I will likely get an offer. Then what do you do? You don't have any other offer, right, so you will probably take it... That happened to me. I got some offers, which were pretty good, but it definitely wasn't my dream job. That advice really stuck with me. I had a pretty good offer from a company. It was a research engineering position. I wasn't super excited about the company or the position, but it was pretty good, pretty good compensation. I didn't have anything else.

So you did not follow the advice?

I interviewed where I could. I didn't want to turn down the interview. It seemed like a good option. Maybe not the top option - I wasn't sure what I could get, so I guess I did not exactly follow [*laughs*], but I wasn't able to land interviews at my top places at the time in 2020. Many places paused hiring. Then Professor Chellappa reached out to my advisor asking if he has any students who are interested in a postdoc. Ashok recommended that I talk to him. I did get the postdoc offer. I had a very challenging decision to make. I have a full-time job offer. I'm on a visa, so I do need some sponsorship eventually. That was a stressful

"If you don't like it, then just quit and give up!"

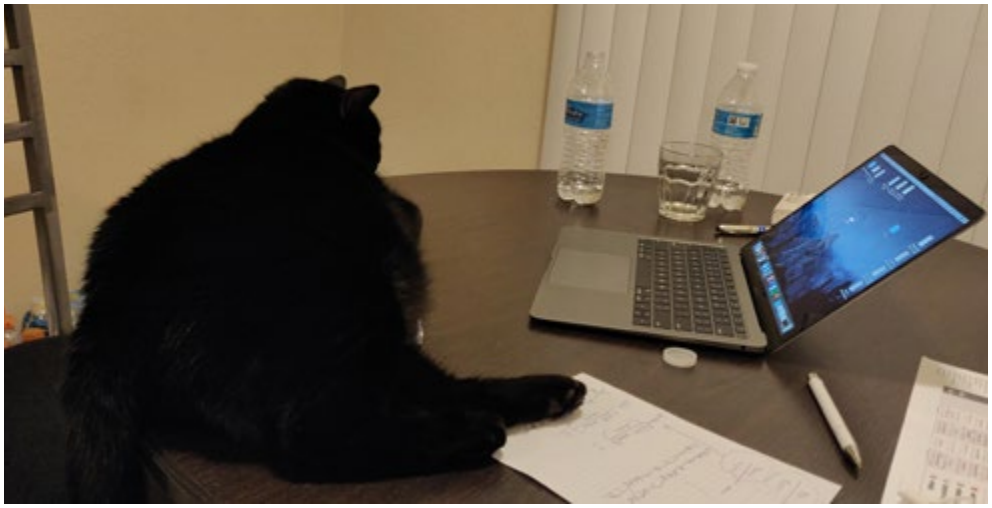

Don’t give up and keep trying!

thing to decide between. Then potentially this opportunity to take this postdoc that might open doors for my dream job, but maybe not, who knows. Maybe I won’t get another offer like that. It took a lot of deciding. I sacrificed some things to be here. I left my family back in Poland. I’m here by myself without family. I spent all of these years going to college here and grad school, trying to do something worthwhile. I shouldn’t settle right now. I don’t have a reason to settle. I should just take a chance and see what happens. I wasn’t sure if that was the right decision because then I have to start all the interviews again. I started the stress of interviewing in this uncertainty, but that also opened some possibilities.

How do you see yourself 20 years from now?

I hope to see myself in a similar position where I can work for a fast-moving company, where I have the resources to succeed and a great team to work with like I do here, and where I can choose what projects are important and still keep working on those. I hope I’ll have the same opportunity to contribute to industrial research.

Did you work with Yann LeCun, yet?

No, I haven’t.

Can you tell me one thing about yourself which is important and that we did not discover yet?

Maybe I’ll just briefly talk about where I came from, or like you know how I ended up where I am. Nobody in my family is doing this sort of technical research in

industry or anything related to computer vision. My mother is a very strong figure in my life. She is a physician, an oncologist, which is a very challenging specialization in medicine, unfortunately, because of the low success rate of treatment. She pursued a PhD as well. So she is an MD PhD. She was a professor for a short amount of time, but mainly focused on her clinical practice. She has always been this very strong figure who has been an inspiration as a mother, and a female, in a male-dominated workplace in medicine, like science, in Poland and the same as here. You can succeed and be a good mom and a good wife and have a normal life. I never felt neglected. I felt like we always spent a lot of time together on vacations or playing games. She helped me with my homework. But, she was definitely aggressive with her career choices. She worked very hard and also took a lot of chances. She always had this attitude that you shouldn't complain. Just work hard. You'll be fine. That has always been very motivating. If she succeeded, I can succeed.

That's very Eastern European to work hard and not complain. In Western Europe, we complain. A lot! [*laughs*]

We're definitely not allowed to complain. When my mom says *"If you don't like it, then just quit and give up"*, she says it with a sarcastic voice! *"Of course, it's hard, otherwise everyone would be doing it. You'll be fine! Stop complaining!"*

Do you have a message that you want to share with the community?

You can accomplish anything you want to accomplish. It might take a long time, and maybe you change your goals or priorities along the way, and that's okay. You'll learn during the process. I think it's important to not give up and just keep trying. Sometimes, it doesn't work out the first time or third time or 10th time. Maybe sometimes it's time to move and try something different. At the same time, you don't need to just work really, really hard and exhaust yourself. It's important to take care of your mental and physical health along the way. Look at it like a marathon instead of a sprint. Take breaks and take care of yourself. Reassess your goals and what drives you, but keep that long-term goal in mind. Don't give up and keep trying!

You'll be fine! Stop complaining!

More than 100 inspiring stories of Women in Science here!

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

The Road to Success is Paved with Rejection Letters

Who said that? One of the most successful people I know, Michael Black. He just wrote "**My other CV - The CV of failures**" to tell us that even successful people go through rejection and how to deal with it. In his own words: "*The road to academic success is full of barriers and when you hit them, they can feel insurmountable. Maybe my stories give a glimmer of hope that helps someone continue through failure. The academic life is worth it.*" And reading Michael's other CV is definitely worth it too! Even **Johannes Kepler** was found not clever enough to be offered a faculty position. And he did just fine all the same. Read More

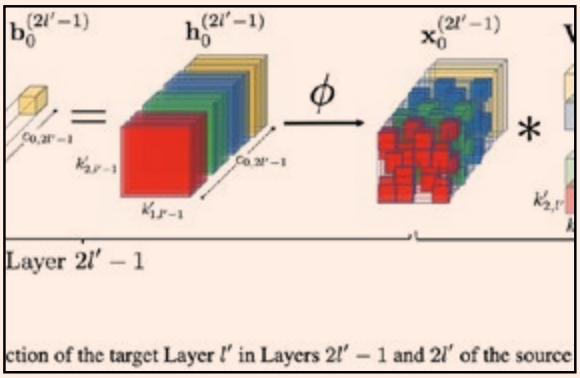

Setting AIs on SIGGRAPH: Top Academic Researchers Collaborate with NVIDIA to Tackle Graphics' Greatest Challenges

NVIDIA co-authors are presenting a record 16 research papers at the year's SIGGRAPH, pushing forward the fields of neural rendering, 3D simulation, holography and more. This nice post on NVIDIA's blog tells us all about them. I was curious to see what's the story at this month's CVPR: and I found out that Sanja Fidler coauthors 8 papers and Jan Kautz 6 other papers! A great choice for CVPR participants. Of course what counts is quality, not quantity. So do not miss all these poster and oral presentations. We probably won't meet at SIGGRAPH in Vancouver, but we'll obviously be at CVPR 2022 in New Orleans! Read More

Davis Summarizes Machine Learning Papers

Every week, there are about **600 new machine learning papers submitted to arXiv**. This gentleman called **Davis Blalock** goes through all these submissions and identifies 10-20 that he thinks are especially interesting, practical, or promising. He then writes a summary of each one, often with some commentary. He admits being biased in favor of papers that improve understanding of **deep learning**, or **show empirical gains on computer vision** or NLP tasks, so there is plenty of great stuff for all the readers of **Computer Vision News**. You can subscribe to his newsletter in just one click! Read More

Artificial Intelligence/Machine Learning Reduces Mistakes and Waste

Apparently, **AI is becoming able to make quality decision at the production chain**: camera systems and deep learning modify process and ingredients to adapt to changing conditions. In this way, they can ensure that the final product is consistent with consumer expectations. The real advantage of this versus an expert technician is that it can be done continuously. New technologies in **vision systems** are on their way to see abnormities that the naked eye cannot see. This very nice article includes also **real-world examples**, like adding decision-support system for a labelling process in a distillery. Read More

Revealed: Intel's Plan to Sell Software Simplifying Computer Vision Training

The folks at **El Reg** claim that they have seen internal documents at **Intel**, revealing that the chipmaker plans to sell a software platform that promises to simplify and speed up the training of **AI models for computer vision**. The name of this computer vision development kit is **Sonoma Creek** and it is supposed to work with **NVIDIA graphics** along with **Intel processors**. Apparently, Intel says that Sonoma Creek brings together five important steps of AI model training across multiple services in a "single workflow": **data collection, data labeling, model selection and training, model optimization, and deployment.** Read More

Explainable AI Can Assist Internet of Behavior Efforts

We all know about **IoT, the Internet of Things**. Many of us probably don't know about **IoB, the Internet of Behaviors**. It is defined as the confluence of IoT devices combined with advanced sensors, **computer vision**, facial recognition, other biometric indicators or location tracking, which, when added to insights from human psychology and data analysis, can **influence, prompt or change human behavior**. It seems that behavioral science and a system based on IoB can achieve user and business benefits, like reducing electrical power consumption or even saving lives during a global pandemic. Read More

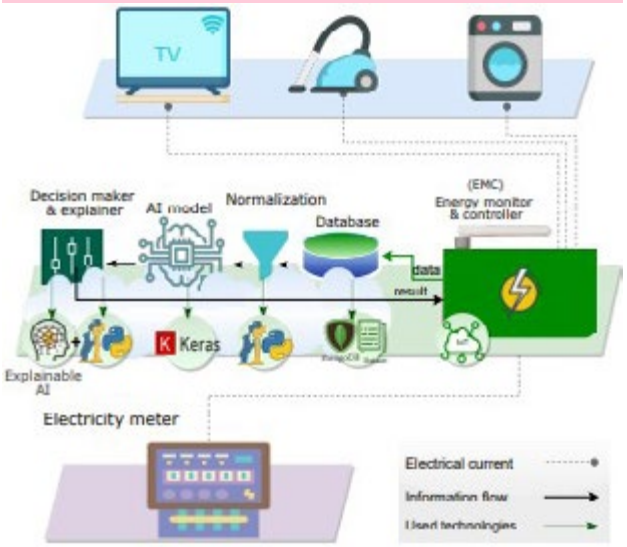

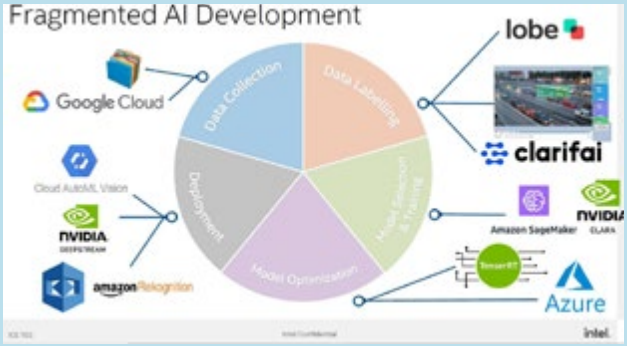

COMPUTER VISION EVENTS

CARS 2022

Tokyo, Japan
7–11 June

CAOS 2022

Brest, France
8–11 June

CVPR 2022

MEET US THERE

New Orleans, LA
19–24 June

Deep Learning World

Las Vegas, NV
19-24 June

BVM

Heidelberg, Germany
26-28 June

Hamlyn Symp. on Medical Robotics

MEET US THERE

London, UK, 26-29 June

MIDL 2022

Zürich, Switzerland
July 6-8

WBIR2022

Munich, Germany
10-12 July

Int. Conf. on Machine Learning & Data Mining

New York City, NY
16-21 July

ICML 2022
Machine Learning

Baltimore, MD
July 17-23

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

JUNE 2022

CRAS 2022 - CONFERENCE ON NEW TECHNOLOGIES FOR COMPUTER AND ROBOT ASSISTED SURGERY

Fanny Ficuciello is an Assistant Professor at the University of Naples Federico II on the tenure track, shortly set to become an Associate Professor. She was also one of the co-organizers of CRAS 2022, the annual conference on new technologies for computer and robot-assisted surgery, which took place in April. Fanny is here to give us the lowdown on this year's main event.

In its 11th edition, CRAS has graduated from its origins as a workshop to become **a fully-fledged single-track conference**. Being focused on surgical robotics fills a space that more generalized single-track events, like ISER, the International Symposium on Experimental Robotics, cannot.

Single-track conferences mean people spend a great deal of time together, with plenty of opportunities to talk and get to know each other. They also get to network with famous professors, like **Paolo Fiorini** and other keynote speakers in a more intimate setting than can be afforded by large multi-track events.

CRAS has a typical attendance of around 70 people, so with **over 100 registrations**, this year was one of its most significant editions yet.

"*I hope that was down to the good organization!*" Fanny laughs.

"*The conference was postponed several times due to the pandemic, but we advertised widely online and through social media. For this kind of event, the way you communicate, prepare, and present it is all essential for its success.*"

Might the stunning location and bounce back from the pandemic have also played a part?

"*Yes, Naples is an amazing city, and people were keen to start traveling and meeting each other again,*" she agrees.

"*Many young researchers took part with great interest. They were all participating and asking questions. If you weren't there, you missed a very friendly environment. We had so much fun staying all together, sharing knowledge, and it's* ***an occasion where professional partnerships are forged*** *because you can talk to everyone.*"

Fanny and her colleagues all agree CRAS 2022 was a big success, but she admits it did not go off without a hitch.

"*The first day was Liberation Day in Italy, so I had my small child with me,*" she recalls.

"*In the afternoon, when I thought everything was under control, I dipped out with her for two hours, and everything that could go wrong happened at that time! We had a power blackout, so we lost electricity and that caused some delays. I wish it were a funny story, but it all worked out in the end at least!*"

As in previous years, **the industry played a vital role in proceedings**, showcasing new technologies, products, and prototypes. Some were invited, but Fanny also received interest from enterprises wanting to participate.

"*It was bilateral communication between us because the industry needs the scientific community as much as we need them,*" she asserts.

"*In the surgical domain, it's crucial to have this interaction between academics and industry so that research is useful and can eventually end up on the market and help patients.*"

Fanny reveals that one of the most important things she has learned from CRAS is that **surgical robotics and other robotics applications require input from different disciplines**.

"*Robotics, in general, is a multidisciplinary science,*" she points out.

"*Even more so in the surgical field because you need this tight cooperation between engineers and doctors. It's essential to create this synergy.*"

Watch in the following pages our review of the Best Paper Award and the Best Demo Award at CRAS 2022.

Awesome **Sophia Bano** *speaking at CRAS 2022*

DUAL-ARM PLATFORM FOR CONTROL OF MAGNETICALLY ACTUATED SOFT ROBOTS

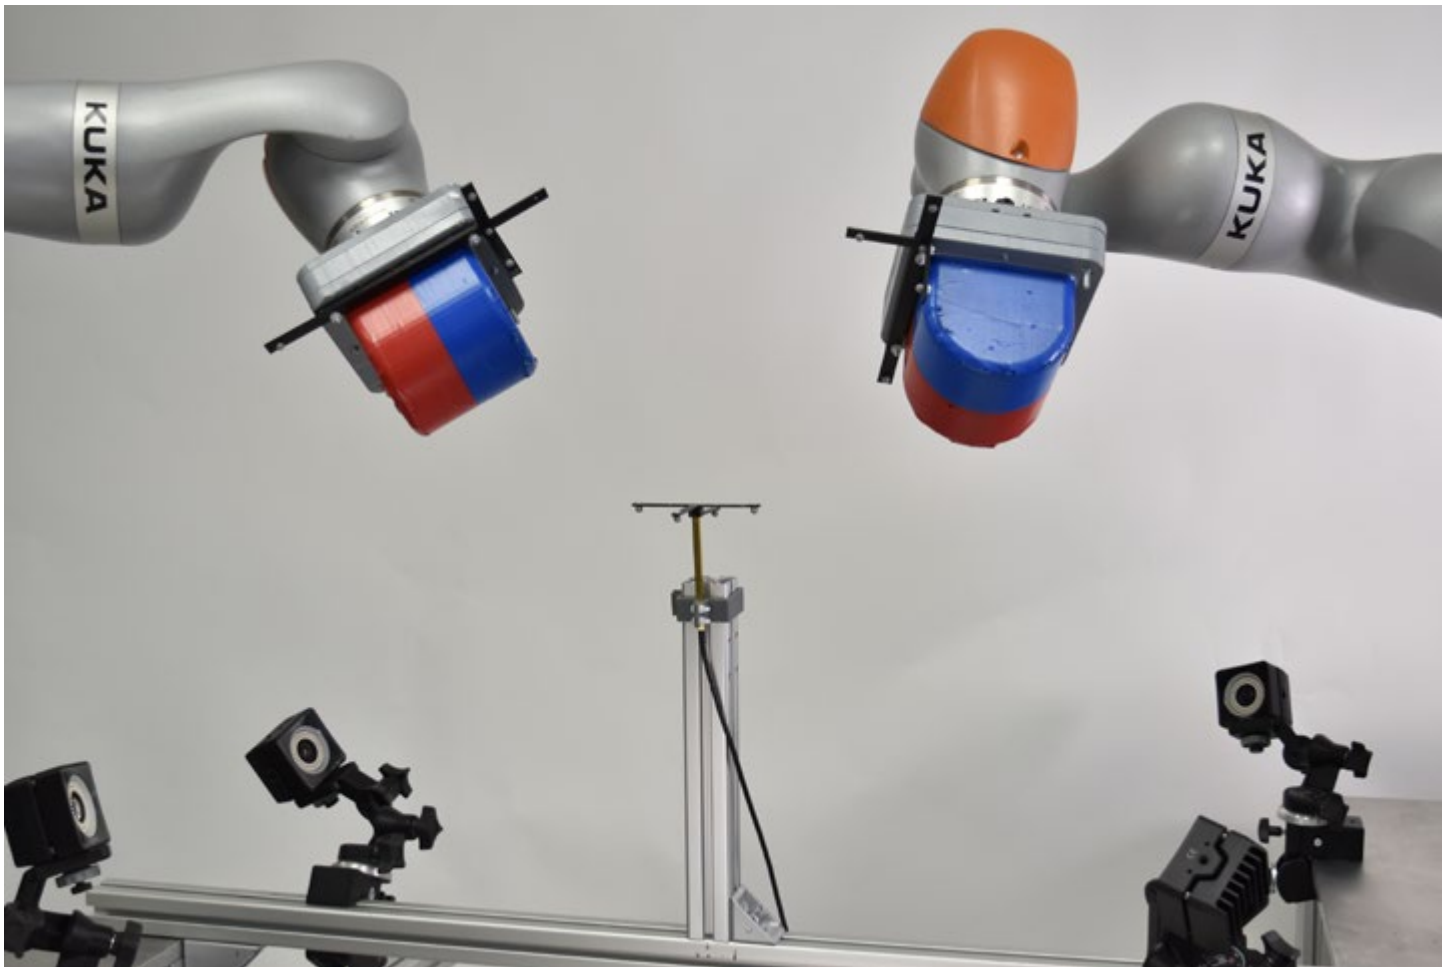

Michael Brockdorff is a second-year PhD student at the University of Leeds. Giovanni Pittiglio is a research fellow at the Harvard Medical School and Boston Children's Hospital. Their paper on a dual-arm platform for control of magnetically actuated soft robots has just won the Best Paper Award at CRAS 2022, and they are here to tell us all about it.

Devices used in bronchoscopy are generally rigid and thick, limiting how deep clinicians can get into the lungs and potentially causing damage and discomfort to patients as they bend the tissue around them. Clinicians use a **pre-computed anatomy map**, often from a CT scan, which is less helpful once they insert the rigid tool and the anatomy moves.

In this work, Michael and Giovanni propose **using soft silicone-based magnetic catheters inside the body controlled externally by two permanent magnets on robotic arms**, allowing for precise control of the magnetic field in the workspace. This contactless form of actuation means devices can be smaller and reach areas of the lungs that are impossible to reach with classic tools. The soft devices will **bend and conform to the body during navigation**, so they will not deform the anatomy, removing any pain or discomfort for the patient and leading to a better-targeted intervention.

"This unique platform combines a large workspace – the human body – with high magnetic fields and a high degree of manipulability," Giovanni tells us.

*"When you actuate the magnetic field within a large workspace, the main difficulty is actuating strong magnetic fields with a high level of controllability. With one permanent magnet, you can control five degrees of freedom. We add **three more degrees of freedom**, up to eight degrees, with two magnets. That is a key part of this platform compared to smaller actuation systems."*

Professor **Pietro Valdastri** from the **University of Leeds** had the initial idea of actuating these devices to conform to

Michael Brockdorff

Giovanni Pittiglio

the anatomy with multiple magnets along the length of the catheter. But the team discovered that if this was not matched with the ability to control the magnetic field in a more convoluted way, the degree of shaping does not improve. The team met with Pietro's colleagues from the **University of Toronto**, who were working on the ability to control the magnetic field in a more convoluted way with eight degrees of freedom, but within a relatively small magnetic workspace, which did not apply to their case.

How did they solve it?

"*Two years of hard work!*" Giovanni laughs.

"*I thought, what if we take that idea of being able to control these degrees of freedom and we scale it up.* ***The way to scale it up usually is to use permanent magnets with systems of coils****, but the energy consumption and need for cooling make that very costly. We started investigating mathematically and theoretically – what can one, two, three, or more permanent magnets do? Luckily, we stopped at two because I don't think Pietro would have paid for ten robots!*"

Michael adds:

"*We're exploring the possibility of using more robots in the future, but we have already maximized how much we can control the magnetic field in our workspace with this platform. More robots might allow the development of new devices that we haven't even thought of yet, but we're focusing on our two-robot system for now.*"

Being picked for a **Best Paper Award** is a huge achievement and no mean feat, but Michael has an idea of what might have influenced the judges' decision.

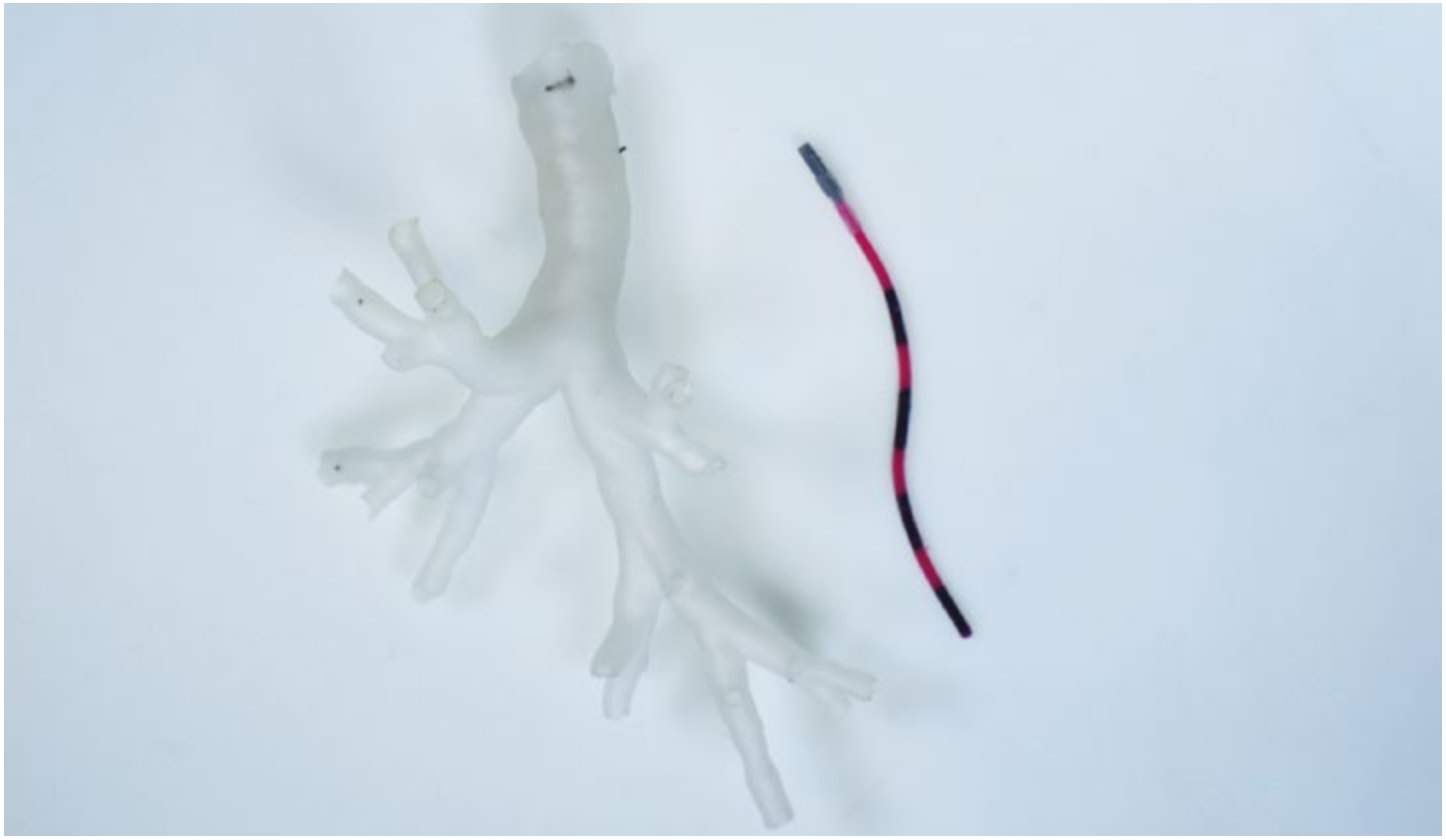

"*I like to think it's the* ***novelty and diversity of our platform,***" he tells us.

"*Compared to most actuation systems, the bigger workspace we offer with our dual-magnet approach makes it easier to step into the medical field but is versatile enough to be used for a whole range of magnetic devices, including in the medical domain and beyond.*"

Looking to the future, Giovanni thinks it is essential they focus on the clinical need.

"*We've done a lot of technical work in the last three years with this project, so* ***having a strong clinical application sooner rather than later****, before moving on technically, could be a good step forward.*"

For more information on the team's work, you can read their recently published paper in the Soft Robotics journal.

Giovanni and Michael with co-author Tomas da Veiga

MULTI-LEVEL-ASSISTANCE ROBOTIC PLATFORM FOR NAVIGATION IN THE URINARY SYSTEM

Martina Finocchiaro and Sanat Ramesh are Marie Curie PhD students and part of the European project ATLAS (or 'AuTonomous intraLuminAl Surgery'). Martina is enrolled in two universities: the Technical University of Catalonia in Barcelona and the Sant'Anna School of Advanced Studies in Pisa. Sanat is jointly supervised by Professor Paolo Fiorini from the University of Verona and Professor Nicolas Padoy from the University of Strasbourg in France. We speak to them fresh from winning the Best Demo Award at CRAS 2022 in Italy for their steerable robotic platform for robot-assisted intraluminal procedures in the urinary tract.

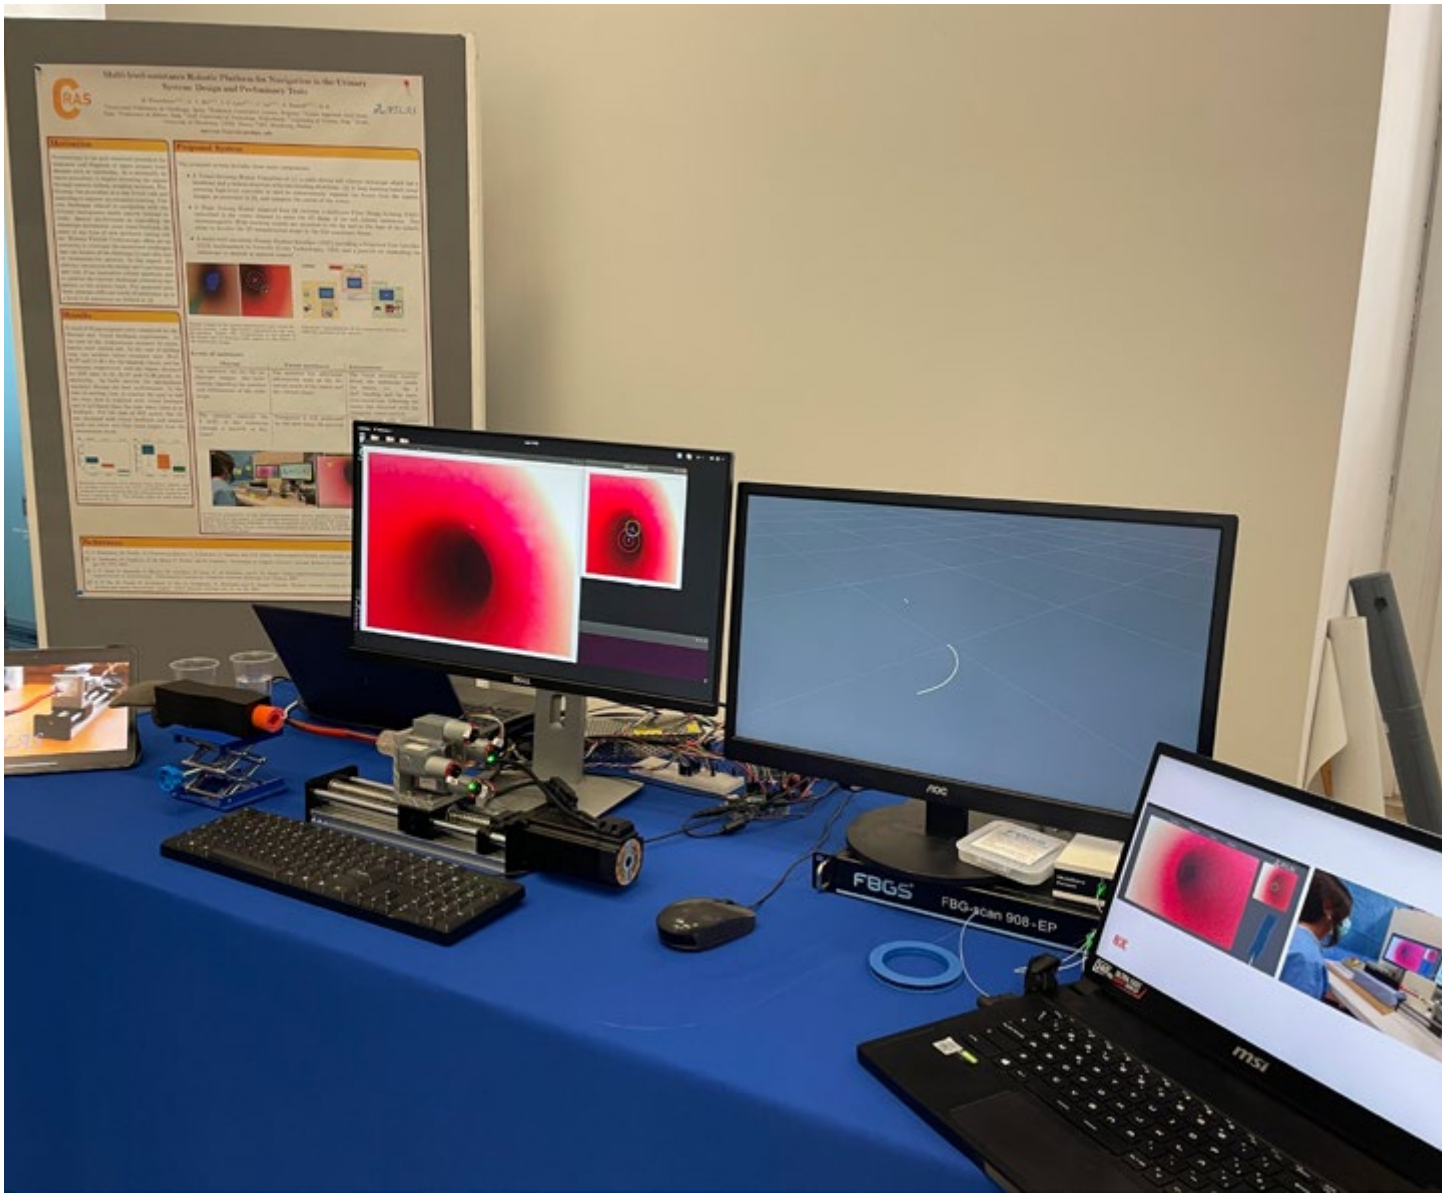

Ureteroscopy is not an easy procedure to perform. **Navigation inside the urinary tract** is complex, and it depends on the user's experience as to whether it will be successful. The operator must avoid touching the walls of the lumen inside the ureter as this could cause lesions. Localization is achieved by performing repetitive X-rays, which can be harmful to both patient and user.

Martina and Sanat's award-winning work includes a **steerable robotic device** – developed by **Chun-Feng Lai**, one of the project's co-authors – which travels along the urinary tract with a camera at the tooltip. It can be controlled autonomously or manually with a joystick.

It has an autonomous visual-motor system that tracks the center of the lumen and moves the device's tip towards it, which helps avoid the risk of lesions. The clinician always has the option to revert to manual control.

"*We use FBG sensors together with EM trackers to localize the whole length of the device and provide this information back to the user,*" Martina explains.

Best Demo Award

Martina Finocchiaro

Sanat Ramesh

"We have a graphical user interface where you can see the endoscopic view from the camera, as well as the 3D shape reconstructed, and the runtime of the device given by FBG sensors registered on the organ. The system offers localization by taking one X-ray at the beginning of the procedure to register with the EM trackers, and then you always know where you are."

The platform reduces the physical load on users with improved ergonomics compared to current endoscopic device controllers. On top of this, **autonomous control reduces the cognitive load**, allowing users to focus on what's most important while supporting less experienced operators. Even under manual control, the information extracted from the camera is fed back to the user to ensure the tip is always directed towards the center of the lumen.

Also, the system saves precious time, which can help patients avoid complications like **sepsis** in certain time-sensitive procedures.

Since the project's outset, the team has been in touch with clinicians, including **Elena De Momi** from the **Polytechnic University of Milan**, who has been key in keeping the PhD students in regular contact with surgeons.

"Initially, we had a video call with surgeons to understand the exact problems they are having," Sanat tells us.

"Everything we're trying to solve comes directly from their input. We interact regularly, they provide feedback, and we go back and forth to see how we should proceed."

Martina and Sanat are modest when we

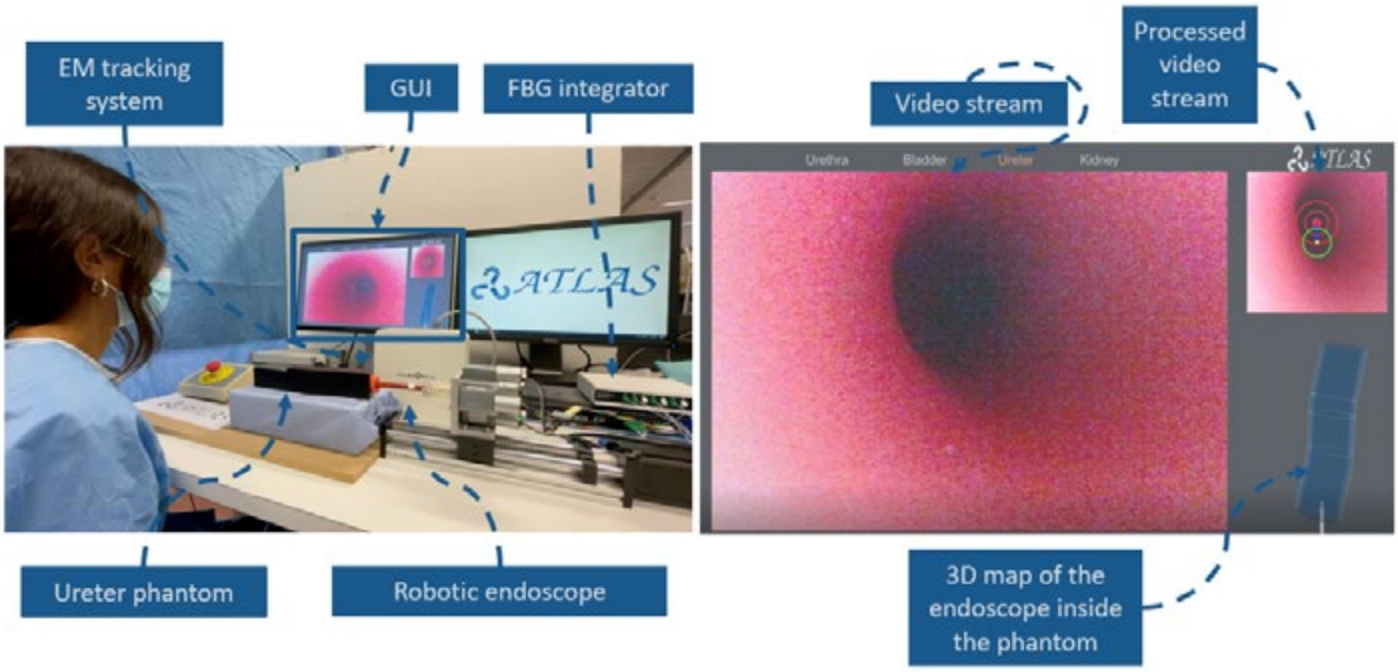

ask what they think convinced the jury to pick them as the winners of the **Best Demo Award**.

"*It was a tough challenge, and there were some great demos in the class,*" Sanat reveals.

"*We had multiple systems working together and could show all the components that we wanted to show together. I think that gave us an edge.*"

Martina agrees:

"*Yes, we had many new technologies and applications combined, including the FBG sensor, the robotic catheter, and the autonomous visual-motor control.*"

Sanat points out that the platform is an **initial prototype**, and it will have to go through many trials and tests before it makes it to humans. Certification is a long process.

Does he have a dream feature that he would like to add?

"*I would say automated lesion detection – detecting polyps, for example. If we get to autonomous treatment and removal, I guess the surgeon could be on a beach somewhere doing this work!*" he laughs.

Finally, with their PhDs coming to an end soon, Sanat is keen to put out a rallying call to other students.

"*I hope this article motivates people to follow this work because we have limited time now,*" he laments.

"*I'd like to see people take it up in the future and make it even more successful!*"

ARTIFICIAL INTELLIGENCE IN BREAST CANCER SCREENING

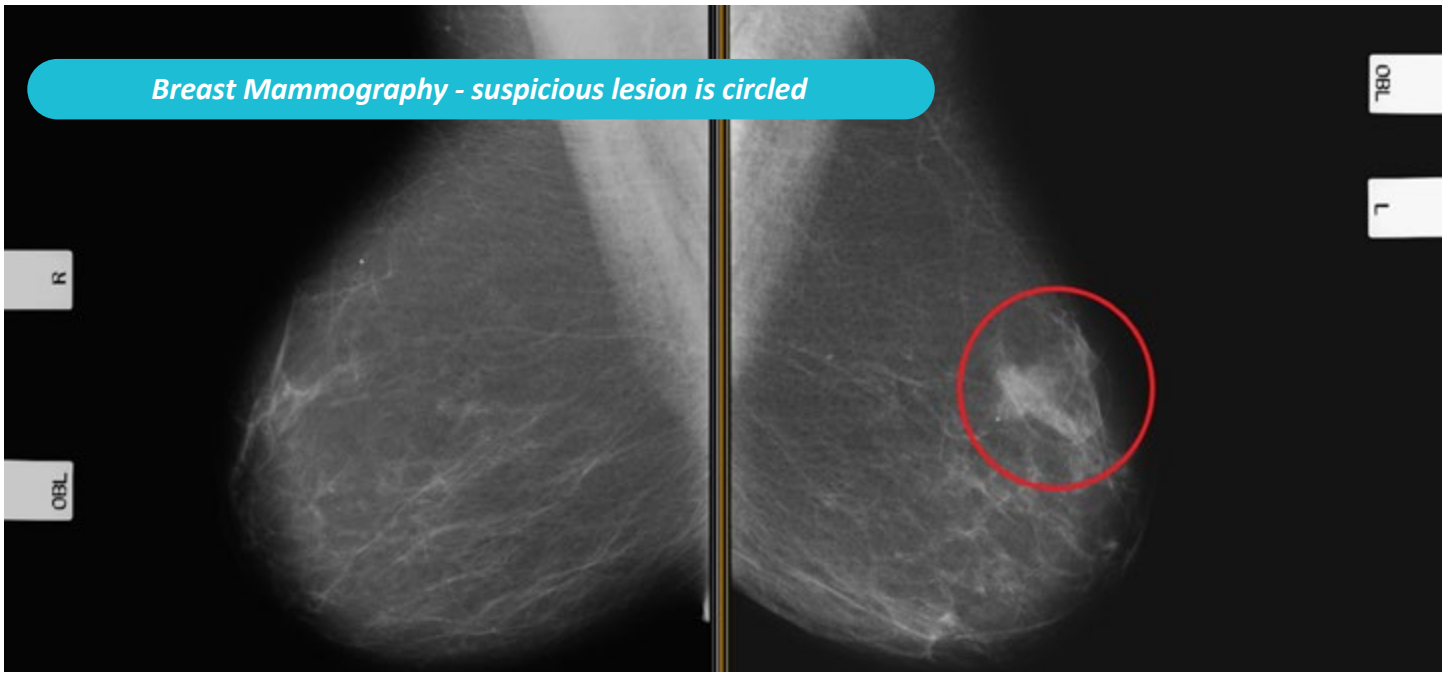

Breast cancer is the 2nd most common cancer in women (after skin cancer). Risk of breast cancer increases with age, and is also affected by family history and genetic factors. Like most cancers, **early detection increases survival rates dramatically**. Therefore, breast screening for women is very common, and multi-imaging modalities are utilized. Current recommendations state that women over 40 should have a yearly mammography examination. Women at higher risk (family history of breast cancer, carriers of BRCA gene, etc.) undergo additional MRI and/or ultrasound exams. It has been shown that a significant portion of radiologist workload originates in breast exams, so any means to assist radiologists in reading these exams can help reduce physician fatigue and error rates. Additionally, some modalities like mammograms use harmful radiation for imaging, so any reduction of radiation dose is beneficial for the patient. Artificial Intelligence applications can be added to any breast screening modality:

Mammography: during this exam, the breast is pressed between 2 surfaces, and a 2D X-ray image is acquired. Alternatively, tomosynthesis acquires a 3D tomography scan of the breast using similar technology.

Applying a lower radiation dose will reduce image quality. However, **neural networks can be trained to reconstruct high-quality images from low-quality ones**. This technique is in use in CT scans, and can further reduce radiation exposure in mammography and tomosynthesis.

Once the image (or tomography) is acquired, the radiologist is required to review it and search for suspicious lesions/sites in the breast for further examination. Both **classic computer vision** and **more advanced deep learning methods** can be used to detect and segment suspicious landmarks.

MRI: a breast MRI exam is also used for

screening and detection of breast cancer. Breast MRI exams are often long - there are multiple sequences to acquire, and breathing causes movement which needs to be accounted for.

This issue can be addressed by **accurate fusion of the acquired sequences**. Deep learning methods can be applied and register each sequence to a predetermined baseline sequence, like Diffusion Weighted Imaging (DWI) to T2, and discard the relative motion between and during acquisition. Additionally, resolution can be enhanced by **training deep neural networks to reconstruct MRI images from suboptimal acquisitions**. Both applications can reduce acquisition time as the need for repeat acquisition is reduced.

In MRI, like mammography, the resulting image needs to be reviewed by the radiologist. Computer vision systems can be developed to support the analysis of breast MRI images. These systems can detect and segment suspicious lesions, calculate their intensity in multiple sequences, and report these results to the radiologist, to be used as the baseline for interpretation and diagnosis.

Ultrasound: Breast US is performed to further examine women with dense breast tissue or with findings in other imaging exams. **Automated breast ultrasound (ABUS)** is a relatively new approach, where a large transducer is placed on the breast and automatically acquires multiple images to cover the entire breast.

Also here, AI can be used for **automatic detection of abnormalities**, and with sufficient data even **classification of malignancy**. This can reduce the miss rate significantly. Additionally, ultrasound is often used for biopsy guidance, and **Artificial Intelligence methods can be applied to provide needle guidance and tracking during the procedure to increase accuracy and reduce procedural time**.

The **Breast Imaging Reporting and Data System (BI-RADS)** was developed to give specific guidelines for reporting findings in breast screening images. These guidelines apply to **mammography, MRI, and ultrasound**, and they account for the lesions size, volume, shape, homogeneity, restriction, and other features. All these parameters can be automatically extracted from the imaging data, and presented to the radiologist for consideration, **saving even more precious time and increasing physician confidence**.

As described above, breast cancer takes a large toll on the medical industry. Numerous scans are conducted daily, and early detection is the main factor affecting survival rate. **Introducing AI modules into the screening and diagnosis pipeline can reduce scan and radiologist review time, lower costs, and increase detection and survival rate**. There are many methods for AI implementation, and **RSIP Vision** can assist in the development of any of the aforementioned solutions.

Breast MRI - MIP with subtraction showing enhanced procedure in left breast

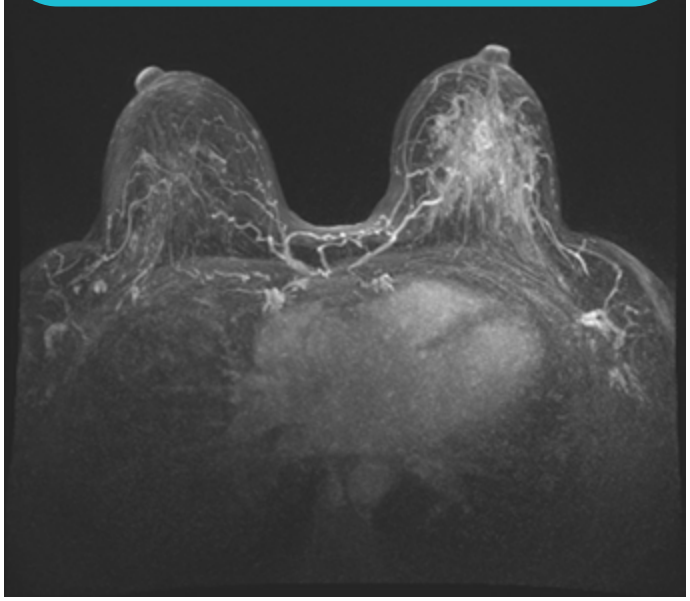

A PREVIEW OF MICCAI 2022

MICCAI 2022 Singapore

25th International Conference on Medical Image Computing and Computer Assisted Intervention

September 18–22, 2022

Resorts World Convention Centre Singapore

Shuo Li is the General Chair of MICCAI 2022 in Singapore.

Shuo, for the first time in three years, MICCAI is back as a primarily in-person event. You must be excited!

Yes, I am! Our community is shaking because people are not meeting each other. Many have said to me, *"I haven't traveled for three years. Please don't have another virtual event. I want to see people!"* We can finally get back to normal because this conference is not just about scientific content; it's a networking event. We come to meet friends and colleagues and learn what they do. The MICCAI community is built on all the connections that we create.

For many new participants, all they know is Zoom presentations. They have never met anybody.

It is a pity that the students, especially those enrolled in 2019, who in 2020/2021 should have been getting out into the community, presenting their work, and participating in social events, didn't get to do so. Some of these students may be graduating this year, and they will have missed out on a key part of their academic training.

How confident are you that enough people will come along in person this year and that habits have not changed?

Many active members in the community have approached us to persuade us to do it, so we're pretty sure. There's always


General Chair Shuo Li

uncertainty nowadays – new Covid variants, the financial burden, and the war even. We know not every individual will be able to come, but most members hope to have an in person event this year. . We can have a virtual meeting anytime, but after three years, we're optimistic that people will return.

I imagine some people may want to attend, but their supervisor will not pay for them to go to a physical event again. How can they convince their boss to send them?

This question has been a recent sticking point in our discussions. The last three years have indeed been burdensome to supervisors with financial challenges. We've come up with a hybrid solution to be more inclusive so that everyone can still attend, but we know it will not be the same.. After discussing it with the board and our sponsors, we've decided to keep our registration fee the same as before Covid. That includes the lunch, the banquet, everything. Hopefully, we can help minimize the burden and make life easier by keeping costs down.

We're also planning to give more free registrations and student travel awards. Student travel awards allow you to travel to the conference and enjoy free registration or accommodation. That's something we're still discussing. Of course, we must consider our financial position in the current climate, and by keeping the registration fee the same and increasing investment, we're pushing the envelope, but it's the right thing to do.

How are the submissions shaping up so far?

This year, submissions are equivalent to 2019 levels. They've fully recovered from Covid, which is positive news and a strong indication that people want to attend. Our program team is working hard to bring a feast to everyone. This year will be the first year we've had two parallel tracks. Traditionally, MICCAI was a single-track event. A few years ago, we briefly tested a dual track, but we will probably have three full days of two tracks this time. There is always someone hanging around in a hallway with a single track because they're not interested in something. Now, we have a diverse range of events to serve the diverse needs of attendees.

On day two, we actually have three tracks! We have two MICCAI main tracks and the CLINICCAI program. CLINICCAI started as a virtual event last year, but we have a new program this year with many physicians attending in person. We have a clinical leader from Europe and a local clinical leader who has assembled a large local team. We expect to set a historical record for clinicians attending MICCAI!

What should we expect in terms of the scientific content this year?

We've just completed the review process, and now the program chair, area chairs, and reviewers are starting to discuss the rebuttals, so we'll have a better idea of what we'll be presenting soon. 30% of papers have already been accepted and are of the highest quality; half have been rejected. The other 20% require further discussion. We're confident in the quality given the high cutting line. Papers are being rejected even though they have a relatively high evaluation from the reviewer compared to previous years due to the competition. It gives us the confidence that we will see some innovative scientific content this year, but we don't know exactly how many papers will be accepted yet.

Can you tell us more about the review process?

We have an open review process for the first time this year, meaning that all the scientific discussions will be openly published. The author will be able to see the details of those comments and the discussion afterward, and everything will be published online later. The reviewers will know that the community can read their comments, encouraging them to provide a thoughtful review. If a rejected paper gets submitted to other places and

a reviewer there says our review was unreasonable, we should accept it. Much scientific evidence indicates that open review increases review quality and the quality of the publication overall.

MICCAI is traditionally a very academic conference, but what is industry's role this year?

Industry hasn't played the strongest part in the main event before. However, there's always strong participation of people from industry. Even when they're not a sponsor or part of the exhibition, many attend in their own right and take part as reviewers. That's the sign of a good conference. They come to learn – it's the source of their knowledge and their next innovation.

In that sense, I always see industry as part of the MICCAI community, and this year it will continue to be part of it. We have generous sponsorship from companies like Google and companies from China like NDI. NDI doesn't sponsor virtual conferences but has shown strong confidence in this one. Even though China is still under lockdown policy, it's a gold sponsor and will do its best to attend.

Do you have a final message for the community?

MICCAI 2022 is a new starting point for many years of growth and development. It will be a chance to reconnect, meet new colleagues, and reboot our society. It presents opportunities for every community member because pretty much everything is starting over. I would strongly encourage everyone working in this area and everyone who wants to work in this area to attend. Remember, we're always here to help. Plus, Singapore is a wonderful place to have a break and reboot!

Vision Transformers in Medical Computer Vision

IOANNIS VALASAKIS, KING'S COLLEGE LONDON
@WIZOFE

Hello again! In this article we are going to have another medical imaging tutorial creating a COVID predictor using chest X-rays!

This month we will look at how Artificial Intelligence applied to the medical domain can have very significant impact. Since COVID-19 attacks the epithelial cells that line our respiratory tract, we can use X-rays to analyze the health of a patient's lungs.

X-ray is a very common modality in most hospitals over the world. Developing a semi-automated analysis system can save medical professionals valuable time and the lives of many patients.

Let's start by importing important metadata and modules.

```
#Create the data for positive samples
IMAGE_PATH = "/content/drive/My Drive/images"

df = pd.read_csv('metadata.csv')
print(df.shape)
```

```
(372, 29)
```

```
df.head()
```

```
  patientid  offset  ... other_notes  Unnamed: 28
0         2     0.0  ...         NaN          NaN
1         2     3.0  ...         NaN          NaN

[5 rows x 29 columns]
```

Load Datasets

Here I will load saved XRAY datasets where the training will happen:

```
TARGET_DIR = "/content/drive/My Drive/Dataset/Train/Covid"

if not os.path.exists(TARGET_DIR):
  os.mkdir(TARGET_DIR)

#Copy COVID-19 images with view point PA from Downloaded directory to Target Directory
cnt = 0
for(i,row) in df.iterrows():
  if row["finding"] == "COVID-19" and row["view"] =="PA":
```

```
    filename = row["filename"]
    image_path = os.path.join(IMAGE_PATH,filename) #IMAGE_PATH + filename
    image_copy_path = os.path.join(TARGET_DIR,filename) #TARGET_DIR + filename
    shutil.copy2(image_path,image_copy_path) #Copy From IMAGE_PATH to TARGET_DIrectory
    print("Moving Image",cnt)
    cnt+=1
print(cnt)
```

```
Moving Image 0
Moving Image 1
Moving Image 2
Moving Image 3
Moving Image 4

….
140
```

```
# Normal chest X rays from Kaggle Data.

import random
KAGGLE_FILE_PATH = "Desktop\\chest_xray\\train\\NORMAL"
TARGET_NORMAL_DIR = "/content/drive/My Drive/Dataset/Train/Normal"
if not os.path.exists(TARGET_NORMAL_DIR):
  os.mkdir(TARGET_NORMAL_DIR)

image_names = os.listdir(KAGGLE_FILE_PATH) #COntains list of all image names
image_names
random.shuffle(image_names) #it will randomly shuffle names in list

for i in range(141):
  image_name = image_names[i]
  image_path = os.path.join(KAGGLE_FILE_PATH,image_name)

  target_path = os.path.join(TARGET_NORMAL_DIR,image_name)
  shutil.copy2(image_path,target_path)
  print("Moved",i)

from google.colab import drive
drive.mount('/content/drive')
```

```
Enter your authorization code:
··········
Mounted at /content/drive
```

```
TRAIN_PATH ="/content/drive/My Drive/Dataset/Train"
VAL_PATH = "/content/drive/My Drive/Dataset/Val"
```

Building Architecture

```
model = Sequential()
model.add(Conv2D(32,kernel_size=(3,3),activation="relu",input_shape=(224,224,3)))

model.add(Conv2D(64,(3,3),activation="relu"))
model.add(MaxPooling2D(pool_size = (2,2)))
model.add(Dropout(0.25))

model.add(Conv2D(64,(3,3),activation="relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Dropout(0.25))

model.add(Conv2D(128,(3,3),activation="relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Dropout(0.25))

model.add(Conv2D(128,(3,3),activation="relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Dropout(0.25))

model.add(Flatten())
model.add(Dense(64,activation="relu"))
model.add(Dropout(0.5))

model.add(Dense(1,activation="sigmoid"))

model.compile(loss=keras.losses.binary_crossentropy,optimizer = "adam",metrics=["accuracy"])

model.summary()
```

This is an exercise left for the reader: feel free to perform data augmentation!

Found 50 images belonging to 2 classes.

Fit The Model

```
hist = model.fit_generator(
   train_generator,
   steps_per_epoch = 8,
   epochs = 10,
   validation_data = validation_generator,
   validation_steps = 2
)

Epoch 1/10
8/8 [==============================] - 179s 22s/step - loss: 0.8587 - accuracy: 0.5261 - val_loss:
```

```
0.6933 - val_accuracy: 0.5000
…
Epoch 10/10
8/8 [==============================] - 9s 1s/step - loss: 0.1225 - accuracy: 0.9522 - val_loss:
0.1918 - val_accuracy: 0.9600
```

```
model.save("Covid19_XrayDetector.h5")
```

```
 model.evaluate_generator(train_generator)
```

```
[0.02655816078186035, 0.991304337978363]
```

```
 model.evaluate_generator(validation_generator)
```

```
[0.19179965555667877, 0.9599999785423279
```

Test Images

```
model = load_model("Covid19_XrayDetector.h5")
```

```
train_generator.class_indices
```

```
{'Covid': 0, 'Normal': 1}
```

Confusion Matrix

By creating a confusion matrix we can get an idea of how well our model performs!

```
from sklearn.metrics import confusion_matrix
cm = confusion_matrix(y_actual,y_test)
```

```
sns.heatmap(cm,cmap = "plasma",annot=True )
```

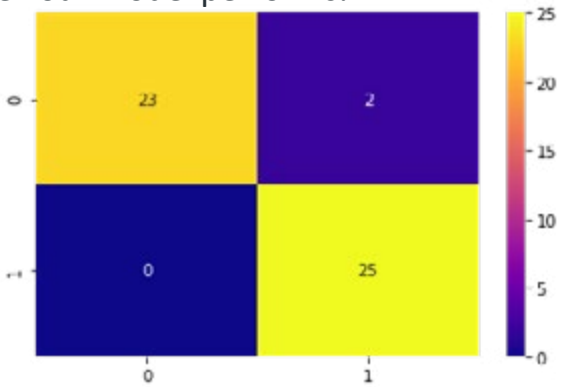


```
import itertools
import numpy as np
import matplotlib.pyplot as plt
from sklearn.metrics import confusion_matrix
class_names = ["Covid-19","Normal"]

def plot_confusion_matrix(cm, classes,
                          normalize=False,
                          title='Confusion matrix',
                          cmap="plasma")
```

Confusion Matrix

```
plt.figure()
plot_confusion_matrix(cm, classes=class_names,
        title='Confusion matrix for Covid-19 Detection')
```

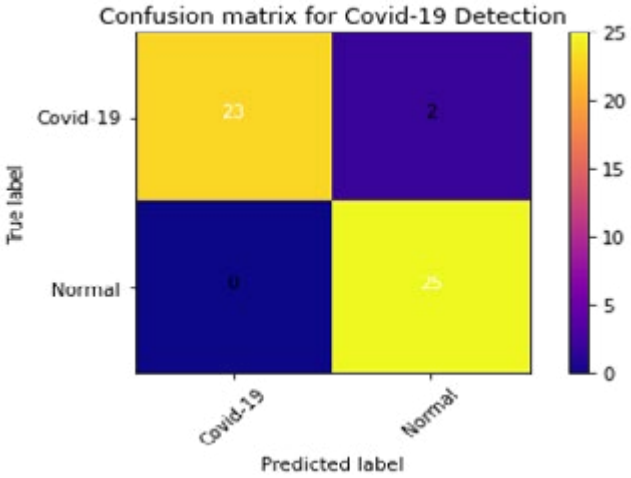


```
Confusion matrix, without normalization
[[23  2]
 [ 0 25]]
```

```
history = hist
print(history.history.keys())
```

```
dict_keys(['val_loss', 'val_accuracy', 'loss', 'accuracy'])
```

Let's plot the accuracy and loss values for both training and test sets:

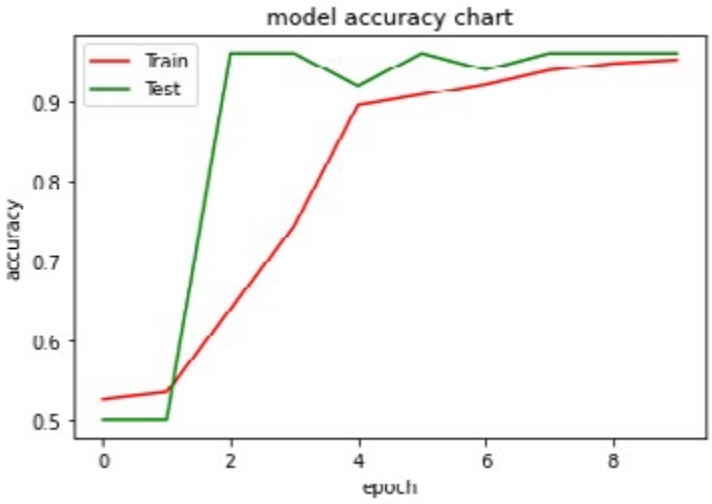

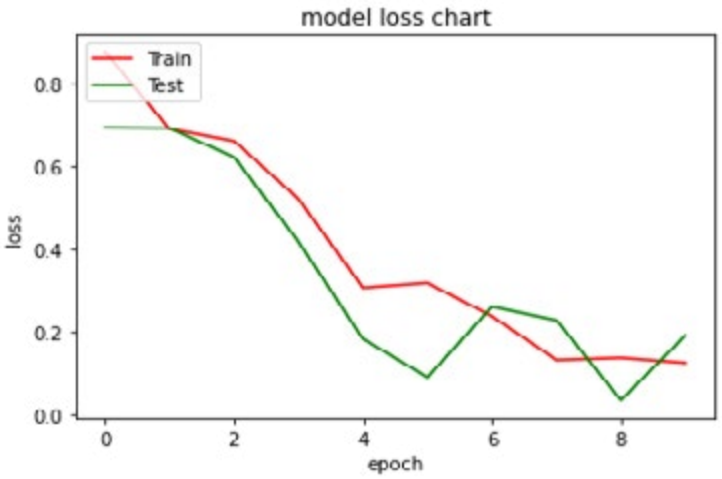

Predictions from X-Ray Images

The following are the predictions for the x-ray images!

```
imggg = cv2.imread('/content/drive/My Drive/Dataset/Predictior_Image/IM-0569-0001.jpeg')
print("Actual: Negative covid-19 ")
imggg = np.array(imggg)
```

```
imggg = cv2.resize(imggg,(400,400))
cv2_imshow(imggg)
print(results)
if results[0][0] == 0:
   prediction = 'Positive Covid-19'
else:
prediction = 'Negative Covid-19'
print("Prediction  : ",prediction)
```

Actual: Negative covid-19

[[1]]
Prediction : Negative Covid-19

Actual:Positive covid-19

[[0]]
Prediction : Positive Covid-19

Next month

As a conclusion to this -hopefully useful- COVID predictor, I want to try the mentioned exercises and feel free to send any of your questions through. Last month, I received in my Twitter two great suggestions from our readers and I am going to have one of them as the next review. Until next month, have a great time and always be curious! ☺

Aziliz Guezou-Philippe has recently completed her PhD at the Laboratory of Medical Information Processing, University of Western Brittany, Brest, France. Her research focuses on biomechanical and statistical shape modelling, statistical analysis and point cloud registration. She aims to help surgeons improve their surgery by developing computational tools that are easy to integrate in clinical routine. She's on the market for a postdoc position and she's awesome. Congrats, doctor Aziliz!

Total hip arthroplasty (THA) has become one of the most common surgeries worldwide as it reduces the pain and restores the hip mobility of around 1 million patients every year. However, it is still at risks of complications, such as dislocation, with an incidence varying from 1% to 5%. Dislocation usually occurs by lever arm effect when the femoral stem impinges the acetabular cup and is directly related to the correct positioning of the implants. The goal of my thesis was to propose a global approach that considers patient-specific functional parameters to optimize the implants positioning in order to minimize the risks of prosthetic instability.

Determine a Functional Orientation for the Acetabular Cup

In order to reduce the risks of dislocation, the acetabular cup must be oriented regarding the pelvis inclination. However, the standard recommendations do not take into account the mobility of the patient: the hip range of motion (ROM) and the pelvis tilting during activities of daily life. Therefore, I developed a functional safe zone, adaptable to the pelvic tilts of each patient in different positions, that would ensure the absence of prosthetic impingement and thus reduce the risks of dislocation (Figure 1). This safe zone is based on the 3D modelling of the implants motion and is simple and fast to compute. The pelvic tilts are acquired using a navigated ultrasound device developed in our lab, which avoid ionizing imaging techniques. Moreover, it can be used to analyze the impact of prosthetic parameters regarding the patient-specific pelvic tilts and ROM, in order to personalize the pre-operative planning and estimate the margin of error for the prosthesis implantation.

Acquire the Patient Reference Plane During Surgery

Once the pre-operative planning is determined, the surgeon must be able to easily position the implants. One challenge is to acquire the patient-specific Anterior Pelvic Plane (APP), which is the plane of reference for the cup orientation, with the patient already installed for surgery. However, most THAs are performed with the patient lying on his side, which makes the APP acquisition difficult as one anatomical landmark is found against the operating table and covered by the patient abdominal apron. In my thesis, I proposed a new method to determine the APP from anatomical landmarks of only one side of the pelvis acquired using navigated ultrasound (Figure 2). To do that, I built a statistical shape model of the pelvis, able to reconstruct the patient's whole pelvis from a few anatomical landmarks. Once the model is fitted to the target landmarks, the APP can be automatically determined. This approach is quick, low-cost and compact, which makes it convenient to use in the operating room.

I am now looking for postdoc opportunities to continue working on medical and translational research. Feel free to contact me at azilizgp@gmail.com. Ken ar c'hentañ !*

* "See you soon!" in Breton

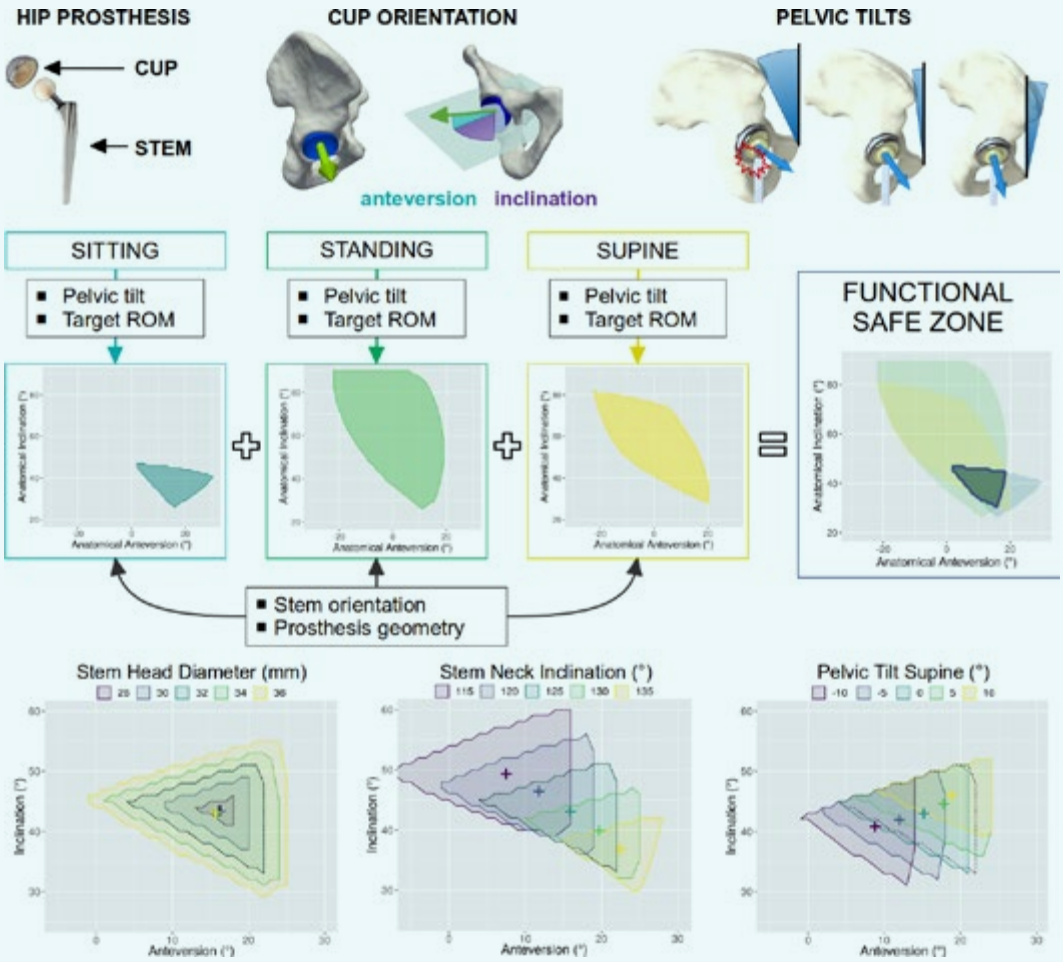

Figure 1. (top) Definition of, from left to right, the implants, the cup orientation and the pelvic tilts. (middle) Computation of the functional safe zone: for each position, the cup orientations safe from impingement are computed and form a "static" safe zone; the functional safe zone is then defined as the intersection of the sitting, standing and supine safe zones. (bottom) Effect of some parameters on the functional safe zone.

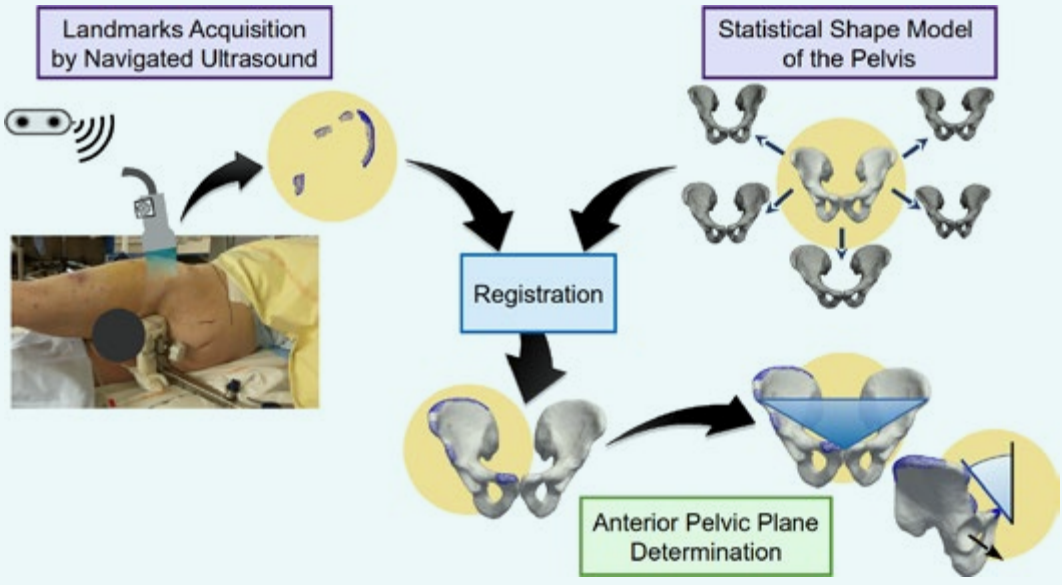

Figure 2. Determination of the patient reference plane: landmarks of one side of the pelvis are acquired using a localized ultrasound probe; then a statistical shape model is registered to the acquired landmarks; finally, the anterior pelvic plane orientation is determined on the fitted model. The acetabular cup can then be positioned and implanted regarding this reference plane.

HECKTOR CHALLENGE AT MICCAI 2022

Vincent Andrearczyk is a Senior Researcher at the University of Applied Sciences Western Switzerland and a co-organizer of the HECKTOR Challenge, which will be presented at MICCAI 2022 in Singapore in September. He tells us more about the competition, which has just kicked off.

Two years ago, the first **HECKTOR Challenge** presented its results at MICCAI 2020. Its task was to **automatically segment head and neck primary tumors in PET/CT images**. After a successful start, the challenge grew last year with more data and a view toward personalized medicine and finding the best treatments for patients. It added a task around predicting patient outcomes and progression-free survival based on PET/CT images and clinical data.

"*This year, for the segmentation task, we propose the segmentation of **primary tumors and metastatic lymph nodes**,*" Vincent reveals.

"*We realized that metastatic lymph nodes are predictive for patient outcomes, so we wanted to include them. We've also increased the number of centers, so we have data from nine hospitals and many more cases in the training and test sets.*"

Combining the segmentation task with the prediction of patient outcomes is something that Vincent tells us separates HECKTOR from other segmentation challenges.

"*I think we were the first segmentation task from PET/CT images,*" he points out.

"*We have the PET image representing the metabolic response of the tumor and the CT modality providing information on the morphological tissue properties. The discriminative information in PET images is important, but **multimodal models***

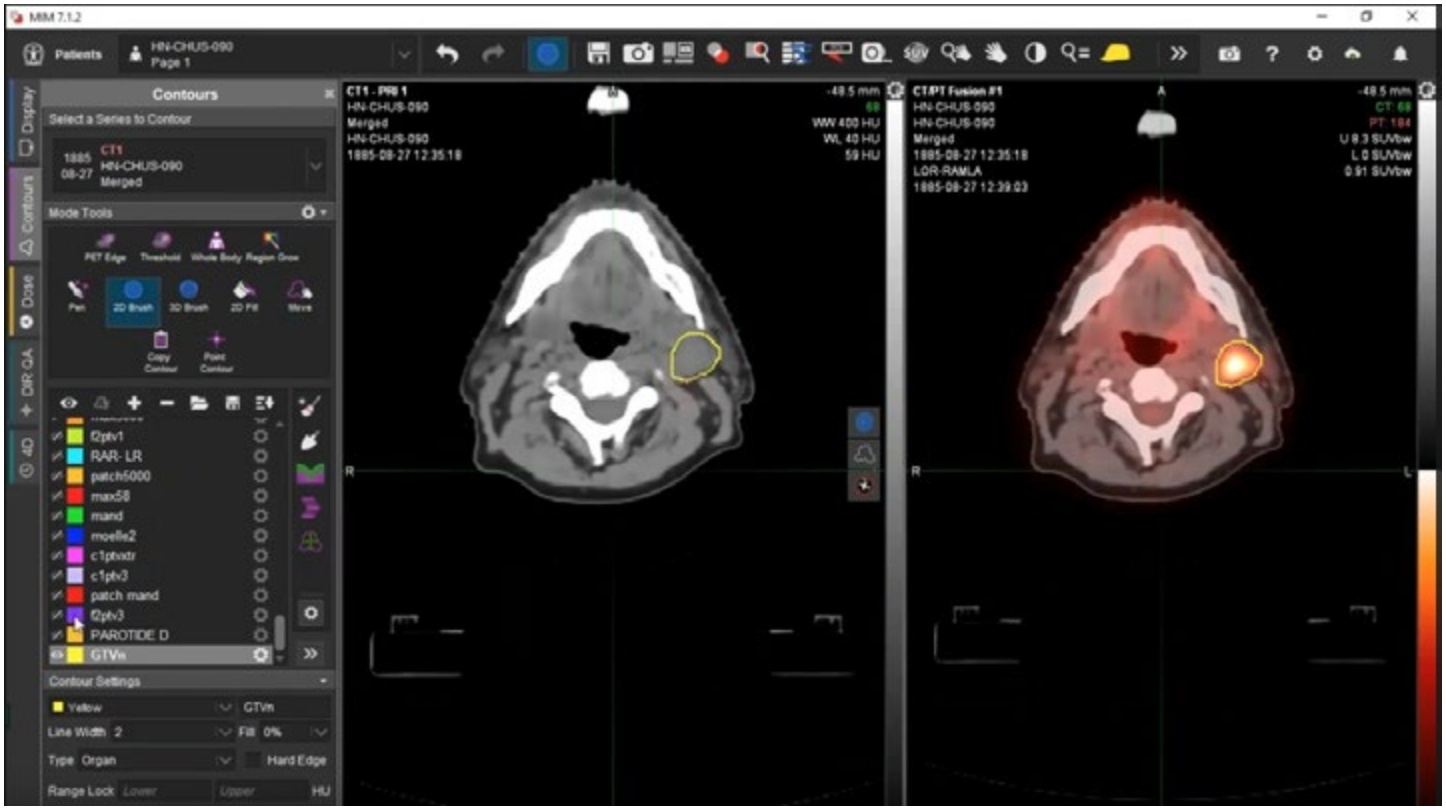

***based on PET and CT images work best**. Combining these two modalities is important for a good segmentation followed by outcome prediction."*

Participants can take part in one or both tasks. This year, organizers are not providing test tumor contours, as they realized radiomics methods based on tumor contours from experts were not always the best, so participants will be segmenting their methods themselves and then predicting the outcome.

Vincent says understanding the data and the problem is key. For segmentation, standard models like a **U-Net** can be helpful and good pre-processing of data and data augmentation. For outcome prediction, it is essential to understand the clinical data and the links between, say, HPV status and how it influences the patient outcome. Also, combining clinical data with the PET and CT images and understanding how that can help **build strong survival models for prediction**.

With this being the third edition of the challenge, what might we see in the fourth, fifth, and beyond?

"We've thought of other prediction tasks, such as predicting HPV status, TNM staging, sarcopenia – these are all relevant predictions for clinical usage," Vincent tells us.

"Also, subpopulation analysis – fitting our populations into HPV-positive patients and predicting local recurrence."

Its consortium of organizers represents the interdisciplinarity of this challenge. Together, they discussed its design and analysis and how to develop guidelines for annotations. The first task seeks algorithmic contributions, but the second is about making decisions on the outcomes and performance metrics that are most meaningful for the clinic in head and neck cancer treatment, ultimately to help clinicians find the best treatment for patients.

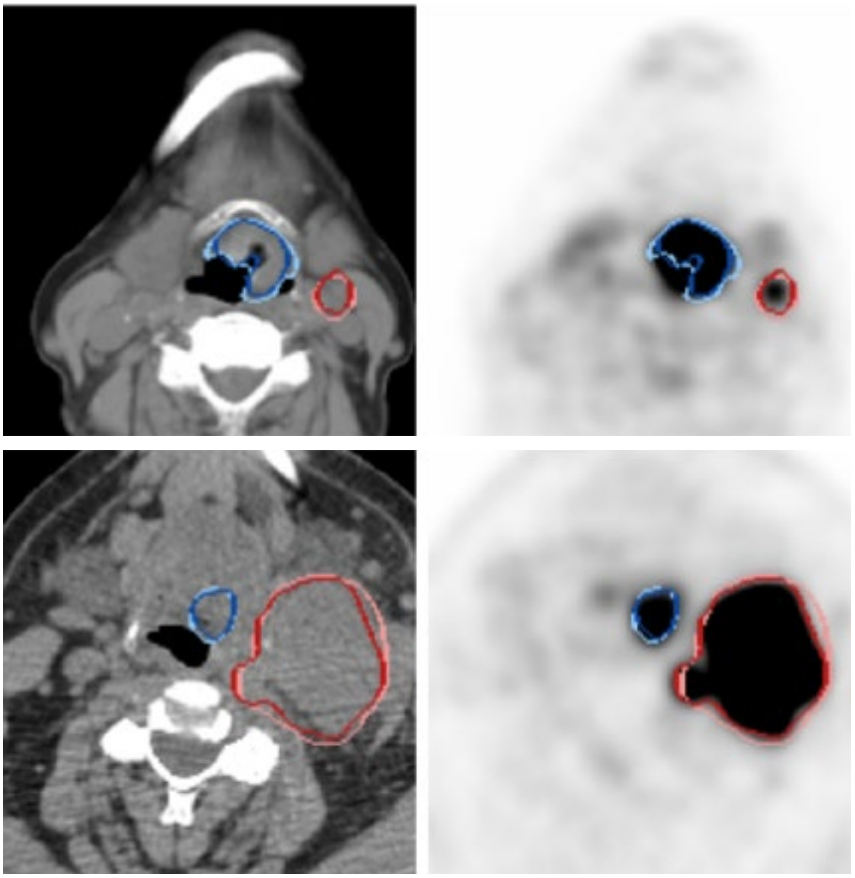

Although the challenge is taking up a great deal of Vincent's time now, particularly this last sprint before releasing the data, his day job focuses on **deep learning and radiomics problems for radiology**, including segmentation and predicting outcomes for different cancers.

"*I also work on the interpretability of the models, which is very important for clinical acceptance,*" he says.

"*The end goal is always clinical interpretability, which is another next step for HECKTOR.* ***Evaluating the interpretability and uncertainty of the models*** *is very important if we want to reach clinical development.*"

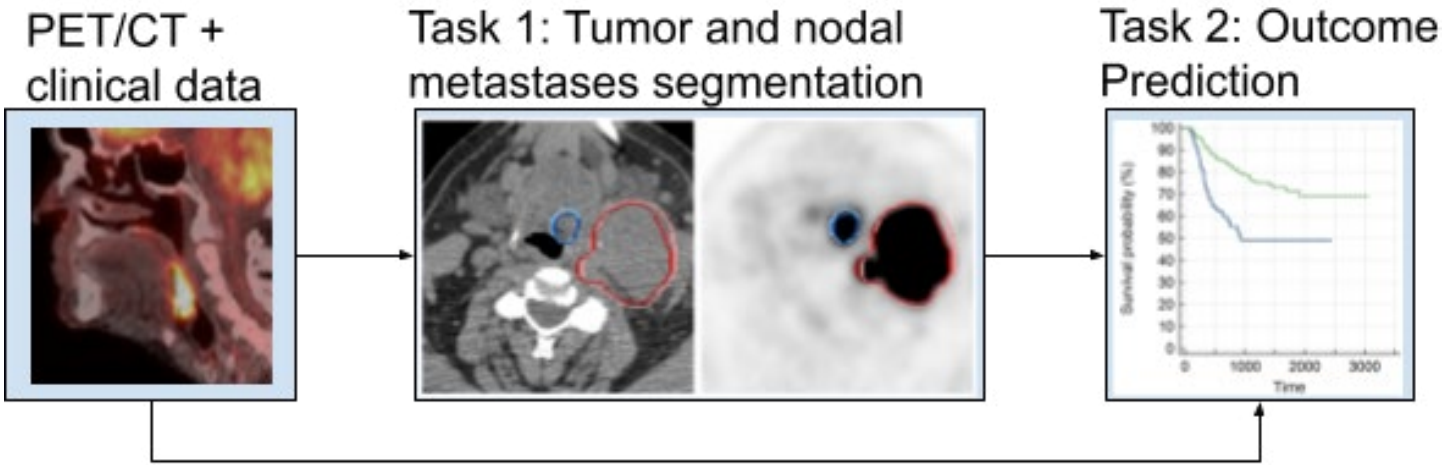

Medical Imaging News has found great new stories, written somewhere else by somebody else. We share them with you, adding a short comment. **Enjoy!**

Nanoparticles: sometimes the smallest things can have the biggest impact

The article we link to is a written by the CEO of **Imagion Biosystems**, so it's a bit self-promotional; nonetheless it is extremely interesting. In brief, **combining nanotechnology with pharmaceutical and biomedical sciences** is enables to re-invent how certain diseases, like cancer, can be diagnosed and treated. Application of nanotechnology in healthcare are getting more and more frequent: 1-100nm **nanoparticles** can be made from a variety of substances including metals, lipids, and polymers. They are naturally metabolized by the body and the iron is recycled to help form hemoglobin. Read More

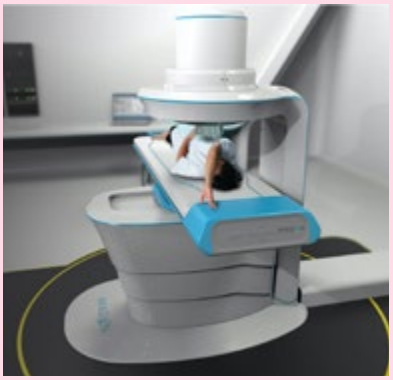

New non-invasive method of risk assessment in liver disease

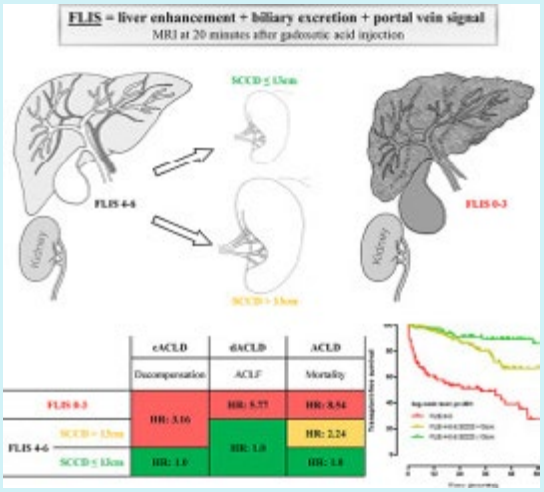

Journal of Hepatology (via eurekalert.com) tells us that a research team from **MedUni Vienna** showed that **functional magnetic resonance imaging (fMRI)** can be used as a non-invasive method for predicting complications in **chronic liver disease**. The scientists combined a simple risk stratification system - the functional liver imaging score (FLIS) - with splenic diameter. This non-invasive combination provided the researchers with complementary data for risk assessment in patients with advanced chronic liver disease and it can be used to complement invasive investigations. Read More

AI Model can Predict whether Crohn Disease will Recur After Surgery

One more proof that **Artificial Intelligence** is bringing – and will keep bringing – immense benefits to our health. A Japanese team of researchers uses an AI tool that emulates how humans visualize and is trained to recognize and classify images: their model predicts **the postoperative recurrence of Crohn disease** with high accuracy by evaluating histological images. This method was already used targeting malignant tumors. This study focused on Crohn disease, in which **postoperative recurrence is a clinical problem**. You can read in the article in what way this can enable **more intensive and successful treatment of high-risk patients**. Read More

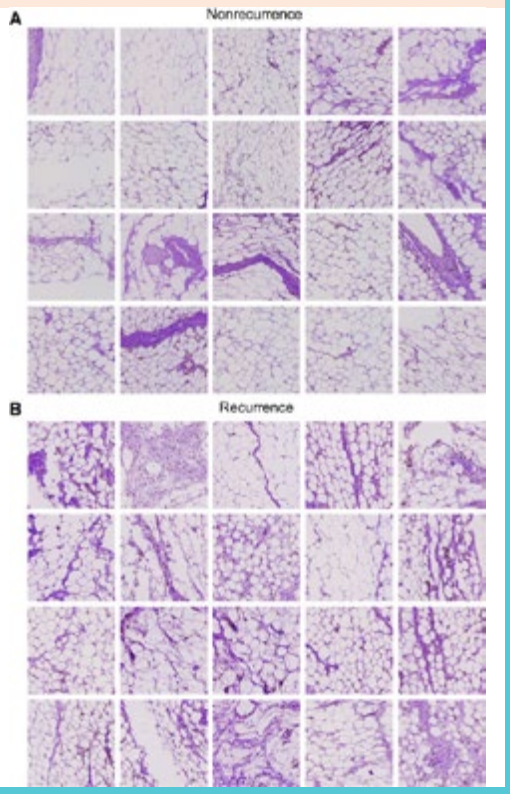

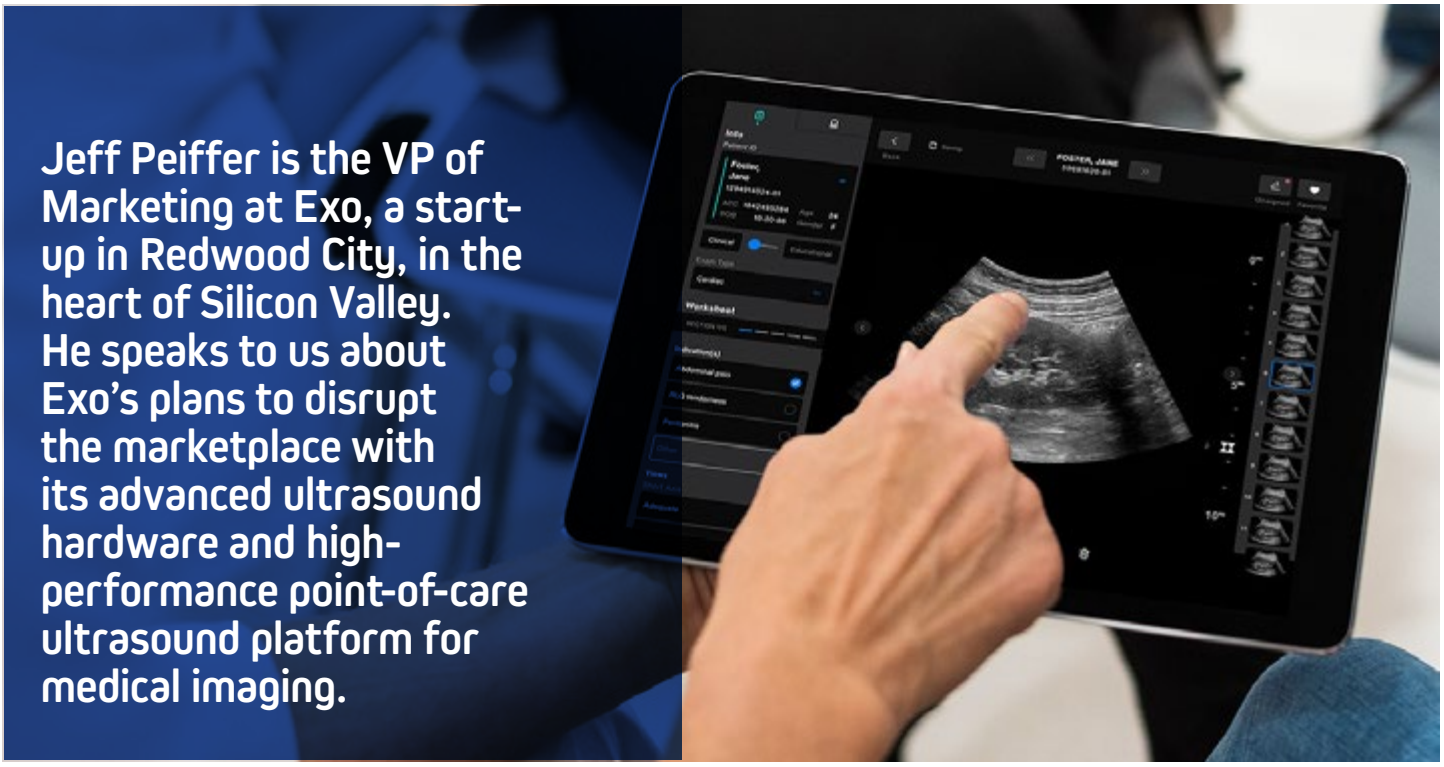

Jeff Peiffer is the VP of Marketing at Exo, a start-up in Redwood City, in the heart of Silicon Valley. He speaks to us about Exo's plans to disrupt the marketplace with its advanced ultrasound hardware and high-performance point-of-care ultrasound platform for medical imaging.

For more than fifty years, ultrasound probes have been made by taking a piece of piezoceramic, dicing it into small elements, and attaching wires to each element. This process is labor-intensive and expensive.

The technology developed by Exo (pronounced 'echo') is completely different from anything on the market today. Described as 'silicon meets piezo,' its nanotech hardware deposits the piezoelectric on a piece of silicon, giving users the capability of a three-dimensional imager in their pocket. It is relatively inexpensive but would offer a much higher performance than other handhelds, and the technology has the potential to be used in many different form factors.

The device itself is a **Piezoelectric Micromachined Ultrasonic Transducer**, or **pMUT**, rather than a **Capacitive Micromachined Ultrasonic Transducer**, or **cMUT**.

"A pMUT uses high-end ultrasound technology and performance capabilities at the onset, whereas a cMUT has some limitations, from the amount of power to the different software algorithms you can use," Jeff explains.

"The P, the piezoelectric, is where the differentiation takes place. The FDA limits the amount of power you can put in the body, but the piezoelectric allows us to do broader frequencies and higher power up to the FDA limit."

That extra power can help image patients for whom it is difficult for the ultrasound to penetrate. For example, someone who is bigger than average, where that extra power afforded by the pMUT is needed to see deep inside of the body.

Jeff joined Exo four years ago, having spent almost two decades at **GE Healthcare**, where his career in ultrasound began.

"What attracted me to Exo is its vision of helping the world with imaging," he tells us.

"Three-quarters of the planet doesn't have access to any imaging, and ultrasound is a

natural modality to provide that imaging globally. The technology we're using to deploy this vision is unique. It's unlike anything I've ever seen before, and I've been deep in the bowels of a large company doing ultrasound for many years."

When Jeff first started working with point-of-care ultrasound, the field had tended to be about very traditional methods, such as radiology ultrasound, cardiology ultrasound, and OB-GYN ultrasound. Then, suddenly, there were physicians in emergency medicine picking up a probe and looking for answers on the spot.

"For me, it was this opportunity to learn about every specialty in point of care that was adopting and using ultrasound at the bedside – it was like going to medical school!" he laughs.

"Almost every specialty can benefit from using ultrasound right there with the patient. I had the privilege of working with the first two medical schools that adopted ultrasound as part of their curriculum. You had all these smart, young med students learning anatomy and physiology picking up a probe like it was nothing!"

When it comes to using artificial intelligence in ultrasound, Jeff says the industry is at the start of the journey but looking at a wealth of opportunities ahead.

"AI, I believe, will have a significant impact on guidance," he points out.

"Ultrasound is a skilled procedure. Users need to know anatomy and almost visualize the body in 3D, which is not easy. AI will help us get to the anatomy quickly, easily, and once we're there, pick the image that best exemplifies what we need to know to make a decision. It could even allow us to do calculations automatically, like ejection fraction for the heart or bladder volume."

Jeff Peiffer

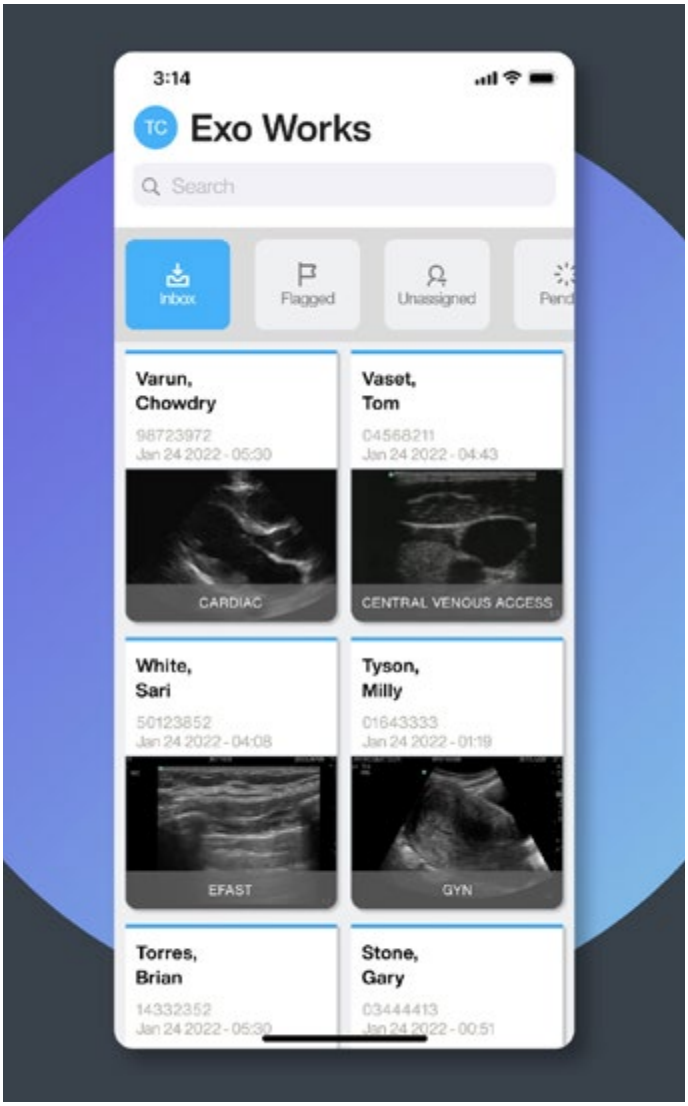

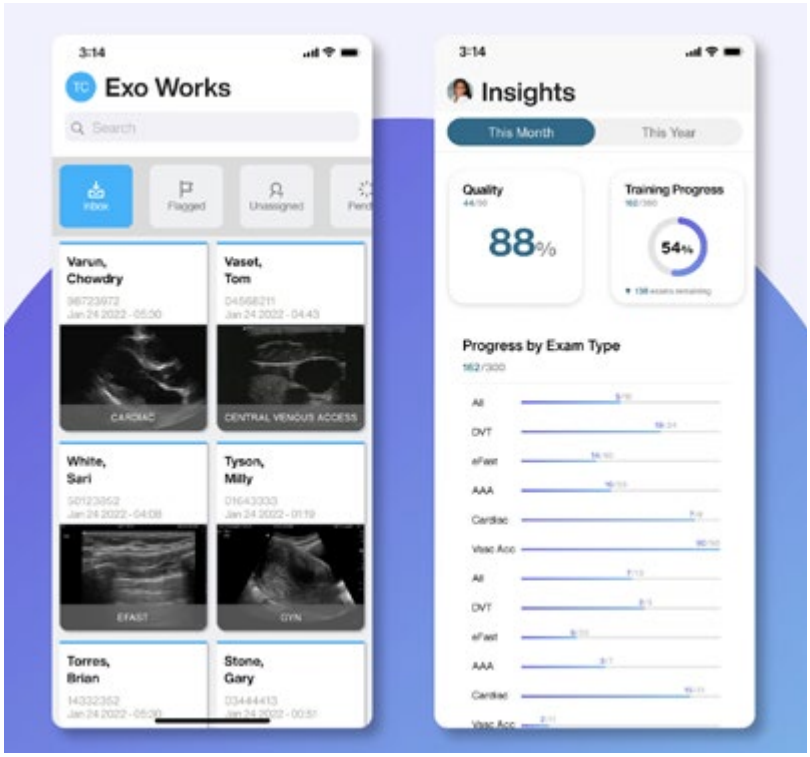

Jeff recalls that last year's **RSNA conference** was the first time he had seen a dedicated AI section – albeit primarily MR, CT, and X-ray. Ultrasound adoption has been slower, but things are evolving, and guidance is at the forefront.

Exo's hardware device is not yet on the market, but its software product, **Exo Works**, is already solving the number one problem in point-of-care ultrasound: workflow.

"*When your GP says you need an abdominal scan, you go to radiology, they know who you are, they know what you're there for, it's all been scheduled, and it's done,*" he explains.

"*If you go to the ED or critical care and need an ultrasound, there's no prescription. It's an encounter. A physician does the scan right there, finds out what the patient needs, and takes care of it. Then they have all this information that must be put back into a system. More times than we know, images get lost, or there are ghost images, and maybe they can't get them into EMR or PACS because things aren't connected. Then they run out of time. What do they do with that information?*"

Exo Works makes documenting exams as easy as picking up a smartphone or a tablet and dialing in. Within seconds, the patient is signed off, and scans are on their way to EMR and PACS for billing and storage automatically. It can hold educational and clinical scans and has built-in QA and credentialing tools. Devices can be added, whether handheld or cart-based. It is an enterprise-wide gateway, bringing the whole platform together as one.

"*I'm so blessed to be working with some of the most talented people in the industry today,*" Jeff beams.

"*Exo is a magnet for talent. We have a dedicated, talented, passionate group of engineers, commercial people, and manufacturing people. Exo has a vision, and it's following it through. It's a little magical, I have to say.*"

Looking ahead to the next five years, he expects the field to be unrecognizably different. He points to the effect the pandemic has had on **lung ultrasound**.

"*Covid accelerated the use of ultrasound to look at the lung by at least five years,*" Jeff reveals.

"*You can use ultrasound to differentiate pneumonia from Covid pneumonia without X-ray. That's remarkable. That's changing the world. That's the trajectory I think we're on. Anyone who's a part of point-of-care ultrasound and technology and doing what we're doing is in a pretty fun place. It's exciting to be even a small part of such a big thing. These opportunities don't come along very often in a lifetime.*"

INTERNATIONAL NEURO IMAGING SUMMER SCHOOL

4TH-9TH JULY 2022

IT'S TIME TO PUT OUR BRAINS TOGETHER

sano

www.neurosummerschool.org

N3S2

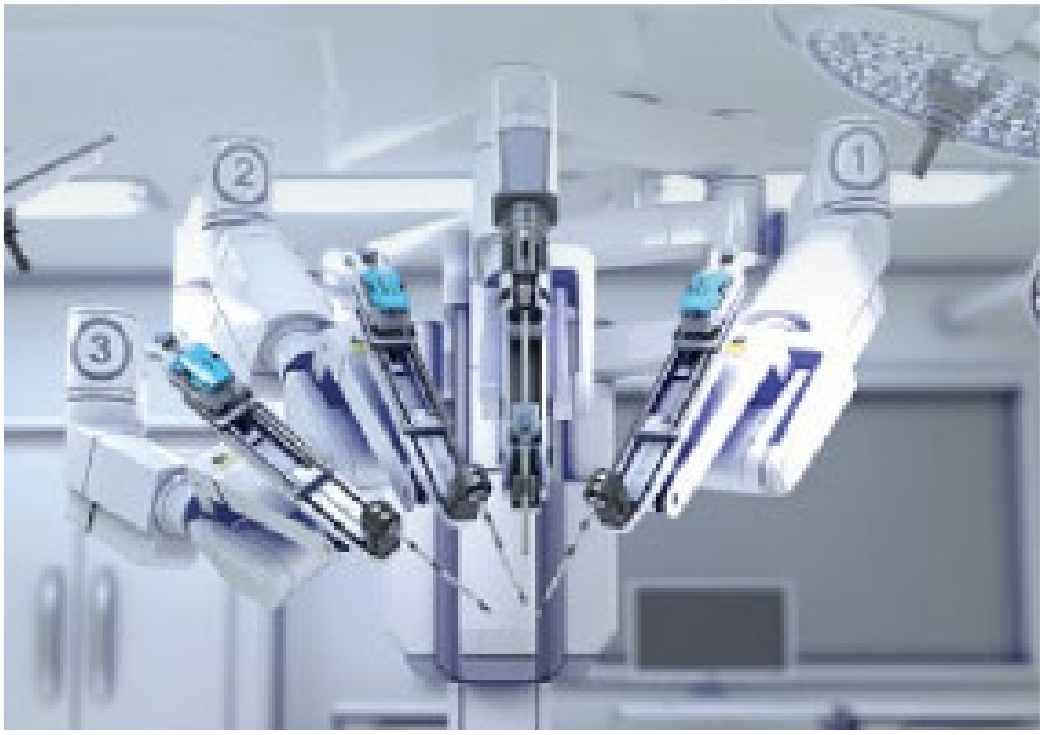

IMPROVE YOUR VISION WITH

Computer Vision News

SUBSCRIBE

to the magazine of the algorithm community

and get also the new supplement Medical Imaging News!

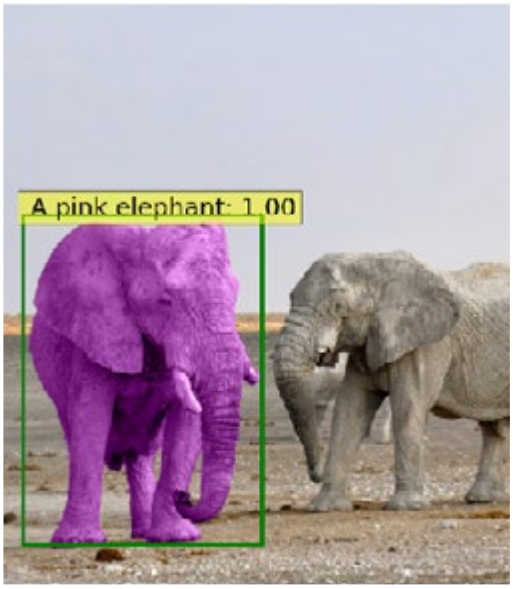

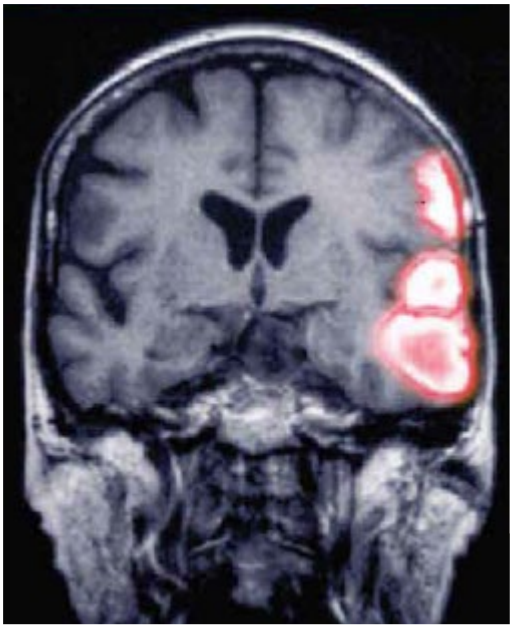

RSIP VISION