Computer Vision News

The magazine of the algorithm community

July 2017

Exclusive Review by RSIP Vision:
Understanding Google's Transformer

Workshop and Challenge at MICCAI 2017:
OMIA4: Ophthalmic Medical Image Analysis

Women in Science:
Ece Özkan
and
Olga Raznitsyna

Project Management Lecture by Ron Soferman:
When deep learning is inferior to other methods

Exclusive Interviews:

Nassir Navab **Leo Joskowicz**

A publication by

RSIP VISION

Research
Google's Transformer

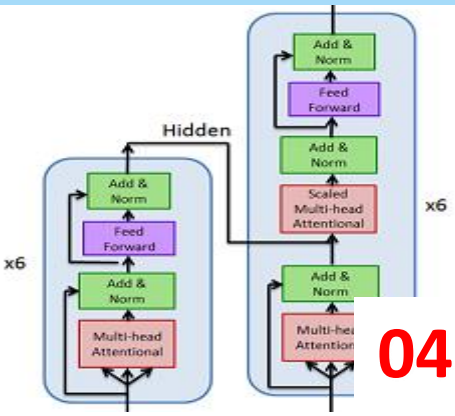

04

Spotlight News

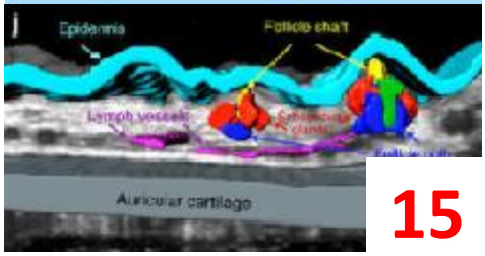

15

Project Management
When deep learning is inferior to other methods

16

Workshop and Challenge
OMIA at MICCAI 2017

18

Guest
Nassir Navab

20

Presentation
Catheter Segmentation in X-Ray Fluoroscopy using CNN

26

Women in Science
Ece Özkan

28

Women in Science
Olga Raznitsyna

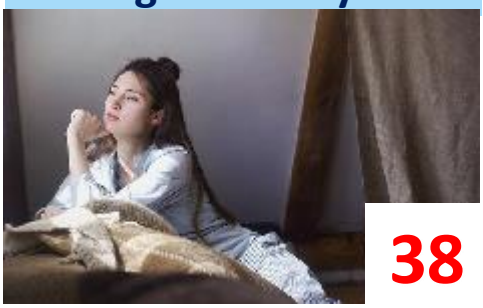

38

Guest
Leo Joskowicz

44

Dear reader,

Computer Vision News, the magazine of the algorithm community published by **RSIP Vision**, has given birth to a new baby: the first ever **CARS Daily magazines** were circulated at **CARS2017**, the Computer Assisted Radiology and Surgery conference (CARS 2017), held in Barcelona only a few days ago. We are very proud of this new partnership with **CARS**, going along similar ones with **CVPR**, **ECCV** and **MICCAI**. We wish to warmly thank **Prof. Heinz Lemke** for fully trusting us, as well as **Franziska** and **Sarah Schweikert** for their precious help and support.

This Computer Vision News issue of July offers many other great reads, in particular **exclusive interviews** with **Nassir Navab** and **Leo Joskowicz**, both collected at CARS2017.

We have mentioned **CVPR**, the venue of our first Daily magazine in Las Vegas, back in 2016. Our partnership is still very strong and they honor us by requesting another series of **CVPR Daily**, later this month in **Honolulu, Hawaii**. Will you be attending? Let us know!

We are also gearing up for **MICCAI 2017**: you will find in this July issue of Computer Vision News a preview of both **OMIA workshop** and **ReTOUCH challenge** in Ophthalmic Medical Image Analysis. Read details by organizer **Emanuele Trucco** of the University of Dundee.

I warmly recommend you give a look at the first article you will find just after turning this page: it is our exclusive review of **a new attention unit by Google called Transformer**; we based our report on three articles published by Google itself only a few days ago.

Please keep sharing this magazine and **enjoy the reading!**

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

Computer Vision News
Editor: **Ralph Anzarouth**
Publisher: **RSIP Vision**

Contact us

Give us feedback

Free subscription

Read previous magazines

Understanding Google's Transformer

Every month, Computer Vision News reviews a research from our field. This month we have chosen to review **Understanding Google's Transformer - the first sequence transduction model based entirely on attention**. In fact, a number of new articles from Google Research came out last week: "Attention is All You Need", "Separable Convolutions for Neural Machine Translation" and "One Model to Learn Them All", all based on an interesting new attention unit called Transformer.

"It took only a few days to add the parsing data-set generator to Transformer and achieve very good results after a week's training"

Transformer (T2T) streamlines the creation of open source models for a wide variety of machine learning tasks, such as natural language parsing, translation, image captioning, etc., making possible much more rapid exploration of a variety of innovative ideas.

In this issue, we shall review the tools and research developments that made the Transformer possible, so even readers with little or no background can get an overview of the field. We'll explore applications and code snippets where appropriate.

The review will be divided into 6 parts: (A) First, we'll go over standard RNNs. (B) Then, we'll turn to LSTM (Long Short-Term Memory) networks. (C) We'll explain the use of RNNs for sequence transduction, that is, encoder-decoder networks. (D) We'll discuss the addition of the attention unit to encoder-decoder networks. (E) We'll discuss attempted improvements of attention units, focusing on the Neural Turing Machine. And we'll conclude with (F) the latest innovation proposed in this area -- the Transformer.

(A) RNN:

Recurrent Neural Networks (RNNs) are a class of deep learning networks created to enable a network to more easily and naturally deal with continuous data, such as text, audio and video, that is, to use the information it has gleaned about a previous word in a sentence or frame of a video to understand the next word or frame. RNN methods have produced impressive initial results in handling a variety of tasks, achieved the ability to understand free text, and even created new original sequences from scratch. Nevertheless, there is still room for improvement.

To deal with the challenge, RNN processing (like human thinking, and unlike regular neural networks) is given persistence -- by building the networks with loops in them. In the diagram on the next page, the loop stretched out. The input x0 and outputs y0 also outputting z0 (the information deemed relevant to handling the next input) in a loop to itself,

then looks at input x1, but this time with the additional input of z0, and outputs y1 and z1, and so on.

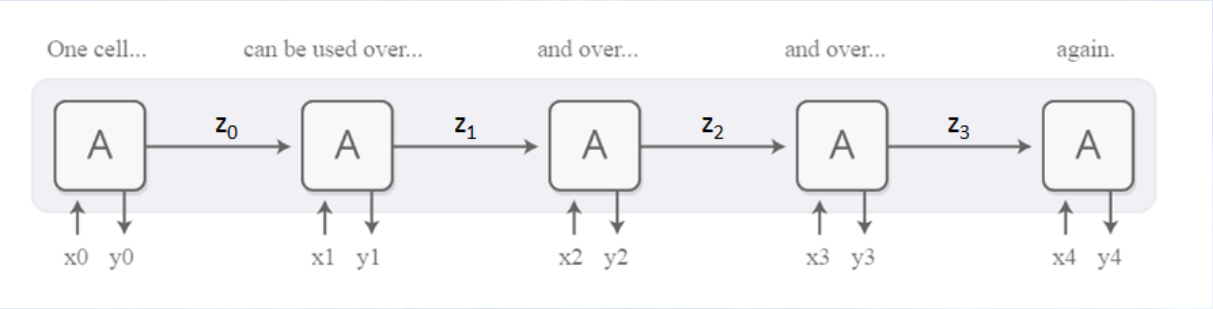

However, basic RNNs had difficulty in making use of information gained early in a long sequence, much later in that sequence, an ability often needed, for instance, in dealing with natural language with its long-distance dependencies. Most of the recent impressive successes of RNNs, in translation, voice recognition, image classification and more, have come from the special class of RNNs called Long Short-Term Memory RNNs (or LSTMs for short).

(B) LSTM:

LSTMs proposed in 1997 by **Sepp Hochreiter** and Jürgen Schmidhuber are designed to remember information for long periods by default, with little or no change, so forgetting must be explicitly handled in this model, rather than assumed. Like in standard RNNs, LSTMs have the structure of repeating identical units of neural network.

There are numerous variants of LSTM units, each unique in some particular feature: we'll describe the basic typical structure of an LSTM unit. The basic difference with respect to an RNN unit is that, rather than a single simple layer (for instance, *tanh*), there are 4 layers designed to produce a clever interaction.

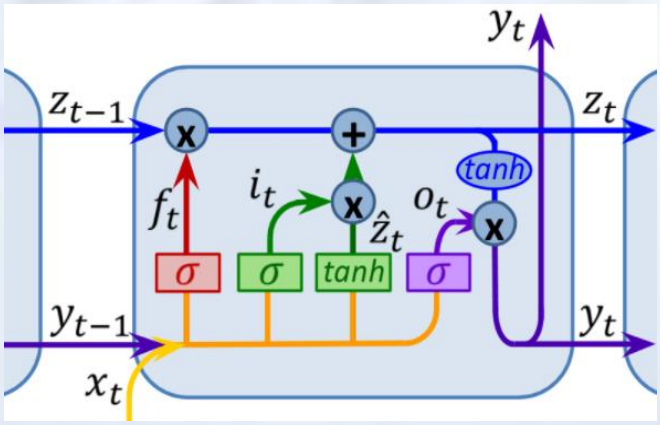

Key to the functioning of LSTMs is the hidden state. The blue line running along the top of the diagram, throughout the chain of (looped) units, only affected by controlled interactions. Hidden state z_{t-1} is received as input by the node from its predecessor. Hidden state z_t , passed on as output by the node, will be changed

from z_{t-1} it received, only to the degree the gates of the node opted to let information through to change it, making it easy for information to proceed along unchanged between iterations.

Three of these gates update the hidden state z of the LSTM at each iteration, and filter what part will constitute output y:

The *forget gate* layer (in red): the first decision is what part of the data of the hidden state z_{t-1} to discard. The forget gate sigmoid layer takes y_{t-1} and x_t and produces a number between 0 and 1 for each value of the hidden state z_{t-1}, dictating whether to completely discard, partially discard or keep that data.

The *input gate* layer (in green): the next decision is what new data should be added to the hidden state. This is comprised of two parallel elements: the input gate sigmoid layer sets which values will be updated, and in parallel the input gate *tanh* layer creates a vector of new candidate values $\hat{z}_t$ that might be added. When the vector of candidate values $\hat{z}_t$ is multiplied by the sigmoid output f_t , the product is that part of the candidate values that will be let through to be added to the hidden state.

Once both the forget gate and the input gate have updated the hidden state z_{t-1}, z_t is ready: $z_t = f_t * z_{t-1} + i_t * \hat{z}_t$

The *output gate* layer (violet): The final decision is what parts of the new hidden state z_t to produce as output. The output gate sigmoid layer acts as a filter setting what will be output and what will remain "hidden". The hidden state is run through *tanh* (to normalize values between -1 and 1) and then multiplied by the sigmoid layer output: $y_t = o_t * tanh(z_t)$

Now that we've understood the basic structure of an RNN/LSTM unit, let's look at using these units to form a network that will be able to perform sequence transduction, also known as: **sequence-to-sequence**.

(C) ENCODER-DECODER:

Neural sequence transduction models have an encoder-decoder structure. The encoder transforms an input sequence x to a hidden sequence z. Given z, the decoder can produce an output sequence y.

The encoder LSTM produces no external output, only the hidden state z is updated at each iteration. Once the encoder has completed producing z_n , the decoder takes that as its initial input z'_0.

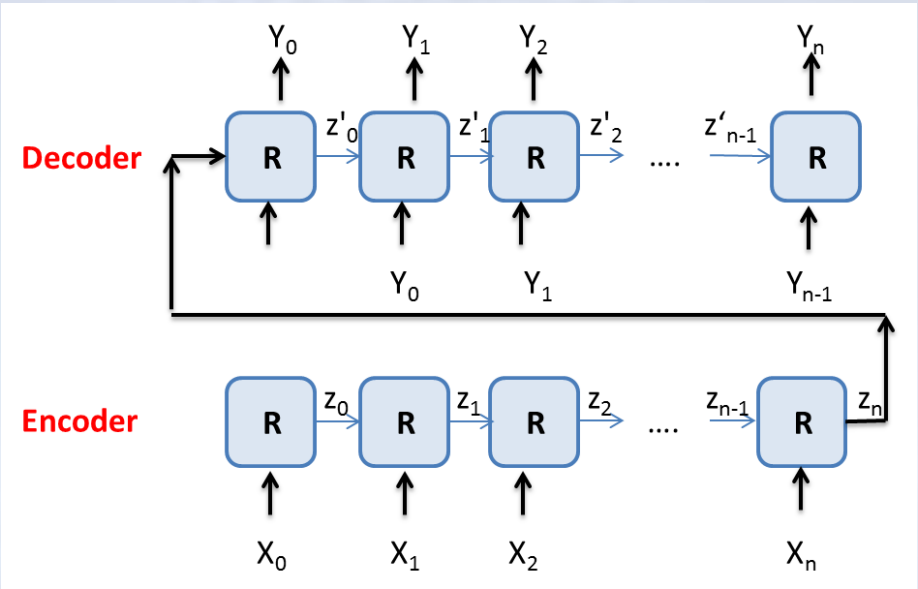

Let's look at language translation as an example (the model involves the use of two LSTMs). The first encodes the input language and outputs a fixed dimensional output z made from embeddings of the input. The final z is fed into another LSTM as its initial z' to output the predicted translated sentence. How does a decoder process inputs and generate outputs? Our decoder inputs are the target language inputs with a "start" token in front and an "end" token in back, followed by padding. The decoder unit takes as its initial input the final hidden state produced by the decoder unit.

(D) ENCODER-DECODER with attention:

Now let's take a look at encoder-decoder models with an attention unit. The difference is an additional input from the attention units received by the decoder, telling it how much attention to give to each of the encoder inputs when predicting the next output.

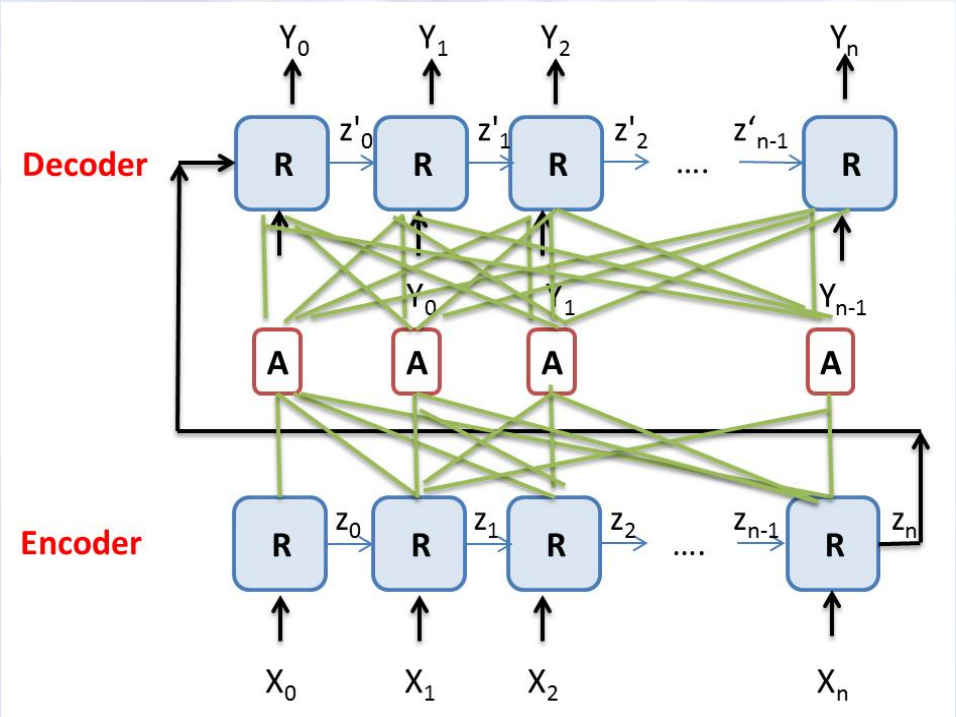

(E) attention units:

Ongoing research on LSTMs with attention units produced many different attempts at improving these units. We will focus on the Neural Turing machine, but will also review two additional types of attention unit to demonstrate their shared attributes:

1. Neural Turing machine
2. Attentional Interfaces
3. Adaptive Computation

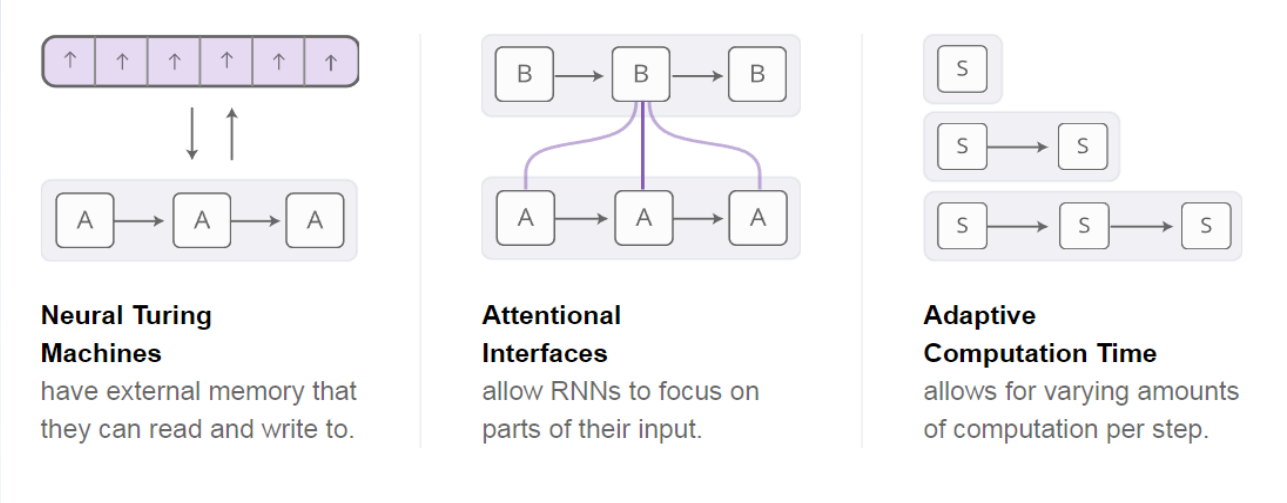

Neural Turing machine

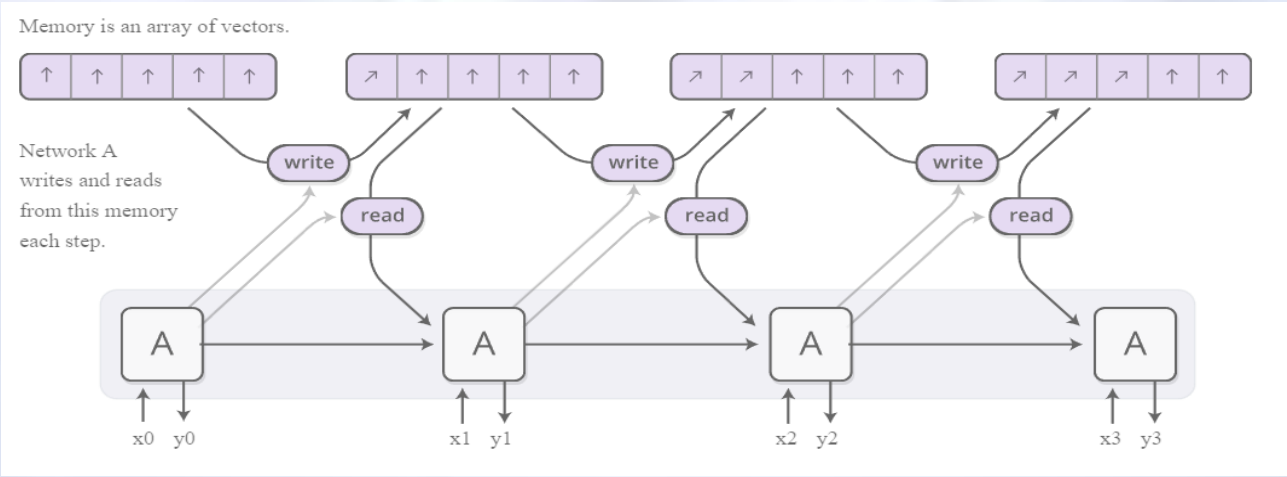

The **Neural Turing Machine (NTM)** adds an external memory to the basic LSTM network, which it can write to and read from, allowing the network to better understand what it is doing (e.g. mimic a lookup table with key-value, learn to store a long sequence in memory and so on). The external memory is a vector array, the question arises where in the vector array should the network write to and read from?

To answer this question and to allow the network to learn the optimal location for writing and reading, we need to represent the problem by a differentiable function (so that an optimizer can be trained to learn its parameters). In fact we want to make the function differentiable with respect to the reading/writing location, so that the network can learn where to read and write. NTM does this by writing and reading anywhere in the vector, but with a different weighting for different locations.

The read result is a weighted sum.

$$r \leftarrow \sum_i a_i M_i$$

That is, rather than assigning a single location the network should write to, the NTM produces an attention distribution and a set of weights designating how much to write at each location, and the read result is the weighted sum of the written values at the different locations.

Two different techniques serve for training the network to learn the weights: (1) content-based attention and (2) location-based attention. Content-based attention enables the NTM to learn to focus on the most relevant areas of the memory vector array, while location-based attention enables more flexibility and movement throughout the memory.

Attentional Interfaces

In learning problems there is an abundance of information and in many cases much of it is not directly related or relevant to the problem at hand. The purpose of the attentional interface is to learn what part of the data is relevant. In other words, its function is to help the network learn to focus on a subpart of the input data. The attentional interface implements an approach resembling that of NTM: we focus on the entire input, but with differing weights.

Network B focuses on a different part of the data than network A, at every stage.

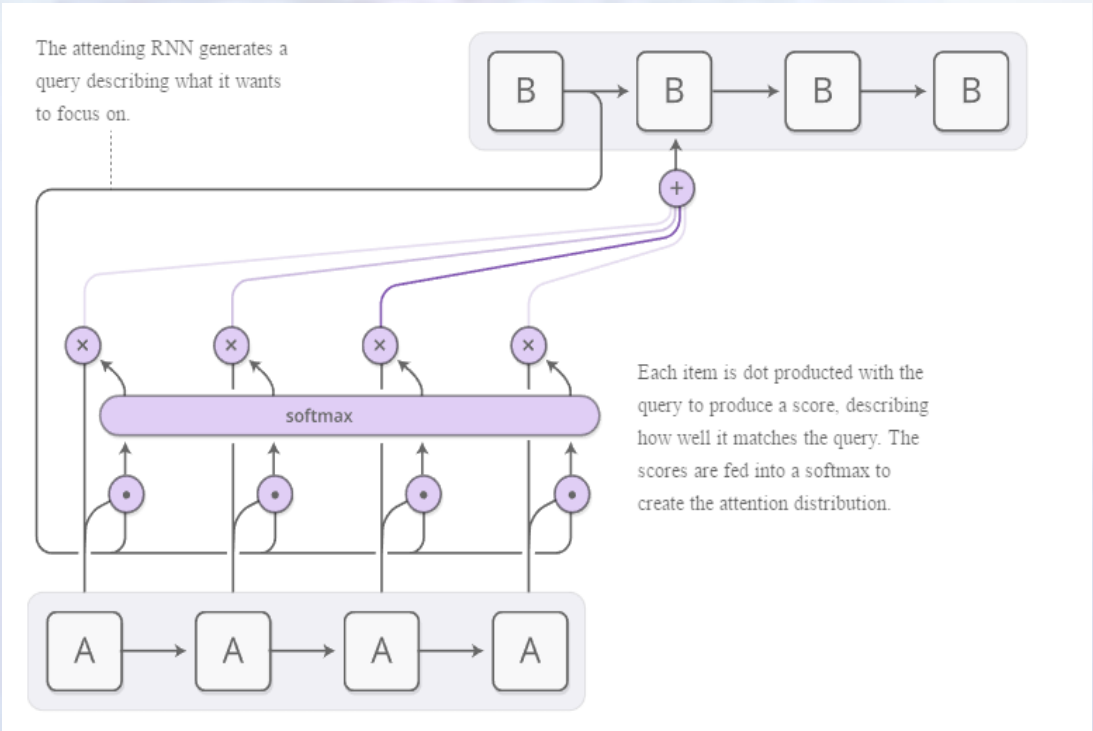

An attentional unit can also function as an interface between networks, for instance between CNN and RNN.

One of the common uses of attentional interfaces is for captioning. First, a CNN extracts relevant regions of an image and their features. Then, an RNN focuses on the regions and features extracted by the CNN to produce a description of the image.

Adaptive Computation Time (ACT)

This is an algorithm enabling RNNs to learn how many computational steps should be performed between receipt of an input and producing an output.

The motivation is the understanding that not every input or every step of processing (correlated to vector locations) justifies or requires the same computational effort. This is intuitive to human thinking: the more difficult the task, the more thought and/or effort it requires.

The idea is simple: for the network to learn how many computational steps should be undertaken, we want the number of steps to be represented by a differentiable function. We do this, yet again, using the same trick; rather than deciding what the number of steps is, we perform the computation on a variety of numbers-of-steps, and the output will be the weighted average of the values (each sub-network performed a different number of steps).

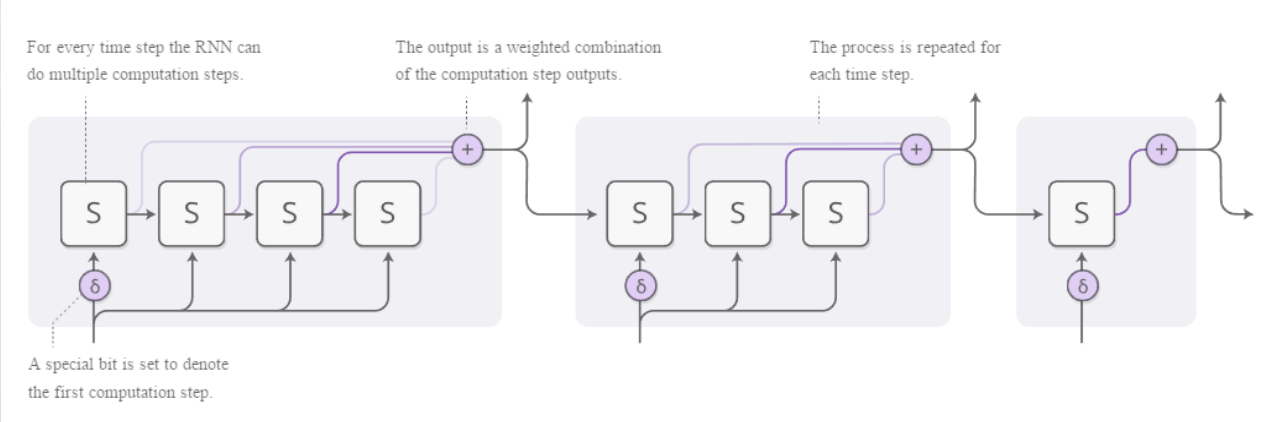

The only open source implementation of Adaptive Computation Time at the moment seems to be Mark Neumann's (TensorFlow).

The main disadvantage of all these attention models is that, because we are computing all the various different options, they require many computations for every step, causing a linear increase in computation costs and massive use of memory.

(F) Transformer - attention only:

Transformer is a new model, hot off the presses from Google. In this model the network relies exclusively on attention units to learn the relationships between input and output, with no use of RNNs whatsoever. The Transformer's architecture allows far greater parallel computation, the effectiveness of which its authors illustrate by training their network for just 12 hours on 8 P100 GPUs to handle the WMT 2014 English-to-German translation task, and achieving the best results to date.

Transformer's attention mechanism is based on key-value pair attention, which resembles the NTM attention approach in that the key is the means of retrieval of the learned values, stored in memory. In a similar manner to NTM, the output is computed as a weighted sum of the values, where the weight assigned to each value is computed by a compatibility function of the query with the corresponding key.

Transformer's attention unit is Multi-Head Attention, which is made up by repeated use of a basic element called a "Scaled Dot-Product Attention". The input consists of queries and keys of dimension d_k and values of dimension d_v .

We compute the dot products of the query with all keys, divide each by $\sqrt{d_k}$, and apply a softmax function to obtain the weights on the values.

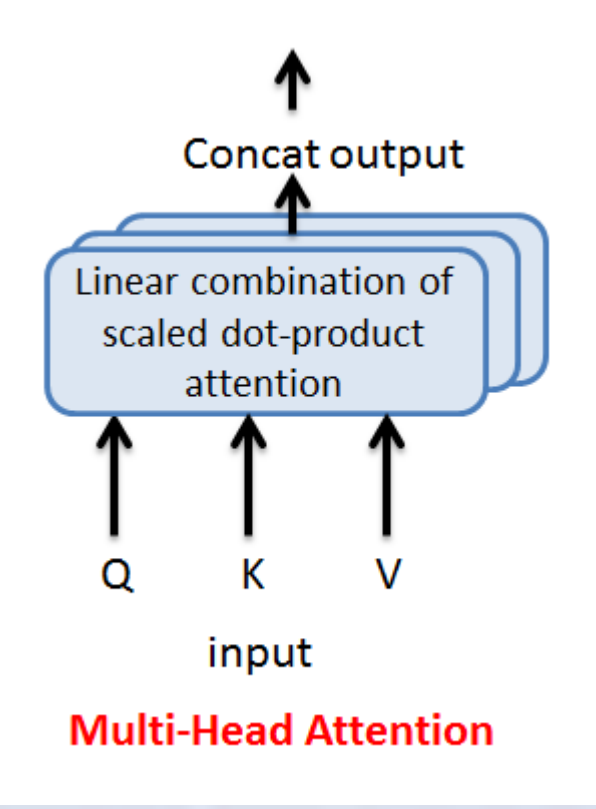

Multi-head attention linearly projects the queries, keys and values h times with different, learned linear projections. Different d_v dimensional output values are achieved by the various projections of queries, keys and values. The final output of the multi-head attention is a concatenation of all the various projections. See figure below.

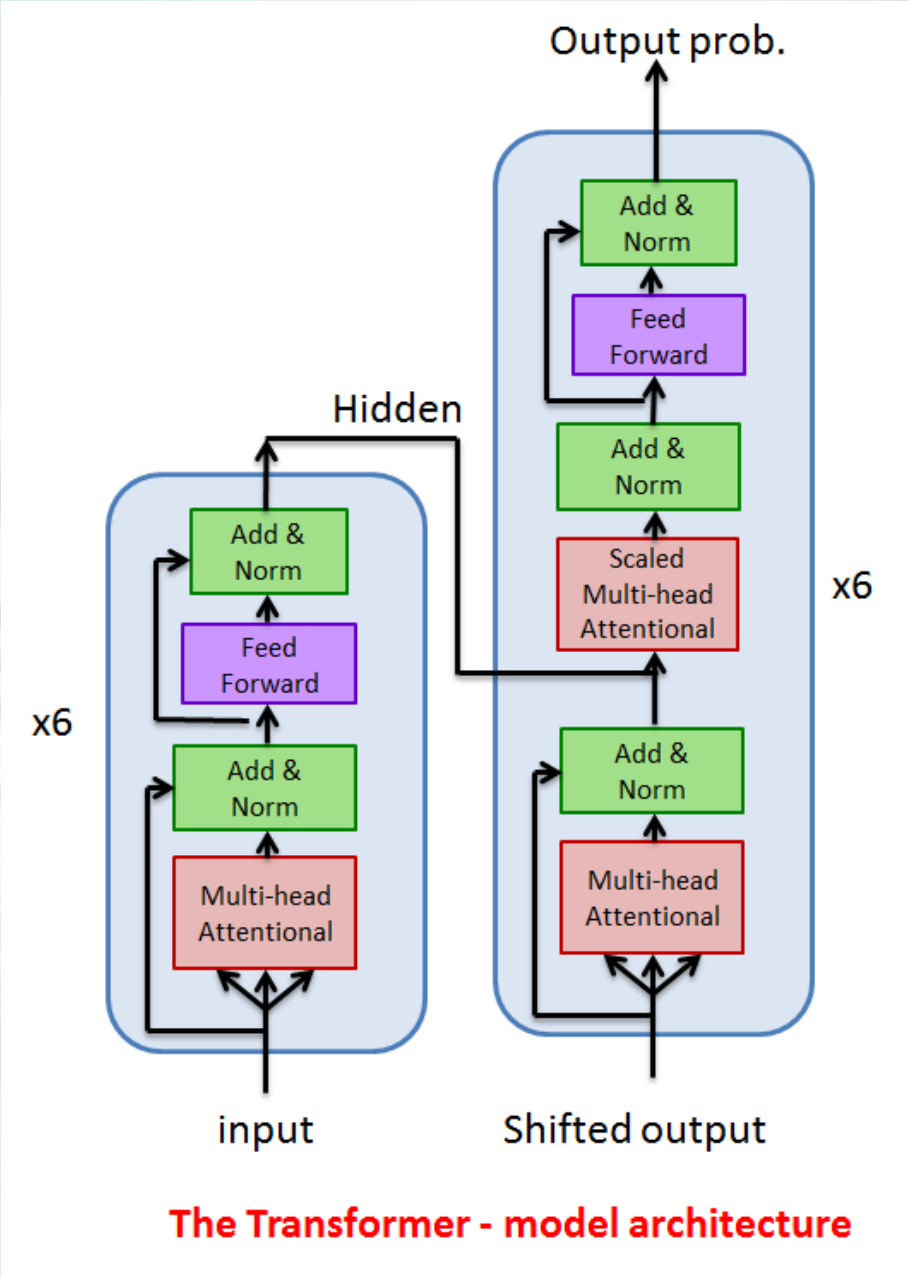

The Transformer - model architecture

Multi-head attention is used at three different locations in Transformer:

1. The first location is the "encoder-decoder attention" layers; it takes the queries from the preceding decoder layer, and the keys and values from the encoder output. This allows each decoder position to attend to every input sequence position, resembling a typical encoder-decoder attention mechanism in sequence-to-sequence models.

2. The second location is at the self-attention layers of the encoder. Self-attention refers to the fact that in these layers all of the keys, values and queries are from the same source: in this case, the output of the preceding encoder layer.

Each encoder position can attend to every position in the preceding layer of the encoder.

3. The second location is at the self-attention layers of the decoder. Each decoder position can attend to every other position in the decoder preceding it and to itself.

The encoder is comprised of 6 identical layers, each made up of two sub-layers. The first is a multi-head self-attention layer and the second is a feed-forward network, positionwise fully connected. The output of each sub-layer is LayerNorm(x + Sublayer(x)), where Sublayer(x) is the function implemented by that sub-layer. Normalization is performed on the residual connections (see the green "Add & Norm" box in the figure at page 12).

The decoder, too, is comprised of 6 identical layers, each made up of three sub-layers: the two sub-layers of the encoder, plus a sub-layer performing multi-head attention over the encoder stack output. Each sub-layer is followed by a normalization layer.

For Transformer to take into account the order of the input sequence, even without any recurrence or convolution, it must somehow be provided with the relative or absolute position of the sequence tokens. This is the function of the "positional encodings" at the bottoms of the encoder and decoder stacks.

Google team reports that they trained Transformer on a wide variety of completely unrelated datasets - and it was subsequently successfully tested on all of them.

As an example, the problem of parsing English sentences has been under research for decades. It can be solved with sequence-to-sequence neural networks, but requires a lot of tuning. The authors report it took them only a few days to add the parsing data-set generator to Transformer and achieve very good results after a week's long training:

Parsing Model	F1 score (higher is better)
Transformer (T2T)	91.3
Dyer et al.	91.7
Zhu et al.	90.4
Socher et al.	90.4
Vinyals & Kaiser et al.	88.3

Transformer is written using TensorFlow tools. Most staples of deep learning are

pre-defined: model architectures, data-sets, hyperparameters, optimizers, etc. Note that it is modular, as it implements a standard tensor-to-tensor interface.

```
t2t-trainer \
  --data_dir=$DATA_DIR \
  --problems=$PROBLEM \
  --model=$MODEL \
  --hparams_set=$HPARAMS \
  --output_dir=$TRAIN_DIR
```

In the command you can set the model, the hyperparameters, the dataset, etc. and run the training to check how it performs.

The installation is very easy, consisting of just 3 lines.

```
# Assumes tensorflow or tensorflow-gpu installed
pip install tensor2tensor

# Installs with tensorflow-gpu requirement
pip install tensor2tensor[tensorflow_gpu]

# Installs with tensorflow (cpu) requirement
pip install tensor2tensor[tensorflow]
```


References:

1. Olah & Carter, "Attention and Augmented Recurrent Neural Networks", Distill, 2016. Link 1, Link 2

2. Github

3. One Model To Learn Them All

We have reviewed the tools and research developments that made the Transformer possible, as well as applications and code snippets where appropriate: now, even readers with little or no background have read an overview of the field.

Computer Vision News lists some of the great stories that we have just found somewhere else. We share them with you, adding a short comment. Enjoy!

Giving robots a sense of touch - let them manipulate objects

GelSight, a sensor technology first unveiled by Ted Adelson's research group at MIT's CSAIL (Computer Science and Artificial Intelligence Laboratory), is now giving robots greater sensitivity and dexterity. In addition to Neural Networks, they used 400 groups of silicone objects with different levels of hardness. Read More and Watch

World-first technology gives birth to next gen digital identification

aiThenticate Computervision Labs, a joint venture of the University of Johannesburg, has announced that they have developed ground-breaking technology that simulates human cognition. The goal is to bring a solution to identity theft, which generates a global loss of about $2 trillion, doubling every year! Read More

No, really. You can see through walls using drones and Wi-Fi

It's not us saying that: theregister.co.uk reports that scholars at University of California in Santa Barbara claim that drones can perform 3D imaging of objects through walls using only Wi-Fi signals. It's not the first time Wi-Fi has been used for imaging. A group from MIT has already used Wi-Fi to capture human motion behind a wall. But this team claims to be the first to do it with drones. Read More...

The Kinetics Human Action Video Dataset, by Google, Inc.

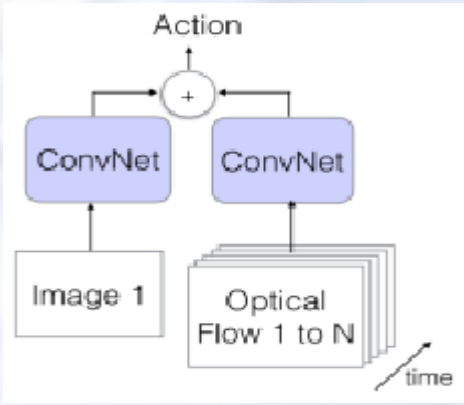

Kinetics is a large-scale, high-quality dataset of YouTube video URLs including a diverse range of human-focused actions. This Google team has released the Kinetics dataset to help the ML community advance models for video understanding. Read More and Consult the 400 classes

Breakthrough Could Make OCT Images Clearer

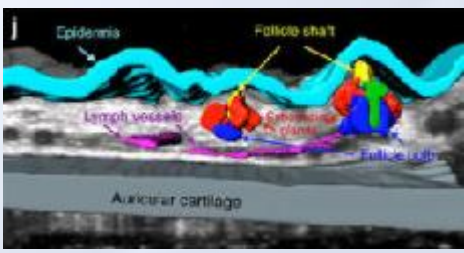

Optometry Today reports that Stanford University researchers claim they have developed a low-cost solution that could increase the resolution of Optical Coherence Tomography (OCT) by several-fold. The successful technique, speckle-modulating OCT (SM-OCT), was found to clarify and reveal previously undetectable structures, with its solution to speckle noise, an inherent noise in OCT images: Read...

Great tool: Essential cheat sheets for Machine Learning and Deep Learning engineers

Nice read by VentureBeat: 5 ways computer vision could impact how we do AI

When deep learning is inferior to other methods

RSIP Vision's CEO **Ron Soferman** has launched a series of lectures to provide a robust yet simple overview of how to ensure that computer vision projects respect goals, budget and deadlines. This month we will report the key points from Ron's lecture at the **Boston Imaging and Vision**, during an evening dedicated to **N3CV (Non-Neural-Network Computer Vision)**.

"In this case, we should not use deep learning!"

We will dedicate this talk to gain some understanding of when to use **deep learning** in a project and when it's preferable to use **non-neural network computer vision technology**.

Let's start by acknowledging the fact that deep learning is a real revolution in computer vision. Like the Gold Rush, we find that many people are going in this direction and trying to find opportunities there. Most of the academic efforts are done in this direction. Between 80% and 90% of the papers that are published in the conferences and all of the oral papers come from the deep learning area. Also, industry is reexamining the current product and their algorithms with this new approach. We can see that there is a big wave of startups that have based their technology on deep learning.

First of all, deep learning is, of course, **the best way to do classifications**. Also, in the last year we found that fully convolutional neural networks are very efficient in segmentation as well, especially when segmentation is more complicated. This means that, far from being a very simple problem, it requires context and a lot of knowledge.

We find that in the medical area, by using detection in CT images, we can find lesions and other problems. It can really give a boost to the radiologist's work. There are more and more applications every day.

We have to take into account that there is a **high computational cost**. The dataset requirements may be problematic especially when we start the development, and for any medical application, we know there is very slow FDA acceptance for this new methodology. In summary, we can say that when there is a task of detection and recognition of a complicated object, we will surely start with deep learning.

On the other hand, when we have some good mathematical grasp of the problem with regards to **physics, optics, or even physiology**, other methods will work much better. We can use it at least in some places where we can start with non-neural network computer vision technology. First of all, for any exact measurements, we will use it. Also, all of the knowledge of the 3D reconstruction including the projection, the camera model, etc. is done solely by geometric calculation and not by deep learning at all.

During registration, it is usually done by other methods. We can use some anchor points that will be based on deep learning; but the registration itself - whether it is rigid body, affine transform, or other transform - will be done with technology based on the physics of the problem. Segmentation can be done with other technology also when we have less data. We can obtain a very good segmentation without training for thousands of images.

Whenever we need to do complete authentication with 100% accurate biometrics identification, we will need to use technologies that are 100% correct and not only statistically working. In this case, we should not use deep learning. Of course, when we perform navigation during surgery, we will prefer to use solid computer vision technology.

The diagram below shows on the left side all the machine learning, classification, segmentation, and detection technology. The leading technology will use deep learning, but of course Support Vector Machine and other technology are very efficient as well.

On the right side, advanced computer vision technology will use graph cut, active contours, mean shift, Dijkstra and Bellman-Ford. All of these are based on solid graph series coming from mathematics. Of course, that also includes 3D reconstruction, optimization methods, compression and color science. Many more examples can be found in any computer vision textbook.

We want to end with a reminder that the best practice is to use both methods because with only one method, you can skip and do very partial work. We find that, especially for the pre-processing and post-processing, you need other tools. As we say, **even the biggest hammer cannot replace a screwdriver**...

RSIP VISION

Deep Learning

- Classification
- Segmentation
- Diagnostics - lesion detection
- Many more, every day
- Computational cost
- Data set requirements
- Slow FDA acceptance

- When you have to detect / recognize complicated objects

N^3 CV

- Exact Measurements
- 3D reconstruction
- Model based reconstruction
- Registration
- Segmentation with less data
- Bio metrics and authentication
- Navigation in medical surgery

- When Mathematics, Physics or Physiology are involved

OMIA - Ophthalmic Medical Image Analysis

Every month, Computer Vision News reviews a **challenge** related to our field. If you do not take part in challenges, but are interested to know the new methods proposed by the scientific community to solve them, this section is for you. This month we dedicate this space to the OMIA workshop and challenge, organized as a satellite event of MICCAI 2017 in Quebec City, Canada. More about it here.

The **OMIA workshop** is a mini-conference, one of the satellite events of **MICCAI 2017**, taking place on **September 14 in Quebec City, Canada**.

It is entirely dedicated to ophthalmic images, which is medical image analysis of images of the eye. That means the retina, cornea, conjunctiva, and all other parts.

This year we can expect a very interesting event, because for the first time in many years, there aren't only presentations, but also a **competition**: The OMIA workshop is organized together with the **OMIA challenge, called ReTOUCH - the Retinal OCT Fluid Challenge**. This happened only once before, several years ago, when it was organized by the **Iowa** group, led by **Michael Abramoff**. It is exciting that after all this time, the community is coming together again around a software challenge.

This first ReTOUCH challenge is about OCT images and, like every challenge, it is organized around submitting software and there are some specific criteria for deciding which the best performing software will be.

The first goal of OMIA is to bring the community together. So, to provide a forum for technical discussion to the community developing software for medical image analysis, specializing on the eye.

The second goal, embodied by ReTOUCH, is to move towards a better assessment and comparison of software doing the same thing in producing pieces around the world. Validation is a crucial aspect, which is not confined only to eye image analysis

but is common to all medical image analysis software development. It is particularly complicated for medical images, because very often you do not have an ultimate ground truth, as you have in metrology for instance. So, part of the goals of international challenges is also to promote the debate about better validation techniques.

We asked organizer Professor **Emanuele Trucco** to tell us his view on the future of the initiative:

Emanuele Trucco is
NRP Chair of Computational Vision
and VAMPIRE Project Director
School of Science and Engineering
University of Dundee

"As far as we know, there are no other challenges for software for retinal or eye image analysis. My personal view is that this should be treated the same way as challenges in computer vision.

Like the famous Middlebury Stereo Vision challenge, many years ago, where a web site is organized not only as a means of reporting results, but also to accept software in a certain format and evaluate results automatically. This means that people can asynchronously load up their software, run it and check immediately how it fares compared with the other ones having been submitted. So, the ReTOUCH challenge becomes both a means of advancing the field and a publicity vehicle for it."

OMIA workshop is organized by Xinjian Chen (Soochow University, China), Mona Garvin (University of Iowa, USA), Emanuele Trucco (University of Dundee, UK), Frank XU (A-STAR, Singapore).

ReTOUCH challenge is organized by Hrvoje Bogunovic, Sophie Klimscha, Sebastian M. Waldstein, Bianca S. Gerendas, Ursula Schmidt-Erfurth (Medical University of Vienna, Austria), and (Freerk Venhuizen, Clarisa Sánchez, Caroline Klaver (Radboud University Medical Center, Nijmegen, The Netherlands).

Frank Zu and Hrvoje Bogunovic sit on both organizing committees to guarantee coordination of the two events.

After the results are known and declared at the OMIA workshop at MICCAI, the plan is to publish them on a web site so that they can be communicated to the wider community.

Guest

Nassir, what in your view is so special at CARS?

In 2009 and 2010, we started to create a new event called IPCAI and at that time, we were deciding, do we do an independent event or do we join CARS? The reason we co-located with CARS was CARS was one of the conferences that had a few characteristics that I like. One was there was a mixture of academia, in terms of engineering, but also a good number of surgeons and radiologists were coming, and there was also some industrial representation. For computer-assisted interventions, which is my focus, having surgeons in the loop and also industry is absolutely necessary. For image processing, which was the focus of MICCAI, this was not absolutely required. You had the imaging data, then you didn't need to be in daily contact. I believe you did, but 90% of the committee didn't have daily contact with clinicians, they were doing their work on their stations with imaging. Then you go to the surgery room and try to do computer-assisted interventions, you can't do it. You have to have the latest industrial solutions, where we'd work with the most advanced technology, while surgeons have to be in the driving seat, they have to define all the requirements. They are the final users, so we needed to have them. We chose CARS because of this nice mixture of these three communities together.

Are surgeons conscious of this responsibility that they should take?

Very good question. It really depends, I would say the majority are not. Not that it's any fault of them, they have been usually looked at as a final customer, which receives a solution and gives his opinion on that solution. Only the big companies will use what they were calling opinion leaders and consultants, to lead them in designing solutions. I worked for Siemens for nine years before joining academia and these very sparse communications in a rare workshop of one week per year, and only working with a selected few opinion leaders who are not representative of the community, because they're only early adopters and they're reactive, it was not the right way. We have to enable the majority of surgeons to know that they have to be the person to do research with us.

Academia should have involved them more, and they were a bit neglected?

They were neglected, but I think it was also their fault, they were more comfortable to work with simulator data, because it was easier than talking to complex surgeons who ask difficult questions. It was a two-sided story. On one side it's easier for an engineer to get a well-defined problem and then go

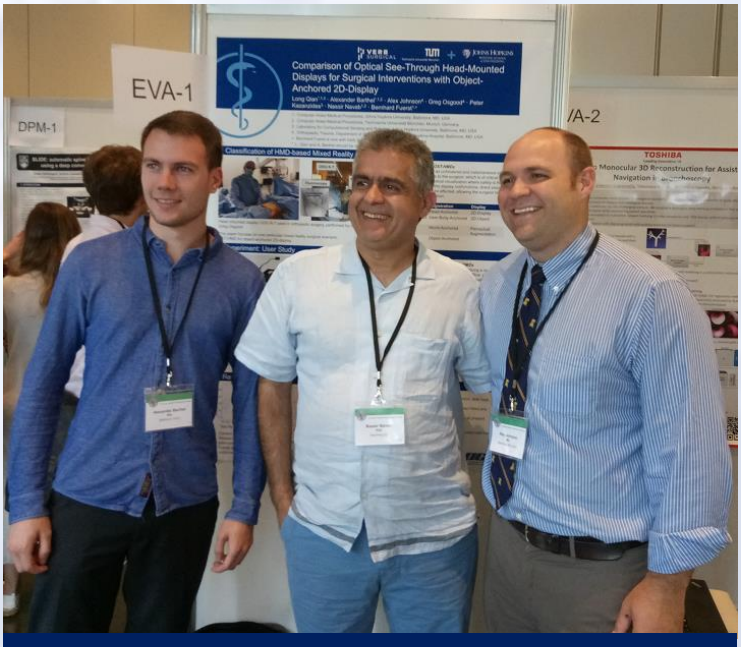

Nassir with Alexander Barthel (left) and Alex Johnson (right)

and work one year on it, than discovering every week that they had to redefine the problem. Clinics, industry and academia should work together on a daily basis.

Do you think that this bottleneck is improving?

Definitely. Two of the presentations of our team in IPCAI, both of them were done by a surgeon with seven years of experience, Dr Alex Johnson from Johns Hopkins. He came and he presented a technical talk, but of course with his flavour of surgeon and surgery, with a lot of bloody images, and with explaining his view. I think this is happening and with more and more surgeons. One of the reasons, as funny as it sounds, is thanks to companies like Apple, Google, etc: they started to provide simple hi-tech user interfaces like iPad or Google Glass. Surgeons are technology freaks, so they started to think: if I have such a nice user interface, why don't I have it during my work? There is a new generation of surgeons who are demanding better solutions.

But these are the main innovators and maybe they do not represent the big mass of surgeons which might not adopt these solutions for some time?

I don't know, 10-15 years ago, this kind of technology, these advanced user interfaces, were only acceptable for surgeons in Johns Hopkins and Harvard and Stanford. Thanks to this mass media, technology like Google Glass and iPad was very cheap and everybody could have access to it. I think these days we have many more young surgeons. We're just asking for more. It's not only key opinion leaders these days, we have many more who are interested, if we can embrace them. That's one of the reasons that, for example, in Munich we have laboratories at three hospitals. I don't expect a surgeon to travel 20 minutes to come to my lab and explain something to me, I have to go to them in their operating room and talk to them.

When did this setting start?

It's 14 years now, since 2003, but it's improving more and more.

Has it been so precious for you?

Absolutely, it has been a key to our progress, in the sense that nowadays, so many accepted us that we have two major projects there about augmented-reality technology, for their teaching. We are completely integrated into the teaching process.

And training also?

Teaching more, training too. The reason is that they teach 800 students who still don't have their medical degree. It's very hard for 800 to organise the cadaver studies, for

Guest

dissection, etc., but doing it with augmented reality and virtual dissection is much easier.

Because it's a simulation?

Because it's a simulation, but we do it on their own body. We have a system called Magic Mirror, where they can look into the system, to a wall display, and see their own body with a lot of detail of anatomy on it. Then they can learn it, they can interact with it. It reduces the need to do a lot of cadaver studies and books. It's better than books and cadavers, it's somewhere in between and they like it.

How do you convince a surgeon that it's in their own interest to cooperate?

This a very, very good question. You have to find out what is their motivation: some of them want to get habilitation or they want to progress in their things. One of them was saying from his point of view he's a handyman, because instead of repairing a car, he is repairing humans. He was saying that even though it's high-level handyman work, he would be bored if he did that his whole life, even though each patient is a puzzle, still he would like to upgrade himself. Some of them are motivated by talent: "I am very intelligent, I was the best of my class. I don't want to do the same thing in 30 years, I want to be the one who brings a new solution." Often the most motivated ones, they're hoping that not only they use our technology, but that they understand our technology, which would be also our holy grail of success. That the surgeons so much understand our technology, they redefine a new way of doing a surgery. This is very rare, the majority of these academicians are providing tools to improve surgery, but changing it completely to say, "Okay, now that I have these things, instead of entering from this part of the body, I can do another thing and then change the whole process and make it more efficient," that would be the holy grail. Some surgeons have that ambition, and those are the most exciting ones, but they are rare.

Is there competition between different labs to attract the most cooperative surgeons?

I would say no. Between companies yes, between labs, no. Usually the doctors work with the local labs. For example, in Munich, I'm only competing with the other labs in Munich, but I did go to Johns Hopkins because I was interested in the medical school most of all and I wanted to collaborate with them. You sometimes go out to do that, but I think there are so many bright doctors that it's not so rare. The fact that you, with your lab, with your students, acquire the culture that you want to interact, convince

Guest

doctors, find excited ones, this is the most challenging part. They are not rare, they have to be discovered. I think in any surgical team there is one person who has these characteristics, but very often you don't get access to them, because you need to go there. We did very funny things in order to attract surgeons to our lab in the hospital.

I would love to hear about that.

The first time we did it was in 2006, I thought we have this cool technology with a mounted display, with augmented reality, but the surgeons don't have time for it. What shall we do? Then we announced all around the hospital that this week, every morning from 6:30-8:30, we provide Weisswurst. Weisswurst is a breakfast of Bayern, which is a white sausage, it's a very fancy breakfast. We announced that during this whole week, you can just come to our lab, we'll give you Weisswurst, we'll give you coffee, we'll give you non-alcoholic Weiss beer, and then we show you demos. The doctors thought: instead of going to a coffee shop at 6:30, I can go to their lab and have a Weisswurst. Suddenly in a week we had over 90 visits. They came to have the Weisswurst and see what's going on. From them we got about 6 of them stuck with our technology. They saw it, they liked it and they came back. 90 of them didn't come back, but among the other 5-6, at least 3 of them stayed with us for about 5 years of research. Now we do it every year, we have weeks of open house, offering breakfast.

All over Germany?

In our hospitals. We also started to find out how each department works, for example, some departments have one hour of continuous education, and we try to find out when it is, and fit ourselves into that one hour. Different ways, different tricks to bring them in. The key is to show them that you have such cool technology, and then you always find 20% or 10% of them will get stuck and fall in love.

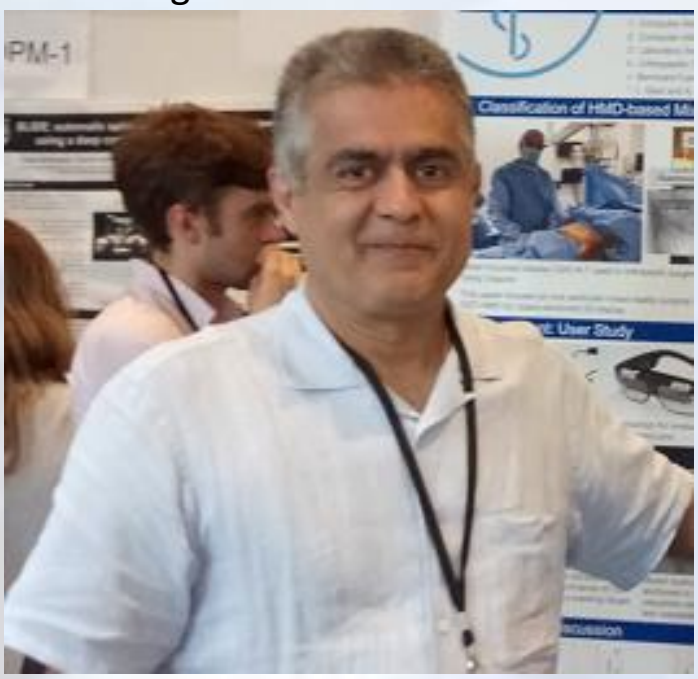

Guest

Things have changed so much for surgeons in the last few decades, but it has also changed a lot for students. Who gained the most from these big changes?

For me, the big gain was to patients. Definitely. For example, in some fields like breast cancer, in Europe, the rate of morbidity and mortality dramatically reduced. There were fields where we had no solutions and now we don't have mortality in those directions. There is a direct effect on that. Those are the big winners I would say, ourselves and our families and everybody in society. On the other hand, some of the big winners are also the countries where they had no access to healthcare and it was too expensive for them. Thanks to telecommunications, thanks to making

things available now, for example, very cheap ultrasounds, very cheap x-rays, this kind of research started to play an important role.

This big variety also gives a lot of uncertainty.

Definitely, you're completely right, there are two kinds of uncertainty. One is uncertainty whether they succeed or not and the second one is whether will the FDA approve and the like. Where should one rely on a solution? These are definitely complex issues.

Guest

"You have to give them freedom!"

You have met and taught hundreds of students, you have met hundreds of surgeons. What is the best thing that you have learned from your students and what is the best thing you have learned from the surgeons that you met?

I learn from good students much more than they learn from me. This is definitely the case, because in this new world, the students have access to information much more than in our time, so they can go much faster and learn things and teach you, and you have to run after them. They challenge us. The students find new things in the literature, I have to read it over weekend to catch up with my students, so I love it! I learn from them that you have to give them freedom. If you tell them exactly what to do, then you are blocking their creativity, from learning more. They have more information access and more information than what I have learned, so I should not limit them to what I know, I should support them in finding how to do science, rather than tell them exactly what to do in terms of science.

From surgeons, what I learned, is that the clinical, real problems are much more complex and exciting than the ones that academia separately tries to solve. The real difficult problems stems from the variability of everything. One of the things we were discussing in the course is, they do many systems to evaluate surgical skill for example: very often scientists define surgical skills in how you move the knife, but for them, it's very difficult to evaluate how the surgeon imagines the physiology of a patient which is 80 years old, 70 years old, 50 years old, they don't move the tools in the same way if the patient under them has 10 different conditions or not. When you work with surgeons, very quickly you find out why you need 30 years of education and 10 years of experience, because intuition of the surgeon is accumulation of 40 years of information.

Understanding the complexity of the biology of human beings.

Exactly. You learn from them that your research team has to look at all the problems in a very holistic way, and not try to simplify it in one equation or one robotic solution, because for them, every single patient is a new huge puzzle. Different anatomy, different genetics, different way of life, different experiences in their life, different weight, different desire, different process. For a doctor, two patients with the same disease are never a couple of each other. What I learned from surgeons is exactly that: that as a scientist I try to turn the problem into a manageable problem with 10 parameters. The real patient has millions of parameters in fact. We learn that we need to be more tolerant to

understand that the problems are more complex. At the beginning, you ask the surgeon, "What do you need?" and the surgeon says, "I need this and this and this and that," and you say, "This is easy, I can do it," but then the next time when you show it, you find out that everything that they told you has itself 20 variations. They say, "I need this," but then each of them is not just an answer, it's a family of possibilities. You learn that you have to be much more curious also, you have to search, you have to find.

And to listen better?

To listen. You have to learn to listen. From surgeons, you learn that you have to listen and I always tell my students, the first time a surgeon tells you something, write it down and don't act on it, because 90% neither you understand it well nor he told you everything: since the surgeons know that their students cannot understand that. They simply write things such that a junior understands it. I tell them write it down, talk to the same person a week later, two weeks later, at lunch, at dinner, at a conference, and then try to accumulate and discover the real problems behind it, because the problems are usually more complicated than it looks like.

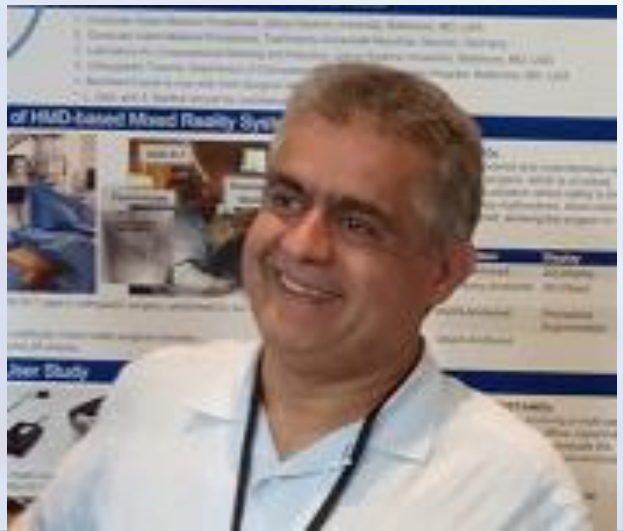

This is very precious for our readers: you are suggesting the way we should listen to surgeons.

Completely. In fact, I asked one of the most famous and most knowledgeable, expert surgeons at the Eye Institute of Johns Hopkins. He does peeling in eye surgery: which is in order of tens of microns, so the majority of the surgeons do that with a contrast agent, and he does it without. I asked him, how can you take this under the microscope, and something that you cannot even see, take it completely out? His answer to me was, an NBA basketball player can close his eyes and he can still hit a target. After 30 years of knowledge and experience, this becomes just an intuition. This is what we have to understand as people, some students when they are very young, they think that in 5 years they can build a robot that will replace a surgeon. If you can make in 5 years a robot to replace Michael Jordan in a real game with another 9 NBA players, you can make a robot that can replace surgeons. It's much more complex, because it's accumulated knowledge into an automatic and very fast decision-making and reaction, in the middle of microsurgery or very delicate surgery.

"Our life depends on it..."

This is not only going right and left and cutting and suturing. It's not an algorithm. It's accumulation of multi-dimensional decision-making, based on multi-dimensional data and a holistic, physiological view of their patient. When they do brain, heart surgery, any surgery, they are doing very complex decision-making, instantaneously. Our life depends on it!

Guest

Theo van Walsum and Pierre Ambrosini

Catheter Segmentation in X-Ray Fluoroscopy Using Convolutional Neural Networks

Pierre Ambrosini

"Previous approaches that exist are either not fully automatic, or if they are fully automatic, they are not real time"

Theo van Walsum

Theo van Walsum has presented at CARS2017 **Pierre Ambrosini**'s work on Catheter Segmentation in X-Ray Fluoroscopy Using Convolutional Neural Networks.

Theo is associate professor at **Erasmus MC in Rotterdam, Netherlands**. He told us that the work is about trying to automatically segment a thin instrument that is visible in x-ray images on video, which is useful for the surgeon during the catheter intervention surgery.

Previous approaches that exist are either not fully automatic, or if they are fully automatic, they are not real time. This new approach would apply this on live imaging, so it must be **fully automatic** and it must be **real time**.

The lack of good algorithms and computational power has hampered this process

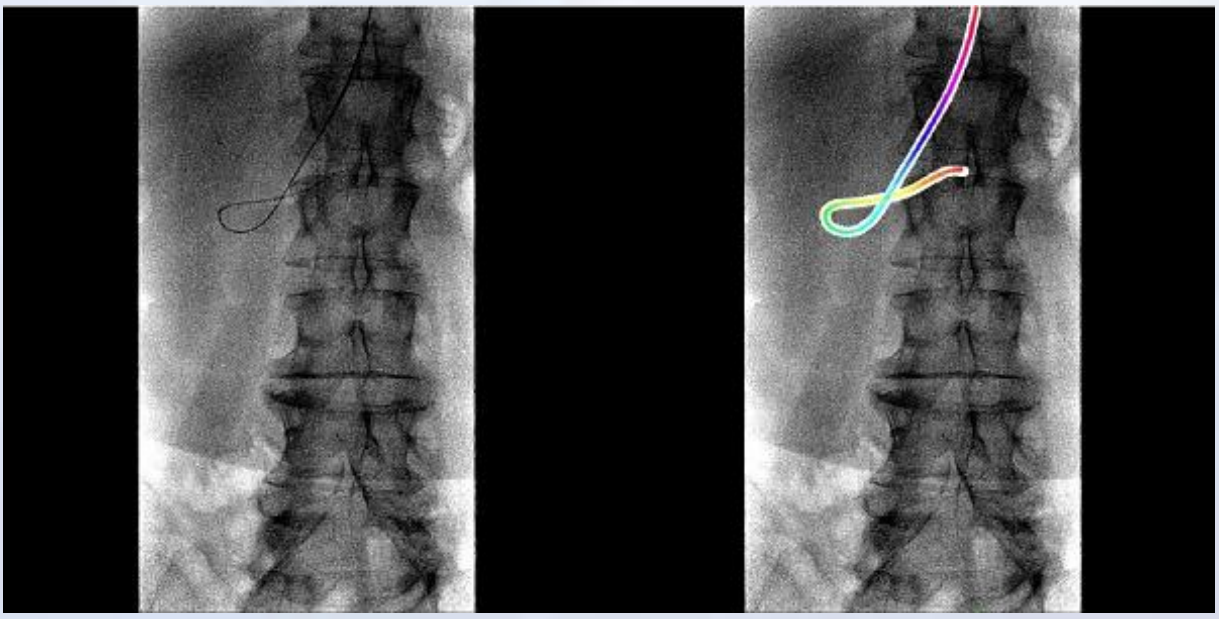

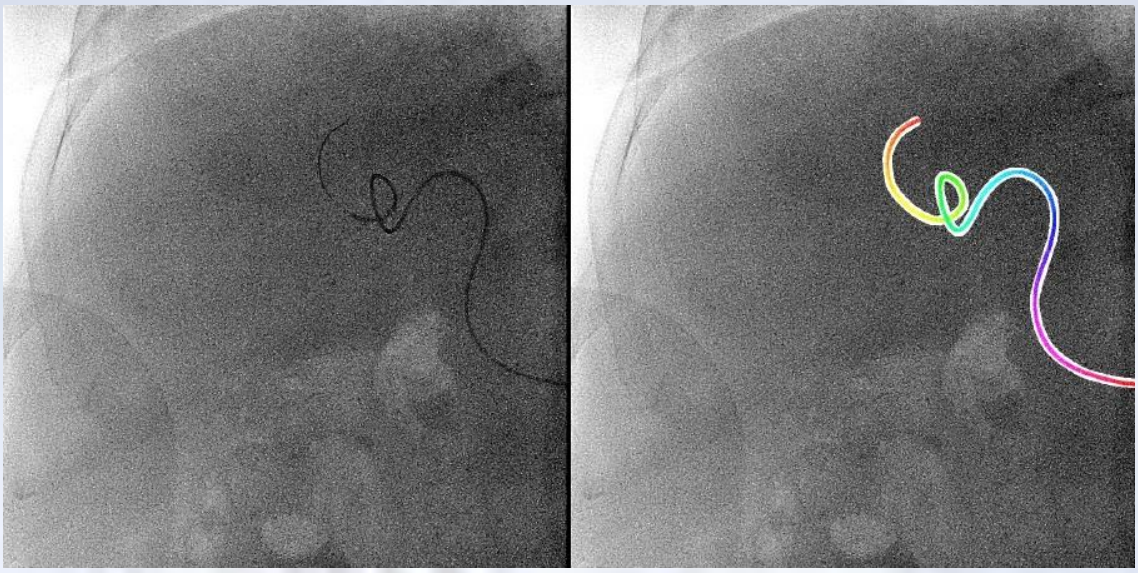

in the past. The new approach uses convolutional neural networks, which is a relatively novel technology, and requires a GPU for the computer issue.

A challenge is that **fluoroscopic images are very noisy** and the **instruments being looked for are very thin**. Using the deep learning approach allows them to segment the instruments with high accuracy. The strength of these deep learning approaches is that they themselves learn the features from example images.

When using deep learning, there are many parameters that can be set in training and it can be a challenge to find the right settings that give an acceptable working result. For the training data, they use a lot of data augmentation. They have 100 sequences more or less with x-ray images, with annotations in a few frames, and then they transform these images to generate extra training data.

Although they have all the components they need, Theo thinks **a much larger training dataset would be of benefit**, with a wider range, different systems and different applications. He gives the example that even with lots of training data, if you have non-fluoroscopic images with more contrast, the detection fails.

Theo tells us about the next steps for this work:

"*This is the last step of a complete approach for guidance in TACE. TACE is a liver procedure that is endovascular. What you previously did was build a 3D model from the vessels, from a 3D image, and align that with the fluoroscopy, providing a roadmap which was live updated based on the motion of the instrument. Up to now, we used minimally segmented guide wires, and this is the final step, so we are currently integrating all steps. That should give us a system that gives a live roadmap during TACE procedures and also allows the physician to see in 3D where his instrument is.*"

Ece Özkan

Ece Özkan is a PhD student at ETH Zürich in the Computer-assisted Applications in Medicine Group. Ece, what did you present yesterday?

I presented and I also had a poster session: my work is about tracking liver sequences recorded with ultrasound. We track the breathing motion using the surrounding information.

Why did you choose that work?

Breathing is important especially for radiotherapy applications. To deliver an accurate dose of the radiation it is important to track the breathing motion.

"I think the competition is with myself"

That is why we do this.

Where are you from originally?

I am from Turkey and I grew up in Istanbul. Until high school I was there. Then I moved to Switzerland to do my Bachelor's and Master's. .

Why did you choose Switzerland?

It is especially because, first of all, I studied in a German high school so I speak German. To go to a university, I did a German high school degree. I wanted to go to a place where they studied in German, either Germany, Switzerland, or Austria.

Are you fully bilingual?

Yes.

Why did you choose a German school when you were in Turkey?

The German school is one of the good ones in Turkey.

Was it your choice or your parent's?

It was both. There were also some American, Italian, and French schools, but the American and the German ones were one of the best in Istanbul. You can learn English anyhow.

Did that brought you to Zürich?

Yes, and it's also because the university is good.

How do you find life in Zürich?

I think it is very nice in Switzerland. Zürich especially is such a nice city.

It is very different from Istanbul.

Yes, it is. People are also very different, but I like it.

Do you miss home?

I do, of course. My parents are there. I miss my parents and my friends.

Is Zürich your new home?

Yes, I would say. I'm building now a new home for myself.

When did you first develop an interest in science? Was it when you were a little girl or did it come over time?

In school, I was always good at math. I was enjoying math and also physics. In natural sciences, I was always good. That's why when I was choosing something to study, I thought that maybe I should study something where I could use math. That's why I chose engineering. Then later I studied electrical engineering.

Did you ever receive any remarks about being a girl that studied math?

Not about math, but about why I studied engineering.

Was that later in your studies?

Exactly - When I started studying, people would say it's weird that I study

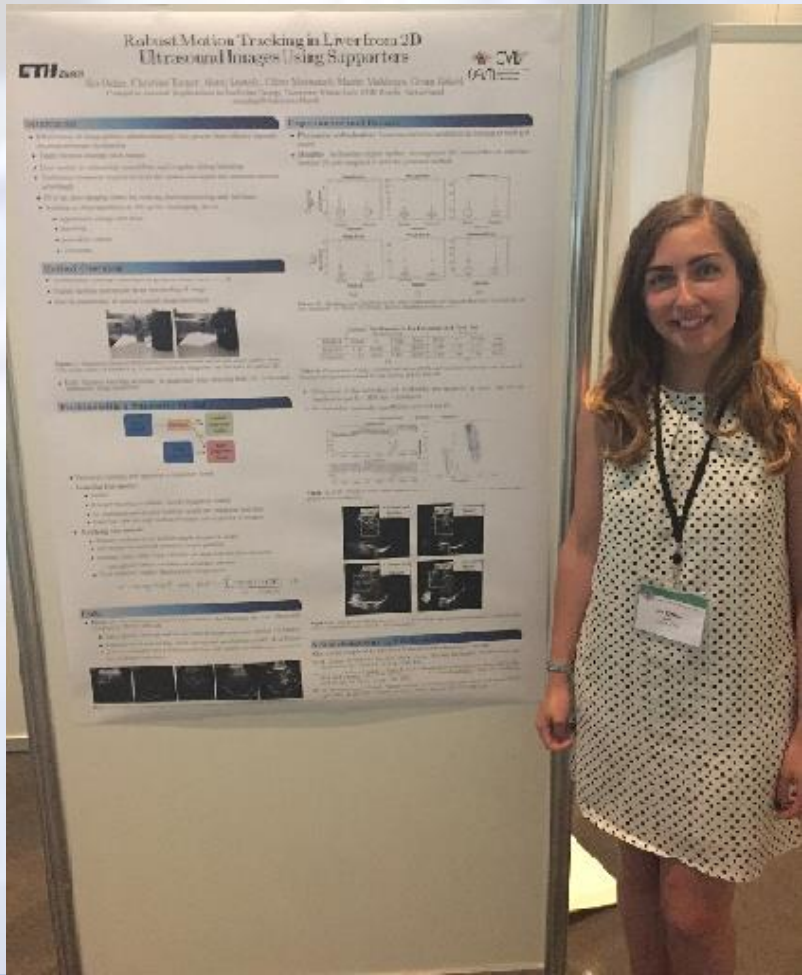

electrical engineering. Also we started as 220 people at ETH, and we were about 20 girls. This was the record.

They never had 20 girls at the same time?

...or never this percentage.

Did you have any support in this choice?

Actually, it was basically my choice, but my parents were always supportive.

Teachers too?

Yes, teachers as well in the school. Since I was always good in the lectures, my teachers were like, "Yeah, sure, why not study engineering?"

Did you have to face any sarcasm about this choice at the time?

From the outside, when I said I was doing my PhD in electrical engineering, people would say: "As a woman, why do you do it?"

What would they want you to do instead?

I don't know... teach? I don't want to say that a teacher is a women's job, but in their minds it might be.

How do you react to that?

I find it actually weird because I think there are no gender jobs. All jobs should be for each gender.

No gender jobs at all?

No, I mean why should it be? Any woman or man has different aspects that they can bring to the job so why not use it?

Is it more girls or boys who told you that?

I think it is more usually girls because I think that boys maybe have something else in their minds. By the way, in Switzerland, there is also a foundation of the city that there are no men or women's jobs. There is a webpage where you take a photo with normal clothes. Maybe there is a male nurse. Then you try to estimate who has that job. There are like 6 photos and 6 jobs, and you need to say who is doing what. This is actually for kids so there is no men's or women's jobs. Like a dancer... why should it be a woman's job?

[*laughs*] I am a better footballer than dancer, I think.

...or fireman or firewoman and so on...

Does it anger you or do you just shake it off?

It angers me of course, because I still wonder why.

That means it touches a nerve.

I think it is because you hear it all your life that girls should behave like this.

Why does it bother you? You could say who cares…

I think it is mostly because women are not seen as capable of many things. Then you try to show it or prove it: "I can also do it…"

If you could speak to all those with that kind of mind frame, what could you tell them to convince them they are wrong?

When you look at the statistics, all girls are as good, maybe even better in school than the boys in mathematical studies. I think that at the end, they are somehow not supported to choose engineering or science. They give up. I think this should not happen. They should not think, *"Oh no, I cannot choose it because this is a man's job."*

They give up, not because it is difficult for them, but because the people do not expect them to be there.

Exactly. I think they also get influenced from that. Of course, it is hard to study, but I think also women give up more easily than men.

Why didn't you give up?

I think I had support from, for example, my parents. If I thought I couldn't do it, they would say, "No you can do it!"

I have another guess: you are very competitive, and you want to prove that you can do it.

Yes, I'm also a competitive person, but I think I'm more competitive with myself. I want to prove to myself that I can do it. Of course, you care what people say, but mostly I think the competition is with myself.

Is it to prove that you are good or is it to prove that, even against other people's expectations, you are able to do great things in this profession?

I would say both. I do it for myself, but I also like to do it.

"Will, passion and talent"

What is the drive for you to take on such a heavy load?

I also like what I am doing. Otherwise, I wouldn't do it. I would say the two important aspects are, I like it, but I am also good at it, and I want to do it.

It's a mixture of three things then?

The will, passion, and talent.

Which one is the strongest?

I think the will. I also think that if you like what you do, chances are better.

Where did you find the strong will to do this?

I think it's just something you have.

You are so lucky! You have the talent, will, and passion. Were you born with these three things that are important to succeed?

I think everybody has that probably, but you can lose it along the way. I don't think that you get it from somewhere. I think you have it, but you

can lose it. I think most people, especially women, think that at the end they are not capable of it. It's not because they are not, but they think they are. They put this idea in their minds. If you want it, you can compensate for the talent.

When you look around, which one of these three components was missing in those who gave up?

I think the will. If somebody, especially a woman, doesn't believe in themselves then they just lose their motivation. If the motivation is gone then they don't see the point in that. I cannot judge in my point of view, but I think that for a woman, it is important to keep the will.

If you talk to a high school girl who is good at science, but she receives conflicting messages from the environment. What can you tell her?

I think she should look back to herself and ask herself if she really wants it and if she thinks objectively that she is good at it. If yes, she shouldn't care what the others say, because at the end it's her life.

How important is it for you to become a good role model?

I think it is very important. That's why I like this interview especially. I think it's important for girls in high school or younger that they should believe in themselves. I think it's good to see a role model that are not the genius ones, but they still continue to do their work.

What are you missing now in order to be the real role model that you would like to be?

I don't influence enough women yet. That's why when I see any opportunity, I want to use it so it can spread.

Do you ever falter?

Of course, you should see me just before a deadline! [*we both laugh*]

I actually have it quite much I would say. Then I see people supporting me. My parents and my boyfriend always support me: "Come on, you can do it. You know that you can do it." They are always supportive. I think without them, I would not be able to do this.

Are they proud of you?

I think so, yes. [*laughs*] My parents told me yesterday that they are proud of me.

They are very much entitled to be proud.

Of course, everyone has their down moments, but I think at the end it is important to get it together again and start again. When you have a bad day, then you just continue. You know you can continue.

What is the most difficult part of the pressure that you put on yourself?

I think the comparison with men, I would say. I always think that I am not good enough and compare myself to others. I would say this is the hardest part.

Do you feel any kind of satisfaction when you are better than a man? If you get a better mark or you win an award?

I have a tendency to make it less of a big deal if I get something. If I get it then I say anyhow, this was easy. What I get is mostly for me. It is OK, but failures influence me much more. That's why I think it gives me more will when I fail and continue. I would say that, especially for the girls, don't give up. You know you can do it. You actually know it, just continue doing it. If you don't, the most important thing is to ask yourself: "Do you want to do it?" Then do it. If you don't want to do it then don't, because you know yourself better. Nobody can tell you what to do.

What will you say to the boys?

I think it is the same for them. It's something they should also ask themselves.

They seem more confident naturally.

Yes, this is what I meant before that I compare myself with boys because they advertise more what they do. I don't think this is something that men should change, but maybe women should learn.

"Don't have prejudice and be nice!"

Would you tell them to behave differently or that they should make life easier for women?

I think for men it is important not to have prejudice. It is not only a man's job.

That's the only thing that we do wrong?

[*laughs*]... and be nice!

In general, we are nice…

That's true, yes!

A few hours after our interview, Ece Özkan won the Best Paper runner-up Award at IPCAI 2017. Congrats for a well-deserved honor!

Daniel Naftalovich

Daniel Naftalovich is a PhD student at the **California Institute of Technology, Pasadena** and also a medical student at the **Keck School of Medicine of the University of Southern California (USC)**. He spoke to us ahead of his poster session at **CARS 2017**.

"I just submitted to the CRAS meeting that will be in September"

Daniel's work is on the use of algorithms for skill assessment, particularly in robotic supply to microsurgery. The work he is presenting today is about a dataset they collected regarding this kind of work. Microsurgery is used as an example dexterous task in order to study skill, skill acquisition and learning curve in this setting, with kinematic and motion data that they recorded and video. They used the da Vinci Research Kit. Unlike previous datasets in this, they did a single user, but many trials in one user and quite substantial trials in each experiment. He made a very long process of training himself in doing a microvascular anastomosis procedure with the robotic assistance in the da Vinci Research Kit, tracking it the whole time both in the video and the kinematic data. The idea is that this can be a useful set of data to experiment with algorithms for skill assessment, and can also become about skill tracking, to track progression along a learning curve.

Daniel in da Vinci console during robotic-assisted microsurgery training

"I was using a non-animal phantom model in the laboratory"

Daniel told us how he thought he fared as a surgeon:

"How did I do? I think I did okay. I definitely started out not okay, and towards the end it got more okay. The question can be answered in many ways. One important thing to mention is that I was using a non-animal phantom model in the laboratory, so that of course, while it's very useful, also has its limitations. To really say, we'll see one day if I continue on to a surgical residency - we'll see for real. But actually, also recently, I did a training course in microsurgery, that I performed

April 2017 - Daniel at the Microsurgery Research & Training Laboratory course at Columbia University's Orthopaedic's department

on rats. Over there we had a more realistic measurement of how I did and at the end of the course I was able to do well, which is great. That fits into the research as well because since we wanted to see a learning curve, we wanted to make sure it's a complete one, so now I'll return to our partner laboratory at the **University of Verona** with **Professor Paolo Fiorini**. I use their equipment and do this research in collaboration with them. I'll return there now after the course and complete the study that I'm discussing at the meeting this week. I'll complete it in the sense that now that we've checked that I do well as a surgeon here, from the course that I did, this should complete the training curve that we're going to be recording."

He told us that one novelty of the work is the performance of microscopic surgery in the **da Vinci Research Kit**. That was exciting on the experimental side, because they had to do a few modifications and to observe some things that were new to them, for example, retuning some of the controller parameters in

"Very quickly, we invented a solution for this"

the robot to reduce the vibrations that they were getting in the research kit. In the clinical mode, the use of microsurgery was already demonstrated more substantially.

He also mentioned the intention of capturing a learning curve process, which is distinct from some of the previous work that used a static snapshot of skills, and the way they distinguish skills by using different people who at a given instant of time have different skills. Here, the approach was very different, it's taking a single person with a dynamic skill, that's changing and increasing through training, and capturing that whole learning curve. That's one additional way it's new.

Another novelty is in the type of data analysis that this dataset supports, he thinks this can move them closer to a regression setting. A lot of the previous work has been doing a skill classification, in two or three classes - so either expert, or maybe beginner, intermediate and expert - and that has some utility, but a nice continuation is to get more of a quantitative type of metric, that can be used for evaluating and tracking skill. This also goes in synergy that when you

External view of the surgical field during Daniel's experiment of performing micro-vascular end-to-end anastomosis using the Konnyaku/Shirataki phantom model at the ALTAIR laboratory in the University of Verona, using the da Vinci Research Kit (DVRK).

are doing a learning curve monitoring, and if you're doing it one day in a real, clinical training setting, like a residency programme that's several years long, it can be nice to track the metric and see how it progresses along the learning curve. For that, a good granularity at the resolution of the metric is better. If it's a number and it's more of a regression, you can see how the number improves and it's not just a step up from today you're beginner, tomorrow you're intermediate, and one day you're advanced.

Daniel told us an interesting story about an experience he'd had doing the experiment:

"When I got there ready to do the experiment, I had a plan of how the experiment should work and we sit down to try to do a microsurgery in the Research Kit, for our first time, suddenly we realised the problem and it wasn't going to work. That problem was on the topic of changing the magnification and changing the motion of the camera and the position in the visual field. When we tried to do it the first time, I realised that I'm going to need some control over the camera motion and the camera field. Very quickly, we invented a solution for this. We had a start-up next door that had a foot pedal plate that mimics the da Vinci's foot pedal plate. We were able to borrow that from them, connect it with USB, add a few lines of code to integrate it into the ROS package, and then we had a new user interface foot pedal that allowed me to control the camera for specifically this microsurgery experiment. I just submitted to the CRAS meeting that will be in September."

Olga Raznitsyna Ольга Разницына

Volha (Olga) Raznitsyna is a first year surgeon from Belarus. In fact, in Belarus surgeons work as an intern for one year after graduating from university. Then they start to work. This is her first year working as a surgeon, in Mogilev City Hospital 1.

Olga, when did you decide to become a surgeon? Did you know you wanted to do this when you were a little girl or is it something you decided later?

I decided to become a doctor at the end of school, and a surgeon during my last years at university. Almost everyone from my family are doctors. I decided that it is the profession that fits me the best because I adore working with people. I like to help them, and I like that this profession is usually outside of anything political. All of us can become ill. When I finished my university, I decided that I'd like to help people, not only with medicine, but also with my hands. I decided to work in the surgical sphere. Of course, it is a very wide sphere. You can be an ophthalmologist. You can be a neurosurgeon. You can be an abdominal surgeon. There are a lot of professions that you can choose. Now I work as a general surgeon. We perform operations on the inner organs of the abdomen. This is what I'm doing now.

There is a stereotype that little girls are afraid of blood. Did this happen to you?

[*laughs*] I am afraid of blood because I understand how important it is. I can't watch movies with a lot of blood because

I think: *"Oh my God, it should be inside the body, not outside!"*

It depends. When I am working, I am not afraid because it's my work. There are a lot of ugly things, but it is my work so I am OK with this… almost. If I see someone bleeding on the street, I might feel afraid. It's a normal reaction. Then I understand that a person needs to help. You concentrate on your responsibility and do what you are supposed to do.

Is it different if the blood is your own or if it's somebody else's?

If I hurt myself or cut myself by chance, I'm not afraid. I think that every surgeon is afraid of blood because in case of profuse bleeding, we have seconds or minutes to save a life. We are also afraid of blood, but in other ways, because we understand how it can end.

Did anyone ever make remarks about you being a girl and wanting to be a surgeon?

My parents seriously asked me if I was sure that I wanted to become a doctor because it's very serious, and I need to dedicate all of my life to it. I need to study constantly without breaks. When they saw me, I was sure about it so they said: *"You are welcome to do it, and we hope that everything works out."* Also, medicine is a very wide sphere. So is surgery. You can choose any speciality you want. Maybe some of them are considered *"more girly"* than others. My parents asked me again if I was sure about surgery. I told them that I wanted to try. It was very interesting for me because you apply all of you knowledge, but also your hands. You can see the immediate result of your work so it's amazing. When it is well done, it brings you so much happiness. Of course, life is long. Maybe after years or sooner, I will think about doing a more specific field.

Are you your parent's only daughter?

There are two girls, but the other one is very young. She is only 6 years old.

You said that the drive to help people is very important. Would I be accurate in saying that you inherited this drive from your parents?

Maybe because all of my life I see good examples of doctors. My relatives really like this speciality. Also, it's not only about helping people. Of course, you help. Everyone helps if they work hard. You can be a scientist, and you can also help people. Moreover, you can

also discover something. It's a perfect speciality to do both things.

We spoke earlier about female stereotypes in science. Of all of these stereotypes, which is the one that angers you the most?

In my field, stereotypes really exist. It's 100% true that there are a lot of stereotypes. There are specialities that are thought to be male fields. Abdominal surgery is one of them. When people say it's not for a woman, maybe they are right just because you need to dedicate all of your time to be successful. It requires a lot of physical effort to stand for many hours in the operation room. That's why my colleagues remind me: "*You are not a woman here. You are a surgeon. Don't complain, just work well!*"

Work like a man… Does this anger you or is it justified?

Really, it doesn't anger me. If you are discriminated against as woman and told to do something else, it's one thing. The other thing is that if you see the operation, study, work hard, and dedicate all of your time, eventually you'll be able to do it.

Do you feel like there is a glass ceiling above you that prevents you from succeeding because you are a woman?

Fortunately, not yet because I just started my work. Everything is possible for me. Maybe in several years, I will change my mind.

Do you think you have to do more work because you are a woman?

Yes - It's because of the stereotypes.

Is it because you feel the stereotypes or because other people make you feel the stereotypes?

Maybe it's because of the other people. Usually, there are few women working as abdominal surgeons. Even if we look at young surgeons, there are a lot of women among them. With experienced surgeons, we can see that the majority of women prefer to choose some specialization rather than stay and perform these long hour operations.

I think that stereotypes usually exist in male fields because it requires certain physical strengths. For example, in my work. Of course, we are equal in terms of the law, but in nature we are not identical. If a woman has children, she has to stop working for a while. Men have that extra time to progress. It's

normal for women to take a break to dedicate time to family. It's more complicated for women. That's why the stereotypes exist.

Also, in male spheres, if a woman makes a small mistake, she will be judged more by colleagues than if it were a man.

Would they be judged more by a male or female boss?

That's a good question. I didn't think about it, but maybe by women also. It depends on the woman. A woman that achieved a lot of success in a male field understands how difficult it was for her.

"My colleagues remind me: *"You are not a woman here. You are a surgeon. Don't complain, just work well!"*

Maybe she would want to motivate other women. Maybe she would judge more harshly to motivate other women.

Do you have children?

No.

Do you think your career will influence whether or not you have children?

Right now, it's not the main goal for me. I would like to have a family and children. I know that I will dedicate my time to my family and children. It doesn't mean that I will give up my speciality. Maybe I will take a break. We have good examples of women and men in this field with families and children. I hope that I will be one of them.

If you had to give advice to your 6 year old sister, would you recommend medicine or would you tell her to do something easier?

I think I would tell her that medicine is an amazing profession. I never regretted my choice. Of course, it requires a lot of work, but it's a very interesting specialization. You will never stop discovering new things.

There are many great fields, and you can find your own way. In general, I would say make sure to choose something you do well. Anything is possible, and never stop trying to prove yourself. It's important for any field.

What if she tries and finds it difficult?

If you are not happy with your choice, it's better to try something else rather than making all of that effort for nothing. A lot of people realize that this field is not for them. It's better to give up. Don't ruin your life. Try something else.

The most important thing is your happiness, and the feeling that you are where you want to be.

Are you happy everyday?

There are moments when I feel upset or sad. There are a lot of difficult moments. There are difficult patients and situations, but in general, I am happy.

What worries you about the future?

I would like to succeed in my profession so I worry about if I will do well enough.

Are you afraid of being inadequate?

Of course!

Is this justified or is it just in your head?

Sometimes I realize that I need to read more about some medical problems and refresh my memory. Other times it's just a fear that I'm not good enough. I think it's normal. It's better to hesitate a little bit because it motivates you to study more.

Are you ever jealous of the self-confidence that your male colleagues have?

Yes - Sometimes I am jealous because I can see how confident they are or how calmly they solve problems. I understand that behind this confidence, they have a lot of years of experience. They have read a lot of books. They have made a lot of mistakes. Nothing is easy to achieve. It's a good motivation.

Can you achieve this confidence too?

I hope so. I hope I achieve it one day.

What else would you like to achieve?

I want to be happy. I put all of my effort in being good at my profession, satisfied, surrounded by friends and family, to learn something new, and feel surprised by life every day.

In my profession, it's very easy. Keep your "doors" open and meet new people. The CARS conference was like a new door for me. I'm very lucky to be here.

"I think that every surgeon is afraid of blood, because in case of profuse bleeding we have seconds or minutes to save a life"

Leo Joskowicz

"Coming to a place like this, listening to lectures and talking to colleagues forces you to think a little more in depth about things"

Guest

Leo, CARS is very close to your heart. Can you tell us why?

CARS is close to my heart because I've been attending for a long time. It has been a forum for a very rich interdisciplinary exchange. The conference is between clinicians and scientists. I enjoy the variety and the interdisciplinarity. I have many colleagues and friends. I am always happy to see the new young students coming along in their careers.

What do you see in this generation of students that fascinates you in particular?

I think on the one hand, this generation has an easier time because things are more accessible, but on the other hand more variety. It's harder to make choices. When you go to a restaurant, and you have 100 dishes, it's difficult to choose one of them.

Having too many choices is a problem for the rich...

Of course, but many choices can also create some anxiety and confusion.

Can you give them guidance on how to overcome this? How can they make a choice that they will not regret?

There are no choices which one never regrets. I don't believe in that. The trick is to make choices and stick to them; then understand when you have to change. There is no formula for that. I think passion and enthusiasm are definitely required. In this field in particular, you need patience. This is not the field in which you do a 100 meter sprint and get success. This is a field with longer runs. You need 10 km, 20 km, a marathon… You need patience. That is why it's not suited for everyone. On the other hand, it is very enriching. The enrichment of interdisciplinarity costs time and effort more than in other fields. This is why you need more perseverance and more patience to let things bloom or change when they don't.

Let's move to the second part of what you said about change. I've read somewhere that half of the jobs that will exist 5 years from now, do not exist today. Whatever the real ratio may be, how can people choose when they don't know where their career will lead them?

This is part of what I was telling you about that when you have more choices and more opportunities. In a sense, it makes things harder. Choosing the right thing is becoming more difficult nowadays. There are changes, and there are industries that change very radically. I think having a good solid basis and being an expert in something specific is helpful. I think academic degrees are still useful for that. Don't overspecialize in one thing, but get a solid foundation. Then you can basically do one thing or another thing. I'm talking mostly about analytics, mathematics, computer science, and bioengineering.

I understand that the young people we are talking about have to deal with uncertainties, more than we did when we studied. How should they learn to deal with that?

I think the uncertainties are always there. I think only experience helps you deal with uncertainties. I don't think that there is a formula of how to deal with that.

It's not a consolation. [*laughs*]

No, it's not a consolation, but on the other hand, as you pointed out, having more choices is a problem of the rich. In the past, 50 years ago, there were 5 careers. You were a lawyer, an accountant, a doctor, an engineer, or a teacher. You had 10 options. You went to one career. Then afterwards you went to some specialization. Nowadays you can have Bachelor's degrees in very, very varied subjects. Sometimes they are overspecialized. Five years later, that specialty goes away. You have more of a problem. I am more

Guest

conservative in the sense of having a solid education in your background.

You told me a moment ago that there are no choices that you can never regret. Did you see people who were at the end of their PhD, and they started to have second thoughts about the huge investment that they made?

Yes, but they don't finish their PhD. Usually, they give up sooner. There are different options when you can stop. You can complete a Master's and such. You can opt out. I also advise my students to make sure that it is really what they want. I can tell you an anecdote when I went to do my PhD at NYU, in the late 80s. I have an uncle in New York that said: "What? You want to do a PhD in artificial intelligence? Two comments: why do you want a PhD when you can hire one? Artificial intelligence? Well, you have no natural intelligence. Why do you need artificial intelligence?" He was a Jewish man living in the Bronx… very practical. He thought there were three worthwhile careers; a doctor, an accountant, or a lawyer. The rest didn't count. A very nice man, but you know, I think he was very right because that basically forces you to question if it's worth the investment. This is a completely personal thing. There are no formulas. That's for sure.

You are very a very international person. You teach in Jerusalem, but you are not originally from Jerusalem. You went through several countries. What can you tell us about the internationality of your career? How has it enriched you?

I was indeed lucky enough to be able to be exposed to different places. I grew up in Mexico. I went to a French school. Then I did my Bachelor's in Israel at the Technion. Then I lived in New York for fifteen years. At a certain point, I felt very much like the wandering Jew, which also comes with some anxieties. In the end, I understood that home is wherever you are and with the people you are with. Therefore, I started to see that this is a great enrichment and privilege. I believe that human relationships enrich you in a great way most of the time. Some cause damage, but then you walk away from them. I think it gives a very good perspective which gives you a better understanding of the different tendencies. I think that it has enriched me incredibly. This is one of the things I like about CARS because I meet here people from all over the world. The diversity is extremely enriching.

Well, this is one of the things that is embedded in the soul of this conference.

Yes, they work very hard for internationality, and the result is here. You see it. That has been very successful.

Since you were a student, you have seen huge progress in technology. Can I ask you, which was the specific one that surprised you the most?

This is relatively recent: if you had told me 10 years ago that cars would drive by themselves, even knowing what is quite going on, I would say, no, this is 20, 30 years away. The past 10 years or even 3 years, we have seen a jump that has astonished me. I must say that when I started seeing the performance, the results, and the hype, of which there is a lot, it was still something that surprised me. I'm dreaming now every time I drive from Jerusalem to Haifa, I dream of sitting in the back of the car, reading, and letting the car drive. I think, wow, this could happen in my lifetime. If you had asked me 10 years ago, I wouldn't have believed that. This is something that has surprised me.

That might save us from road rages one day. We will not be as angry as we are today when we drive.

Well, that's to be seen because technology always gives rise to many unexpected behaviors and side effects that we have no clue about. Believe me, I will be very happy sitting in the back of the car with the car driving.

Can you also tell our readers which technological development will probably surprise them in the future by arriving unexpected?

[*laughs*] Well, that's a difficult question. I can tell you what probably will not happen, and it's likewise dramatic advances in healthcare and business. If, for some reason, it's somewhat happening in the field of autonomous driving, I think that biology is much more complex. You have to respect it much more. I would not necessarily see this type of spectacular advance in any field in medicine because of biology.

Even with the incredible success of technical research?

Yes, but again, after all, you have to go back to biology. Biology is much more complicated than anything we know. Certainly, it is more complex than autonomous driving. We don't really understand it fully. We're making progress, but that's where I don't see or expect a revolution there. Of course, I have yet to be surprised.

Which is the organ that fascinates you the most?

I can tell you that it's not necessarily the brain. That's the obvious answer. It's certainly fascinating but as I tell my colleagues at the brain center in Jerusalem (I'm also a member of that)… what's happening below the neck. I think that things that have to do with orthopedics are fascinating.

Guest

Guest

Do you think that there is a bottleneck that prevents some of the great research which is done here and at other places from reaching the real world between the clinical phase to industry?

I think that it's a naturally inefficient process. I talk about that a lot, and in the end, 1 out of 1,000 succeeds. I think it's very difficult to increase the rate. I think that by definition, research is extremely inefficient in terms of value and revenue… in terms of you do something, and you get something back.

Money complicates the equation.

I think by definition, research is that way. Sometimes you need 10,000 teachers to have one that invents the right formula. By nature, it is very inefficient. Of course, you can try to make it more efficient and have the technology transfer. Our field is particularly difficult. I think that in a sense, we have to accept the inefficiency of a system in which it is not very clear how the innovation and transfer is done. I think we're not doing that bad. My opinion is actually that many more things get transferred or tried than they used to be 40 years ago. That's the good news. On the other hand, for the perspective, I think that it is an inherently inefficient system.

I would like to conclude pointing out my belief in the importance of deep thinking and taking time to pause and reflect on more basic issues like the kind of activity that happens at academic conferences like this one. You can live fast and in the moment with very superficial things. I think coming to a place like this, listening to lectures and talking to colleagues, forces you to think a little more in depth about things. I think that's probably the main reason why I am here. I always look forward to coming back.

Upcoming Events

ICIAR - International Conf. on Image Analysis and Recognition
Montreal, Canada July 5-7 Website and Registration

EECVC - Eastern European Conference on Computer Vision
Odessa, Ukraine July 14 Website and Registration

ICPRS Internat. Conference of Pattern Recognition Systems
Madrid, Spain July 11-13 Website and Registration

MICCVTA Int. Conf. of Computer Vision Theory and Applications
Toronto ON, Canada July 15-16 Website and Registration

ECPA 2017 European Conference on Precision Agriculture
Edinburgh, Scotland July 16-20 Website and Registration

CVPR 2017 - Computer Vision and Pattern Recognition
Honolulu HI, USA MEET US! July 21-26 Website and Registration

InfoAg Conference - Precision Agriculture
St. Louis MO, USA MEET US! July 25-27 Website and Registration

ECVP - European Conference on Visual Perception
Berlin, Germany August 27-31 Website and Registration

BMVC 2017 British Machine Vision Conference
Imperial College London, UK Sep 4-7 Website and Registration

MICCAI - Medical Image Computing & Computer Assisted Intervention
Quebec, Canada MEET US! Sep 10-14 Website and Registration

ICIP 2017 IEEE International Conference on Image Processing
Beijing, China Sep 17-20 Website and Registration

RE•WORK Deep Learning Summit
London, UK Sep 21-21 Website and Registration

Did we miss an event? Tell us: editor@ComputerVision.News

FREE SUBSCRIPTION

Dear reader,

Do you enjoy reading Computer Vision News? Would you like to receive it **for free in your mailbox** every month?

Subscription Form
(click here, it's free)

You will fill the Subscription Form in **less than 1 minute**. Join many others computer vision professionals and receive all issues of Computer Vision News as soon as we publish them. You can also read Computer Vision News on our website and find in our archive new and old issues as well.

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

FEEDBACK

Dear reader,

How do you like Computer Vision News? Did you enjoy reading it? Give us feedback here:

Give us feedback, please (click here)

It will take you only 2 minutes to fill and it will help us give the computer vision community the great magazine it deserves!

RSIP VISION | Global Leader in Computer Vision

JOIN THE AI AGE IN OPHTHALMIC IMAGING

Be part of the AI revolution in ophthalmic imaging. RSIP Vision, a global leader in computer vision and image processing, has completed numerous ophthalmology projects involving development of advanced algorithmic software. In order to achieve state-of-the-art results, the company has invested in breakthrough technology like AI and deep learning.

Your applications can be first-class too! Talk with us about how to gain a competitive edge which will give your device a top-notch performance and a huge advantage in the marketplace. New technologies can benefit all eye care fields in all their components, from telemedicine to diagnosis, from data collection to cloud operations.

Consult our experts now!

Ophthalmology software Research & Development:

- Machine Learning
- Artificial Intelligence
- Image Processing
- Deep Learning
- Computer Vision
- Ophthalmic Imaging

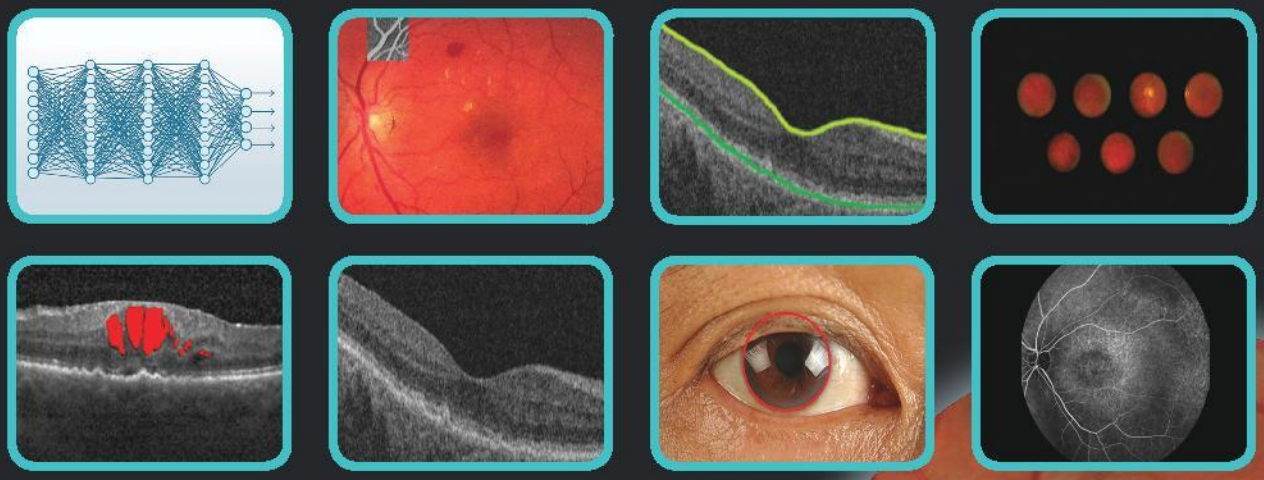

Proud sponsor of ARVO The Association for Research in Vision and Ophthalmology

www.rsipvision.com