DECEMBER 2021

Computer Vision News

The Magazine of the Algorithm Community

MEDICAL IMAGING NEWS

New supplement on page 21

RSIP VISION
Visual Intelligence for MedTech

DILBERT

Computer Vision News

Editor:
Ralph Anzarouth

Engineering Editors:
Marica Muffoletto
Ioannis Valasakis

Publisher:
RSIP Vision

Contact us

Free subscription

Read Previous Magazines

Dear reader,

As a busy year for the community draws to a close, I think we should all take a moment to reflect on just how much has been achieved. Despite the challenges we continue to face, we've seen some awesome events and so much innovation in 2021. I hope we all get the chance to meet again in person at one of the many planned events in 2022. There is so much to look forward to!

Computer Vision News has continued to grow this year, and we are now two magazines in one! We have been receiving great feedback about our new supplement, Medical Imaging News, and this month we have the very first edition of **MedTech Spotlight News**, building on the success of AI Spotlight News in the main magazine. You can find it on page 44.

Also this month, we speak to **Professor Russell Taylor**, one of the founders of MICCAI and a pioneer of Robotic Assisted Surgery. He talks to us about the state of the art in computer-enhanced interventional medicine and his perspective on the future of the field. You don't want to miss his exclusive interview on page 30!

We have an exclusive preview of the upcoming NeurIPS workshop **ImageNet: past, present, and future**, also featuring Olga Russakovsky, one of the original creators of ImageNet. Find out more on page 12.

Last but not least, you're all invited to an exciting webinar called **Emerging Technologies in AI for Surgical Imaging** on December 8. Head to page 45 to get your individual access link.

We have much more to offer this month, so enjoy the reading, and remember to share our link with your friends and colleagues so they can **subscribe for free**!

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

Follow Us

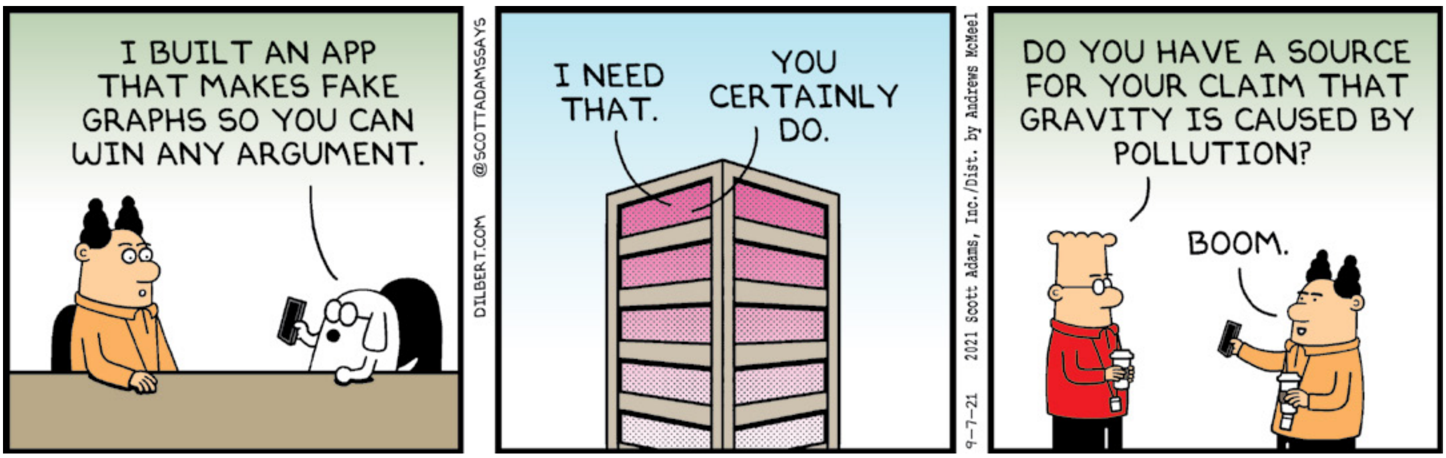

Computer Vision News

04
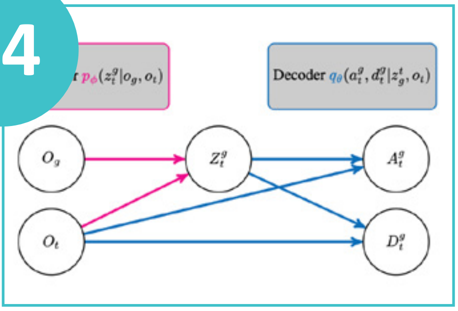

12


14
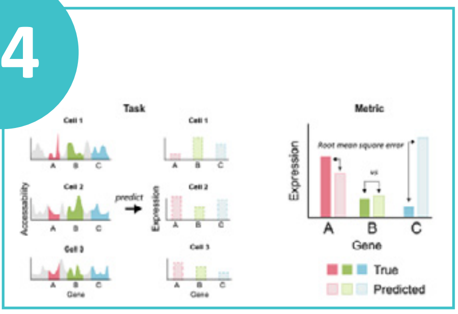

18
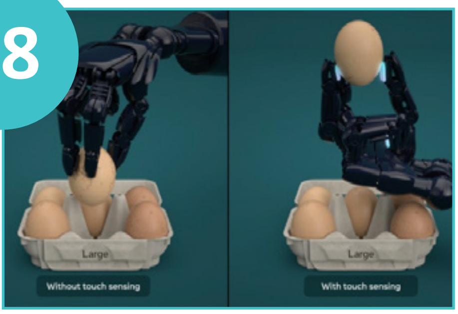

Medical Imaging News

22
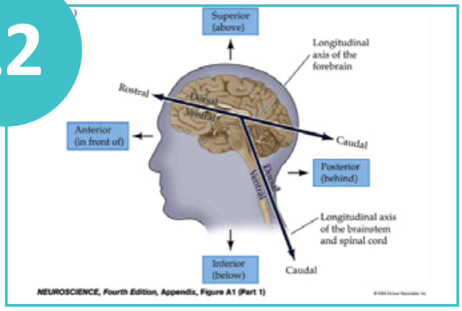

30


36


46


RAPID EXPLORATION FOR OPEN-WORLD NAVIGATION WITH LATENT GOAL MODELS

by Marica Muffoletto (Twitter)

Dear readers, this month we review a paper from UC Berkeley and Carnegie Mellon University which was recently presented at the 5th conference on Robot Learning in London!

We welcome this new computer vision research, entitled **Rapid Exploration for Open-World Navigation with Latent Goal Models** (link here). We are indebted to the authors Dhruv Shah, Benjamin Eysenbach, Nicholas Rhinehart, Sergey Levine for allowing us to use their images to illustrate this review.

Topic

This work focuses on the problem of goal-directed **exploration** for visual **navigation** in novel environments. Fundamental to this work is robustness and adaptiveness to visual distractors. Who needs a robot that can reach a goal only when it's sunny or when there is no obstacle in the way?

A compressed representation of perceptual inputs and goal images is core to this research. As similar state-of-the-art methods, it's based on a topological map learning a distant function and a low-level policy. The main difference with other methods lies in the efficiency and robustness of RECON. While recent works reason about the novelty of a state only after visiting it or need high sample complexity and hence a simulated counterpart, RECON makes use of previous experience in different environments to accelerate learning in the current one using an information bottleneck.

Data

The authors carefully approach the problem of using previous experience to build a robust method, by curating their own dataset made of over 5,000 self-supervised trajectories. This accounts to 9 hours, collected over 9 distinct real-world environments of varying complexity (see Figure below d) Training Environments, e) Unseen Test Environments). The dataset contains measurements from a wide range of sensors, including more and less accurate ones, following the assumption that learning-based techniques coupled with multimodal sensor fusion can provide a lot of benefits in the real world. This dataset is also made available online, contributing to open-source and replicable research.

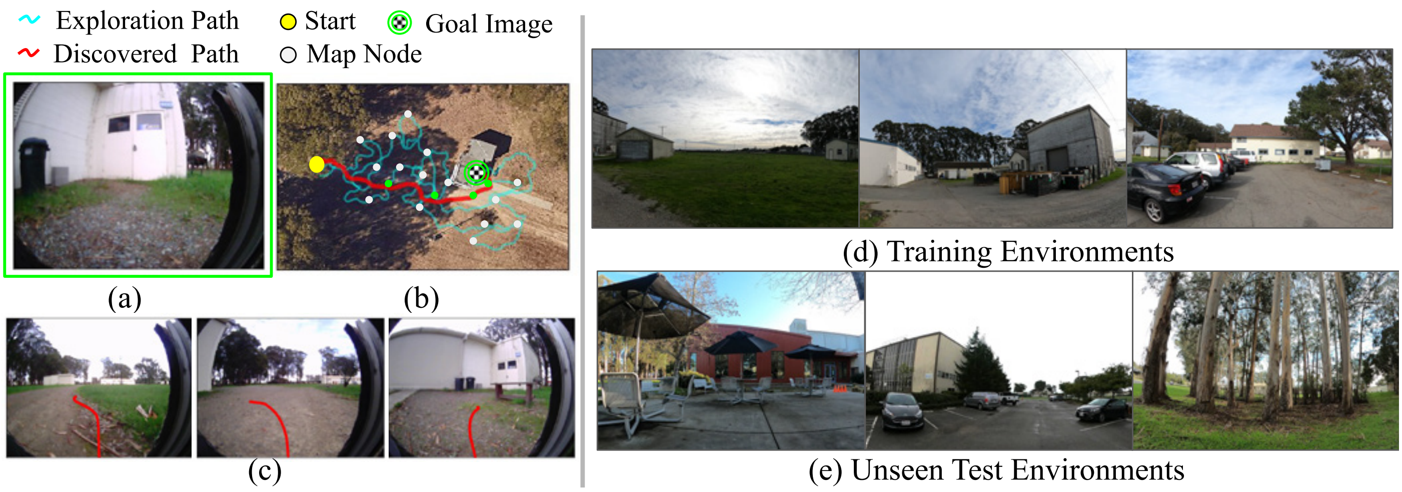

To gather these data, the authors use a time-correlated random walk and a mechanism to detect if the robot is in collision or stuck and then an automated backup manoeuvre that drives the robot out. The collision detectors are further used to generate event labels for the collected trajectories.

Implementation

The method introduced in this paper is called **RECON** (**R**apid **E**xploration **C**ontroller for **O**utcome-driven **N**avigation). It is based on two novel elements:

- An **uncertainty-aware, context-conditioned representation of goals** that can quickly adapt to novel scenes. This is also referred as the latent goal model, which encodes prior knowledge about perception, navigational affordances, and short-horizon control.
- A **topological map** with memory of the target environment.

These serve a two-fold purpose: exploring a novel environment which uses a combination of frontier-based exploration and latent goal sampling with the learned model, and navigation of an explored environment using the topological graph and the learned model. All of this is made possible through an initial step of supervised training using prior experience on previously visited environments provided by the acquired dataset.

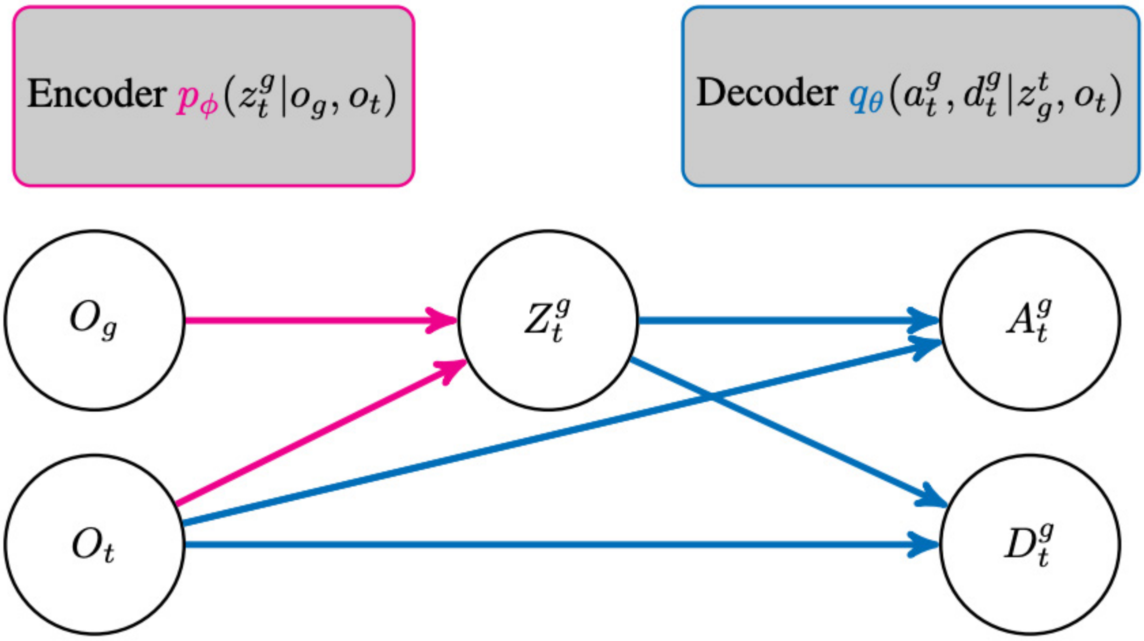

The first element is based on previous work called InfoBot: Transfer and Exploration via the Information Bottleneck which postulates an information bottleneck architecture. The graphical model on the right explicates the architecture: the encoder extracts a representation z_t^g of the goal image o_g, conditioned on the current observation o_t. The encoded representation, which is built such as to only retain relative location of the goal from the context, is then decoded into a prediction of the best action a_g^t and the temporal distance d_t^g to the goal.

The objective of the model consists into maximising the model's ability to predict actions and distances from the encoded representation, and the model's compression of the incoming goal image.

This architecture is used to pre-train a latent goal model on the offline dataset described before, which is used as building block for the second component of RECON: the topological map. This is created during the exploration phase and then employed again in the navigation one, to quickly navigate towards the goal o_g. The map's topological memory is built incrementally with exploration by looking for subgoals, which are represented by the latent variables in the model. Given a subgoal, which is chosen based on the robot's estimate of the goal reachability and its proximity to the frontier, the algorithm executes actions and its distance to the goal is used to construct edges of the topological graph.

More implementation details (hyperparameters and architecture) and algorithmic components which are at the core of the RECON method (exploration techniques) are also described by the authors in great details as part of the supplementary and provided as pseudo codes.

Results

The authors experiment RECON on a mobile ground robot in open-world exploration scenarios. The results are quite exceptional given that they usually outperform state of the art methods in both time and robustness.

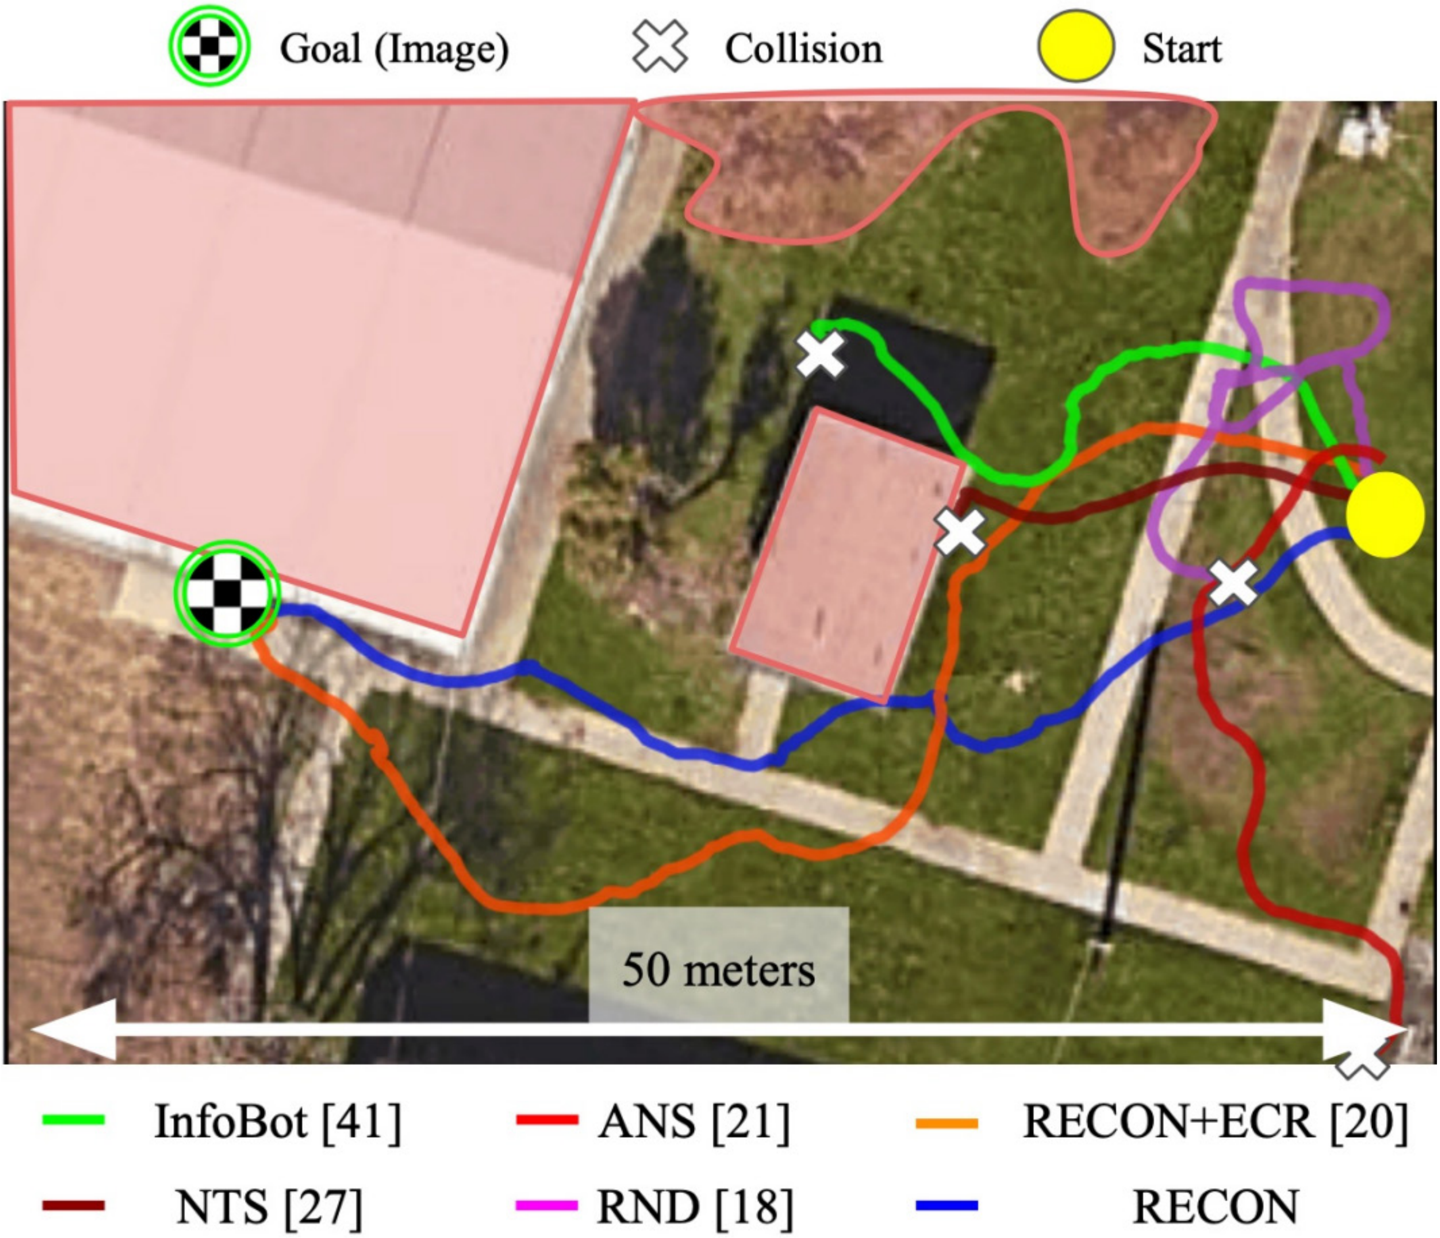

They compare RECON to five baselines, trained on the same 20 hours of offline data, and finetuned in the target environment. RECON outperforms all baselines both in time spent discovering the desired goal and in time spent navigating to the discovered goal using prior exploration. As can be observed below, RECON achieves to navigate to the checkerboard goal, and uses a shorter path taking 30% less time than the only other successful method.

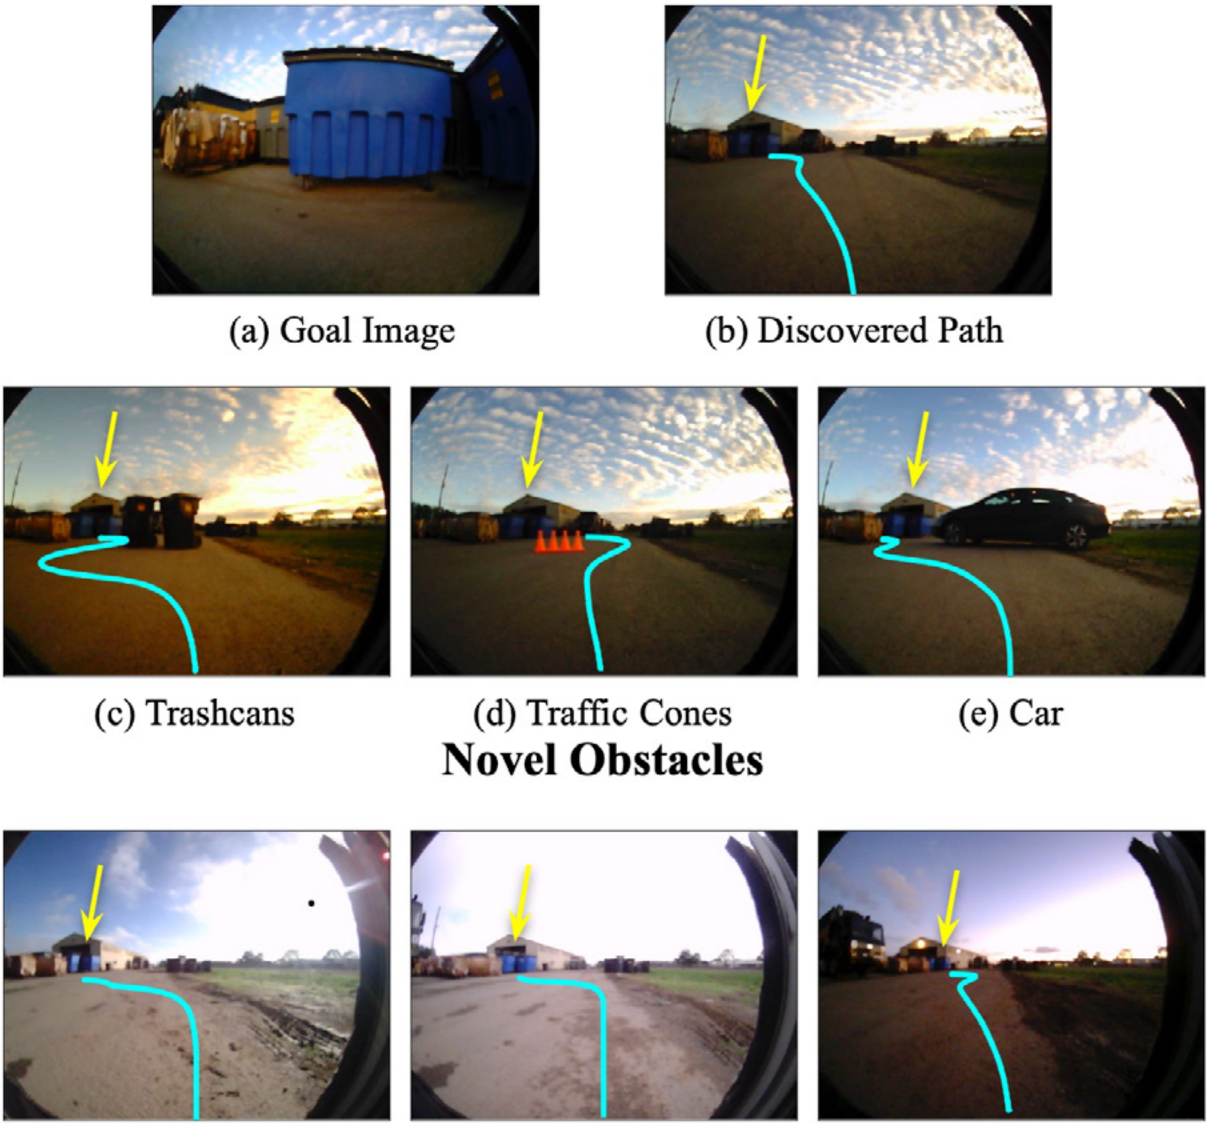

(a) Goal Image (b) Discovered Path

(c) Trashcans (d) Traffic Cones (e) Car

Novel Obstacles

(f) Sunny (g) Rainy/Overcast (h) Twilight

Varying Lighting Conditions

RECON is also proven to be stable amidst previously unseen obstacles and weather conditions. This is warranted by the invariance of the learned representation to such factors. To test this, the authors made RECON first explore a new "junkyard" to learn to reach a goal image containing a blue dumpster, and then they evaluated the learned goal-reaching policy when presented with previously unseen obstacles or lighting conditions. In the figure above, under Novel Obstacles the trajectory slightly changes to avoid the obstacles, while under varying lighting conditions the trajectory is not affected.

Finally, some ablation studies were also led to ensure that the building blocks of the RECON method (the topological graph for memory, rollouts to sampled goals, the information bottleneck) are fundamental to its robustness and performance. This was confirmed by experiments showing that the full algorithm substantially improves on the other variants' timings (1.58-4.58 minutes for exploration time, 2.9-11.4 seconds for navigation time).

Conclusions

This paper represents a great example of research in the robotics field which leverages on recent developments in the computer vision and reinforcement learning fields, to output a perfectly viable and practical application. The authors claim that **RECON can discover goals that are up to 80m away in under 20 minutes** and given its proven robustness to obstacles and outdoors conditions, seems like a reliable algorithm for open-world navigation! We would also love to try it out and test exploration and navigation abilities in a new environment. Meanwhile, we can enjoy the short video below, which shows the method deployed on a Clearpath Jackal ground robot and exploring a suburban environment to find a visual target.

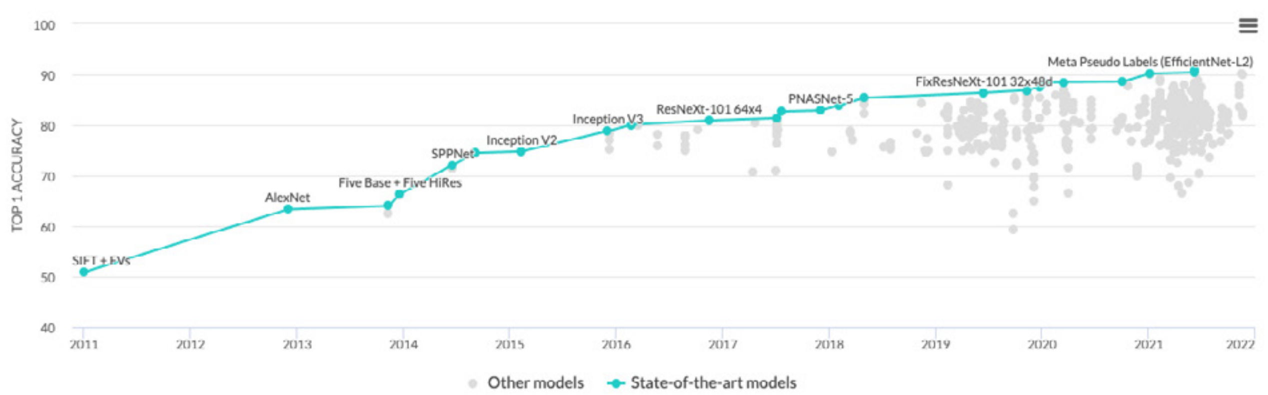

Is ImageNet still competitive? Will it be in the next few years?

The answer on page 12!

Congrats, Doctor!

Muxingzi Li recently completed her Ph.D. at Inria Sophia Antipolis-Méditerranée. Her research interest lies at the intersection of geometry processing, computer vision and graphics, with the goal of enabling compact representations of objects from low-budget image data. She received her B.A. degree from University of Oxford and her MSc degree from KAUST. She has also spent time at Alibaba DAMO Academy. She will start as a researcher at miHoYo. Congrats, Doctor Muxingzi!

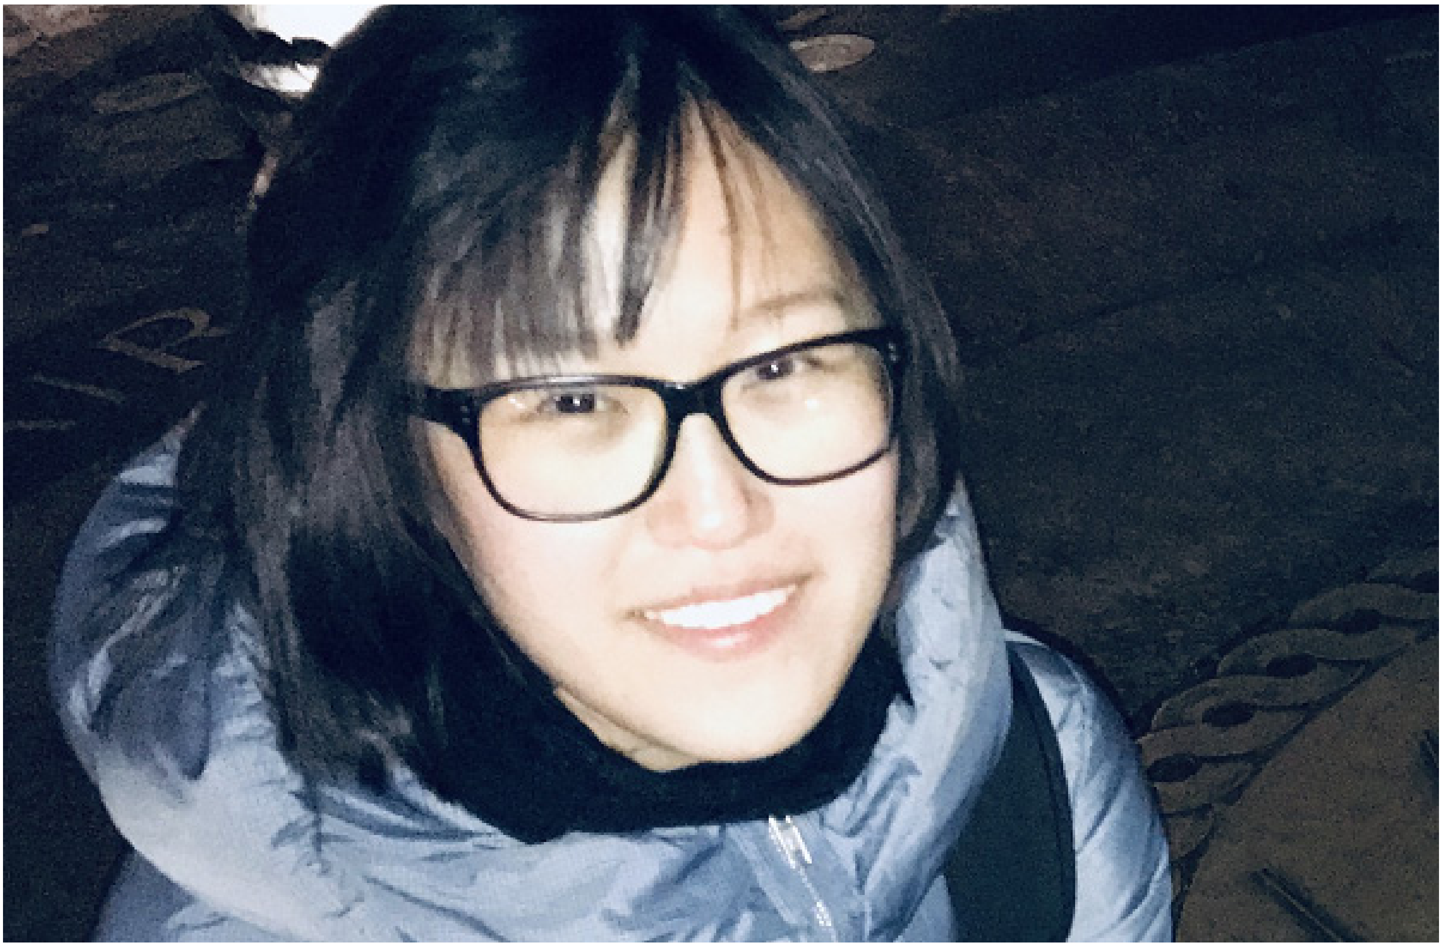

Geometric approximation of urban objects with compact and accurate representation is a challenging problem that concerns both computer vision and computer graphics communities. Existing literature mainly focuses on reconstruction from high-quality point clouds obtained by laser scanning which are too costly for many practical scenarios. This motivates the investigation into the problem of geometric approximation from low-budget image data. Dense reconstruction from a collection of images is made possible by recent advances in multi-view stereo techniques, yet the resulting point cloud is often far from perfect for generating a compact model. In particular, our goal is to describe the captured scene with a compact and accurate representation. We propose two generic algorithms which address different aspects of image-based geometric approximation. The proposed algorithms could be used sequentially to form a pipeline for geometric approximation of an urban object from a set of image data, consisting of an overhead shot for coarse model extraction and multi-view stereo data for point cloud generation.

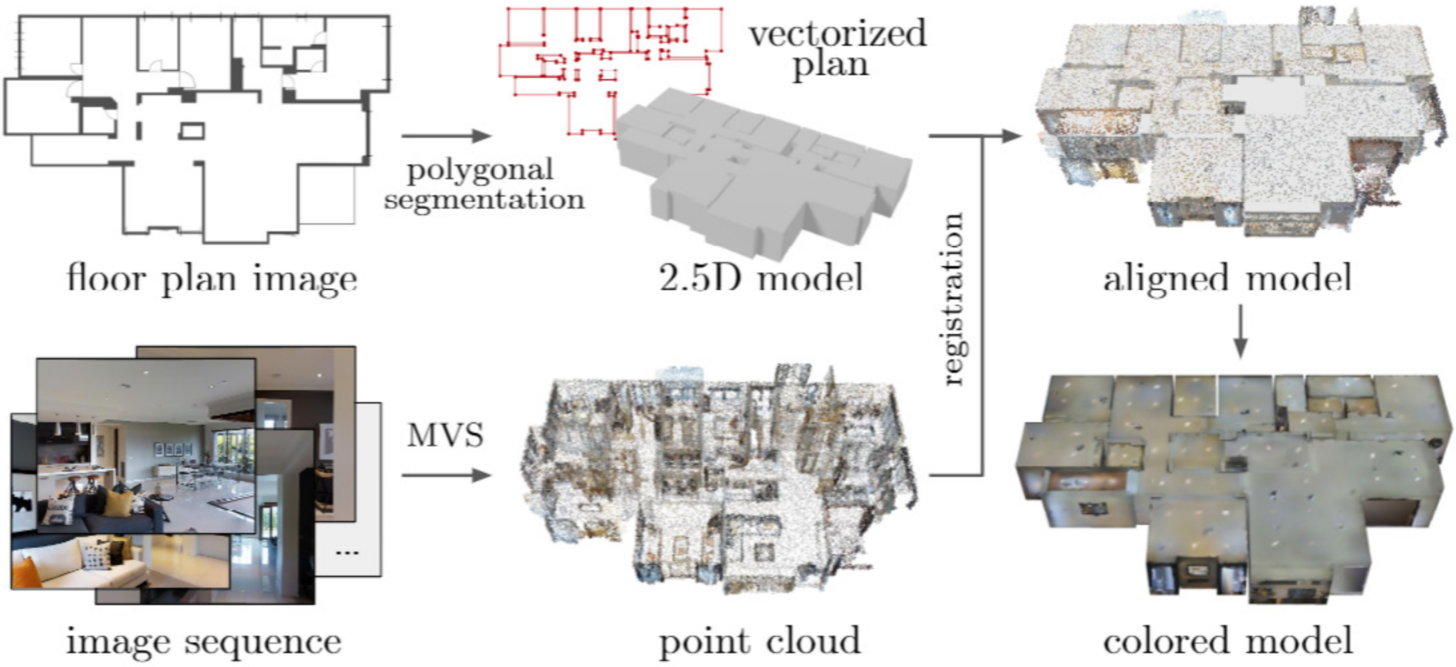

Polygonal image segmentation

We present an algorithm for extracting and vectorizing objects in images with polygons. Departing from a polygonal partition that oversegments an image into convex cells, the algorithm refines the geometry of the partition while labeling its cells by a semantic class. The result is a set of polygons, each capturing an object in the image. The quality of a configuration is measured by an energy that accounts for both the fidelity to input data and the complexity of the output polygons. To efficiently explore the configuration space, we perform splitting and merging operations in tandem on the cells of the polygonal partition. The exploration mechanism is controlled by a priority queue that sorts the operations most likely to decrease the energy. We show the potential of our algorithm on different types of scenes, from organic shapes to man-made objects through floor maps, and demonstrate its efficiency compared to existing vectorization methods.

Alignment of multi-modal geometric data

We present a global registration algorithm for multi-modal geometric data, typically 3D point clouds and meshes. Existing feature-based methods and recent deep learning based approaches typically rely upon point-to-point matching strategies that often fail to deliver accurate results from defect-laden data. In contrast, we reason at the scale of planar shapes whose detection from input data offers robustness on a range of defects, from noise to outliers through heterogeneous sampling. The detected planar shapes are projected into an accumulation space from which a rotational alignment is operated. A second step then refines the result with a local continuous optimization which also estimates the scale. We demonstrate the robustness and efficacy of our algorithm on challenging real-world data. In particular, we show that our algorithm competes well against state-of-the-art methods, especially on piece-wise planar objects and scenes.

IMAGENET: PAST, PRESENT, AND FUTURE

Sangdoo Yun is a research scientist at Naver AI Lab in Korea and is a co-organizer of the first “ImageNet: past, present, and future” workshop at the upcoming NeurIPS conference. He speaks to us about what we can expect from the event.

For more than a decade, **ImageNet** has played a crucial role in advances in computer vision, deep learning, and artificial intelligence. Originally created to train image classifiers, over the years it has become the go-to **benchmark for modern architectures and training techniques**. It has been used to maximum advantage for tasks such as few-shot learning, self-supervised learning, semi-supervised learning, and many more.

Sangdoo and his team at Naver work on fundamental computer vision and machine learning and have been focusing on training a strong and robust vision model using ImageNet.

“We’re very familiar with the ImageNet dataset and have been discussing it amongst ourselves for a while now and thinking about its future,” he reveals. *“We thought, why don’t we organize our own event to get everyone together to talk about it? So many people work on or around ImageNet, but they have no place to meet and discuss ImageNet itself. This workshop will be a great opportunity for all of us to do that!”*

The workshop will be a chance to reflect, regroup, and look to the future. Did we solve ImageNet? What have we learnt from it? What are the remaining challenges? What should the next generation of ImageNet-like benchmarks look like?

Based on these questions, Sangdoo and his team have lined up a diverse range of invited speakers, interesting posters, and spotlight presentations. He hopes participants will gain a new perspective.

Speakers include a great friend of our magazine, **Olga Russakovsky**, one of the original creators of ImageNet, who can speak about its origins. Olga was so kind to give us a hint about her talk: *"PhD student **Kaiyu Yang** and I are talking about our joint work with **Jacqueline Yau, Klint Qinami, Fei-Fei Li** and **Jia Deng** on 'fairness and privacy aspects of ImageNet.' We will reflect on the history of the dataset and the field of large-scale visual recognition and discuss some of our recent work on filtering and balancing the ImageNet dataset (more details here)."*

There will also be a talk by independent researcher **Ross Wightman**, who is very famous on GitHub for creating the **timm package**, which has 15,000 GitHub stars. He is an expert on training ImageNet from the engineering side, so there will be a lot to learn from him.

"All the ingredients of our workshop are fascinating, but the speakers in particular will be so exciting for everyone," Sangdoo tells us. *"We'll also have a panel discussion with all the speakers together, which will be extra special. The audience is sure to learn many insights from these field-leading researchers."*

Sangdoo's fellow organizers are a diverse bunch too, each with a unique perspective on ImageNet. Five, including Sangdoo, are from Naver; two are from Google; and two are from the University of Tübingen in Germany.

Looking to the future of the workshop, Sangdoo has one thing to say: watch this space!

"We've put all our ideas into this year's workshop, so haven't given a lot of thought yet to another one, but we're talking about the possibility of making it annual," he teases. *"We'll be discussing next steps afterwards. The most important thing is the future of ImageNet. Maybe we can hold a **next generation ImageNet workshop** next year... maybe!"*

NEURIPS CHALLENGE - MULTIMODAL SINGLE-CELL DATA INTEGRATION

Daniel Burkhardt is a machine learning scientist at Cellarity, a Boston-based biotechnology company. He is also co-organizer of the Multimodal Single-Cell Data Integration competition at this year's NeurIPS conference. He is here to tell us all about it.

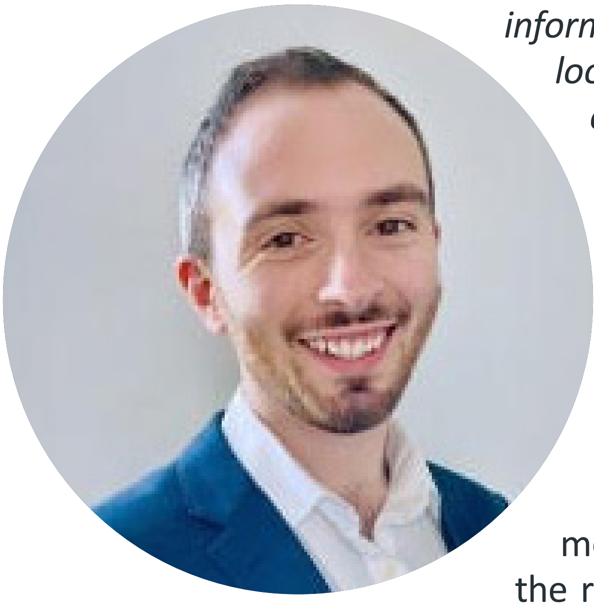

Daniel is part of an organization called **Open Problems in Single-Cell Analysis**, with **Smita Krishnaswamy** from Yale, **Fabian Theis** and **Malte Lücken** from Helmholtz Munich, and **Angela Pisco** from Chan Zuckerberg Biohub. It is a collaboration that started when he was at graduate school. The goal of the organization is to **standardize benchmarking in a domain of life sciences research called single-cell analysis**.

"The fundamental problem we're trying to address is an emerging data modality called ***single-cell transcriptomics****, or* ***single-cell omics*** *more broadly,"* he tells us. *"There's been a massive influx in data over the past six years creating an incredible opportunity to look at how individual cells function. This data exists at a scale unheard of before in the life sciences, particularly in terms of genetic information. We can start to look at the expression of tens of thousands of genes and hundreds of thousands to millions of cells."*

Alongside this increase in scale and resolution, there is a need for new methods to analyze the data, and a concomitant need to compare these methods and understand the relative merits of one over another.

Open Problems has organized this competition to take some of the standard machine learning approaches that you can compare a whole bunch of different methods to, such as centralized benchmarking and formalized metrics, into this field of single-cell analysis, where for the most part those kinds of comparisons are performed in a one-off manner, but they want to create a living benchmark.

"One of the things we really struggle with in the field of single-cell analysis in general is understanding what our ground truth is," Daniel explains. *"There are different ways you can do that, but one of the things we're taking advantage of is a new kind of technology called* ***multimodal single-cell data****. The idea is you can measure both the DNA and RNA, or both the RNA and the*

protein in individual cells, at this 100,000 million cells scale. One of the nice things is multimodal data has a built-in ground truth because you can say I measured one part of this system and then in a paired manner I measured this other aspect."

This throws up questions such as: If I only expose you to one, can you predict the other? What if I give you many sets of both, and I want you to match them? What if I have a set of annotations I like about the data and I want you to learn a joint embedding into a lower-dimensional space where certain properties are preserved?

The goal of the competition is to introduce some of these tasks for the first time for multimodal single-cell data, and to incentivize broad-scale competition from statisticians, computational biologists, and machine learning practitioners to bring these communities together in a way that has not been done before. The team have aimed to provide as straightforward a task as possible to enable people who do not necessarily have a life sciences background to participate. They came up with tasks that were intuitive for someone with a machine learning background, while introducing some of the background knowledge to explain what people are looking at and what the different tools are.

The idea for the competition formed a couple of years ago when a paper called **'Eleven grand challenges in single-cell data science'** came out.

"I really liked the paper," Daniel tells us. *"One of our co-organizers was a co-author of it, so there's no bashing here, but it was very generic and not formalized, so it wasn't clear what the problem was you needed to find a solution for. In the field of mathematics there are well-defined problems for which we don't have effective solutions. How do you mathematically quantify what you would like people to solve?"*

Part of the goal of this effort was to take some of that formalization and mathematical rigor and apply it to something which might just be called a challenge. However, it is more than that – it is a huge open problem that takes a lot of work to solve, so the team wanted to reflect that in the name of

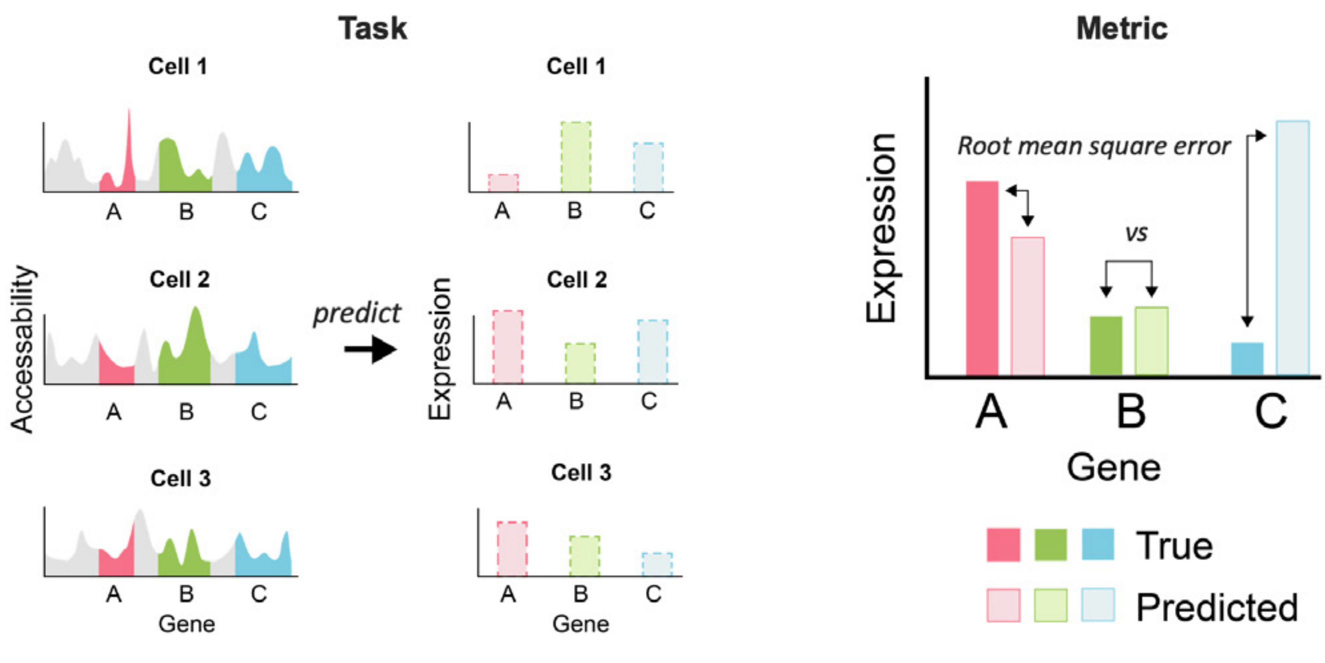

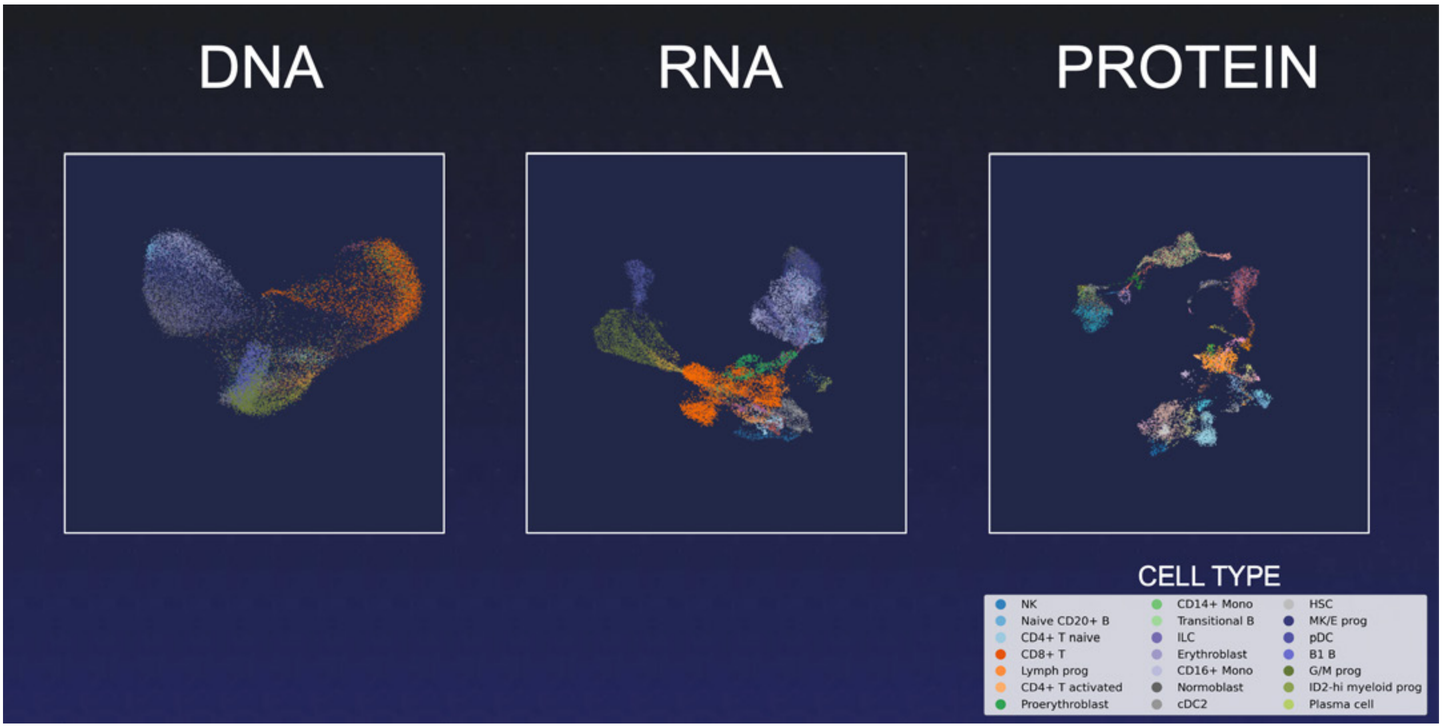

the initiative: Open Problems.

Daniel says this work excites him because he uses these kinds of methods every day, and it can be frustrating trying to figure out what tool to use for a given task. Also, he just loves bringing people together.

"Academia can get very competitive – it's funny I'm saying academia is competitive and one of the ways I'm going to solve that is through a competition!" he laughs. *"But there's a lot of good-natured effort by people who want to solve these problems and want to know if they're doing well or not compared to some other tool. It makes it a lot easier if you've already come up with the task you're trying to solve because then people have a clear understanding of how well they're performing and in what context. When are they better than another tool? When are they not? There's a lot of interest from many kinds of people to want to come together."*

Generating the dataset for this competition was a huge team effort that took place across four different academic institutions around the world with 37 people involved with generating the data, processing it, analyzing it, and making sure it was properly formatted. Part of what makes this so special is that the effort brings together wet lab biologists who have no experience with machine learning and machine learning practitioners who have little prior experience with single cell biology. Just generating a benchmarking dataset is a landmark effort. How did they manage to get so many people motivated towards a common goal?

"A lot of people just love science and they're really interested in biology!" Daniel responds. *"People are just innately curious about what is going on with this brand-new kind of data and they all understand the value of having these methods properly benchmarked."*

The biggest barrier to this work has been the funding. The team managed to secure financial sponsorship from both **Cellarity** and the **Chan Zuckerberg Initiative**.

"It'd be remiss if I didn't acknowledge that

the Chan Zuckerberg Initiative has been a huge sponsor of this effort so far," Daniel points out. *"Their support towards this kind of open scientific effort is non-zero-sum. They have a long track record of doing these kinds of events that bring people together. Continuing to get to be involved with them is a huge privilege as an organization."*

Daniel is already looking at planning other competitions in the future. The team are in preliminary talks about an event next year around a different set of tasks. They very much hope they will get to **run the competition in person next time**. Like many other events in recent times, **NeurIPS** is virtual this year.

When he is not doing this open-source community work, Daniel works as a machine learning scientist at Cellarity, a young biotech company under the Flagship Pioneering umbrella. Flagship has been in the news recently as the investor behind Moderna.

"What we're really trying to do is take these kinds of high-dimensional datasets that give you a holistic view of what's happening at the level of cells and tissues, to understand how we can focus on the effect of drugs on biological systems," he explains. *"We do this by not focusing on a single protein, which is what the traditional pharmaceutical pipeline does, but rather looking more holistically at the cell as a complex dynamic system. I spend a lot of time thinking about methods to understand what's actually happening when that system is perturbed in a certain way. How can I quantify the changes the cell undertakes and predict what kind of intervention would allow us to go from a diseased state to a healthy state to enable a patient to go about living their life?"*

Daniel and his co-authors will also be presenting a paper at NeurIPS called **A sandbox for prediction and integration of DNA, RNA, and proteins in single cells**.

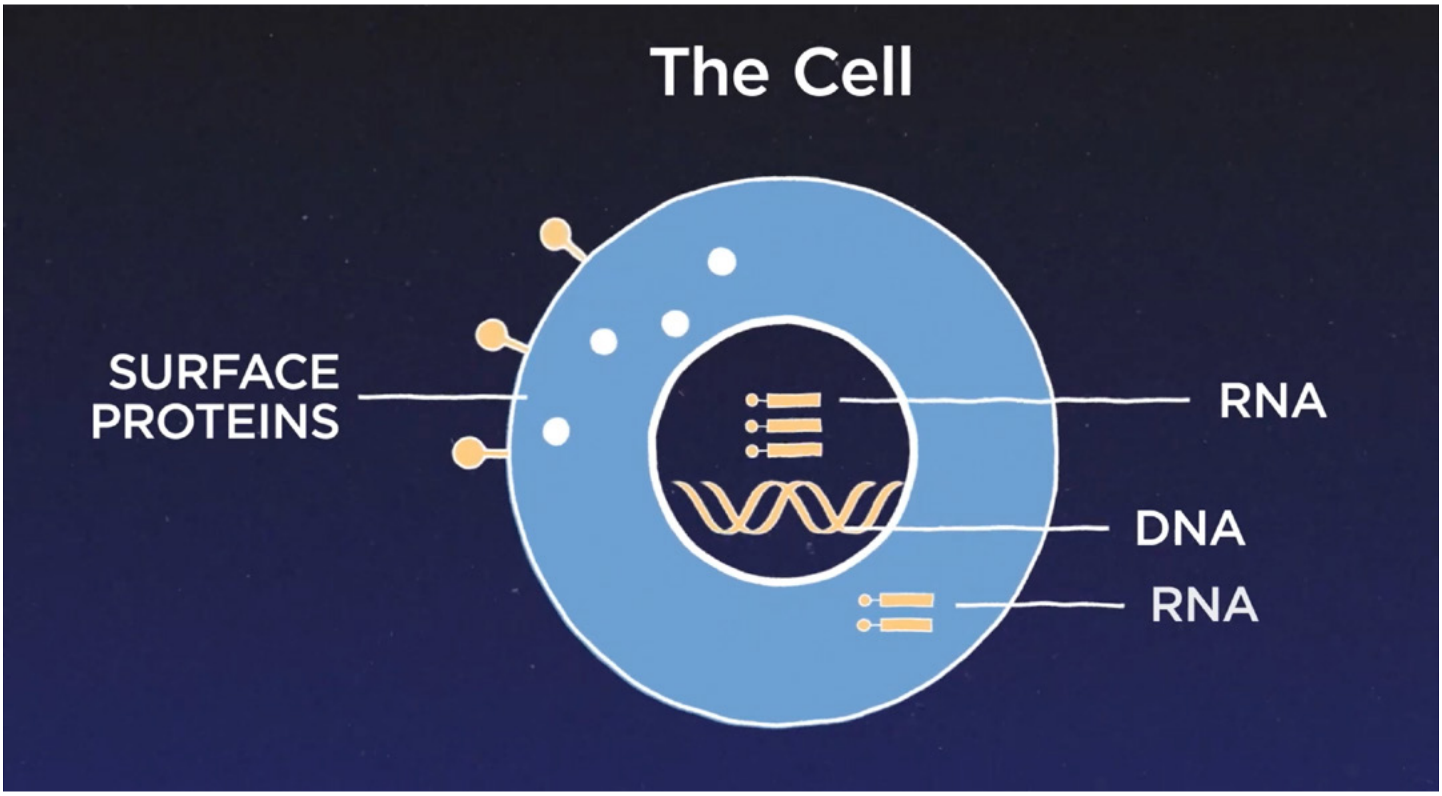

Computer Vision News has found great new stories, written somewhere else by somebody else. We share them with you, adding a short comment. **Enjoy!**

'Cooperative Perception' Could Help Autonomous Vehicles See Pedestrians Around Corners

New technology developed by researchers in Australia could help autonomous vehicles steer clear of pedestrians and cyclists without a direct line of sight. Three researchers from the Australian Centre for Field Robotics at the University of Sydney claim they can "substantially improve the efficiency and safety of road transportation." One of the "intelligent roadside units" was even able to "see" a "visually occluded pedestrian behind a building" seconds before it would have otherwise. The image here shows a cooperative-perception enabled vehicle detecting a cyclist behind a bus. Read More

A New Camera Can See Through Almost Anything, Including Human Tissue and Bones

A different camera: this too has the ability to "see" through things. Developed by Northwestern Engineering researchers, this new HR camera can see around corners and through human skin and even bones. It also has the potential to image fast-moving objects such as speeding cars or even the beating heart. This non-line-of-sight (NLoS) imaging camera is so precise that it even captures the tiniest capillaries at work. It intercepts the scattered light coming from an object to reconstruct the inherent information about its time of travel and reveal the original object. Read More

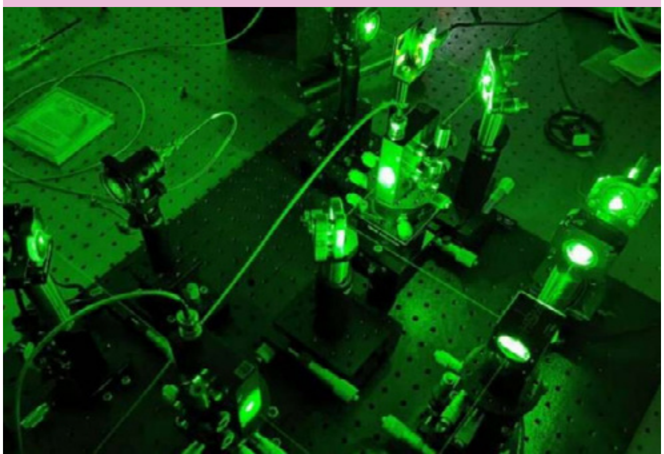

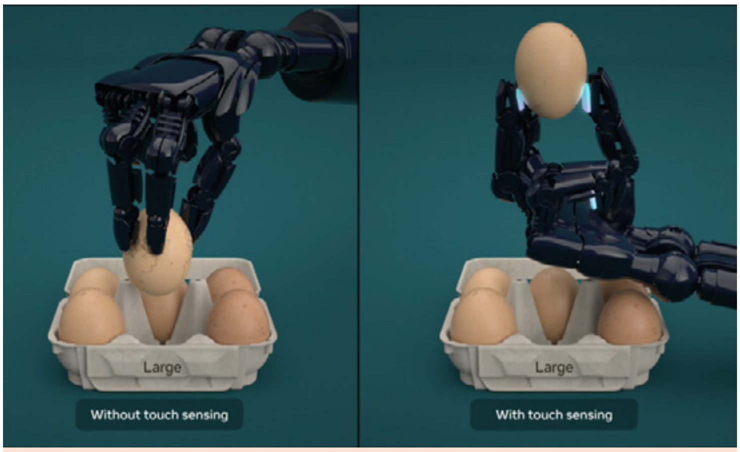

Facebook battles the challenges of tactile sensing

Facebook (or Meta if you want) wishes to lead the rush for touch sensing and has developed a full set of hardware and software that provide this function, an emerging need in the robotics field: it aims to understand and replicate human-level touch in the physical world, providing information about temperature, texture, weight and more. Together with the Carnegie Mellon University, they presented a synthetic "skin" called Reskin, which will be an inexpensive and replaceable solution for long-term use, employing an unsupervised learning algorithm to help auto-calibrate the sensor. Read More

If you missed the last month's Computer Vision News

Exclusive Interview with Raquel Urtasun: "We are actually going to change the world!"

Best Paper Award at ICCV 2021 - Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows

Honorable Mention Award at ICCV 2021 - Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields

COMPUTER VISION EVENTS

Hamlyn Winter School Online 6-10 December	NeurIPS Virtual-only 6-14 December	Bay Vision Webinar Online 8 December
The AI Summit New York, NY 8-9 December	Face and Gesture Recognition Virtual Event 15-18 Dec	AAAI Conf. on AI Virtual 2-9 Feb
AI Ethics Summit S.Francisco, CA 17-18 February	Deep Learning Summit S.Francisco, CA 17-18 February	
ICARA Aut., Robotics and Apps Prague, Czech Rep. 18-20 February		
SPIE Medical Imaging S.Diego, CA 20-24 February		

SUBSCRIBE!

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

FREE SUBSCRIPTION

(click here, its free)

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

Fill the Subscription Form

it takes less than 1 minute!

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

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

MEDICAL IMAGING NEWS

DECEMBER 2021

I am a researcher in Medical Image Processing and Artificial Intelligence.

Find out about me on page 46!

Introduction to Neuroscience Image Processing using Jupyter Notebooks

IOANNIS VALASAKIS, KING'S COLLEGE LONDON
@WIZOFE

Hello again and welcome!

Are you excited about neuroscience and coding with brains? This month it's all about this. Let's dive in. Without further ado, I am presenting you NIFTI, NIBABEL but also a short introduction (to be continued in the next issue) of General Linear Models (GLMs).

We'll start with the required setup, introduction and presentation of data and how to use a Python Jupyter notebook to access and process them.

There's also an amazing package introduced here, it's called BData from Kamitani lab (where amazing info were taken from) and here's used as a machine learning data format which you can utilize in your later models. Enjoy and feel free to let me know if you use it in your project and how!

Setup

```
!pip install numpy matplotlib
!pip install nibabel
```

```
!mkdir data
```

```
!curl https://s3-eu-west-1.amazonaws.com/pfigshare-u-files/27995283/sub01_sesanat-
omy_T1w.nii.gz -o data/sub-01_ses-anatomy_T1w.nii.gz
!curl https://s3-eu-west-1.amazonaws.com/pfigshare-u-files/27995286/sub01_sesper-
ceptionNaturalImageTest01_taskperception_run01_bold.nii.gz -o data/sub-01_ses-per-
ceptionNaturalImageTest01_task-perception_run-01_bold.nii.gz
```

```
!ls data
```

Introduction

In the neuroscience community, MRI images are often saved, processed, and shared in **NIfTI-1** format. The file extension is `.nii` (uncompressed) or `.nii.gz` (gzipped). A single NIfTI-1 file can contains either 3-D (spatial) or 4-D (spatial + temporal) MRI image.

More information: https://nifti.nimh.nih.gov/nifti-1/

Visualization of NIfTI images

- Use **nibabel** to handle NIfTI images: https://nipy.org/nibabel/
- `nibabel.load` load an NIfTI image as a nibabel image instance.
- To obtain the image as a Numpy array, run np.asanyarray(<img>.dataobje).

```
import os

import nibabel
import numpy as np
import matplotlib.pyplot as plt
```

Anatomical MRI image (3-D NIfTI image)

First, we are going to see an anatomical MRI image.

```
# Load an anatomical MRI image (3-D NIfTI image)
img_t1 = nibabel.load('data/sub-01_ses-anatomy_T1w.nii.gz')

# Extract the image as an array
data_t1 = np.asanyarray(img_t1.dataobj)
data_t1.shape
```

The loaded image should have shape of (256, 256, 208).

You can check the spatial direction of each dimension of the image array by the following code.

```
nibabel.aff2axcodes(img_t1.affine)
```

N-th letter represents the incremental direction of N-th dimention of the image array.

- **A/P**: Anterior/Posterior
- **S/I**: Superior/Inferior
- **L/R**: Left/Right

See also: [Anatomical terms of location - Wikipedia]

Thus, the dimension-direction correspondense of the loaded image is:

- 1st dim: anterior to posterior
- 2nd dim: superior to inferior
- 3rd dim: right to left

Terminology on brain anatomy

- Anterior vs posterior
- Superior vs inferior
- Dorsal vs ventral

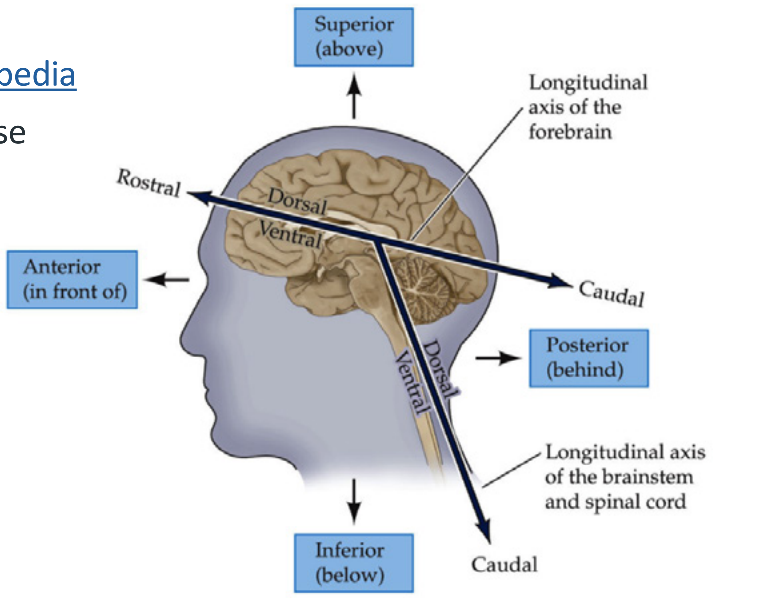

***NEUROSCIENCE, Fourth Edition*, Appendix, Figure A1 (Part 1)**

Image from "Neuroscience, 4th ed." Sinauer, 2007

You can check the size of each voxel (each element of an MRI image) by the following code (unit: mm).

```
img_t1.header.get_zooms()
```

Now, let's visualize a sagital slice of the loaded MRI image by display the slice at the center of the left/right axis.

Sagittal, coronal, and axial slices

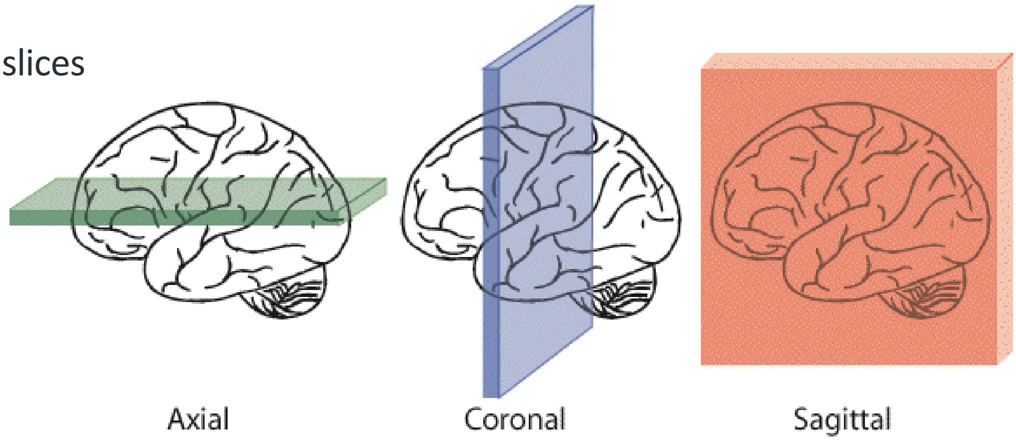

Image from https://users.fmrib.ox.ac.uk/~stuart/thesis/chapter_3/section3_2.html

```
# Visualize the middle sagittal plane
plt.imshow(data_t1[:, :, data_t1.shape[2] // 2].T, cmap='gray')
```

Exercise 1

Visualize axial and coronal planes of the loaded image.

```
# Axial slice
plt.imshow(<Complete the code>, cmap='gray')
```

```
# Coronal slice
plt.imshow(<Complete the code>, cmap='gray')
```

Functional MRI image (4-D NIfTI image)

Next, let's take a look at a 4-D MRI image.

```
# Load functional MRI images (4-D NIfTI image)
img_fmri = nibabel.load('data/sub-01_ses-perceptionNaturalImageTest01_task-percep-
tion_run-01_bold.nii.gz')

# Extract the image as an array
data_fmri = np.asanyarray(img_fmri.dataobj)
data_fmri.shape
```

You can see that the loaded image is 4-dimensional. The first three dimensions are spatial and the 4th is temporal. So, the image contains whole brain fMRI data from 239 timepoints. The fMRI image at each timepoint is called volume.

Let's check the dimension-direction correspondense and voxel size.

```
nibabel.aff2axcodes(img_fmri.affine)
```

Thus, the dimension-direction correspondense of the loaded image is:

- 1st dim: right to left
- 2nd dim: anterior to posterior
- 3rd dim: inferior to superior

```
img_fmri.header.get_zooms()
```

This indicates that the voxel size is mm and each volume was collected at each 2 sec (*).

* Strictly speaking, it takes 2 seconds to collect the data of one volume.

The image size (i.e., the number of voxels) are much less than the anatomical image we investigated above because the spatial resolution of fMRI images is lower than that of anatomical images (i.e., the voxel size of the fMRI image is much larger than the anatomical image).

Let's visualize a sagittal slice of the first volume in the loaded MRI image.

```
# Sagittal plane
plt.imshow(data_fmri[data_fmri.shape[0] // 2, :, :, 0].T, cmap='gray', orig-
in='lower')
```

Exercise 2

Visualize axial and coronal slices of the first volume in the loaded fMRI image.

```
# Axial slice
plt.imshow(<Complete the code>, cmap='gray')

# Coronal slice
plt.imshow(<Complete the code>, cmap='gray')
```

In 4D fMRI, each voxel is a time series of BOLD signals. Let's select one voxel and display its time series.

```
# Time course of a voxel

voxel_index = [33, 78, 42]

plt.imshow(data_fmri[voxel_index[0], :, :, 0].T, cmap='gray', origin='lower')
plt.plot(voxel_index[1], voxel_index[2], 'r+', markersize=12)

resp = data_fmri[voxel_index[0], voxel_index[1], voxel_index[2], :]
plt.plot(resp)
```

Exercise 3

Calculate temporal mean of the fMRI responses of each voxel, and plot the temporal mean as a brain image (in either sagittal, coronal, or axial plane).

```
# Complete the code

data_fmri_mean =
plt.imshow()
```

GLM

This notebook provides brief instruction describing how to run general linear model (GLM) analysis on fMRI data with Python.

Setup

Check if you have not installed the following. If not, you can safely skip them.

```
!pip install numpy
!pip install nipy
!pip install nilearn
!pip install git+https://github.com/KamitaniLab/bdpy.git
```

Download data

```
!mkdir data
```

```
# Subject 2
!curl https://ndownloader.figshare.com/files/28089525\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-02_task-localizer_bold_preproc_native.h5
!curl https://ndownloader.figshare.com/files/28089570\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-02_anatomy_t1.nii.gz
!curl https://ndownloader.figshare.com/files/28089528\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-02_template_native.nii.gz
```

```
# Subject 3
!curl https://ndownloader.figshare.com/files/28089534\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-03_task-localizer_bold_preproc_native.h5
!curl https://ndownloader.figshare.com/files/28089582\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-03_anatomy_t1.nii.gz
!curl https://ndownloader.figshare.com/files/28089537\?private_link=3bd9a1c29f-19649c8c0d -o data/sub-03_template_native.nii.gz
```

Import modules

```
import os
from itertools import product

import bdpy
from bdpy.mri import export_brain_image
import matplotlib.pyplot as plt
import numpy as np
from nipy.modalities.fmri.glm import GeneralLinearModel
from nipy.modalities.fmri.experimental_paradigm import BlockParadigm
from nipy.modalities.fmri.design_matrix import make_dmtx
import nibabel
from nilearn import plotting
```

fMRI data

In this example, we run GLM analysis on fMRI data collected in the **higher visual areas' localizer experiment**. The aim of the experiment is to identify visual areas related to processing of complex visual information such as objects, faces, or scenes. During the experiment, a subject was required to look at an image presented in the scanner. The image was either of them.

- Object images (IntactObject)
- Scrambled version of the object images (ScrambledObject)
- Face images (IntaceFace)
- Scrambled version of the face images (ScrambledFace)
- Scene Images (IntaceScene)
- Scrambled version of the scene images (ScrambledFace)

Example images:

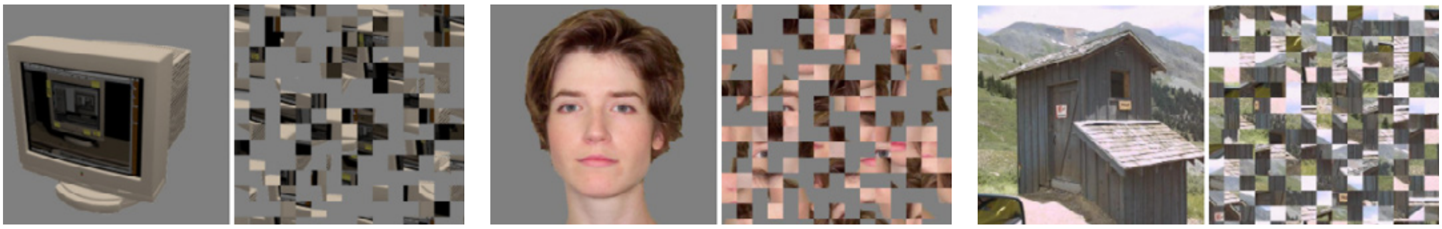

By taking contrast between intact and scrambled images of a particular domain (e.g., IntactObject vs. ScrambledObject), we can get a brain region related to the visual processing of the domain (e.g., lateral occipital complex for object vision).

Here is a list of typical brain regions identified with the experiment.

- LOC (lateral occipital complex; related to object recognition; Malach et al., 1995)
- FFA (fusiform face area; related to face perception; Kanwisher et al., 1997)
- PPA (parahippocampal place area; related to scene perception; Epstein & Kanwisher, 1998)

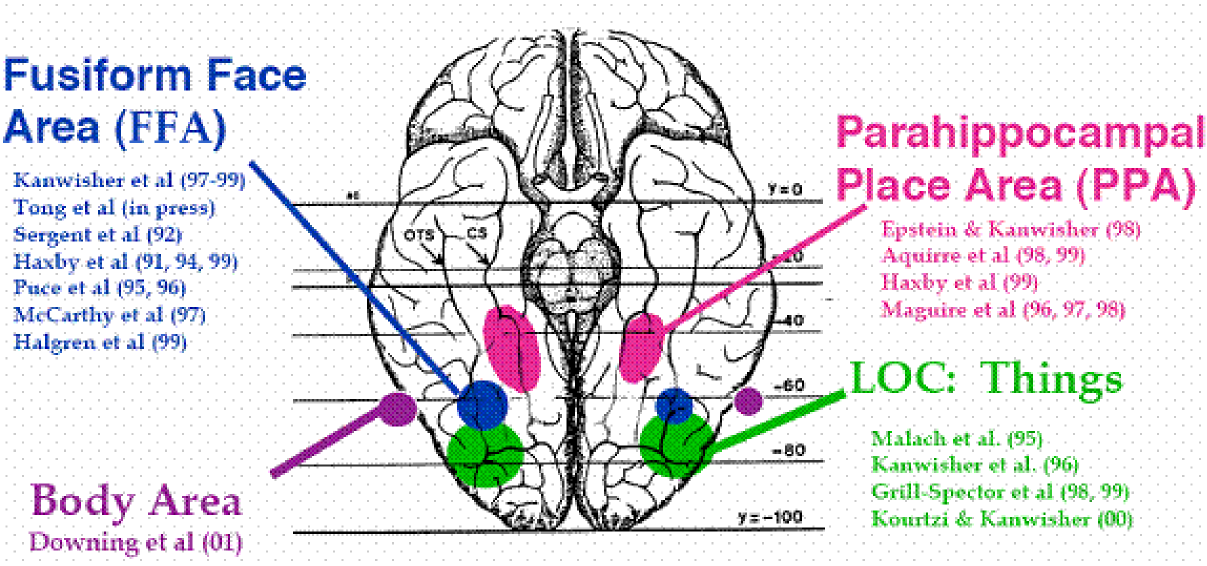

From: Prosopagnosia | Fusiform Face Area Functions and Impact https://prosopagnosiabytmayo.wordpress.com/

An experiment that identifies a brain region (region-of-interest; ROI) based on the function of the brain (e.g., how voxels are activated by certain stimuli) is called a *functional localizer experiment*.

Experiment information

- TR: 3 sec
- 100 volumes/run (300 sec/run)
- 8 runs

```
Blank (24 s)  Stim. (15 s)  Stim. (15 s)  Blank (15 s)
Blank (6 s)
+----------+  +----------+  +----------+  +----------+
+----------+
|          |  |  Intact  |  |Scrambled |  |          |
|          |
|          |  |  Object  |  |   Face   |  |          |  ... (repeated 6 times) ...
|          |
+----------+  +----------+  +----------+  +----------+
+----------+
```

- Stimulus presentation: different images in the category were flashed every 500 ms with blank.

fMRI data format in this example

The fMRI data are saved as *BData*, machine-learning analysis oriented data format developed in Kamitani Lab. [bdpy] is required to read the BData.

```
bdata = bdpy.BData('data/sub-02_task-localizer_bold_preproc_native.h5')
```

```
# Get fMRI data
fmri_data = bdata.select('VoxelData')
fmri_data.shape  # volumes x voxels
```

So the fMRI data are composed of 800 volumes and 32028 voxels.

```
# Get labels
labels = bdata.get_label('block_type')
np.unique(labels)
```

The data have nine events. Three of them ('PreRest', 'InterRest', and 'PostRest') are 'rest' event in which no visual stimulus was presented. The rest event is not inc luded in the GLM model.

In the other six events ('IntactFace', 'IntactObject', 'IntactScene', 'ScrambledFace', and 'ScrambledObject'), a visual stimulus was presented. These 'task' event should be included in the GLM model.

The labels are used to create *task regressors*, which model the signal changes caused by the experimental conditions (e.g., stimuli).

```
# Get run numbers
runs = bdata.select('Run')
runs.shape  # samplex x 1
```

This is an array including run numbers, used for creating run regressors, which explicitly model run-wise fluctuations of the fMRI signals.

To be continued...

Wrapping up!

Keep looking at the answers and feel free to contact me if you have any questions till next month! Amazing work from Kamitani lab where tutorials and the bData package were written! Thanks a lot

Take care for yourself and the people around you! See you next month.

If you missed the last month's Medical Imaging News

...ents in
...l Surgery:

...tely measure
...ocedures can
... outcomes. It
...e to use sizing
... in orthopedic
...plant fitting is
...mo-implants, in
...zing is achieved

...cedures use video
...g a 3D module –
...g – is standard in
...he 3D perspective
...ension to the image
... spatial orientation
...ocedure.

... is the baseline for
...ate measurements
...ng well established
...nera calibration (e.g.,
... coordinate system can
...n the image, providing
...ng of the field-of-view.
...e of stereo camera is
...dly in many different
... is integrated in many
...s. The ability to view and
...depth and size of objects
...ssibilities during surgery.

There are several benefits for this application.

1. **Standard Measurements:**

a. many procedures require **accurate tissue excision**. In orthopedics, it is vital to saw the bone accurately to fit the implant without causing limb length discrepancy. During soft-tissue excisions, to avoid perforation of nearby organs or blood vessel it is essential to know how deep into the tissue the tools are. Gastric bypass procedures often require tedious measurements of the gastrointestinal tract, which can be executed efficiently using stereoscopic measurements.

b. Diagnostic procedures often rely on measurements to **assess severity**. Cancer staging is determined by the size of the tumor. This can be accurately measured, provide adequate staging, and reduce complications due to misdiagnosis.

c. Another aspect where accurate measurements surpass current practice, is **surgical precision**. If all incisions are performed using a defined measurement it ensures that surgeries are performed in a similar manner, making the procedure repeatable.

2. **Buffer Zones:**

a. When removing tumors, healthy tissue from around the tumor is also removed to verify a "negative margin" and prevent recurrence. Verifying that excision occurs at a constant distance from the tumor is challenging, therefore an **objective measurement** tool can assist in this task. Additionally, this buffer can be overlayed on the image given a manual segmentation of the tumor.

3. **Temporal/Spatial Measurements:**

a. During stitching, applyin... too little or too much tension o... the sutures can have devastating effects. Assessing tension based on images only is difficult, and there is no tactile feedback in RAS. Strain is defined as:

$$\epsilon(t) = \frac{L(t) - L(0)}{L(0)}$$

b. Where L is the length of the suture as it is continuously measured throughout the stitching process. Positive strain means the suture is under tension, whereas negative strain means the suture is loose. **Upper and lower strain limits can be defined** to provide proper feedback to the surgeon.

b. Other procedures can als... benefit from detection of di... **changes over time**. A... ...uire total ...

...implem... the lack of ... **improving procedu... increasing patient safe...**

Read more about Robotic Assisted Surgery: https://www.rsipvision.com/surgical-robotics/

Medical Imaging R&D - 3D Measurements in Robotic Assisted Surgery

AI Research Paper - Anomaly Detection in Medical Imaging with Deep Perceptual Autoencoders

Anomaly detection in medical imaging with deep perceptual autoencoders

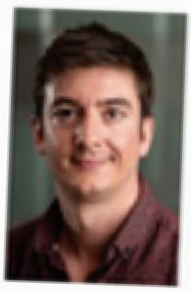

Hi everyone, how are you doing? Let's have another article to read. This month's review is "Anomaly detection in medical imaging with deep perceptual autoencoders" by Nina Shvetsova et al.

This is a pre-print which was uploaded this month, but it already has 15 citations and you can easily find it on Arxiv, to read it in full.

Anomaly Detection

Anomaly detection is a task with significance, especially in the deployment of machine learning models. The knowledge of a "a normal" data sample would be used to compare -in a sense of a ground truth- to an "abnormal" one. To identify less often occurrences is another application where anomaly detection is useful and with the method proposed in this paper, the efficacy of autoencoders for anomaly detection is utilized.

has been extensively studied. Domains including fraud detection, cyber-intrusion detection, as well anomalies in videos and financial analytics are some of the latest fields where research is active. Distribution based methods try to predict if the new example lies in the high-probability area or not.

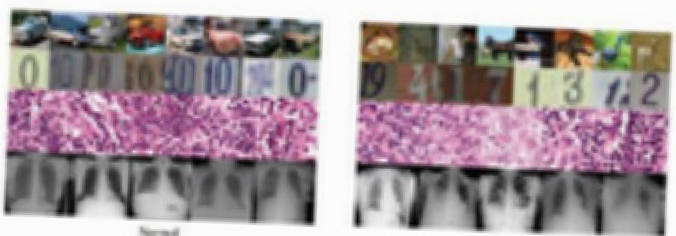

Examples of dataset with abnormality vs normality comparisons. CIFAR10 dataset images, with SVHN digit dataset in the row below, followed by histopathological and chest X-ray images from the NIH dataset.

DEEP PERCEPTUAL AUTOENCODER

Autoencoder-based approaches rely on the fact that they can learn shared patterns of the normal images and then restore them correctly. The key idea of the ...

PROFESSOR RUSSELL TAYLOR

Professor Russell Taylor is the John C. Malone Professor of Computer Science at Johns Hopkins University, with joint appointments in mechanical engineering, radiology, and surgery. He is also director of JHU's Laboratory for Computational Sensing and Robotics. With a career spanning over 50 illustrious years, he speaks to us about his experiences to date and his hopes for the future.

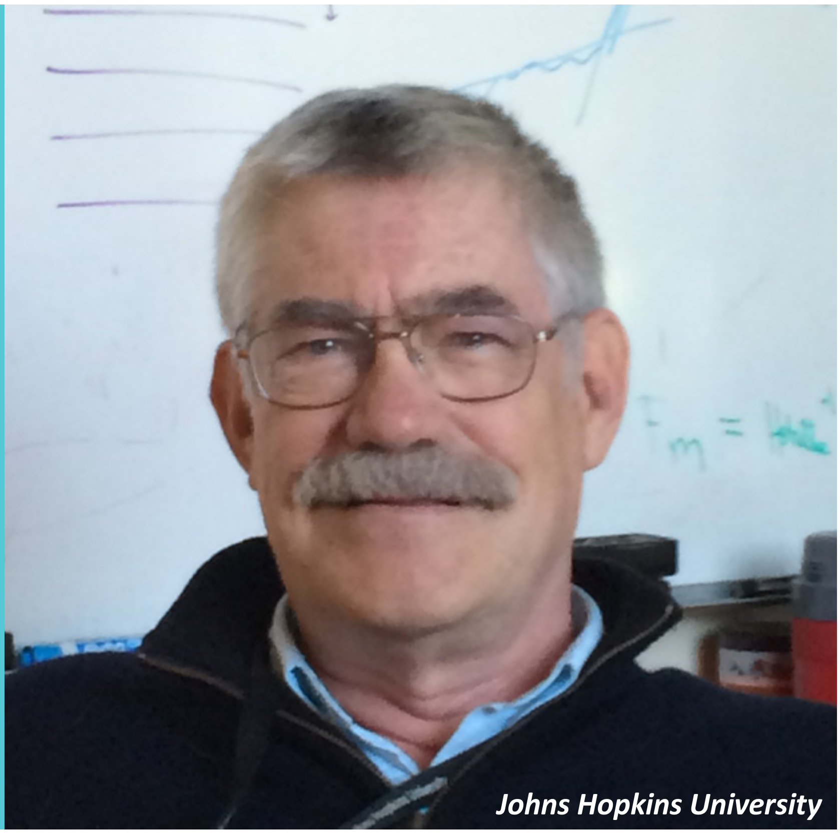

Johns Hopkins University

It is a pleasure to speak to you Professor Taylor. Can you tell us more about your work?

For over 30 years, both here at Johns Hopkins and before that at IBM Research, I have been focused primarily on developing a three-way partnership between humans, technology – machines, robots, sensors – and information, to improve surgery and interventional medicine. Through this partnership, you can improve the intervention for an individual patient by making it safer, more efficacious, and less invasive. You can save that information and relate what you did to outcomes, using statistical methods to improve your treatment processes for the next patient. It really is the synergy between this patient-specific control loop, and this broader machine learning-based use of the information that I think will be driving interventional medicine and the progress of medicine in the coming years.

What is the current status of works in this field?

In the last 10 years, there has been an exponential increase in the amount of research and commercial activity in this area. Particularly surgical robotics, which seems to get the most publicity. Personally, we are very active. We have a large center here at Johns Hopkins that grew out of a National Science Foundation funded Engineering Research Center that we established shortly after I moved to JHU. There are very, very large centers in this broad area of computer-enhanced interventional medicine around the world in Europe and Asia, as well as in North America now.

Are there any advancements that have made you go “wow” in the last few years?

What has been most impressive has been the strength of machine learning methods very broadly to improve the use of information in this field. We are beginning to see something I thought for a long time we would be seeing more of, which is that machines are beginning to be more actively involved in interventional procedures. The first autonomous medical robots were radiation therapy machines that moved around the patient and shot beams of X-rays to kill a tumor – and you hoped not do too much damage elsewhere. The system we developed at IBM Research and which was then commercialized, Robodoc, did the machining step to prepare bones for orthopedic implants.

With surgical robots there are really two key questions. The first is, how can I tell the machine what it is supposed to do in a way it can understand? Second, how can you be sure it can do what you have told it to do safely and properly and it is not going to do something else? At extreme ends of the systems, these are fairly straightforward. The stereotactic systems, like radiation therapy and bone preparation systems, have planned everything out from a CT image or other set of medical images, and there may be levels of autonomy in the planning process – it is typically a human-machine partnership there – but the actual doing of it is very straightforward. The only thing you have left to do is to register all the images and the coordinate systems of these plans to the physical patient. Then the system just does it, perhaps with feedback during the process in case of patient motion or changes.

The dominant paradigm today in medical robots is teleoperation. The machine moves its instruments in the same way I am moving my control handles. The very successful da Vinci surgical robots are an example of that. The commanding and what the robot is supposed to do is very straightforward. There is still a lot that has to be done and the machine has to make a bunch of decisions internally, but they are very straightforward. Now, if I want the robot to take over and do a step while I am doing interactive surgery, then the robot and the surgeon both need to have an agreement as to what the environment of the patient is. You need richer situational awareness. You need things like real-time modelling, real-time computer vision, much better understanding of surgical tasks and how you interact with humans and all of that.

“… how can I tell the machine what it is supposed to do in a way it can understand?”

This is probably the first human clinical case of the Robodoc system in 1992

How can we achieve this?

This is where the machine learning comes in, and just generally the much greater power of computers, making it possible to have much richer environmental models for this situational awareness. There is a danger with that as well that machine learning methods need to learn more humility. They do not know what they do not know. One of the challenges is, how can the machine recognize when it is in a situation it has not been trained for? We are beginning to see some research in that area, but I think it is a key barrier.

What can you see happening next?

I see there being a gradual increase in the capabilities of these systems either remotely, like teleoperated systems, or another way of controlling these robots which is more like hand-over-hand control. If the robot and the surgeon are holding a tool and the surgeon pushes on the tool, the robot can feel that and move accordingly. But because it is a robot, and it does not have hand tremor it can also refuse to move in a certain direction. It can enforce a safety barrier, or control can be shared in various other ways between the surgeon and the robot.

In some cases, the robot might be adding a palpating motion to some guidance a surgeon is using, say with a joystick or hand-over-hand control. In other cases, the surgeon might guide the robot to a place where it can start and take over and do something on its own, which is what happened with the system we developed at IBM all those years ago. The surgeon used hand-over-hand guiding to help get the robot into the position where first it could do some registration, but then they would hand-over-hand guide it to a position where it was safe for the robot to begin machining, and then it would take over and do that task on its own while the surgeon supervised.

I think you are going to see a lot of examples

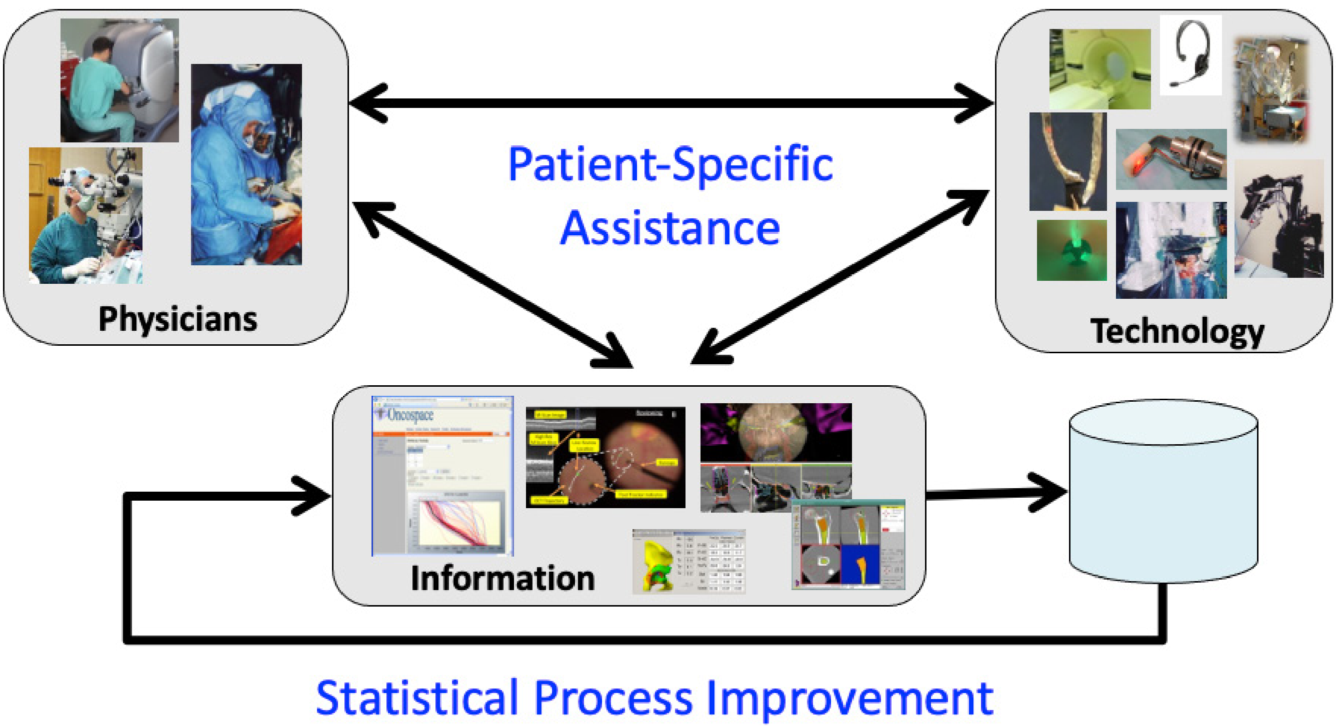

Copyright Russell Taylor - used with permission

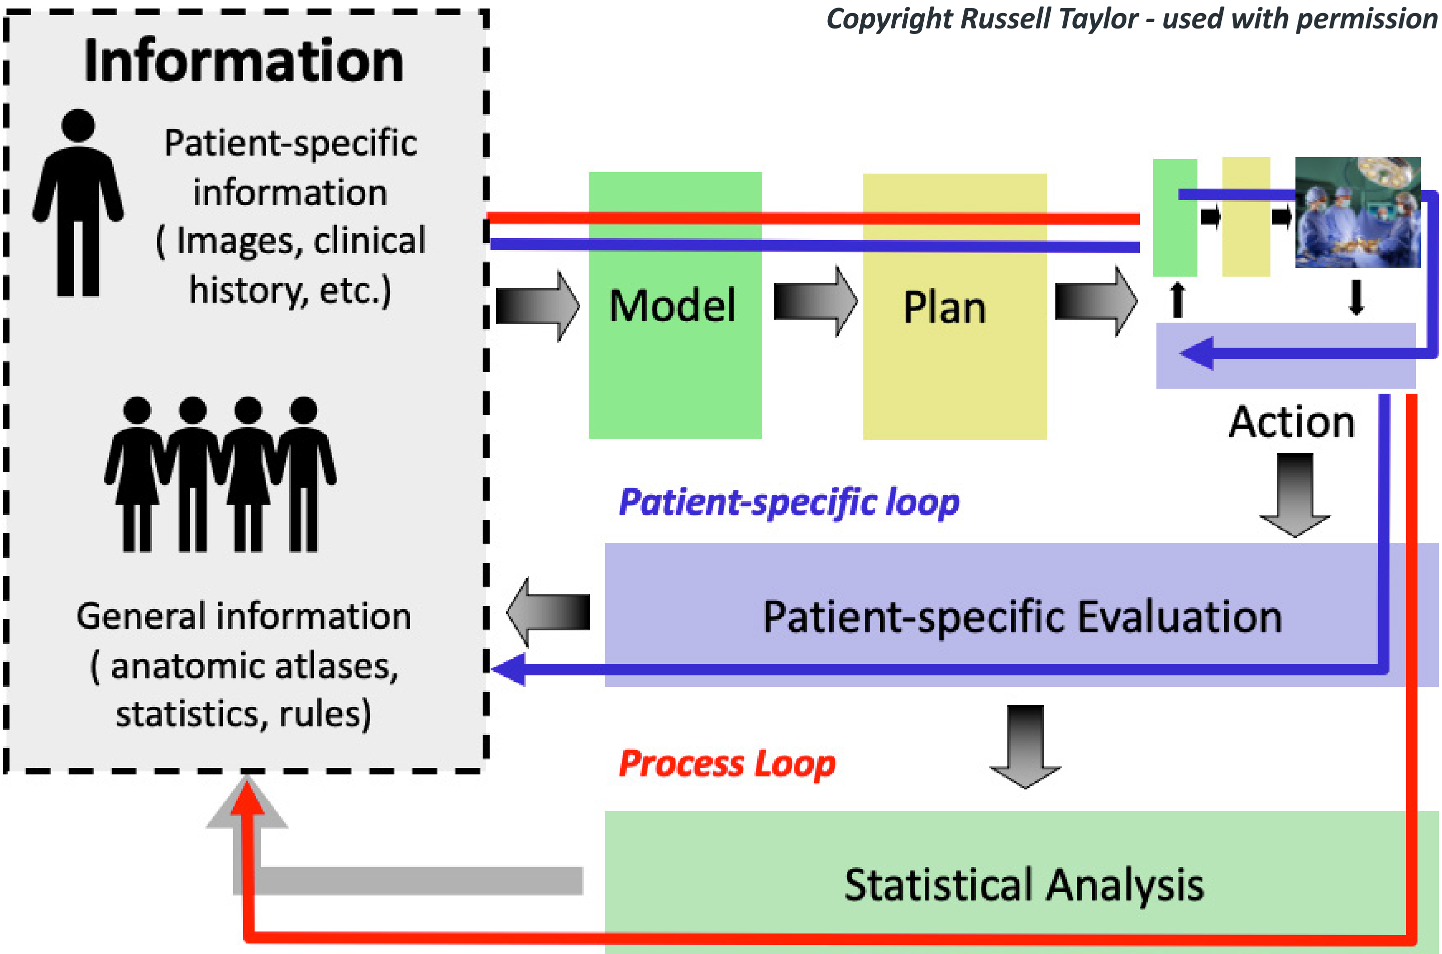

of that over time. There is a lot of interest in suturing, but personally I am a little less excited by that than with some of the other things you might ask a robot to do for you.

What would the other tasks be?

There are a number. Some are maybe easier. Retraction, placement of therapy, tissue resection and debridement, there are a whole bunch of things like that. Suturing is a complex skill, so if you can do that, you can do a lot of other things. There are also tools that help there. The obvious one is a stapler, where you do not suture, you staple. Intelligent retraction might be another one.

I suppose there still is a lot of space for advancement in R&D in this field...

Yes, of course. Some technology is just now beginning to be mature enough to be commercialized, but there is a huge development gap to go from something that is proof of concept to clinical deployment. There are also things we do not know how to do and that is across the board. From the point of view of academics, what I was just telling some students the other day is it is very important to understand the whole problem. To understand in this three-way partnership what is the advantage you are trying to provide to a surgeon, or to a patient? Then to achieve that there are some things that are simply good engineering that we must do. There are some things frankly we do not yet know how to do. It is at that boundary between not knowing how to do it and it being good engineering to do it that I think the most productive academic research occurs.

What is the biggest advancement that

you hope in our lifetime we will be able to attain?

In my lifetime? Well, I am an old guy, but there are a number! If you could get a surgical robot that was able to be a really effective assistant, that would be one thing. Another would be where a robot could do tasks that currently you would expect a skilled resident to do. That is something I would really hope to see. Also, more on the engineering side, I would like the system to be able to look over a surgeon's shoulder or another robot's shoulder and be able to follow along what is being done, detecting when they are about to get into trouble and raising the alarm.

How far are we from that?

Oh, at least four Einsteins! [*he laughs*] I would not want to speculate. When I was a graduate student at the Stanford AI Lab – and that was 50 years ago – I remember we were beginning to do research on robotics, and we had a hand-eye project. People were already talking somewhat derisively about the Summer Vision Project, where all of computer vision could be done in a summer, but we also had the notion with what we knew at the time that we could program a robot to assemble a chainsaw engine. That practical problem drove lots of research. Some of which led us to have a much better appreciation for manufacturing engineers and led to advances in understanding how to do programming to take advantage of sensing to make a precise, but not geometrically accurate robot able to do some of these tasks. That experience and many others since then have led me to be careful about predicting the future. Obviously, the past is much easier to predict!

What advice would you give to a young student now, like you were at Stanford 50 years ago, who is starting his or her career and wants to have an impact?

If you want to get into my field, there are a few things. You obviously need to develop deep knowledge in an area of computer science, which is itself becoming somewhat more specialized, but it is also very important to develop broader knowledge as well. We are dealing with systems that have mechanical components, electronic components, sensors, information, programming, human interfaces. Also, get yourself into an environment where you can work with end users. Ideally, a place where you can go into the operating room fairly frequently to talk to physicians and try to understand their problems. You do not need to go to medical school in my opinion. You will quickly

"I have been focused primarily on developing a three-way partnership between humans, technology – machines, robots, sensors – and information, to improve surgery and interventional medicine."

Working with former student Kevin Old's head-and-neck robot that he developed as part of his PhD thesis.


Johns Hopkins University

learn a lot of medicine is about vocabulary, but it is important to get some domain knowledge as well, if you can, and be open to asking questions.

The other thing is what are sometimes referred to as the Heilmeier questions. At IBM Research, Ralph Gomory had a similar set of questions he would ask. What do you want to do? What is your goal? How will you know when you have succeeded? How will you measure success? Also, who cares? Why is that important? Then you can begin to worry about how you are going to go about it and what you do not know. That is general advice whatever field you are going into but understanding that is really important.

Finally, after a challenging couple of years where we have all been kept apart, do you have a message for the MICCAI community in advance of our planned reunion in Singapore next year?

I was one of the founders of the MICCAI Society and the growth has been remarkable. We have to be careful not to go off in being so enamored of one new technique, such as deep nets, that we lose sight of what our larger goals are. We need to begin to look at the complete path of information in medicine. It is medical image computing and computer-assisted intervention and often you see this tension between MIC and CAI communities (Medical Image Computing and Computer Assisted Interventions), but they are beginning to come together more. The CAI folks need to understand the information they can get from imaging and the imaging folks need to focus on what more they can do with information other than get a nice picture or a nice analysis or something. I am a great believer in MIC and CAI together. That was the whole point. MICCAI was formed from the union of three conferences, and it is wonderful that we are continuing that legacy.

Challenges in Biopsies

There are various types of urological cancers, the most common ones are **prostate and kidney cancer**. In cancer, early and accurate detection significantly improves survival rate and reduces treatment cost. Similar to other cancers, detection is performed using a biopsy – inserting a needle to the suspicious tissue and sampling it. Often there is uncertainty regarding exact positioning, so the tissue is penetrated several times to verify adequate sampling.

The biopsy is performed with ultrasound guidance – real-time imaging for needle tracking and tissue visualization. Ultrasound limitations include poor image resolution and high variability among users, making it very difficult to ascertain the exact position of the needle and the correct scatter of samples.

Current diagnosis protocol when suspicion of cancerous tumors in the urological tract arises, requires a CT or MRI scan. These scans provide high-resolution 3D data of the anatomical regions, which allow visualization of the lesions within the urological tract.

By simply adding AI you can get a superior procedure

Segmentation and 3D reconstruction from CT/MRI

Recent developments in computer vision and AI allow accurate segmentation of desired organs and lesions within. Neural networks can be trained to segment the organs within the urological tract, and even sub-organs within them. Also, lesions or tumors can be segmented, resulting in a 3D model of the patient, with the interesting regions segmented and highlighted. This model is a map of the patient.

Biopsy Planning

The biopsy requires prior knowledge of where the organ is located (e.g. prostate), where the suspected lesion is within the organ, and what is the least harmful direction to get to the lesion (without damaging nearby structures). In real-time the physician uses experience and ultrasound guidance to perform this task accurately, but all this information exists in the CT/MRI data!

Based on the scan and the segmentations, the best path from the skin to the lesion can be determined, automatically or by the physician. The system can then extract the exact point of penetration on the skin, the ideal angle of the needle, and the distance from skin to the lesion which the needle needs to penetrate. The physician can use these directions as a baseline for the biopsy and be certain that the needle is pointing in the right direction.

Additional modules can use Ultrasound or X-ray for needle tracking. The probe position can be tracked using stereo cameras or 'sensors on the probe, and this position can be used to verify needle position. Needle distance from skin can be calculated using the Ultrasound or X-ray calibration and needle tracking, so that the physician knows exactly how deep the needle is penetrating. All these tools can improve the procedural accuracy and reduce time and error-rate.

All this information exists in the CT/MRI data!

The benefits of this module, aside from the obvious clinical advantage, is the relatively simple implementation. It requires minimal hardware additions for needle tracking, and it utilizes the available means – a PC - to improve the procedure. By simply adding AI you can get a superior procedure. **More projects in AI for urology here.**

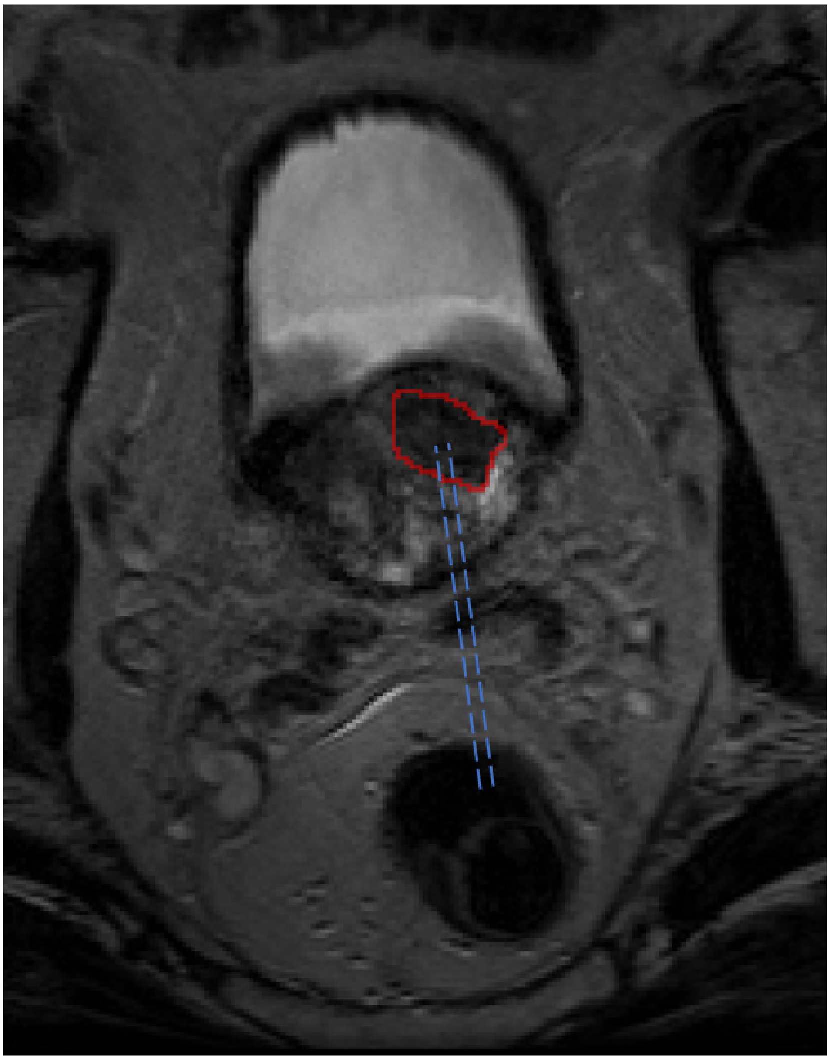

Jannis Fischer is co-founder and CEO of Positrigo, a MedTech startup that has been developing an exciting new technology for the early diagnosis of Alzheimer's disease. He tells us more about this revolutionary system that is set to take the medical imaging industry by storm.

Alzheimer's disease affects millions of people worldwide and costs the world hundreds of billions of dollars each year. Work on medication continues, with several new treatments on the horizon, but early diagnosis is key to ensure medication is given at the first opportunity to ensure the maximum effect.

The best way to diagnose Alzheimer's is via a brain scan using a technology called **Positron Emission Tomography (PET)**. Amyloid plaque deposition is a very strong risk indicator for Alzheimer's. If you do not have any of those plaques in the brain, you have no risk of Alzheimer's in the next 10 years. If you do, then it should be monitored because those plaques are always there when Alzheimer's develops, and **they are there 10 to 20 years before the first cognitive symptoms**.

Currently, PET scans are very expensive and require large devices that take up a lot of space. A few years ago, Professor **Alfred Buck** from **University Hospital Zurich** approached **Positrigo** about the need for a smaller, simpler, more affordable system. He suggested they work on a solution.

"PET imaging is the gold standard for the

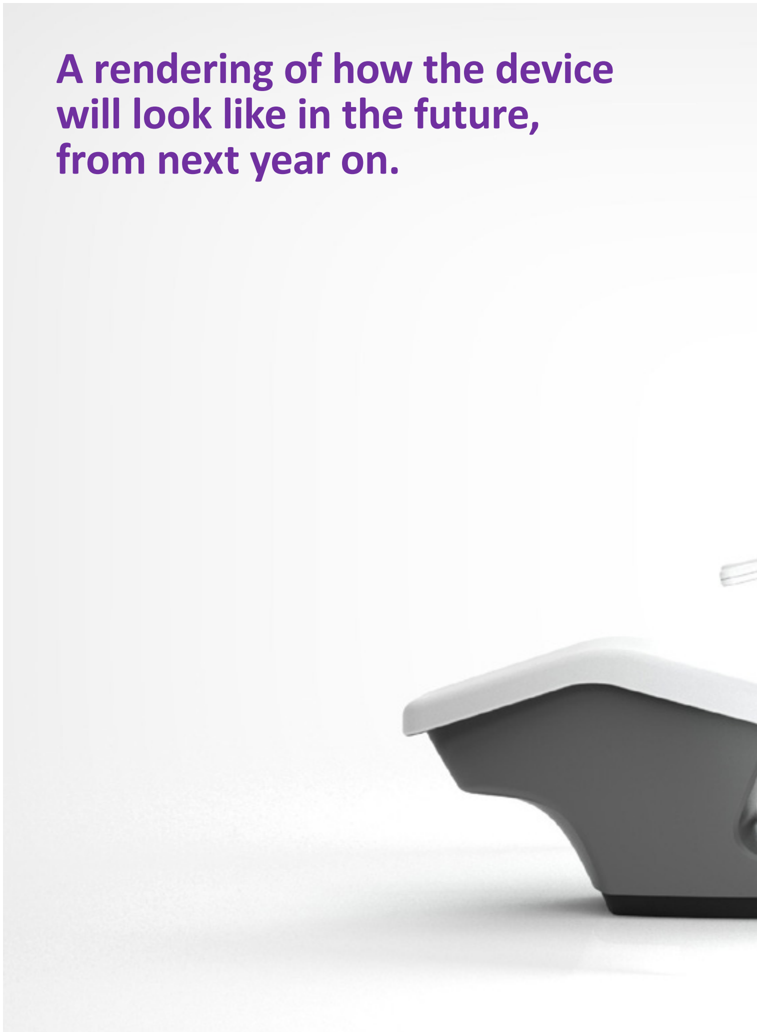

A rendering of how the device will look like in the future, from next year on.

early diagnosis of Alzheimer's, but at the moment it's very expensive and not widely available," Jannis tells us.

"Our vision is to make this possible for everyone and bring it to more hospitals, more doctors, and more patients than ever before."

The 14-strong team at Positrigo have developed a new PET system called **NeuroLF**. It fits the brief of being more compact and significantly more affordable than traditional scanners. The diagnosis itself only takes **10 to 15 minutes** and there is no need for any extravagant hospital setup. All the device needs is a small room to operate in.

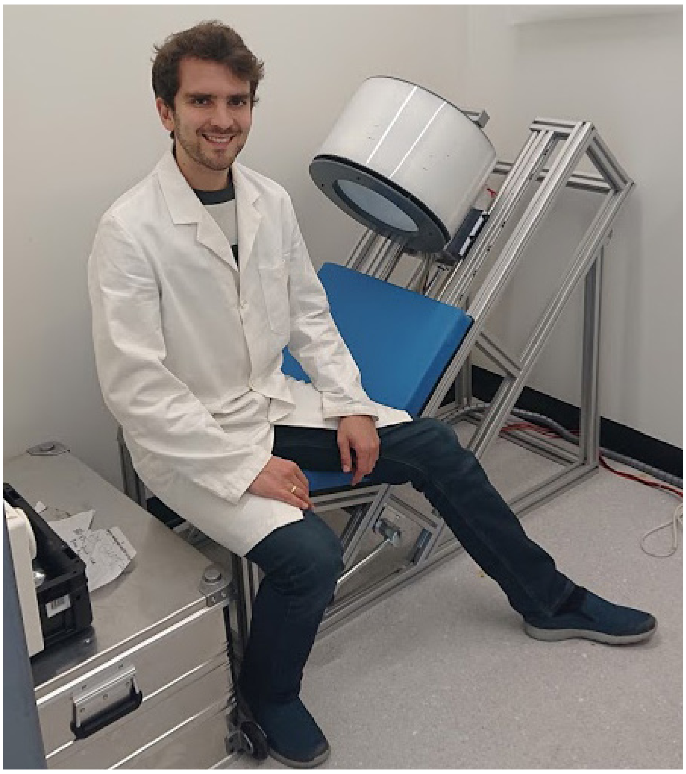

Co-founder and partner Max Ahnen on the very first prototype

This has been a group effort between **ETH Zurich**, where Jannis graduated, **University of Zurich**, and **University Hospital of Zurich**, with advisors including **Alfred Buck**, **Günther Dissertori**, **Werner Lustermann**, and **Bruno Weber**, working alongside Jannis and co-founder **Max Ahnen**.

"We were in the right place at the right time with the right people," Jannis discloses.

*"We were at ETH Zurich when they felt they needed this device, and we had the expertise to build it. It's come from **particle physics** and not many people can do that. Also, we're in Zurich close to the developers and the hospital that did studies on medication against Alzheimer's. It all came together perfectly!"*

Positrigo's vision is that **doctors will screen for Alzheimer's in the same way they do now for colon cancer**. Once people hit a certain age, they will be able to go to their doctor for a regular scan, with the outcome being that they either do not have to worry about Alzheimer's for the next five to 10 years, or that they need closer monitoring, and perhaps medication.

The only other alternatives to this are blood tests, which are not very precise, or cerebrospinal fluid, where they tap into the spinal cord and take out some of the liquid, which can be very unpleasant and invasive.

"As we are in the pandemic, I think a good analogy is that the blood test is like the antigen test for Alzheimer's and PET imaging is the gold standard PCR," Jannis explains.

"There will be a lot of blood tests, but then we'll always have people that need confirmation of blood test results, and this is where imaging comes in."

With NeuroLF, people will have **an**

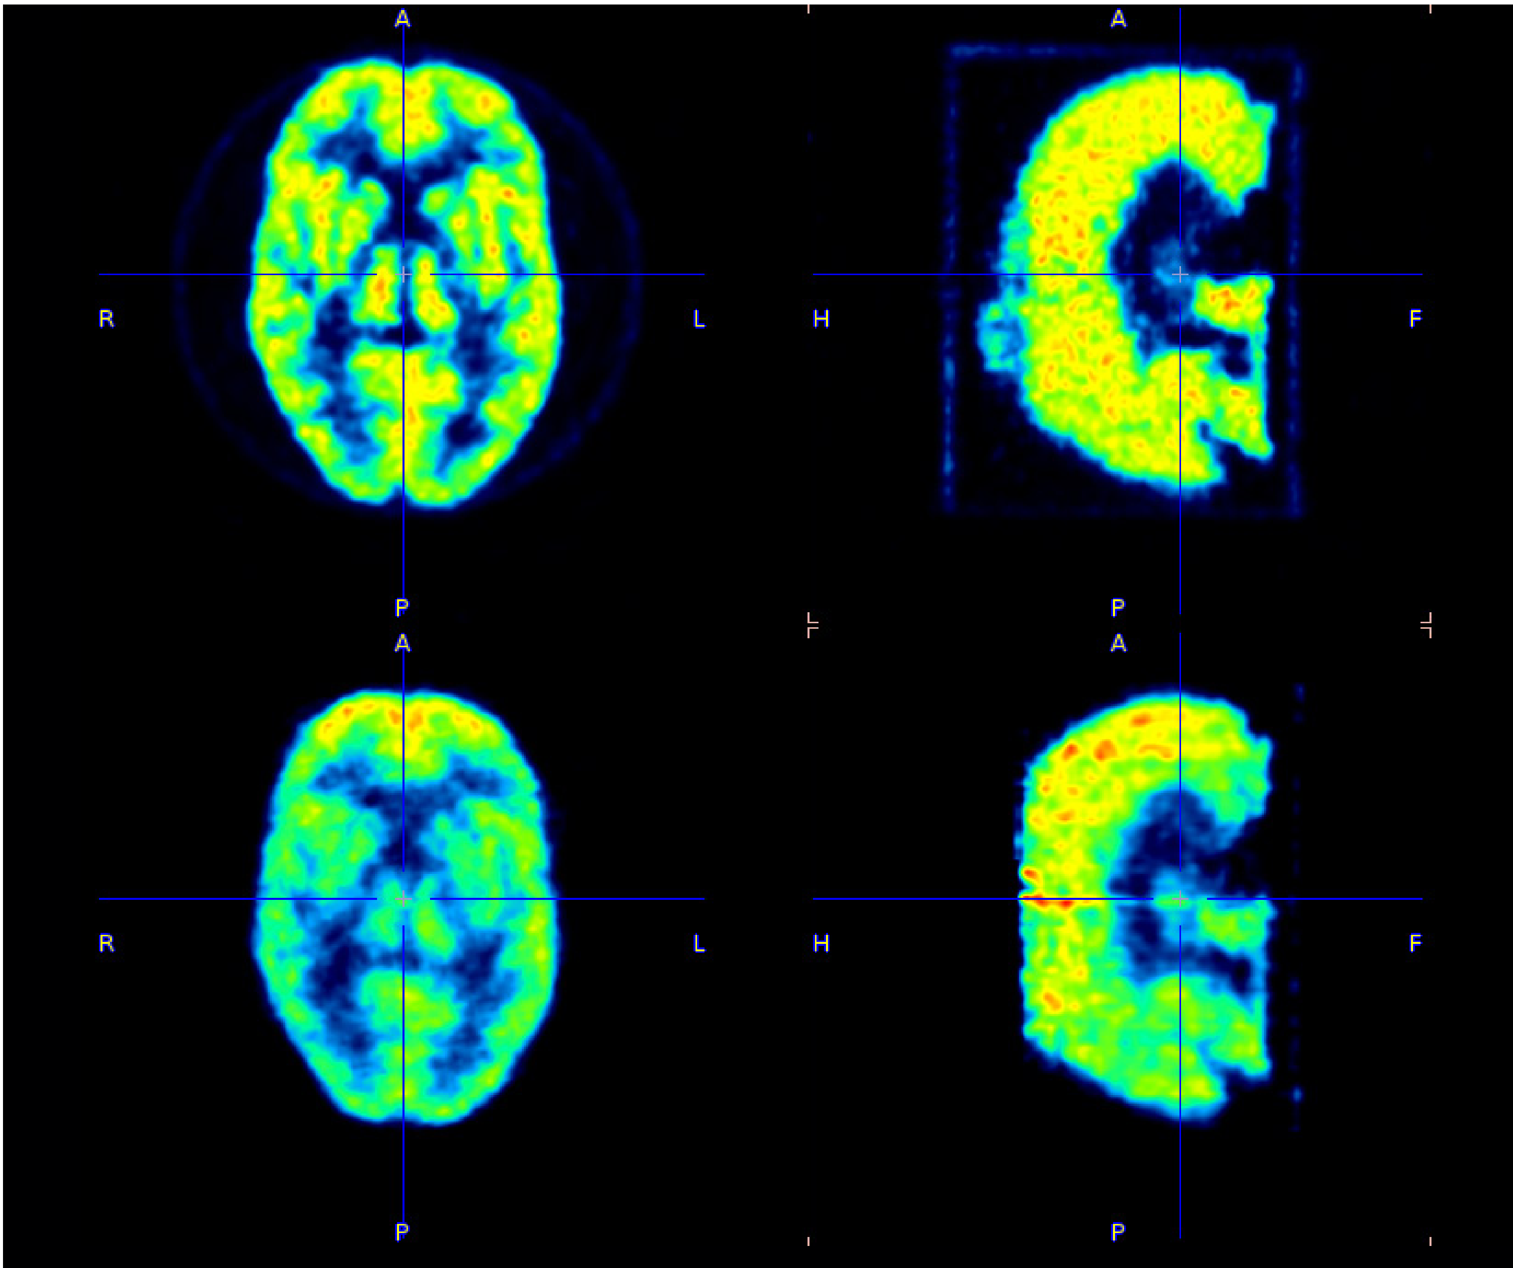

An image of a brain phantom: bottom row is that prototype, top row is a reference image from a big scanner.

intravenous injection with a small dose of radioactive tracer. The radiation dose is no higher than is needed for current systems, which is already no more than someone would be exposed to on a transatlantic flight. The brain scan can also identify **other brain-degenerative diseases, such as Parkinson's and ALS**, to get an individual assessment of someone's risk of brain function decline.

"Early diagnosis is vital because if you're too late, it's harder to treat," Jannis points out.

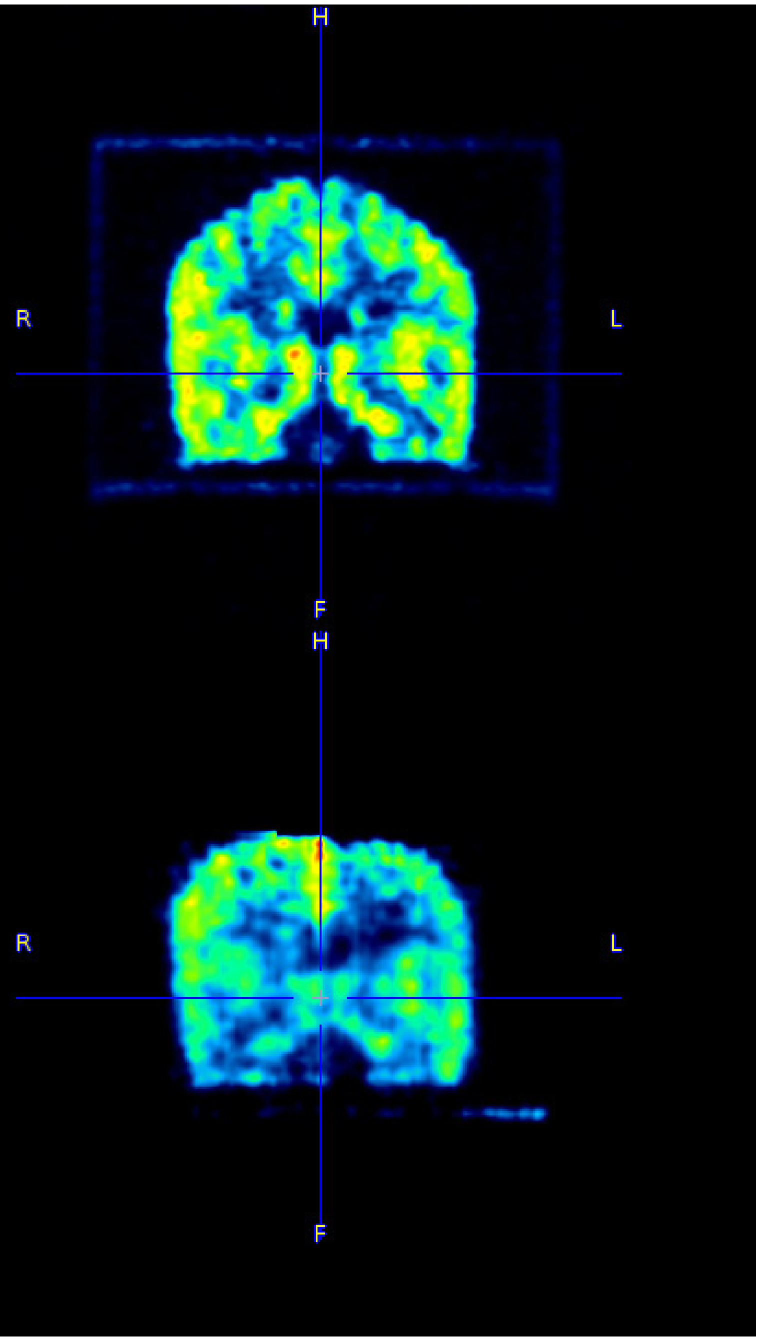

"This is a game changer for an ageing society because it makes it possible to diagnose early and intervene early. There are hundreds of medications in different phases of clinical trials."

He is confident there is now a clear path to market. Positrigo has worked closely with clinicians and engineers to validate the technology and aims to make it available in the next one to two years. It is entering a new series B funding round next year to support roll out into its first markets in Germany, Switzerland, and the US. The hope is that in a decade's time it will be widely available for everyone at the point of care.

Positrigo doesn't intend to stop here. This hardware is at **the core of a software platform** it is developing, and with the images it gets it is devising new models with AI-driven data science to help doctors make better diagnoses.

*"Compared to current software, this will have an **improved user experience and usability**,"* Jannis reveals.

"If you make it nice and easy to use, that's also part of the road to making this exam more accessible and available."

Positrigo is hiring!
Visit its website for current openings and look out for new software engineer positions next year. If you have other engineering disciplines, keep an eye on their page for future roles.

Introduction

Chen (Cherise) Chen recently completed her Ph.D. at Imperial College London (ICL). Her main research interest lies in machine learning and computer vision, and their use in medical image analysis, with the ultimate goal of improving machine learning algorithms' robustness and reliability in real-world healthcare applications. She will continue to work as a postdoctoral researcher in her current group: BioMeDIA. Chen would like to thank her supervisors Daniel Rueckert and Wenjia Bai for their continued support and guidance during this long and exciting journey. Congrats, Doctor Chen!

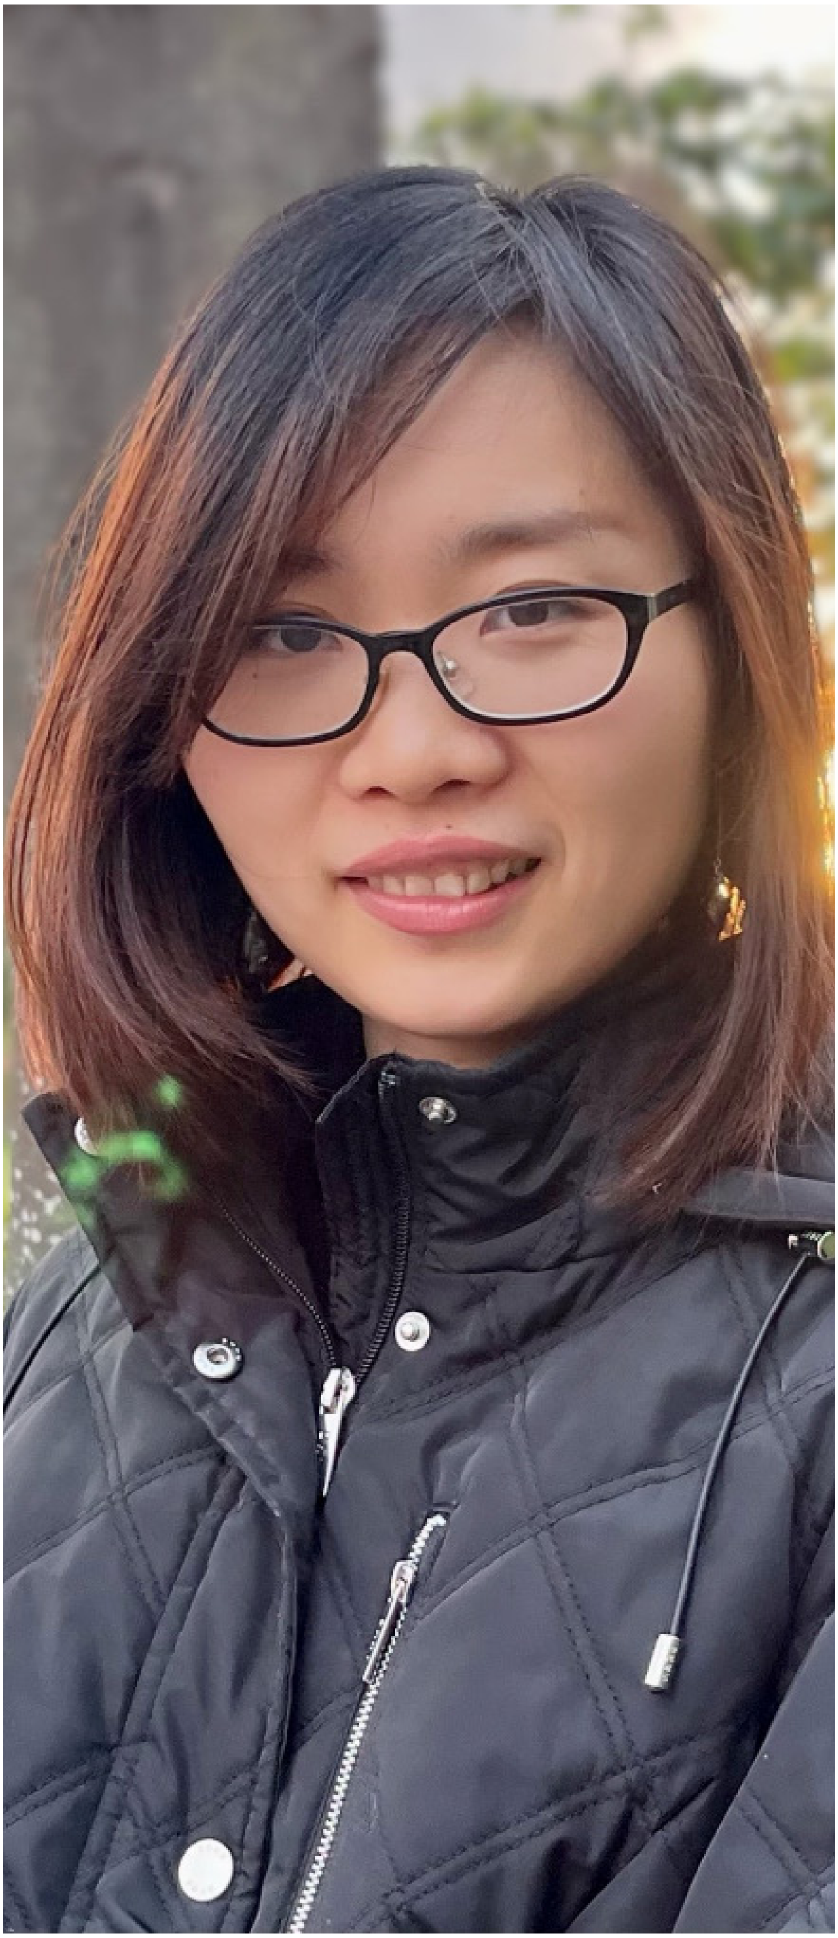

Deep neural networks have become a powerful tool to process and analyze medical images, with great potential to accelerate clinical workflows and facilitate large-scale studies. However, in order to achieve clinically acceptable performance at deployment, these networks generally require massive labeled data collected from various domains (e.g., hospitals, scanners). Such large labeled datasets are rarely available in practice, due to the high collection and labeling costs as well as data privacy issues. So, how can we improve the domain generalization and robustness of neural networks for medical imaging using limited labeled data?

Learning with unlabelled data

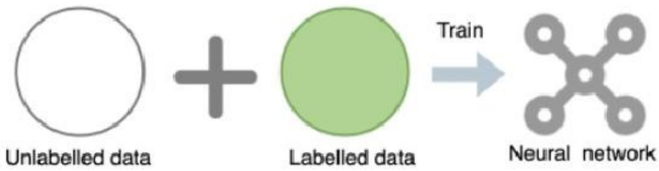

We consider exploiting the value of unlabelled data, which is *cheap* and *often widely available*. We develop a novel **adversarial data augmentation (DA) method with realistic image transformations** (adv bias, adv chain) to augment labeled and unlabelled images for more effective consistency regularization. This adversarial DA can be applied as a plug-in module in both supervised and semi-supervised medical image segmentation frameworks.

We further propose an **unsupervised style transfer method with a cascaded U-net for unsupervised domain adaptation**, where there is a distributional gap between source domain images (labeled) and target domain images (unlabelled). We first train an image-style translation network by utilizing a set of labeled bSSFP MR images (source domain) and unlabelled LGE MR images (target domain). We then employ the trained network to translate labeled bSSFP images to be LGE-like images, which are then used to train an LGE segmentation network. Our approach does not require labeled LGE images for training, yet it can surprisingly achieve high segmentation accuracy on LGE images, ranking 1st in the public multi-sequence cardiac MR segmentation challenge [More].

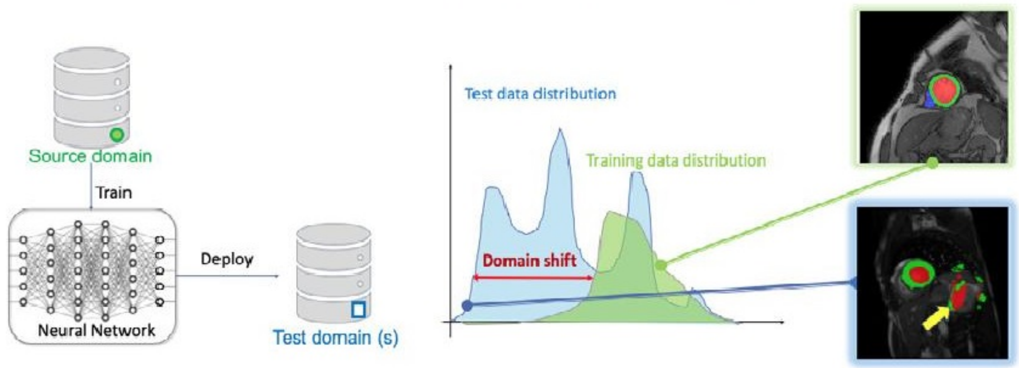

Acknowledgment: The research was supported by the SmartHeart EPSRC Programme Grant (EP/P001009/1).

Learning from limited, single-domain data for improved cross-domain performance

In the worst scenario where we *only* have access to limited labeled data from a single domain for training, the model can easily fail to predict images from *unseen test sets* due to the domain shift problem (see Figure above). A natural solution is data augmentation, which simulates the domain shifts at training. We first design a general training and testing pipeline with **data normalization and augmentation** to improve cardiac image segmentation performance across various unseen domains. We further propose a **latent space data augmentation method with a cooperative training framework** to enhance model robustness against unseen domain shifts and imaging artifacts. By performing random/targeted code masking in the latent space, a set of challenging examples are generated to strengthen the network capacity. This latent-space masking-based data augmentation method is effective and practical, as it does not require domain knowledge and expertise to design a particular image augmentation/corruption function.

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

Her Machine Learning Tools Pull Insights from Cell Images

A very nice interview with awesome **Anne Carpenter**. Not as good an interview as ours, of course! Anne explains how computers can mine images of cells for patterns that identify their cell type and disease-associated traits. This kind of image-based profiling is **speeding up drug discovery** by improving screening for compounds that desirably modify cells' characteristics. What she does (this is how she puts it) is turning *"messy, qualitative biology into precise, quantitative numbers"*. Anne also expands on the **CellProfiler**, her free, open-source software that enable biologists to do just that! Read More

Artificial Intelligence in the Intensive Care Unit: UF Researchers Developing Novel Solutions

The pandemic has brought many previously healthy people to Intensive Care Units all over the world. **University of Florida Health** researchers are developing **an intelligent ICU**, an autonomous patient-monitoring system driven by **artificial intelligence**. Capturing, analyzing and acting on information provided by cameras and other sensors should improve patient care and facilitate physicians' decision-making. Ideally, the AI-driven system will predict whether an ICU patient will improve or decline, giving the docs precious time to intervene before a patient's condition changes. Read More

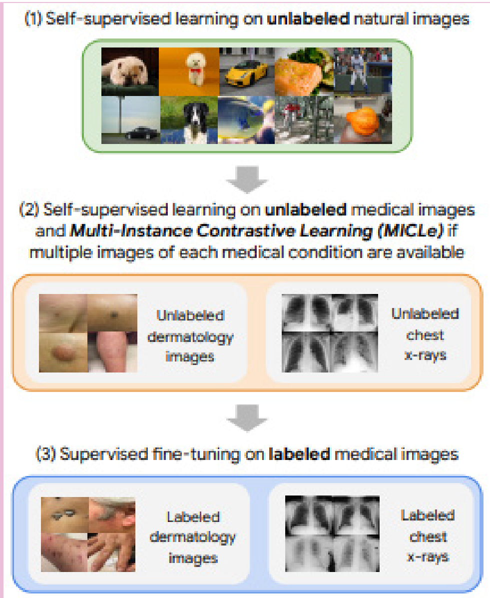

Google Research Changes the Game for Medical Imaging with Self-Supervised Learning

The guys at VentureBeat did not fail to notice this exceptional new work by a team at Google Research, led by awesome **Shekoofeh Azizi**: researchers suggest **a new technique that uses self-supervised learning** to train deep learning models for medical imaging. Early results show that the technique can reduce the need for annotated data and improve the performance of deep learning models in medical applications: *"Using self-supervised learning, we show that we can significantly **reduce the need for expensive annotated data** to build medical image classification models!"* Read More

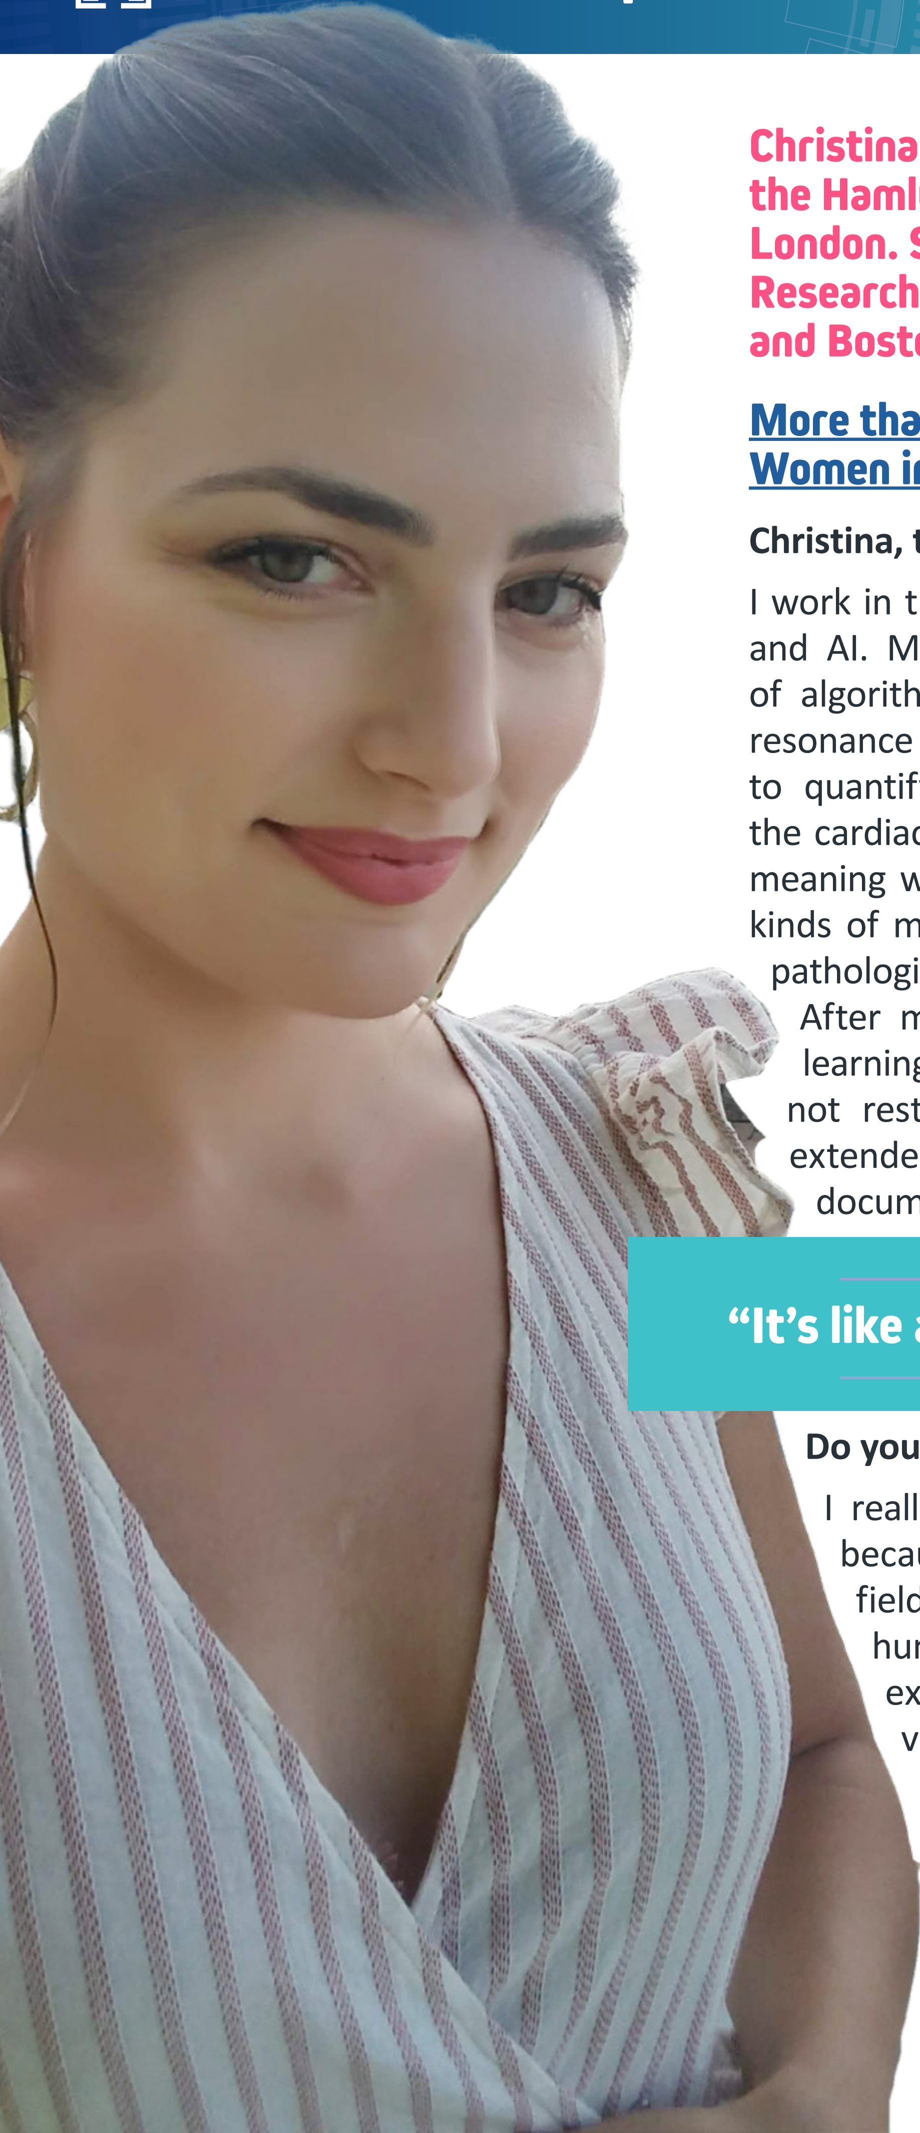

Christina Koutsoumpa earned a PhD at the Hamlyn Centre, Imperial College London. She is a former Postdoctoral Researcher at Harvard Medical School and Boston Children's Hospital.

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

Christina, tell us about your background.

I work in the field of medical image processing and AI. My PhD was about the development of algorithms for the processing of magnetic resonance images from the heart. The goal was to quantify and visualize the contractility of the cardiac muscles during cardiac contraction, meaning when the heart beats. Creating these kinds of maps will assist doctors in diagnosing pathological areas of the cardiac muscles earlier. After my PhD, I moved on to more deep learning and computer vision applications not restricted to the medical field only. I extended this work to the processing of legal documents mainly.

> **"It's like a kind of magic!"**

Do you prefer medical or non-medical?

I really like the medical field, first of all, because I have most of my expertise in this field. Secondly, I get this feeling of helping humanity. But I find other fields equally exciting in deep learning, AI, computer vision and imaging.

Where did you study?

I completed my Bachelor's from the National Technical University of Athens in Greece. It was a five-year degree, equal to a Master's. The National Technical University

is supposed to be the most prestigious technical university in my country. Then I proceeded to my Master’s studies. I studied biomedical engineering in Greece. After my Master’s, I moved to London. At Imperial College London, I studied my PhD in the Department of Computing at the Hamlyn Centre. Shortly, after my PhD, I moved to the US for a postdoc at Harvard Medical School at the Boston Children’s Hospital for robotic surgery.

I would like to share with our readers that Christina is currently open to finding her next job. She is a great catch! Maybe some of our readers have something to offer her! Tell me something Christina, what have you enjoyed the most in your studies up until now?

I really like computer vision and deep learning. I really like creating models that detect features on images that are not visible to the human eye. It’s like a kind of magic! [*laughs*] You extract information that you cannot easily see or easily quantify. I find this really exciting to extract the useful information from images and show it in a comprehensive way, especially medical images which are not easily understandable due to not only the lack of background for most people but also due to high corruption of artifacts. You help the medical community and doctors to perform a more accurate diagnosis. You also create visualizations and quantitative measures that most people can understand. I hope in the future that we can help patients better understand their pathologies. I find this really, really exciting, this combination of technology and medicine. It’s something that was not easy to combine in the past. For years, we thought these were completely two separate fields. It turns out that they can be very complementary to each other. I really like this.

You spoke about a lot of interesting research. But you didn’t know about these fields before you started. What inspired you as a young Greek girl to take this path?

This is a really nice question. Thank you for asking me! I rarely get asked this question. When I was a young student, I really enjoyed math. I would spend hours solving problems. I didn’t know what I wanted to become exactly. I knew that I really enjoyed math, physics, and solving problems. Later on, I decided to become

Until that point, it just sounded like a crazy dream

an engineer because I wanted to see the science applied to real applications. Again, I wasn't sure if I wanted to go too technical. I had this internal intention to work on something that would benefit society, humanity, and people. I was not sure what this was exactly until I came across a video on the news about robotic surgery. It was completely unknown in Greece at the time. It was many, many years ago. I found this very exciting at the time. It seemed really important and also sci-fi! [*laughs*] That was a very important stimulus for me. When I graduated from the National Technical University of Athens, I decided that I would leave Greece. I finally felt like I could pursue this dream. Until that point, it just sounded like a crazy dream: this field did not exist in Greece at all, and there were no jobs.

Sometimes dreams come true!

That's true! Then I moved to Imperial. I did a PhD. Then I found myself at Harvard, which was another super crazy dream.

Why did you decide to go to Hamlyn in the UK and why did you decide to go to Harvard in the US?

My main drive for Imperial College London at the Hamlyn Centre was to work on medical robotics. At the time, there were not many groups working on this specific subject. From my research back then, I concluded that the Hamlyn Centre was probably the best one in the field, so I contacted Professor Guang-Zhong Yang. I explained my background and research interests, went through a technical evaluation and eventually got admitted. I was assigned a project, which was not in robotics but medical imaging instead. Although this was not my initial goal, I found the project particularly exciting. I was delighted to do that!

Who was your supervisor?

Prof. Guang-Zhong Yung along with Duncan Gillies later. I kind of went with the flow. During this journey, I discovered a lot of interesting aspects and interesting sub-fields of biomedical engineering and computing that I didn't know about. Imperial College London and London, in general, is a very stimulating environment. I consider myself super lucky to have spent five and a half years there. Then Harvard came along very smoothly, I would say. Then the post-doc at Boston Children's Hospital/Harvard Medical School came along very smoothly, I would say. I didn't spend a lot of time preparing an application or completing exams. Apparently, my research project overlapped with the research focus of the CHAI group at BCH/HMS, the Cardiac Surgery Department. The PI of this group

Dr Douglas Perrin, who became my PI later, showed interest in extending together my denoising methods for cardiac MRI to novel denoising methods for echocardiography. He contacted me, we had an interview and...

....You started working for him?

Yea! [*laughs*] That was it! I had paid a visit to Harvard Medical School a few months earlier, during which I presented my PhD work to a small group of the Cardiac Surgery Department at BCH/HMS. I presume this presentation stimulated the discussions, which led to my interview with Dr Perrin and admission later on.

What advice do you have to give a young student starting down the same path? Would you tell your younger self to go with the flow, as you did?

The advice I would give to my younger self is to keep dreaming and set long-term goals. But during the journey, try to be flexible. Don't let yourself down over the difficulties or obstacles. Whenever you see a slightly different path, consider it. See if you like it. Be flexible. Go with the flow, but never forget your ultimate goal. Overall, do something that you like, because this is the only way to become successful, to put effort into it... Or, if not successful, a happy person [*laughs*]!

You certainly made a lot of sacrifices to come this far. You dedicated many hours to this. Was it worth it?

That's a really good question! [*laughs*] I think of my favorite poem that accompanied me during this journey. It's called Ithaka from Constantine Cavafy. This poem says something like as you set out for Ithaka, you will encounter many obstacles along the way, many scary creatures and things that will discourage you from reaching your goal. Don't give up! One day you will reach your goal, in this case to go back to your hometown, your island. When you reach your island, you will realize that the important thing wasn't reaching the island in the end, but the journey and how it

If it were super easy, there would be no need to discover it!

I didn't know how to fail!

matured you along the way to become the person you are now. I get this! Especially during my PhD, I would read this every single day. I feel very lucky that I went through this journey because it shaped me. I would go through it again, despite the sacrifices.

One day you will reach your goal!

Of all the things that you acquired during this journey, does one lesson stand out?

The most important lesson was to not give up. People who decide to do a PhD are usually very good students. They don't learn how it feels to fail. I was the best student at my school during all the years of my studies. I didn't know how to fail. That's not realistic because life is full of failures. Research, by definition, is like 80 percent failure. Of course, you will fail. If it were super easy, there would be no need to discover it! Someone would have done it already. We think that these are failures, but these are lessons to develop and learn to become better. We reject methods and move on to other methods. This is what I learned: don't give up. Take every little failure as a lesson and an experience to continue. I'm still working on it! I'm not super thrilled when something doesn't work but still, I have made progress.

Although the readers cannot see that, you gestured the word failures with inverted commas: "failures". This is a great lesson. Is there anything else you'd like to share with our readers?

Pursue your dreams! Do what you like. Value all factors in your life. Career is very important; the fulfilment of achievement is very important, but not the only thing. Try to become happy. This is very personal and different for each person. This is very generic advice that I would give to all people, especially girls since this is about women in science. Please don't get discouraged when you feel like your environment does not expect you to succeed. When you feel like you may not make it or that someone will look down on you, just ignore it. Keep trying. Keep doing your work. When you feel like it's difficult for you, imagine yourself in a man's shoes. Think of what he would do and how he would see his fault, if there was a fault. Take courage from this. Just ignore all the noise and do what you like!

I spoke to [Jessica Sieren](#) some time ago. She said, "*When you do a PhD, if you don't have a crisis, you're not doing it right!*"

Yea! [*laughs*] For a PhD student, everyone goes through a tough time. If you don't go through a tough time, maybe you're not pushing yourself enough.

We spoke a lot about the past. I want my last question to be about the future. What is your dream?

My dream is to work in my field and make an impact in my field. Maybe one day, I'll start my own company and put in my own ideas. I'd love to see my ideas become a useful product

Let me remind my readers that awesome Christina is available for work, and she is a great catch! Don't miss out on her!

[More than 100 inspiring stories of Women in Science here!](#)

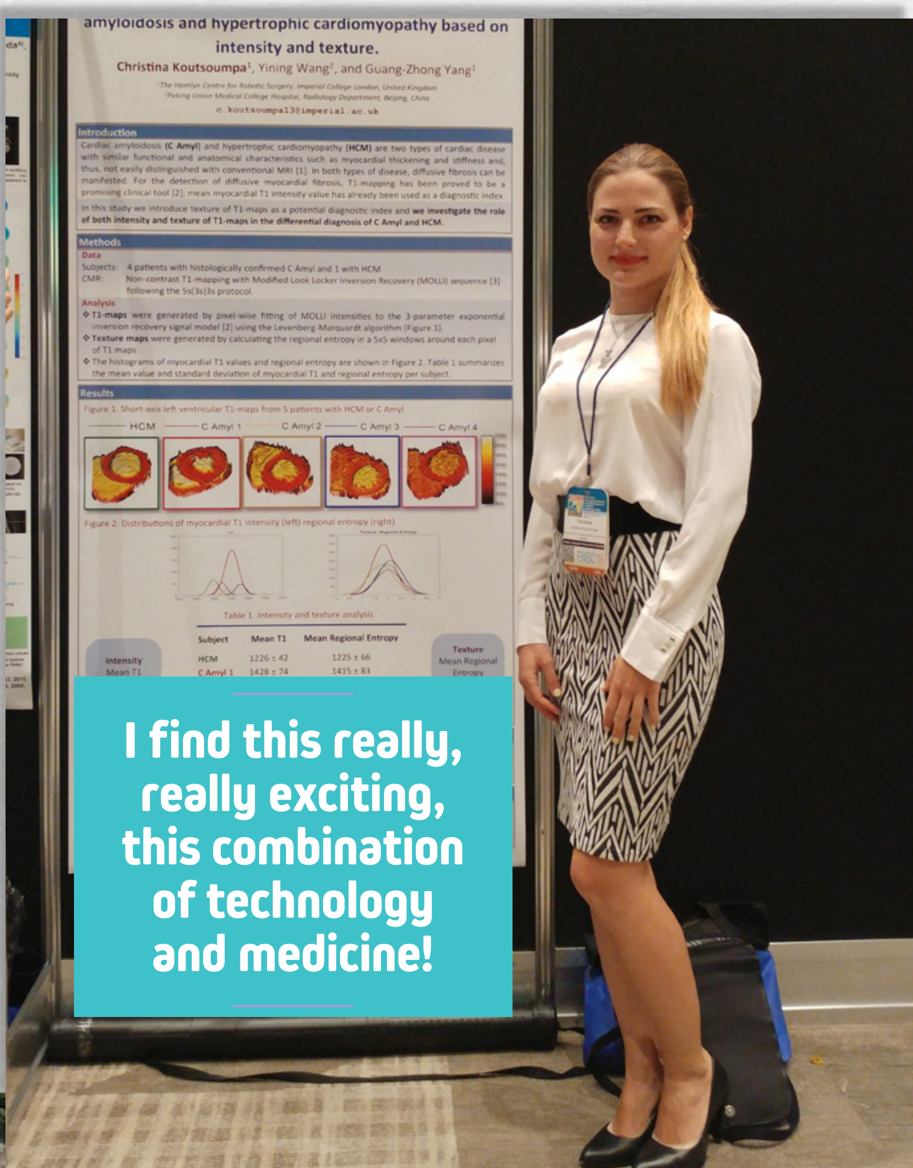

I find this really, really exciting, this combination of technology and medicine!

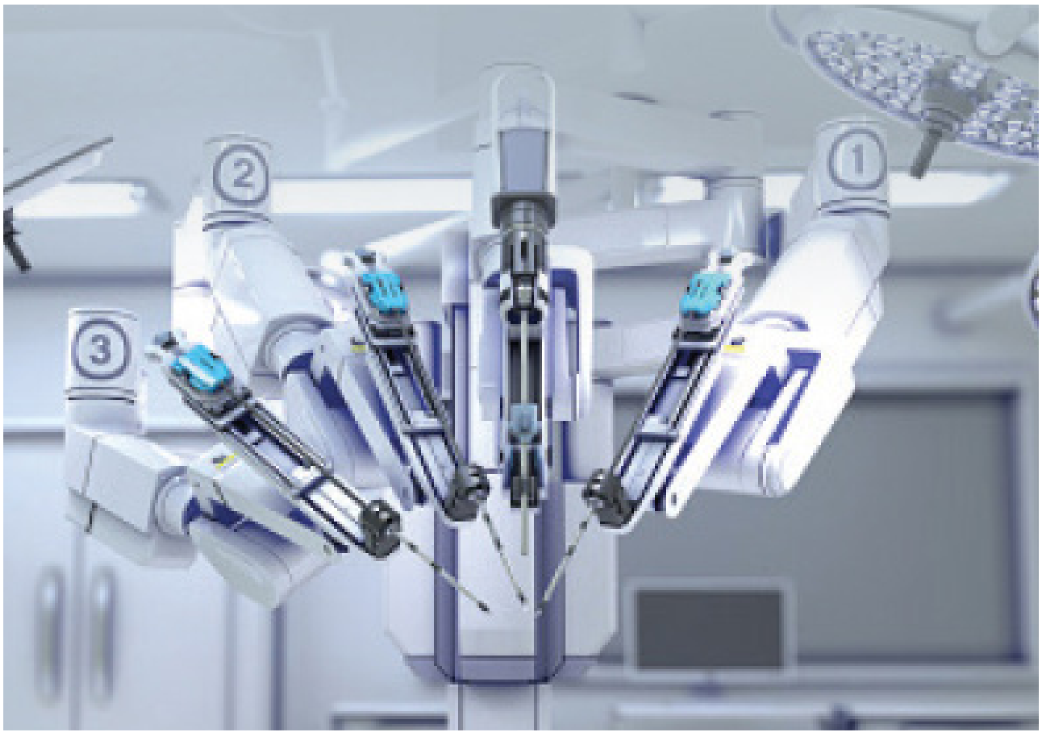

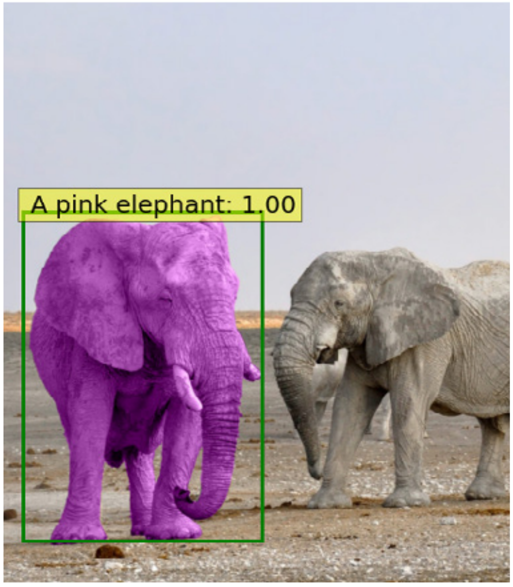

A pink elephant: 1.00