April 2021 Computer Vision News The Magazine of the Algorithm Community

ANNIVERSARY

HAPPY CELEBRATION

Iris Automation

Congrats, Doctor!

Unsupervised Domain Adaptation from Axial to Short-Axis Multi-Slice

Recognise Traits of ASD

Women in Computer Vision

Metal Implant Segmentation

Open-Access Journal

AI Spotlight News

Upcoming Events







2 Editorial



Computer Vision News

Editor: Ralph Anzarouth

Engineering Editors: Marica Muffoletto Ioannis Valasakis

Head Designer: **Rotem Sahar**

Publisher: **RSIP Vision** <u>Contact us</u> <u>Free subscription</u> <u>Read previous magazines</u>

Copyright: RSIP Vision All rights reserved Unauthorized reproduction is strictly forbidden.

Follow us:



This is a special moment for us at **RSIP Vision**: our flagship publication **Computer Vision News** was launched in April 2016, marking our **5th anniversary** this month! We went through many changes during these last 5 years: from the Deep Learning revolution to the advancement of the autonomous vehicles and the unprecedented transformations brought by AI to healthcare and many other fields. All this persists through a global pandemic that has changed every aspect of our lives, including all conferences going virtual.

Despite the challenges, Computer Vision News has always come to you at the beginning of each month. We never failed to bring you **CVPR Daily, MICCAI Daily** and all the other Dailies, every single day of their respective events, in Canada, Hawaii, Korea and everywhere.

We want to celebrate this milestone together with the algorithm community, in particular with the **AI**, **Computer Vision** and **Medical Imaging** community, which we are proud of serving as a loyal voice and pioneering R&D partner: many friends of the magazine, from featured scientists and professors to regular readers, have sent us their anniversary wishes. With their permission, we have decided to publish the most touching of them on pages 28 to 39.

Many thanks to the thousands of scientists who have shared their knowledge and work on almost **4,000 pages** of free and accessible-for-all scientific content. Many thanks to our affectionate readers, who have viewed more than **4,500,000 pages** during the last 5 years. And many thanks also to our RSIP Vision family, in particular to **our CEO Ron Soferman**, for making this dream come true, month after month. We are very excited and full of energy for the next 5 years and more! **Enjoy reading this festive April 2021 issue!**



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



© Scott Adams, Inc./Dist. by UFS, Inc.

Summary 3









- 40 Emma Robinson Women in Computer Vision
- 46 AI Spotlight News
- 48 Iris Automation Computer Vision Application
- 52 Frontiers in Computer Vision Open-Access Journal
- 56 Bay Vision Meetup
- 57 Upcoming Events



- 24 Fetal Brain (FeTA) Challenge
- 28 Happy 5th Anniversary Greetings from the Community

4 Research

Unsupervised Domain Adaptation from Axial to Short-Axis Multi-Slice Cardiac MR Images by Incorporating Pretrained Task Networks

by Marica Muffoletto @maricaS8

Every month, Computer Vision News selects a research paper to review. This month is a special one: spring has finally arrived and the rise in vaccination rate is giving us much needed hope for the coming summer. Filled with all these positive vibes, I decided to celebrate by choosing to review a paper close to my own research field, called Unsupervised Domain Adaptation from Axial to Short-Axis Multi-Slice Cardiac MR Images by Incorporating Pretrained Task Networks, written by several authors located mainly in the German centre for Cardiovascular Research (Heidelberg).



We are indebted to all of them (Sven Koehler, Tarique Hussain, Zach Blair, Tyler Huffaker, Florian Ritzmann, Animesh Tandon, Thomas Pickardt, Samir Sarikouch, Heiner Latus, Gerald Greil, Ivo Wolf, Sandy Engelhardt) for allowing us to use their images to illustrate this review and especially to Sven and Sandy who provided us with extra beautiful illustrations. I hope this review makes you as enthusiastic about their work as I am after reading their paper, which can be found <u>here</u>.

This paper addresses a common problem in medical imaging: the presence of distribution shifts among imaging scans. Of particular relevance, but not very much discussed, is the problem of scans acquired across different axes. In particular, Cardiac Magnetic Resonance, which offers a non-invasive non-radiating option to image the heart, is vastly used to assess measurements of left ventricular and right ventricular volumes for diagnosis and treatment of several cardiomyopathies. This is generally acquired in short axis (SAX) view as



Figure 1: CMR Images in SAX view on the left with corresponding segmentation of RV, LV and myocardium - CMR Images in AX view on the right.

Unsupervised Domain 5 Adaptation from ... 5

this can be easily related to the heart anatomy, oppositely to axial images that are oriented according to the thorax, making them subject to variation between individuals. However, these types of slices remain fundamental to investigate right ventricular quantification measurements for diseases that specifically involve right ventricular (RV) geometry, such as patients with congenital heart defects (CHD) e.g., tetralogy of fallot (TOF). These data (AX view) are uncommon and their corresponding annotations very scarce, making it a real challenge to train deep learning algorithms that can perform high quality tasks on them.

Unfortunately, the problem is particularly complex because AX and SAX volume stacks describe a different physical volume, which means that each of the images contains information which is not present in the other image, and a rigid relative transformation between both images varies heavily from patient to patient. The high variation - due to difference in resolution and field coverage - makes it impossible to simply get the transformation from each image and then just apply the average rotation to all AX images in order to close the domain gap.

To ensure that this is really the case and further research is required to fill in the domain gap, the authors preliminarily try to directly apply a segmentation model trained on SAX images on axial slices (experiment 1). The experiment actually indicates a large distribution shift, similarly to an experiment (experiment 2) which shows that direct regression on the transformation parameters is also not sufficient.





Figure 2: VIDEO showing 3D spatial relation between SAX view in red and AX view in yellow. Example contours for left ventricle.

Hence, this study focuses on leveraging the knowledge and availability of SAX slices from two datasets to get optimal segmentation results on AX images of a TOF dataset with no ground truth labels. This is done by using an "unsupervised domain adaptation approach, which is able to align anisotropic image slice stacks with significantly different fields of view and small overlapping image regions".





The authors refer to their domain adaptation pipeline as an input adaptation type because it addresses the domain gap at the input level. It aims to alter the samples from the target domain (AX images) to bring the distribution of the target closer to that of the source (SAX images) and transform the result back into the target domain.

The main DA approach is carried out by a **Spatial Transformer** which learns the 3D geometric relations between the two domains extracting a set of transformation parameters. These are used to transform the slice orientation to a short-axis view. Once this is completed, a task with pixel-wise predictions (Segmentation Module) can be performed on the SAX representation and the resulting image

Unsupervised Domain 7 Adaptation from ... 7

is transformed back to the target domain. The authors first present a baseline transformation module (experiment 3) which is then expanded into an extended version (experiment 4) using two main techniques: the cycle consistency on forward and backward transformations, so that predicted transformations are invertible (SAX2AX - AX2SAX); the task-specific focus which guides the network in optimizing the transformation such that regions relevant to the segmentation task are included in the target grid.

The final expanded network architecture, as employed in experiment 4, is described in the comprehensive figure below.



Figure 4: Proposed final network with extensions (experiment 4).



What does Alyosha Efros wish Computer Vision News for our 5th anniversary? Find out on page 36!

8 Research

A 2D U-Net architecture is employed for the segmentation module, with specific choice of hyper-parameters. Modifications to the original network, include the employment of Exponential Linear Unit (ELU) activation, a batch normalization layer after each non-linear activation and a combined loss function (Cross Entropy + Soft Dice Loss) with a weighting factor of 0.5 for the CE loss.

The transformation module is built with a **localisation network**, which predicts the parameters of a transformation matrix, M; the **grid generator**, which implements the transform, T; and the **sampler**, which is responsible for the interpolation. The localisation network outputs three rotation parameters ϕ , θ , ψ and three translation parameters $\vec{t} = [tx, ty, tz]$, that are fed to the Euler2Affine layer which converts them into an affine matrix representation to allow valid inverse transformations, such that cycle constraints can be imposed. The MSE loss is always employed to compute the output of the transformation with the original AX/SAX images.

The cycle consistency, added to the network to improve stability and ensure correct rotational parameters, comes in the shape of a $Loss_{Cycle}$ composed of two MSE losses that measure the quality of the target to source transformation and the corresponding inverted output transformation, both output by the Eurler2Affine layer. The other novel element, the **task-oriented guidance**, is added to fine-tune the translation parameters because most of the predicted transformations excluded relevant lower slices of the volume. This is created by feeding the domain transformed image into a branched out sub-network (the Segmentation Module) whose gradients are back propagated into the same localisation network. This allows this network branch to maximize the number of predicted foreground voxels ($Loss_{Cycle}$) and to therefore reduce cutting of relevant structures due to sub-optimal translation. Finally, the overall loss of the network results in:

$$Loss = Loss_{Cycle} + 0.1Loss_{Focus}$$

Two main datasets are employed to train this network: the TOF dataset, made up of two sub-cohorts, and the ACDC Dataset, including healthy individuals and some affected by four different conditions. The segmentation module was trained using

Unsupervised Domain 9 Adaptation from ... 9

SAX images only, from the 1st sub-cohort of the TOF dataset, including 193 sets (144 used for training – 39 for testing) and the ACDC data (75-25). The transformation module was trained in a four-fold cross-validation using only the 2nd cohort of the TOF dataset - 81 patients with pairs of AX and SAX images, shuffled and split into 60-21 for training and validation.

To summarize, the authors use the method described above to lead four main experiments, whose performance is measured using the Dice score, Volume difference and Hausdorff distance. These are:

- 1. Domain gap when SAX model is applied on axial slices without domain adaptation
- 2. Upper limit with respect to a direct regression on the transformation parameters
- 3. Baseline domain adaptation approach
- 4. Complete domain adaptation approach where Dice and HD are considerably improved with respect to all three.

According to the DICE score, results show that there is an increase in performance of 8%(LV) - 2%(MYO) - 25%(RV) on the 4th experiment compared to the 1st; of 2-4% on all classes from 2nd to 4th and of 24%(LV) - 18%(MYO) - 22%(RV) from 3rd to 4th.



What does Tal Arbel wish Computer Vision News for our 5th anniversary? Find out on page 28!

10 Research



Figure 5: Results comparing experiment 4 to 1. This graph shows improvement in performance according to all metrics.

One might wonder why the volume difference appears among the metrics. It is extremely important here as the authors take care to mention that in AX images, the average RV volume appears 51 ± 15 ml larger than in the corresponding average SAX volume. A significant difference between the two is found performing the T-test between RV ground truth volumes (AX-SAX). This variation between the shape and volume of RV in the two axes might be more than what would be expected to be found due to the different grid orientations. And in conclusion, this means that the network might systematically predict smaller RV regions than the expert contour in axial view, because it is trained only on native SAX ground truth contours.

As promised, this paper delivers a full pipeline to segment axial images just leveraging the much bigger dataset of short-axis scans and labels. It is really worth looking at one example where the complete model has been applied to a held-out AX image stack. For this image, the obtained Dice scores for each volume are 0.9274 (LV), 0.8707 (RV) and 0.6605 (MYO).



Figure 6: Final prediction of the segmentation module over the axial and short-axis slices.

Unsupervised Domain 11 Adaptation from ... 11

The paper opens up to huge opportunities in the field. Transforming images acquired in one axis to acquisitions in another axis could be a very powerful resource; this is taken a step further by performing specific tasks on both scans, even when ground truth segmentations are not available in one of the domains. If you are not yet fascinated by this work, you should definitely watch the video below, another great example of transformation and task prediction which fully shows the excellent and novel work conducted by this group. This can also be found on their github page which I highly suggest visiting if interested.



Figure 7: WDEO showing transformation from SAX to AX view and segmentation task

<text>

12 Medical Imaging Projects

Metal Implant Segmentation for Arthroplasty

When planning many orthopedic computerized procedures, а (CT) tomography is taken scan in advance in order to get clear three-dimensional and accurate visualization of the bones. This threedimensional bone model assists the orthopedic surgeon in non-invasive decision making and planning of procedure. Accurate surgical the visualization is particularly essential arthroplasty procedures (joint in replacement), such as knee, hip, or shoulder replacement surgeries, as well as intricate spinal procedures. The modelling processes involves segmentation of the joint bones, osteophytes (bone spurs) and various anatomical landmarks. The resulting high-resolution model is crucial for accurately planning and designing the shape, size, location, and placement of the new implant.

One of the most common problems in bone segmentation occurs when a CT scan includes the presence of metals in the bones. These metals appear in CT scans when a patient has undergone either previous orthopedic procedures (such as hip or knee replacements) or previous surgical corrections after traumatic

injuries (such as spinal fusions). In these cases, the CT images become limited and challenging because the metallic presence creates "glare" (or streak) artifacts which blur and distort the image and complicate the bone segmentation process. These artifacts appear in the scan as bright and dark high intensity streaks that originate from the metal object and often block the surrounding anatomical structures and prevent accurate image processing. In other words, the presence of high intensity metal objects in CT scans produces an "aura" around them which obscures valuable information about the surrounding anatomy.



A knee CT scan in the presence of metal. The star-shape streak artifact ("aura") can be seen originating from the metal implant on the top.

Metal Implant Segmentation 13

The reason for these artifacts is that CT images consist of multiple X-ray images, taken using a rotating X-ray generator and detector. The images are acquired at multiple stages and from multiple angles along the rotation axis. This rotation process causes the metal, which almost completely absorbs X-rays, to repeatedly create the artifacts, resulting in the starshaped streak artifact. In general terms, metal artifacts are caused by nonstandard absorption values (Hausdorff) due to crosstalk between absorbing pixels and additional artifacts. These metal artifacts are caused by a variety of reasons, from beam hardening and scatter effects,

which result in the surrounding bright streaks and dark streaks between the metals, to Poisson noise and edge-related artifacts, due to under sampling, subject movement, cone beam, and windmill artifacts. These metal artifacts are particularly pronounced in iron and platinum, which have a high atomic number, and less pronounced in titanium and other low atomic number metals. The result is a challenging CT image for standard segmentation tools, which result in inaccurate segmentation that limit the ability of the orthopedic surgeon to properly plan the arthroplasty.

RSIP Vision has launched a new



Knee joint segmentation 1: A knee CT scan with image artifacts due to an existing metal implant (left) segmented to all bone and metal components (right)

14 Medical Imaging Projects



Knee joint segmentation 2: A knee CT scan with image artifacts due to an existing metal implant (left) segmented to all bone and metal components (right)

utilizes state-of-themodule that art deep learning (DL) algorithms designed and trained in advance to deal with the presence of metals in CT scans. This module accounts for the possible presence of metals by using an accurate, physical-based, and robust simulation tool of these metal artifacts. By adding synthetically produced data to the training process, the degrading effect of metals in CT scans has been dramatically reduced.

"This module has a vital and important contribution to the treatment of the patients undergoing orthopedic procedures", says RSIP Vision CEO Ron Soferman, "since many patients undergo additional or follow-up procedures throughout their lifetime". The result is accurate discernment of the location of metals in the bones and exact delineation of their outlines from the bones. Using **RSIP** Vision's metal segmentation toolbox, accurate bone modelling is made easy even in the presence of preexisting metallic implants. The result is improved orthopedic surgical planning, safer and more accurate procedures and overall better patient treatment and outcome.

What does Georgia Gkioxari wish Computer Vision News for our 5th anniversary? Find out on page 28!

What does Marc Pollefeys wish Computer Vision News for our 5th anniversary? Find out on page 35!





What does Vittorio Ferrari wish Computer Vision News for our 5th anniversary? Find out on page 35!

16 Congrats Doctor!



Grace Kuo recently completed her PhD at the University of California, Berkeley, where she worked on designing computational cameras, microscopes, and displays. By jointly designing the optics and processing algorithms, she demonstrated imaging systems with simplified hardware (for example, no lenses) and expanded capabilities (for example, 3D). After graduating, Grace began working at Facebook Reality Labs where she develops new display technology for virtual and augmented reality applications. Congrats, Doctor Grace!

Historically, cameras have been optimized to capture high quality photographs on film with no post-processing. However, today's images are almost entirely digital, and post-processing is ubiquitous. This raises the question: Is the current camera architecture best for all applications?

For instance, lenses are an integral part of most imaging systems, but they are typically heavy, expensive, and do a poor job of recording depth information. With my collaborators^{*}, we replaced the lens of a traditional camera with a diffuser -- a bumpy piece of thin plastic that refracts light -- to create a **compact lensless camera** which we call **DiffuserCam**. The diffuser scrambles the incoming light making the sensor measurement unintelligible to a human, but we developed an algorithm to recover the image from the raw data using a physically-based model of the system.



In addition, our lensless architecture naturally records 3D information. By applying techniques from compressed sensing, we demonstrated recovery of an entire 3D volume from a single 2D measurement taken with DiffuserCam.

Grace Kuo 17

Capturing 3D content in a single acquisition enables faster 3D video, which is particularly useful for biological samples. With this in mind, we redesigned DiffuserCam for **fluorescence microscopy** and demonstrated 10x higher resolution than our original prototype.

The ideas of DiffuserCam are not limited to image capture; they can also be applied to display systems. For example, **holographic displays** are a promising technology that can generate 3D content but suffer from limited



3D reconstruction of a tilted USAF resolution target

fields-of-view. We overcame this challenge by adding our diffuser into the system. The diffuser scatters light to higher angles, **expanding the field-of-view with simple, static hardware**. As before, we model light propagation through the system and use optimization-based approaches to compensate for the diffuser when computing the hologram.

For more information, read the full publications on <u>3D DiffuserCam</u>, <u>fluorescence microscopy with DiffuserCam</u>, and <u>diffuser-based holographic</u> <u>displays</u>. We also have a great tutorial on how to build your own DiffuserCam</u> with a Raspberry Pi camera and scotch tape as the diffuser!



* I'd like to acknowledge my advisors <u>Laura Waller</u> and **Ren Ng**, who worked with me on all of these projects, and my co-authors, **Nick Antipa** (DiffuserCam) and **Andrew Maimone** (holographic displays). This work wouldn't have been possible without them.

Single-shot 3D Imaging with DiffuserCam

18 Computer Vision Tool

Screening toddlers for Autism using the dataset from New Zealand Manukau Institute of Technology



IOANNIS VALASAKIS, KING'S COLLEGE LONDON



I hope everyone is well and safe and the same for the people you love and your families. As more vaccination efforts are rolled out in the world, it proves how effective science has been. For a critical view on vaccine safety, you can read a very interesting article on Nature 590, 538-540 (2021).

In continuation to the theme of last month's research, autism

is again what we are going to be looking at. Specifically, using the dataset curated by **Fadi Fayez Thabtah, Manukau Institute of Technology in New Zealand**, we are going to explore whether it is possible to recognise traits before an Autism Spectrum Disorder (ASD) for toddlers. This would be a very valuable tool for parents, as it can provide a hint on their children's condition and react early.

Autism Spectrum Disorder

The dataset is available from many sources online. Specifically, a few different approaches have been followed on Kaggle, but in this article we will present an approach developed on the basis of those Quantitative Checklist for Autism in Toddlers (Q-CHAT) data.

The Q-CHAT could detect forms of autism spectrum disorders for the age of 18 to 24 months. It contains 25 likert-type items (quick questions) and describes behaviours a child may or may not engage in.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
import seaborn as sns
%matplotlib inline
df1 = pd.read_csv('Autism_Data.arff', na_values='?')
df2 = pd.read_csv('Toddler Autism dataset July 2018.csv', na_values='?')
```

This will display the first rows of the dataset described earlier. Now let's get the info for the table values

dfl.info()

which returns

<class 'pandas.core<="" th=""><th>e.fra</th><th>ame.DataFr</th><th>ame<mark>'></mark></th></class>	e.fra	ame.DataFr	ame <mark>'></mark>		
RangeIndex: 704 ent	ries	s, O to 70)3		
Data columns (total	21	columns):			
A1_Score	704	non-null	int64		
A2_Score	704	non-null	int64		
A3_Score	704	non-null	int64		
A4_Score	704	non-null	int64		
A5_Score	704	non-null	int64		
A6_Score	704	non-null	int64		
A7_Score	704	non-null	int64		
A8_Score	704	non-null	int64		
A9_Score	704	non-null	int64		
A10_Score	704	non-null	int64		
age	702	non-null	float64		
gender	704	non-null	object		
ethnicity	609	non-null	object		
jundice	704	non-null	object		
austim	704	non-null	object		
contry_of_res	704	non-null	object		
used_app_before	704	non-null	object		
result	704	non-null	int64		
age_desc	704	non-null	object		
relation	609	non-null	object		
Class/ASD	704	non-null	object		
<pre>dtypes: float64(1), int64(11), object(9)</pre>					
memory usage: 115.6 + KB					

What does Rita Cucchiara wish Computer Vision News for our 5th anniversary? Find out on page 31!



Now let's get the info for the second dataframe (you can also do run the head command like before, to get an idea of what's in there).

df2.info()

which returns

< class <mark>'pandas</mark>.core.frame.DataFrame'>					
RangeIndex: 1054 entries,	0 to	1053			
Data columns (total 19 col	lumns)	:			
Case_No	1054	non-null	int64		
A1	1054	non-null	int64		
A2	1054	non-null	int64		
A3	1054	non-null	int64		
A4	1054	non-null	int64		
A5	1054	non-null	int64		
A6	1054	non-null	int64		
A7	1054	non-null	int64		
A8	1054	non-null	int64		
A9	1054	non-null	int64		
A10	1054	non-null	int64		
Age_Mons	1054	non-null	int64		
Qchat- 10 -Score	1054	non-null	int64		
Sex	1054	non-null	object		
Ethnicity	1054	non-null	object		
Jaundice	1054	non-null	object		
Family_mem_with_ASD	1054	non-null	object		
Who completed the test	1054	non-null	object		
Class/ASD Traits	1054	non-null	object		
dtypes: int64(13), object(6)					
memorv usage: 156.5 + KB					

Exploring the data further

Let us see the percentage of adults and toddlers having ASD.

```
sns.set_style('whitegrid')
adults = df1[df1['Class/ASD' ]== 'YES']
toddlers = df2[df2['Class/ASD Traits '] == 'Yes']
print("Adults (%): ", len(adults)/len(df1) * 100)
print("Toddlers (%):", len(toddlers)/len(df2) * 100)
```

Which returns

```
Adults: 26.84659090909091
Toddlers: 69.07020872865274
```

Recognise Traits of ASD 21

Hmm...! From the general population statistics, we do know that ASD is in about 1-1.5% of the whole. Then why is this so off?

The reason is that the features of the test parameters define only the qualitative properties of the ASDs. Let's run a visualisation for the missing data.







What does Gérard Medioni wish Computer Vision News for our 5th anniversary? Find out on page 30!



22 Computer Vision Tool

Now let's explore some of the traits included in this dataset. From the earlier exploration, we noticed jaundice, included in those. Jaundice is a yellowing of the skin and the whites of eyes that happens when the body does not process bilirubin properly. The common name is icterus and bilirubin is the yellow-colored waste material that remains in the bloodstream after iron is removed from the blood and it might signal a liver problem. But what's the relation with ASD? Let's find out.

```
fig, ax = plt.subplots(1, 2, figsize=(22,8))
sns.countplot(x='jundice', data = adults, hue = 'gender', ax = ax[0])
ax[0].set_title('ASD positive Adults born with jaundice based on gender')
ax[0].set_xlabel('Jaundice while birth')
sns.countplot(x = 'Jaundice', data = toddlers, hue = 'Sex', ax = ax[1])
ax[1].set_title('ASD positive Toddlers born with jaundice based on gender')
ax[1].set_xlabel('Jaundice while birth')
```



The ASD prominence is about six-fold in adults and 2-fold in toddlers comparing to non-jaundice born individuals. According to research, this can be up to 10-fold! Children born with jaundice have a weak link with ASD and it is more common among the males.

Now how about the distribution in the population?



Recognise Traits of ASD 23

The interesting fact is that for the adult group, the ASD positive lies around 20-30 years. For the toddlers that is ~36 months. There is a linear negative relation between the ASD and the age for the adults, while the reverse is true for the toddlers. One reason is that there are strategies developed to assist adults with ASD (and here is why this is very important from a social point of view and policies to be created). For children, the autism shows in about 3 years of age. Again, this is very important in the sense that as the prognosis of ASD (or even better the diagnosis) can be achieved in earlier stages, more strategies can be developed to help those children, adapt, learn, interact and thrill; much better than today.

Wrap up and future...

I do wish this was a valuable approach to using an existing dataset and understanding if it's possible to get a prognosis for ASD for a toddler. In the June magazine, my article will continue to explore the dataset and reveal more features that can be extracted and predicted.

Approaches such as this are currently developed into tools and frameworks, that are tried on clinical settings to infer the possibility of using this in a clinical workflow. This would add agility and will provide with hints for the practitioners to better understand such developments. Till next month, enjoy ③



24 Challenge

Fetal Brain Tissue Annotation and Segmentation Challenge

Roxane Licandro works at the Vienna University of Technology (TU Wien) and at the Medical University of Vienna (MedUni Wien) as a Project Assistant and University Assistant. Kelly Payette is a PhD student at the University of Zurich and is currently working in the Center for MR **Research at the University Children's** Hospital in Zurich. Roxane and Kelly are co-organizers of the Fetal Brain **Tissue Annotation and Segmentation** (FeTA) Challenge, which takes place this year during the 6th annual Perinatal. Preterm and Paediatric Image Analysis (PIPPI) Workshop at **MICCAI 2021.** They speak to us about the competition and why it is so vital that we understand more about fetal brain development.

Fetal MRI has begun to emerge as a powerful tool for investigating the development of fetuses with congenital disorders, which are one of the leading causes of infant mortality worldwide. Typically, when a pregnant mother has an ultrasound, she is referred on to an MRI if a potential issue is found. In the majority of these cases, **a brain disorder is the concern**.

The FeTA Challenge is about





Kelly Payette

Roxane Licandro

finding automatic solutions for the segmentation of fetal brain tissue into different individual tissues, including white matter, gray matter, cerebellum, and brainstem. This will improve the diagnostic process over more time-consuming manual methods and help us to understand more about fetal brain growth in both normally developing cases and those with congenital disorders.

"The brain is a very important organ to observe," Roxane tells us.

Fetal Brain (FeTA) 25

"It controls just about everything we do, including our motor control and our breathing. We want to assess its growth patterns as early as possible so that if surgical or other interventions are needed, they can happen at the right time, leading to **better outcomes for the children** concerned."

This is not a typical segmentation challenge because the fetal brain changes drastically throughout gestation. The challenge introduces a dataset of around 70 cases – with more to be added by the time it is released – aged between 20 to 33 weeks.

The training set aims to bebalanced across that gestational age range as the structural changes that occur for all fetuses within that time frame are vast.



"This is a challenging task to solve," Roxane asserts.

"What we have here is a highly variable dataset, particularly in terms of the pathological cases, and in clinical settings we must be able to segment these highly pathological



26 Challenge

Fetal Tissue Annotation Challenge 2021



cases. We have brought this challenge to life at **MICCAI** to create a platform for all researchers working in the segmentation field to propose new approaches and techniques to help us solve these problems."

There aren't many centers currently using these kinds of imaging techniques nor with access to a high volume of cases, so any algorithms must be able to handle a relatively low number of datasets. The hope is the challenge will bring **new data to the community** and solutions and approaches which can ultimately be transferred to the **clinical routine**.

"We want this challenge to encourage research groups from different fields beyond just developmental imaging to pull together their expertise in segmentation to create these algorithms," Kelly says hopefully.

"Then other groups and clinicians can use the algorithms to analyze their own data in their own clinics. It's a problem that needs to be solved in many different hospitals and this is a great opportunity to come up with solutions." Once the challenge has completed, the organizers are planning a **joint publication** in a journal paper, so as well as being able to compare results with fellow competitors, participants will be able to propose their resultsto the whole community. This will lead to much greater visibility.

As a new challenge on the block, how easy was it for FeTA to be selected for inclusion in the MICCAI programme?

"Our reviewers were generally really positive," Kelly responds happily.

"In the fetal imaging community, there is a lack of open-source data available. It's time consuming to create it and you need highly specialized people to segment fetal brains. There are also ethical issues around releasing it. It's incredibly challenging. This really is a unique dataset that we're offering to the MICCAI community. Everybody has been very supportive."

Good luck to everybody taking part in the FeTA Challenge this year. We very much hope to find out the results, if all goes well, in person in Strasbourg in September! What do Cristian Canton and Laura Leal-Taixé wish Computer Vision News for our 5th anniversary? Find out on page 29!



What does Wiro Niessen wish Computer Vision News for our 5th anniversary? Find out on page 29!

What does Nicu Sebe wish Computer Vision News for our 5th anniversary? Find out on page 28!

28 Happy 5th Anniversary

Happy Anniversary to Computer Vision News!

First of all, let's give a big thank you to Ron Soferman, CEO of RSIP Vision, who backed the entire 5 years of publishing Computer Vision News, a free magazine with no advertising. May he continue offering this precious contribution to the community for many years to come!

We also want to thank all the friends who sent their personal wishes: we value the many years of scientific work, the warm friendships, as well as the awesome talent and energy behind each message. Thank you. It feels great!

Congratulations on 5th your anniversary, Computer Vision News! You have been an integral part of the CV community, always there to follow the latest trends and report the highlights of computer vision, from Venice to Seoul, from Munich to Long Beach. I wholeheartedly enjoy the personal touch you bring to conferences, which would otherwise be dull and overwhelming! I hope you stay with us for another 5 years and many more after that! Thank you Ralph for all the hard work and the joy you bring to our computer vision lives!

Georgia Gkioxari - Facebook Al Research

Dear Ralph, great to realize that your magazine has been accompanying the CV community for 5 years already. Thank you to the whole team for your wonderful, impressive initiative and indefatigable work. Here's to 20 years more!

Emanuele Trucco - University of Dundee

Congratulations on hitting the 5-year mark. Thank you for your fantastic coverage of the amazing work being performed in the field(s). The magazine is so well-written and enjoyable, with articles that provide insight into the research as well as into the researchers as individuals. Well done!

Tal Arbel - McGill University

In my opinion, Computer Vision News has brought all these years fresh air to the community. It has become a standing point at our conferences and I am looking forward to seeing the next editions. Happy anniversary!!! **Nicu Sebe - Università di Trento**

Congratulations to RSIP Vision on its 5-year anniversary! Ralph's hard work has paid off in creating the leading computer vision magazine.

Su-Lin Lee - WEISS Centre, University College London

For 5 more years of CVPR dailies! With fun, insightful interviews, helping us discover the B-side of CVPR!

Laura Leal-Taixé - Technical University of Munich

Congrats on this huge milestone: 5 years! I can't conceive a conference without your magazine which shares relevantinsights, personal perspectives and a touch of humor. Wish you many more years to keep informing the CV community!

Cristian Canton - AI Red Team, Facebook

Computer Vision News has been a true gift to the computer vision community, helping us to recognize and celebrate many of the talented individuals who help our community be what it is! Danna Gurari - The University of Texas at Austin

"a true gift to the computer vision community"

I met Ralph, the editor of Computer Vision News, 5 years ago at MICCAI 2016. I was impressed by the enthusiasm plus the gift he and Computer magazine showed Vision News communicating finding in and genuine personalities and stories from the attendees of the event. Congratulations to the whole team and I trust they keep up the good work! Gozde Unal - Istanbul Technical University

Congratulations on 5 year, Computer Vision News!! I remember how at MICCAI 2016 in Athens we had your daily editions for the first time. By now, they have become an integral part of a MICCAI meeting. Thanks for your contributions to the field!

Wiro Niessen - Erasmus MC, TU Delft and Quantib

Congratulations, Computer Vision News! If I would've known that Ralph would become so famous when I met him in Amsterdam in 2016, I would've asked him an autograph! Jordi Pont-Tuset - Google Research

Congratulations to Ralph Anzarouth and the Computer Vision News team for five years of outstanding work dedicated to inform and keep up to date with the latest developments in our community and for presenting its human part as well. This is definitely an accomplishment to be proud of! It is a great service to our community and we are all very thankful for it -keep going! I look forward to reading and contributing for many years to come!

Leo Joskowicz - The Hebrew University of Jerusalem

Congratulations on all your efforts in bringing the magazine to where it is today!

Fırat Özdemir - Swiss Data Science Center

30 Happy 5th Anniversary

Computer Vision Happy Birthday, News! Thanks for being with us, no matter where our conferences go :) Zeynep Akata - University of Tübingen

Huge congrats to RSIP Vision on a five year anniversary -- wow, time flies! **Olga Russakovsky - Princeton University**

Heartfelt congratulations to Computer Vision News on its 5th year anniversary and for the extraordinarily work you have done and will continue to do. Over these years, your publications have forged an impressive communication tool for the algorithm community. I hope your work keeps connecting passionate readers as it has been doing with each issue.

Nour Karessli - Zalando SE

"Time flies!"

Many congratulations to five years of Computer Vision News! This is a great achievement and shows that dedication and determination pays off! I wish you many more years of successful writing and hope to see you again soon next to one of our posters!

Bernhard Kainz - Imperial College London

Ralph love how makes L every interview turn out to be so relaxed and conversational. Hope we chat in person one day!

Arsha Nagrani - Google Al

Time flies! I cannot believe it has been 5 years since we first met at ECCV 2016 in Amsterdam. The Computer Vision News magazine is really a special thing for the computer vision community. It has accompanied us for five years! I really enjoyed the stories and shares by colleagues from the community when I was attending conferences. Wish it continues as an exciting place for sharing news and ideas in the future!!

Miaomiao Liu - Australian National University

Whenever I attend CVPR (always) or ICCV (mostly), I look forward to opening Ralph Anzarouth's daily news update. In particular, it brings the to life the people behind the science.

Gérard Medioni - Amazon

Happy quintanniversary to Computer Vision News, and both congratulations and thanks on becoming an essential component of the MICCAI experience, linking and bridging together our community!

Simon Duchesne - Université Laval

The best way to start a conference day - a huge cup of coffee and the morning issue of Computer Vision News! Happy Anniversary, Computer Vision 5th News!

Roxane Licandro - Medical University of Vienna and TU Wien

Through the efforts of all those dedicated to Computer Vision News, the magazine has been reflective of the developments of the Computer Vision society and has promoted and communicated consistently highquality research to the profession. It is an honor to extend my congratulations on its 5th year anniversary. Keep up the good work!

Faezeh Tafazzoli - Mercedes-Benz

"The best way to start a conference"

It is my pleasure to wish the wonderful Computer Vision News - Congrats 5th anniversary! Though it seems challenging, I hope that the next 5 years will bring us even more breakthroughs in our fascinating field of research! Thanks to you for these tireless efforts! **Tobias Würfl - Siemens Healthineers**

Happy 5th Anniversary, Computer Vision News – and many happy and informative returns! Keep up the great work, Ralph (3) All best wishes from Julia Schnabel - King's College London

Congratulations on your 5th anniversary! Your contributions to computer vision community are greatly appreciated. I wish you continued success and I hope we all can gather at a traditional in-person conference once again!

Shekoofeh Azizi - Google

Dear Computer Vision News Happy Birthday!!!

Without you, Computer Vision events would never be as interesting as they have been in the past, and as I hope they will be in the future! Interesting news, deep overview of what is emerging, faces of friends and new faces to know! Thanks a lot to exist!

Rita Cucchiara - Università di Modena e Reggio Emilia

For the past 5 years, Computer Vision News has been introducing not only our scientific discoveries, but also our human side to the community, from its unique perspective. Thanks CVN and Ralph for the always supportive, unbiased, and timeous news! To many more 5 years ahead!

Ilke Demir - Intel

"Interesting news, deep overview"

Thanks to Computer Vision News bringing the community closer during every MICCAI for the past years, even a remote one :-)

And none of it would be possible without the energy and drive of Ralph. Happy anniversary!

Lena Filatova - Philips Image Guided Therapy Systems

Congratulations Ralph! Keep digging into the core of computer vision! **Muxingzi Li - INRIA, interning at Ali Baba**

32 Happy 5th Anniversary

Warmest congratulations to Computer Vision News magazine for this great achievement! My best wishes for many more, as successful magazine issues to follow!

Stamatia Giannarou - Imperial College London

I cannot believe it has been 5 years since Computer Vision News has come into existence. I remember how excited I was to be interviewed for my first CVPR in 2017, and what interesting picks and "behind-thescenes" info it had about some allstar papers & researchers. Of course, the majority of those cases I couldn't attend the presentations for those papers because the CVPR venues are huge and there are usually parallel sessions happening, but thanks to Computer Vision News, I became aware of them and marked them to look up later. I also remember you were pulling all nighters to make sure it was published on time every day the next day. Well done, and congratulations on your 5th birthday!

Derya Akkaynak - Harbor Branch Oceanographic Institute, Florida Atlantic University

Great service to the vision community!

Matthias Niessner - Technical University of Munich

It has been a challenging year for all of us. Fortunately, the Internet and technology have connected us far more closely than ever before. I am amazed to see how computer vision and machine learning technologies transformed the wav have we communicate, think, and live in the past year. Moving into a new chapter, am looking forward to seeing Computer Vision News magazine to share more exciting news of people and technology in this community and bring us all together.

Xin Lu - Adobe

Congratulations Computer Vision News! Another year, another data point, and your model keeps getting better and better! Thank you for all your fantastic work and knowledge sharing.

Dan Stoyanov - WEISS Centre, University College London

Happy five years anniversary to Computer Vision News! Thank you for connecting the community and for uplifting women in computer vision. Looking forward to many great conversations to come, both virtual and in person post-pandemic. Best wishes for the next five years and beyond!

Amy Bearman - Facebook Reality Labs

Computer Vision News 33

Dear Ralph and the Computer Vision News team, a big congratulations on the 5-year anniversary of the magazine! Always look forward to reading the magazine at conferences. Here's to another 5 exciting years! Marta Kersten-Oertel - Concordia University

"It has been 5 years"

It has been 5 years since you started this initiative. It seems like ages as the Magazine reaches out wide and deep into the multiple communities liked to computer vision and its numerous application domains. Computer Vision News, for instance, has brought the work in MICCAI to the attention of the wider computer vision community and vice versa. This interaction was timely over the past 5 years, where both got benefits, I believe, in interacting and learning from each other.

Yet it has been only 5 years! I can't believe it: nearly 4,000 pages of scientific knowledge shared for free, reaching over 4,500,000 total pageviews.

We cannot thank you enough, and with you to RSP Vision, for your leadership and fantastic work. I am joining many other colleagues in wishing you all the best in the next 5 years and beyond! **Alejandro Frangi - University of Leeds** Thanks to Ralph, year after year Computer Vision News amazes us, makes us smile, helps us discover innovations and people. The magazine has become an essential and much awaited partner of our MICCAI weeks. Thank you Ralph and happy birthday to Computer Vision News!

Caroline Essert - Université de Strasbourg

Wow it has been five years! Congratulations! I really enjoyed my interview with Ralph -- he asked such great questions -- and I can't wait to see more awesome content from RSIP Vision.

Chip Huyen - Stanford University

"You're the best !! "

A big congratulations to Computer Vision News Magazine for their 5th year anniversary!! What many may not know is that this magazine is organized by Ralph and RSIP Vision and it is a completely independent process from the conference organization. I know this is an immense amount of work and your passion shines through. Thanks so much Ralph for your dedication to the field and fostering our community. Your articles bring humanistic aspects to our colleagues that I enjoy reading greatly. Looking forward to the 10th year anniversary :)

Angjoo Kanazawa - UC Berkeley

34 Happy 5th Anniversary

Congratulations for your five year anniversary of Computer Vision News, empowering our research community with much needed information. May the next five be even more successful and thrilling both for the research community and the magazine.

Alon Baram - TAU and Biosense Webster, a J&J Company

Happy CVNews-anniversary, 4000 pages is quite an impressive achievement... glad to be part of this adventure! Congratulations!

Aïcha BenTaieb - Tempus Labs, Inc.

Congratz Computer Vision News for your 5th anniversary. Keep publishing and giving us awesome and cuttingedge news from the CV world. You're the best!

Amnon Geifman - Weizmann Institute of Science, Deci AI and former CVN editor

Many congrats on 5th anniversary, Computer Vision News! Thanks a lot for your inspiring interviews and bringing our ever-growing CV community closer together!

Hyo Jin Kim - Facebook Reality Labs Research

Cheers to your special day! Happy 5th Anniversary! Wishing more success in the upcoming future!!

Qi Dou - The Chinese University of Hong Kong

Congratulations to the creative team of a great magazine - and especially to the marketing manager with the biggest smile (Ralph A.)

Lena Maier-Hein - German Cancer Research Center (DKFZ)

4,5 Million page views! What a colossal achievement! Big congratulations! Thanks for connecting us to our community and making us united in our diversity!

Islem Rekik - BASIRA Lab, Istanbul Technical University

Congratulations on 5 outstanding years of Computer Vision News, Ralph! As our conferences grow in size, it is always a pleasure to learn about all of the exciting new developments in our field through your entertaining and informative magazine.

Serge Belongie - Cornell University

Happy 5th Birthday Computer Vision News!

Congratulations on the amazing job that you do at presenting, spreading and celebrating great work and great people in the computer vision community. Thank you also for having the opportunity to read about inspiring Women in Computer Vision. I look forward to reading more of your editions in the future! Jelena Frtunikj - ArgoAI

Computer Vision News 35

Dear RSIP Computer Vision News team,

Congratulations on the 5 year anniversary of your magazine! The updates and stories from the magazine are of great relevance and interest to our community, as exemplified by the >4,5 million page views. Look forward to more exciting updates from your team!

Kristin McLeod - GE Healthcare

"4,5 Million page views!"

Congratulations to Computer Vision News for the 5th anniversary! Thank you for all the hard work in bringing us, readers the latest updates in computer vision and in being a platform for the community to keep in touch. Wishing you many more successful years to come.

Adriyana Danudibroto - Agfa

Happy 5th anniversary to Computer Vision News! With my best wishes of success for the coming years and an ever friendly coverage of our conferences!

Carole Sudre - University College London

Happy Birthday Computer Vision News! Always looking forward to researcher portraits and research highlights in your pages.

Marc Pollefeys - ETH Zürich and Microsoft

Since its inception, Computer Vision News has been an exciting, thoughtful, and informative publication driven by a kind of genuine enthusiasm that can't be found outside of the computer vision community. Thank you for 5 years of excellence!

David Chambers - Southwest Research Institute (SwRI)

I would like to send my best greetings for the 5 year anniversary of Computer Vision News and deep congratulations for the amazing work. Through the special issues at CARS and MICCAI international conferences, for instance, Ralph and the magazine created a media that knew how perfectly reflect the scientific to excellence of the presented research and the human atmosphere. He and it were always relevant, useful and more than this did an important effort on gender parity with the women in science section. Above the media, Ralph and the magazine brought a very positive and warm and hearty spirit. I wish to see him again for years.

Pierre Jannin - Université de Rennes

Happy birthday Computer Vision News!

You are the glue that binds us together!

Vittorio Ferrari - Google

36 Happy 5th Anniversary

I've always loved the fact that our discipline, despite its rapidly growing size and popularity, managed to retain much of the friendliness and collegial character of its early days. So, when I heard that we are getting our own magazine, at first I was quite apprehensive. Is this going to be some shiny, glossy thing like MIT Tech News or Gizmodo, or crazy selfpromotional rumor mill like Reddit, Hacker News, or, god forbid, academic twitter? But once I read the first few issues of the Computer Vision News, I was happy to see that my fears were totally misplaced. The magazine has the feel of a small town newsletter, with research news sharing the pages with casual interviews, photos of group hikes and people's pets, profiles ofup-and-coming young researchers, advice columns on how to enjoy grad school, etc. Indeed, this is exactly what we need to retain the friendly, no-drama culture of our discipline while making sure that the new-comers are welcomed and don't feel overwhelmed.

Happy birthday, Computer Vision News!

Alexei (Alyosha) Efros - UC Berkeley

"Happy birthday, Computer Vision News!"

To Ralph and all members of the Computer Vision News magazine, congrats on celebrating your 5th anniversary!!

It has been a while since we had our in-depth interview 5 years ago, we've both grown during this time, and today I feel fortunate to have Nanit featured on one of your first magazine editions, in August 2016 ;)

Keep on the good work of providing quality news to the computer vision community!!

Assaf Glazer - Nanit

"Congratulations to the RSIP Vision team and

especially to Ralph!"

It has been a while since we met in Athens and once again in Shenzhen! Indeed, RSIP Vision is among a very few magazines which highlight the news of scientific events! I wish you and RSIP Vision a great success!

Shadi Albarqouni - Helmholtz AI and TUM

Congratulations to the RSIP Vision team and especially to Ralph! Thanks for your interest in the computer vision community and also for your enthusiasm and dedication. In bocca al lupo!!

Clara Fernández - ETH Media Technology Center

Computer Vision News 37

Wow, five years! Thank you to all the RSIP Vision crew that creates and brings Computer Vision News to us. It's a remarkable contribution to our industry by wonderful people. Wishing you all the best for many more years,

Jason Knapp - Riverain Technologies

Congratulations on the 5 year anniversary! I first came across Computer Vision News when they did the "CVPR Daily" in 2016. Since then, it has been a delightful part of my conference experience, highlighting a variety of interesting events, topics and people I would have likely missed otherwise. Thank you for what you do!

Devi Parikh - Georgia Tech and Facebook Al

Happy anniversary to the Computer Vision News magazine, and many thanks and congratulations to Ralph for his job, his enthusiasm and ability to transmit positive energy!

Looking forward to meeting again in live for celebrating together this anniversary and all the other important events that have happened in the past 12 months and still waiting to be celebrated in an appropriate manner!

A warm greeting,

Silvia Tozza - Università di Napoli Federico II

Many thanks to Computer Vision News for creating quality content for years and highlighting daily updates at major conferences for the community.

Gül Varol - Ecole des Ponts ParisTech

"must read every morning!"

Congratulations to Computer Vision News on its 5th anniversary, the superfantastic easy-to-read magazine for computer vision researchers and practitioners!

Óscar Déniz Suárez - Universidad de Castilla-La Mancha

Congratulations Computer Vision News! You have been keeping us well informed during some of the most dynamic times in our field. Thank you for your outstanding coverage of the important advances and the researchers behind them. Looking forward to many more years of learning with you!

Stephen Aylward - Kitware

Congrats to all the Computer Vision News team for delivering the latest news as well as giving a voice to underrepresented groups in the computer vision community for the last 5 years. During conferences, the daily edition has always been a must read every morning!

Sergi Caelles Prat - Google Research

38 Happy 5th Anniversary

Congratulations on 5 years of Computer Vision News! Great work you've done over the years :) 4,5 Million page views is impressive but not surprising! Keep up the good work, and I'm looking forward to the next 5 years of Computer Vision News! Luisa Zintgraf - University of Oxford

Happy anniversary, keep up the inspiring work!

Stefanie Speidel - NCT Dresden

"Congratulations on the five years of your super magazine"

Congratulations on the 5 years of Computer Vision News! It feels like yesterday, when we were chatting in Amsterdam and Athens! Time indeed flies! Thank you for promoting our work, you are amazing! I am always impressed by how many details you remember for each person you meet in the hundreds of conferences! **Kevis-Kokitsi Maninis - Google**

Your news magazine is great. I look at it from time to time. Nice to stay connected with the world that way. The content is up-to-date, interesting and detailed enough to be able to make me curious and go to the related reference. Keep up the great job! Dana Cobzas - MacEwan University

"It is already 5 yearsss!! So niceee!!!"

Many congratulations on the 5th anniversary of Computer Vision News. Thank you for supporting our community of medical imaging and computer assisted interventions by writingabout our innovations, and helping highlight the many achievements of the people in this community. I look forward to grabbing a coffee and chatting in person sometime soon.

Parvin Mousavi - Queen's University at Kingston

As our field has grown huge, it could also have grown impersonal. RSIP Vision has played a critical role in helping us keep a strong sense of community that is inclusive and welcoming to everyone.

Michael J. Black - Max Planck and Amazon

I can hardly believe that it was almost 5 years ago in Las Vegas that we first met. Many congratulations to Computer Vision News on the five year anniversary! I am looking forward to more Computer Vision News, and to the conference Daily News in person.

David Fouhey - The University of Michigan

Computer Vision News 39

Happy 5th birthday Computer Vision News! Keep the nice work to bring the community together with interesting scientific and social news! :) It is already 5 yearsss!! So niceee!!! Maria Vakalopoulou - CentraleSupélec, Paris

Congratulations on five years! I hope you keep making scientific content sohuman and relatable for many years. You do an amazing job!! María Victoria Sainz de Cea - IBM Watson Health Imaging

Happy Anniversary! It's a great magazine that presents things in an accessible and engaging way. I wish we had more publications like this! Adriana Kovashka - University of

By focusing on the human aspect behind our work, Computer Vision News has done an invaluable service to our community. Thanks Ralph for what you're doing, and happy anniversary!

Marcello Pelillo - Università Ca' Foscari Venezia

Congratulations on your milestone of sharing 5 years of computer vision knowledge! I sincerely wish you and the team continued success.

Modar Alaoui - Eyeris

Pittsburgh

Congratulations to Computer Vision News and to Ralph for editing it over the last 5 years! Thanks for keeping us together even when COVID doesn't allow us to meet personally! Sandrine De Ribaupierre - University of Western Ontario

"Thanks for keeping us

together even when COVID

I always find your vision conference summaries really helpful. Please keep up the good work and congrats on reaching this milestone.

Mohammad Musa - Deepen Al

Congratulations Computer Vision News, we hope you continue everybody's make computer to special. conferences more vision From the moment met we **CVPR** MICCAI) Ralph, our (and experiences have been enriched! **Amy Zhao - Facebook Reality Labs** Adrian Dalca - Harvard Medical School and MIT Guha Balakrishnan - MIT

Congratulations on the five years of your super magazine - I am so happy for you!!! You have changed our community and the vision landscape! Wonderful months and years and centuries to come! Vicky Kalogeiton - École Polytechnique

40 Women in Computer Vision

MOME

Toilets

Emma Robinson Emma Claire Robinson works at King's College London as lecturerforthe Department of Biomedical Engineering, Division of Imaging Sciences & Biomedical Engineering.

DRINK I

Emma, tell us about your work!

My work is focused on the development of new machine learning and advanced techniques processing for image analyzing medical imaging data of the brain. Nowadays, I'm much more focused on cortical modeling, which is trying to understand the outer surface of the brain: the area of the brain which is responsible for complex thought and cognitive processes that are implicated a wide range of neurological in psychiatric disorders. I am also starting to be more interested in cognitive neuroscience, modeling of how the brain works.

How are things going?

We've had some success. My biggest area of focus, at the moment, is the fact that traditional approaches for analyzing this data make simplified assumptions, which make it very difficult to pick up on subtle differences between individual datasets. Historically, people make inferences at the population level: "On average, people with Alzheimer's have the following brain properties" and "On average, people with schizophrenia have the following brain properties". But a lot of these diseases that affect cognition originate in highly complex parts of the brain and are highly variable across individuals. So, they tend to be confounded by individual variability in brain shape and brain organization, which nobody really ever takes into account. We had a recent

"You cannot treat brain data in a data-agnostic way!"

NeurIPS paper about deep generative modeling on MRI brain datasets, which very specifically looked to disentangle this issue. It has one latent space which models natural brain variability and one latent space that models the variability that is specific to the disease or the phenotype. So that works quite well actually.

I have interviewed a couple of scientists who work in that area in Canada: <u>Sandrine de Ribaupierre</u> and <u>Marta Kersten-Oertel</u>. I know you worked with <u>Carol Sudre</u>.

Yes, Carol is on that paper. She helped us with all the Alzheimer's data.

Marta taught me that some areas of the brain can be touched by the surgeon, while other areas shouldn't.

I don't know those specific areas as well as she would. A lot of the data stuff that I've done has been on modeling healthy brains to develop new methods which we can use for disease. We're starting to look at trying to detect the signs, the origins, the seizure point for epileptics, things like that. Do I know which areas of the brain are more vulnerable than others? Not so much... [*laughs*] It's a difficult question!

42 Women in Computer Vision

Did you choose the brain because of the Human Connectome Project or did you decide it independently and you joined the project after that?

I was working on the brain beforehand, so the whole time I've worked on the brain. Originally, I joined a lab at Imperial College modeling baby brains. I did my PhD between there and the computer department at Imperial. I've kind of come full circle. I've always worked with brains. I don't have any interest really in working with other organs. Everyone tells me to diversify. I think it's a big enough problem in itself. It's really important. You need a lot of specialist knowledge. You cannot treat brain data in a data-agnostic way. You need to have some understanding of the constraints of image acquisition and neuroscience to build informed models. You have to have sensible constraints. You've got very high dimensional data

space, very complex multimodality data. Knowing something about the underlying neuroscience can really help.

Is the reason for learning about the fetal brain to be able to one day intervene before the birth of the child?

In my opinion, this is not about trying to detect whether someone might have some condition. We're never going to be able to be that certain. Also, this is not a purely genetic thing. Just because someone has a set of genes, the basic initial blueprint that may indicate that there might be a possible disability would develop, that's that never going to be set in stone. A lot of these complex conditions have significant environmental influence. My interest is trying to understand, from cognitive neuroscience, the brain and maybe help identify early signs to help more people at the earliest possible stage. If you're

"You must absolutely understand neuroscience!"

METRICS Lab at King's College London

talking about fetal screening, absolutely not. It's an incredibly difficult question. It's a balance. There's a huge amount of good that could be done, but as with so many things in science, you have to make sure you use it for the good.

I've known you for only a few minutes, Emma, and you told me several times that the goal of your work is to understand things. When I asked you about their practical application, like fetal surgery, it sounded like that is not exactly what you are aiming at. Can we go beyond understanding and find some practical applications for your study?

You're talking about the difference between work that has an immediate path to the clinic and work that has a long-term path. In terms of immediate impact, we're looking at tools for detecting signs of epilepsy for surgical planning for epilepsy. Are we anywhere near being in the position to detect really early signs of dementia from MRI? No. When you start getting into complex psychiatric heterogeneous neurological disorders, then I think the question becomes much more complicated. Do I have a practical goal in mind? Yes. Am I tackling more difficult problems than can be immediately delivered to the clinic? Yes.

You talk about detection as a practical application, but in a non-scientific mind like mine, the only practical

"We want to propagate that information using a combination of novel, portable technologies, apps, wearables, and all sorts of things..."

application is something that you intervene on. Detecting doesn't seem like a practical application, to someone like me.

Let's go back to the basics, for example. Let's say that through our analysis, we are able to detect, at an early age, which babies are more likely to be at risk of long term cognitive impact. Then you have a set of biomarkers from which you can use in clinical trials of potential treatment. The same for neurodegenerative diseases. The idea is that you want to be able to detect early signs of disease that can subsequently be used as biomarkers for clinical trials and interventions. Also, to develop causal models of disease. The direction that we're going in now is to build a model of how a brain might change over time and create, what they call, counterfactuals, from which one might be able to pose different questions, and through that, understand different disorders.

You are dedicating the best years of your career to this subject. What would be your dream achievement?

What we want to be able to do is to detect an individual with a high level of

44 Women in Computer Vision



accuracy: what are the chances for them to go on and develop some problem later in life?

Is there anything specific that you would like to achieve, even if it would be difficult?

Can you really reliably detect at birth which babies will go on to be at significant risk of longneurological impairment? term That's really difficult. There are environmental factors. There are compound behavioral measures. You've got to marry many high dimensional datasets. You've got real problems trying to robust and generalize this. Can we reliably detect which babies are most vulnerable?

What else should our readers know about your work?

One is you can't treat brain data agnostically. Two is that you have to account for significant heterogeneity, brain's natural heterogeneity and heterogeneity in conditions. Brains are hard! That's my message. Brains are hard. [laughs]

You worked at Imperial and now you are working at KCL. Can you tell us how it is to work at these institutions? What are the advantages of working in that kind of setting?

King's is a really exciting place to be right now. There's a huge investment in healthcare technology research with some really great minds. If somebody wanted to develop training or healthcare technology, or machine learning for healthcare, it's absolutely the right place to be right now. They've got so much investment, huge GPU clusters, great people, diverse skills, a great marriage between clinical knowledge and technical knowledge, everyone's working collaboratively. It's part of the golden triangle. We still work with Imperial. We work with Oxford. Everyone's working together. It's a great place to be.

Emma Robinson 45

thing in your work right now, what would it be?

We don't have enough genetic samples.

Not enough data?

It's data, yes, and that's the problem. The problem is the datasets are too small. We need to solve that somehow.

How will you get data?

I think we need to move in the direction of new technologies where it's cheaper to acquire similar information.

Cheaper technology like X-ray and ultrasound?

Yeah, ultrasound, EEG...



Why has that not happened yet?

It is moving in the direction of portable technologies. I mean you can't scan everything, but we could do things hierarchically. There is the UK Biobank, which is this enormous adult dataset. We can build models on that. We can then propagate those slightly smaller baby datasets. Then what we want to do is build informed models, more complex models of what's going on, into the brain. We want to propagate that information using a combination of novel, portable technologies, apps, wearables, and all sorts of things that I don't know anything about, to be honest, but I would like to be able to make this more useful.

What do you think of that? When I was younger, I wanted to work in banking but my uncle told me, "Never work in banking. Banking is a bad career!" [both laugh] Finally I followed his advice. So if a young student decided today to go into medical imaging and study the brain, what would your advice be?

My advice is you must absolutely understand neuroscience. You do your best to understand, as far as possible, neuroscience. That will point you in the right direction.

Find out 100 more interviews like this with Women in Computer Vision!

46 Artificial Intelligence

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

Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks

Yes, you are right, the object in the image is not a **coffeemaker**. It is most probably a teapot. Researchers at **MIT** have found that popular test sets in machine learning are NOT immune to labeling errors. They made an analysis of 10 test sets from datasets used to train countless computer vision algorithms (including **ImageNet**) and compared that to their guess at what the correct label might be and to the consensus label among 5 mechanical Turk human raters. They estimate an average of 3.4% errors across the 10 datasets, where for example 2916 label errors comprise **6% of the ImageNet validation set**. Here are their <u>blog post</u> and <u>paper</u>. **Read More**



ImageNet given label: coffeemaker

Watch the winners of this year's 'Dance Your Ph.D.' contest

Let's dance! This is really fun: not-your-typical nerds have joined Science's annual "Dance

Your PhD" contest. If you didn't know that yet, it's 14 years that this goes on and the teams invest much effort in their dances. The overall winners (but not only them) wanted "to show nonscientific muggles that science can be fun, silly, and exciting". By doing that, they beat out 39 competitors for \$2000 - and **eternal geek fame**. The level of the productions is amazing! Our favorite is the winner of the

Biology category **"Fragmentation of plastics"**. Which one is your favorite? To all participants: congrats, doctors! **Read More**





Smart Safe Keeping: Blending Artificial Intelligence with Sea Turtle Conservation

Princess Aliyah Pandolfi says that "AI has advanced exponentially while drone development has expanded to many conservation studies" and she knows because she is the Executive Director of Kashmir World Foundation (KwF) and what she does there is worth of a Princess: she created the Fly for Conservation workshop to educate researchers and biologists on the value of custom drones embedded with AI on the edge. That scientific team builds, programs, and operates drones to fly over the Yucatan Peninsula and survey sea turtles and find patterns with a powerful CNN. By the way, the diver posing with that friendly sea turtle is truly yours, in Hawaii! **Read More**

Spotlight News 47

Can Artificial Intelligence Replace Human Therapists?

The WSJ article deals with the opportunity of relying on algorithms for our mental health - potential benefits and issues. It deserves a read, but it is behind a login wall. If you do not care subscribing, here is at least one of the key takehome thoughts of the feature: AI would be very powerful in analyzing large amounts of data generated from socialmedia posts, smartphone data, electronic health records, therapysession transcripts, brain scans and other sources to identify patterns that are difficult for humans to discern. However, the most frequent mental diseases do not require such a complex system to diagnose, let's say, depression. **Read More**





Reddit: When taking pictures goes horribly wrong!

Readers who follow us since the beginning, 5 years ago, know that we would have liked to have a blooper images section including the bugs of the month, submitted by our readers. We sometimes published a few of those that we found particularly funny, but we never made it a regular thing. No worries: there is actually not much need for that section, since Reddit is already doing that better than we could! Have a good laugh by watching many bug shots submitted by Redditors. Like this one called "Kitty's head is gone...". Yes, cat images are popular also on Reddit. **Find out many more bugs**



What does Pierre Jannin wish Computer Vision News for our 5th anniversary? Find out on page 35!

48 Computer Vision Application



Alejandro Galindo is VP of Research & Development at Iris Automation. Iris Automation builds detect-and-avoid systems for industrial and commercial drones. He speaks to us about their innovative work which is bringing together the conventional worlds of computer vision and geometry with the latest advances in artificial intelligence and deep learning.

To integrate safely into civilian airspace, drones must be able to detect obstacles and get out of their way as early as possible to avoid collision. This means being able to see other aircraft that may be a kilometer or more in the distance. Detect-andavoid technology powers unmanned aircraft with the ability to understand what is in their surroundings and to take the safest course of action to keep the airspace safe.

Practically, this is complex for drones because of their **size**, **weight and power**. In the past, there have been attempts to develop solutions using RADAR, but for drones that is too power-hungry and not aerodynamic. LIDAR is also unsuitable because it only provides short-range accuracy. Some have tried to use sound, but that won't detect certain aircraft, such as balloons or gliders.

Recent advances in science and technology have allowed this work to progress at a pace. Products such as **NVIDIA Jetson devices** have been a huge enabler for the robotics and autonomous vehicles industry, while **deep learning** has led to giant leaps forward in **computer vision**.

Alejandro Galindo asserts there is a way to merge more traditional techniques with the cutting edge. His team have developed a solution which uses visual spectrum imaging with cameras and fuses geometry and conventional computer vision techniques with the state of the art in deep learning. High-resolution imagery is acquired in real time, at least 10 frames per second, and processed on an embedded computer

Iris Automation 49

platform to give accurate and precise estimations of where things are in space.

"I am personally against the ideas of using deep learning to completely replace certain things that could be modelled in a general but accurate way by geometry," he tells us. "Geometry models the world so well, why completely replace it and try to learn things end to end? I understand the motivation, but from a practical point of view I think it's better not to do it."

The following video is a presentation of a live test scenario of the Iris Automation Collision Avoidance System.



A Cessna 162 aircraft makes three passes across the drone's field of view (FOW)

"Potential economic savings are huge, but they can also save lives..."

"The way deep learning operates right now is not going to solve this problem. The number of false positives would be too high!"

"You're trying to detect a moving object that is very far away and pretty small in the image," Alejandro explains. "If you give that problem to a recent graduate, they will probably throw deep learning at it and what you end up with are some good detections but a ton of false positives. The way deep learning operates right now is not going to solve this problem. The number of false positives would be too high."

But why are false positives such a bad thing? If there is an obstacle, it will be avoided, and if there isn't, won't it be a case of no harm, no foul?

"Imagine if that is happening 10 times per minute," Alejandro counters. "All those triggers will end up executing an avoidance manoeuvre. This could be an aggressive manoeuvre and if you're doing that all the time then your missions are totally disrupted. It might mean that a manual pilot needs to take over."

50 Computer Vision Application

Drones share the same airspace as crewed aircraft, so Iris Automation works with regulators around the world to make them safer and more accessible as airspace regulation continues to evolve. Detect-and-avoid technology is a vital part of this.

"We add eyes and brains to drones," he explains. "While they are flying by themselves, we will keep them safe, looking around to make sure there is nothing else in the same airspace, and preserving the required safety buffer. If there is something in the vicinity, we can classify what it is and get out of its way."

These capabilities bring a host of new possibilities for drones, such as package delivery and railway inspection. They can support surface mining, precision agriculture, and search and rescue. **Potential economic savings are huge, but they can also save lives**, both indirectly and directly. A number of people every year lose their lives during power-line inspections, for example, which is something drones can perform autonomously.

The team has extensively tested the technology on more than **12,000 real-world encounters**, but this is costly and not without risk. To complement this, they have built their own simulator which has created more than **50,000 synthetic encounter** scenarios, allowing them to test very dangerous events like head-on collisions.

"This is just vanilla not your simulator," Alejandro tells us proudly. "Our Simulation Architect worked for NASA and Industrial Light & Magic in the movie industry, amongst other places, and he developed a simulator that can model all the things that affect computer vision, such as lens distortion, noise, and illumination. The simulator can try to understand what a scenario will look like and how the system would sense what is happening. We've been developing this for four years and now we're at the point where we can generate data at a large scale."





() IRIS AUTOMATION™

Looking to the future, although they have tested on a huge variety of scenarios, Alejandro says they won't stop there, and will keep generalizing so that they can guarantee performance in many different locations and conditions. He is keen to be able to fly the drones at night by modifying the system to use **thermal camera** instead of visual spectrum ones. Also, due to the use of geometry, there are certain conditions that present challenges he would like to solve, such as **flying over water where the ground is moving**.

50 Automation Iris has around employees, with 10 people in Alejandro's perception team. Customers are already using one version of their detect-and-avoid solution, and they have other products coming out later this year, so watch this space.

Finally, we ask Alejandro what he finds most exciting about his work.

"... now we're at the point where we can generate data at a large scale."

"There are so many things that I could mention," he smiles. "One of them is that the solution is currently being used. It's so exciting that I've put something out into the world and our customers are using it in real operations. Also, this industry is less served by academia, so a lot of what we do is novel. We don't just take a paper and implement it. A lot of things don't work in the way they have been designed. There is a lot of discovering and a lot of inventing. That is extremely exciting!"



52 Open-Access Journal

Frontiers in Computer Vision

Marcello Pelillo is a Full Professor at the University of Venice leading the Computer Vision and Pattern Recognition group. He has been working in the field of computer vision for 30 years now, including serving as General Chair at ICCV2017 in Venice. Marcello shows no intention of slowing down just yet. In fact, he has just taken on a new position as chief editor of the Computer Vision section of the Frontiers in Computer Science journal. He speaks to us about his exciting new role and how open access is changing the landscape of scientific publishing.

Frontiers in Computer Science is one of over 100 journals published by leading Switzerland-based openaccess publisher and open-science platform Frontiers. Each journal has a number of sections covering a range of academic disciplines.

Led by **Kaleem Siddiqi**, Frontiers in Computer Science has eight sections of which **Computer Vision** is one.

"I've been on the editorial board of this journal for many years," Marcello tells us. "Originally, the section was focused on computer image analysis. When Kaleem and the editorial staff approached me about taking the computer vision lead, I was a little



reluctant at first because there are already excellent journals in the field, such as **PAMI** and **IJCV**, and excellent conferences, including **CVPR**, **ICCV** and **ECCV**. I wanted to understand if they shared my ambitions and vision, and in the end, we agreed our common goals and I was happy to accept. These journals are popular and very good in terms of impact factor. I know that it is going to be challenging, but I said let's give it a try. I like challenges!"

One of the first things Marcello did was

Frontiers in Computer Vision 53

set up a high-profile editorial board featuring some of the most eminent names in computer vision, including Octavia Camps and Laura Leal-Taixé, whom our readers know well. I'd also add a few more names such as Sven Dickinson, Edwin Hancock, Dimitris Metaxas and more ...

"I called on a lot of friends!" he laughs.

When people submit their papers, they are assigned to an associate editor who will find at least three reviewers to read them and provide comments.

"You will have your paper reviewed by experts in the field," he reveals. "In my humble opinion, the effort I'm making to set up **a great editorial board** is unique in this sense, and among open-access journals it's one of our biggest strengths."

Reviewers work on a voluntary basis, which is standard practice.

"Reviewers typically agree to review a paper in the first place because they want to do **a service to the community**," he explains. "By reviewing a paper in your area of expertise, you can learn new things and gain new ideas. As a reviewer myself, I have had this experience. It

"These journals are popular and very good in terms of impact factor."

can take up a lot of time though, so I do sometimes have to say no!"

The review process itself differs to what people might be used to - it is faster, and authors can interact directly with reviewers, who stav anonymous until papers are accepted, so you aren't waiting months to read and respond to a review of your work. It is more of a conversation. process This interactive means things move much faster and being an open-access journal, once work is published, anyone anywhere in the world can read it, which is great for visibility.



54 Open-Access Journal

"I don't know of any other publisher or journal who has implemented this idea," he says enthusiastically. "It's really exciting, especially for young researchers. It's a completely transparent system. When a paper gets published, you will have the name of the editor who handled the paper and the name of the reviewers. That's really important, but it's not commonplace."

Currently the Computer Vision section has a number of special issues – journal jargon for research topics – including Video Story Understanding, Perceiving Humans, and Differential Geometry; Marcello says they are always on the hunt for new ideas. Frontiers is a **gold open-access publisher**. This means authors (or rather, their institutions) help to pay for publications, as opposed to classical journals where the reader pays to subscribe.

"I know there are arguments that if authors have to pay to publish their papers it will affect the quality of the journal," Marcello remarks. "My reaction to that is to say: think about the major conferences in our field. As an author, you have to pay to register for those, but nobody says that the quality of CVPR, for example, is bad. If you have a good, high-profile editorial and reviewer board, the quality of the journal does not depend on the fact that you have to pay."



Frontiers in Computer Vision 55



For now, both models survive, but many in the community are pushing for more publishers to move to open access. Marcello thinks that eventually they will succeed, but he is up for the challenge of standing out in a relatively crowded space.

"I want this journal to become a venue where high-quality and possibly off-the-beaten-track papers get the visibility they deserve," he says passionately.

"That's my ambition and of course I ask my colleagues in the computer vision community to help contribute to this aim!"

We wish Marcello good luck as we

" ... the effort I'm making to set up a great editorial board is unique..."

say our goodbyes. Here at **Computer Vision News**, we share the same spirit. We don't believe that people should have to pay to read our content, which is free and accessible to all.

"As I told you once, I like very much what you do because you show the human aspect of our work."

Thank you, Marcello, that is high praise indeed!

56 Bay Vision Meetup

You are invited to follow an awesome Virtual Meetup with Dr. Ester Bonmati, senior research fellow at the Wellcome / EPSRC Centre for Interventional and Surgical Sciences (WEISS) at University College London:

Endoscopic ultrasound guided interventions current approaches and challenges

Join us on April 22!

Click here to register now and get your personal access It's a Bay Vision Meetup organized by RSIP Vision

Bay Vision Virtual Meetup

Endoscopic ultrasound guided interventions: Current approaches and challenges

Apr 22nd

Thursday | 10 am PT 1 pm ET





Guest Speaker: Ester Bonmati Senior Research Fellow at the Wellcome / EPSRC Centre for Interventional and Surgical Sciences (WEISS) at University College London



Hosted By: Moshe Safran CEO, RSIP Vision USA







Upcoming Events 57

	HI
T	
	IP
Ĭ	
ľ	
ł	
	L.

IEEE ISBI 2021 Full virtual conference April 13-16	Al Applications Virtual Summit Online April 14-15	Ai4 Healthcare Summit 2021 Taking place digitally May 5-6	
Digital MedTech Conference Online May 6	CIARP Virtual May 13	SSVM 2021 Virtual May 16-20	
ICCP Haifa May 23-25	Women in AI Online June 15	SUBSCRIBE! Join thousands of Al professionals who receive Computer Vision News as soon	
CVPR 2021 Virtual June 19-25 Meet us there	FREE SUBSCRIPTION (click here, its free) Did you enjoy reading Computer Vision News?	as we publish it. Yo can also visit <u>ou</u> <u>archive</u> to find new and old issues a well.	
FIMH 2021 Stanford University (virtual or hybrid) June 21-24	Would you like to receive it every month? <u>Fill the Subscription Form</u> - 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 to go virtual or to be held at another date. Please check the latest information on their website before making any plans!







A PUBLICATION BY

IMPROVE Your Vision With Computer Vision News



The magazine of the algorithm community



