


Computer Vision News

The magazine of the algorithm community

June 2017



Research Paper:
Mask R-CNN by Facebook AI Research - FAIR
Another outstanding work by Kaiming He et al.

Project:
Augmented Reality Applications

Trick:
parfor in MATLAB

Women in Computer Vision:
Deepti Pachauri

Application:
Arterys

Spotlight News

Events:
MICCAI 2017 Preview
With Simon Duchesne


Project Management:
Communication in Computer Vision Projects

Challenge:
Data Science Bowl 2017
Turning Machine Intelligence Against Lung Cancer

A publication by



Research - Facebook AI Mask R-CNN



04

Data Science Bowl Challenge Lung Cancer Detection




20

Women in Computer Vision Deepti Pachauri



26

Application Arterys



10

Project Augmented Reality Applications



22

Project Management Tip Communication in Comp. Vision Projects



32

Trick parfor in MATLAB

```
d = rand(5,5);
A = zeros(1,5);
parfor ii = 1:5
    A(ii) = myMax( d(:,ii) );
end

function res = myMax( data )
    myStruct.max = max( data );
    res = myStruct.max;
end
```

14

Spotlight News



25

Upcoming Events and MICCAI17 Preview



35

- 03 Editorial**
by Ralph Anzarouth
- 04 Research Paper**
Mask R-CNN - by Facebook AI
- 10 Application**
Arterys
- 14 Trick**
parfor in MATLAB
- 20 Challenge**
Data Science Bowl 2017 - Kaggle
Lung Cancer Detection

- 22 Project - RSIP Vision**
Augmented Reality Applications
- 25 Spotlight News**
From elsewhere on the Web
- 26 Women in Computer Vision**
Deepti Pachauri - 3M
- 32 Project Management Tip**
Communication in CV Projects
- 35 Computer Vision Events**
Calendar of June-August events
MICCAI17 preview with Simon Duchesne



Did you subscribe to Computer
Vision News?
It's free, click here!

Computer Vision News

Editor: **Ralph Anzarouth**

Publisher: **RSIP Vision**

[Contact us](#)

[Give us feedback](#)

[Free subscription](#)

[Read previous magazines](#)

Copyright: **RSIP Vision**

All rights reserved

Unauthorized reproduction
is strictly forbidden.

Dear reader,

The month of June arrived and with it a new issue of **Computer Vision News**. But this is far from being the only thing to be pleased about: the other good news are that this Summer we will have many opportunity to meet.

In fact, as you can see in our list of upcoming events at page 35, we will be present for the first time at **CARS**, the annual congress for **Computer Assisted Radiology and Surgery**, which will be held on June 20-24 in Barcelona, Spain. You will not be surprised to discover a new daily publication at that event...

After CARS, we have our usual meeting at **CVPR**, the **Computer Vision and Pattern Recognition** conference. It will be held in Honolulu, Hawaii at the end of July and (of course) we will publish a [Daily CVPR magazine](#).

MICCAI 2017 will be in September and we can witness that preparations are proceeding at full speed: Program Chair **Simon Duchesne** tells us all about the International Conference on **Medical Image Computing and Computer Assisted Intervention** at page 36 of this magazine.

The first article you will read after turning this page is our report about an outstanding paper by **Kaiming He, Georgia Gkioxari, Piotr Dollar** and **Ross Girshick** from **Facebook AI Research**. We are fascinated by their research and we recommend you do not miss it.

This Computer Vision News issue of June offers many other great reads. Please keep reading, enjoying and sharing them with colleagues and friends. We hope to see you at least at one of the big Summer events!

Enjoy the reading!

Ralph Anzarouth

Marketing Manager, **RSIP Vision**

Editor, **Computer Vision News**

Mask R-CNN - by Facebook AI Research (FAIR)

Every month, Computer Vision News reviews a research from our field. This month we have chosen to review **Mask R-CNN**, an outstanding paper offering a simple and flexible framework for **instance segmentation**. We are indebted to the authors (**Kaiming He, Georgia Gkioxari, Piotr Dollar and Ross Girshick from Facebook AI Research - FAIR**) for allowing us to use their images to illustrate this review. The paper is [here](#). The authors promised that they will publish their original source code (Caffe 2). In the meantime, here are third parties' implementations using [TensorFlow](#) and [PyTorch](#).

Without bells and whistles, Mask R-CNN surpasses all previous state-of-the-art single-model results on the COCO instance segmentation task, including the heavily engineered entries from the 2016 competition winner.

Background:

Results of object detection and semantic segmentation have dramatically improved over the last few years, mainly owing to Fast/Faster R-CNN (for object detection) and Fully Convolutional Networks - FCN (for semantic segmentation). These frameworks are relatively simple, flexible and fast to train. Mask R-CNN is a model that offers the same advantages in the area of object instance segmentation with improved accuracy.

Specifically, object instance segmentation is challenging because it requires both: (1) correct detection of all objects in the image, and (2) precise segmentation of each object detected. Instance segmentation combines two techniques: (1) object detection -- which estimates bounding boxes per object and then classifies the object within them; and (2) semantic segmentation -- where the goal is to classify each pixel into a fixed set of categories without differentiating an instance of what object the pixel is being assigned to.


The first impression is that a very complex method is likely needed to combine the tasks of object detection and semantic segmentation. Yet, Mask R-CNN offers a simple yet flexible and fast system that can surpass previous state-of-the-art instance segmentation results.

Mask R-CNN is a deep learning method for instance segmentation of objects in an image. It efficiently detects objects in an image, while creating a high quality segmentation mask for every object in the image. This system extends Faster R-CNN by an object mask prediction mechanism, in addition to the existing mechanism which estimates a bounding box per object.

Mask R-CNN is a relatively easy network to train and adds only a small overhead to the training-time of Faster R-CNN, while its rate of image handling is 5 fps.

In addition, the method is easy to extend to further needs and tasks: authors demonstrate how it may be used to predict the pose of a person in an image.

Mask R-CNN outperforms all existing, single-model entries on every task, including the COCO 2016 challenge winners. It achieves the best results on all three COCO challenges: (1) segmentation, (2) bounding-box object detection, and (3) person keypoint.



Previous work:

Mask R-CNN extends Faster R-CNN by adding a branch to the network for predicting a segmentation mask per region of interest (RoI), alongside the existing branches predicting bounding boxes and classification.

Readers who are not up to date on Fast R-CNN and Faster R-CNN can refresh their memory in the [March issue of Computer Vision News](#).

Reminder: R-CNN is primarily made up of three parts: (1) extract regional proposals; (2) extract features for each region – this part is performed by convolutional layers taken from a standard ConvNet network (such as AlexNet, ResNet or VggNet); and (3) classify the extracted features per category and region, using SVM and LR, respectively.

Generally speaking, Faster R-CNN is composed of two parts: (1) Region Proposal Network (RPN) extracts regions from the image that potentially include relevant objects, these potential regions being marked with bounding boxes; (2) RoIPool extracts features from every region of interest previously selected, and it classifies and performs regression on the bounding box. This second stage

implements a network called Fast R-CNN, a neural network developed in 2015. Faster R-CNN uses the combined features from both stages for faster processing. The following table summarizes recent developments in R-CNNs:

	Paper Title	Author	Year
R-CNN	Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation	Girshick et al.	2014
Fast R-CNN	Fast R-CNN	Girshick et al.	2015
Faster R-CNN	Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks	Ren et al.	2016

Method:

Mask R-CNN adopts the same two two-stage approach: the first stage is identical to the one used in Faster R-CNN. At the second stage, alongside prediction of bounding box and classification, a third segmentation Mask branch is added, predicting a binary mask for each RoI.

In practice, every Mask branch uses an $m \times m$ mask per RoI to predict the object, using FCN. This technique allows each of these branches throughout the network to maintain their spatial layout, without representing it as a unidimensional vector, lacking spatial context. This pixel-to-pixel relation requires the RoIs to be well-aligned to preserve the spatial relation of each pixel. This led to the development of a pixel-to-pixel alignment element -- the main piece missing in Faster R-CNN for Mask R-CNN -- the RoIAlign layer.

Traditionally, R-CNN networks used RoIPool for extracting small feature maps. RoIPool quantizes a floating-number RoI $[x/16]$ to the discrete granularity of the feature map. This has a dramatic negative impact on predicting pixel-accurate masks, unlike on classification. In response, the authors developed RoIAlign, properly aligning the extracted features with the input. The proposed change is simple: it obviates the need for any quantization of the RoI boundaries (i.e., we use $x/16$ instead of $[x/16]$). In addition, the approach uses bilinear interpolation at four regularly sampled locations in each RoI bin and it aggregates the result.

The Mask R-CNN has a per-pixel sigmoid and binary loss: this approach decouples mask and class prediction, with masks across classes not competing. This is a unique approach because usually neural networks have a per-pixel softmax and a multinomial cross-entropy loss, with masks across classes competing. This formulation is key for the performance of Mask R-CNN.

The Network Architecture:

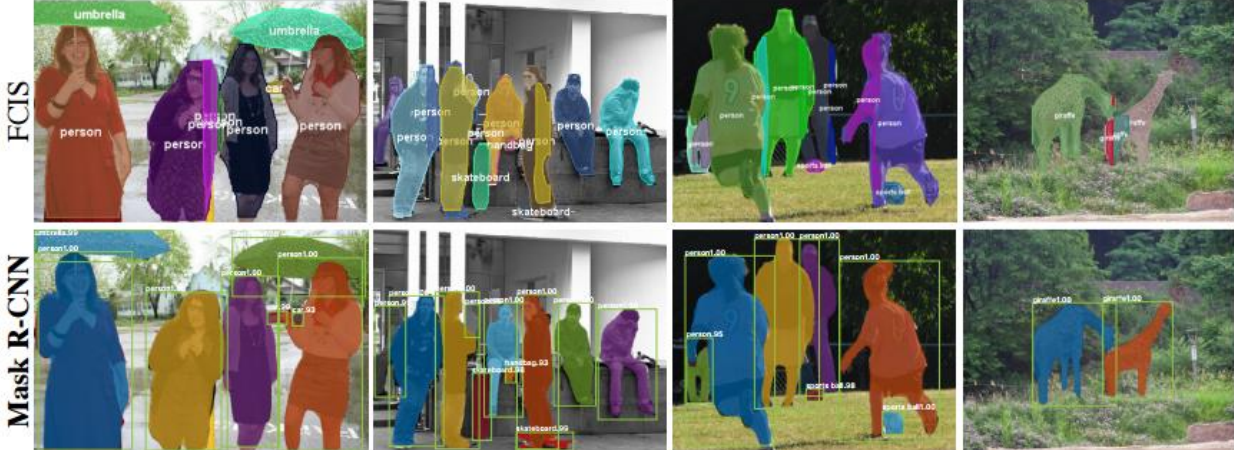
The authors evaluated several convolutional backbone architectures, used for feature extraction over an entire image: ResNet (ResNet-50-C4 - convolutional layer of the 4-th stage) and ResNeXt networks of depth 50 or 101 layers and Feature Pyramid Network (FPN) recently proposed by Lin et al.

During training, Mask R-CNN defines a loss function comprised of three elements: $L = L_{cls} + L_{box} + L_{mask}$. The first two are the classification loss L_{cls} and bounding-box loss L_{box} , identical to the loss defined in Faster R-CNN. The third loss is L_{mask} -- defined as the average binary cross-entropy loss: for an ROI associated with ground-truth class k , L_{mask} is only defined by the k -th mask (other masks not contributing to the loss).

Run Times:

As mentioned, Mask R-CNN runs at about 200ms per frame on a GPU, and training on COCO takes one to two days on a single 8-GPU machine. We believe the fast train and test speeds, together with the framework's flexibility and accuracy, will benefit and ease future research on instance segmentation.

Results:



	backbone	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
MNC [10]	ResNet-101-C4	24.6	44.3	24.8	4.7	25.9	43.6
FCIS [26] +OHEM	ResNet-101-C5-dilated	29.2	49.5	-	7.1	31.3	50.0
FCIS+++ [26] +OHEM	ResNet-101-C5-dilated	33.6	54.5	-	-	-	-
Mask R-CNN	ResNet-101-C4	33.1	54.9	34.8	12.1	35.6	51.1
Mask R-CNN	ResNet-101-FPN	35.7	58.0	37.8	15.5	38.1	52.4
Mask R-CNN	ResNeXt-101-FPN	37.1	60.0	39.4	16.9	39.9	53.5

Instance segmentation mask AP on COCO test-dev. MNC and FCIS are the winners of the COCO 2015 and 2016 segmentation challenges, respectively. Mask R-CNN outperforms the more complex FCIS+++.

	AP	AP ₅₀	AP ₇₅
<i>softmax</i>	24.8	44.1	25.1
<i>sigmoid</i>	30.3	51.2	31.5
	+5.5	+7.1	+6.4

Mask R-CNN has a per-pixel sigmoid and binary loss: this decouples mask and class prediction, with masks across classes not competing. The table above provides a comparison between a per-pixel softmax and sigmoid loss, showing the efficiency, especially in training, of the approach predicting a binary mask without concern for categories.

	align?	bilinear?	agg.	AP	AP ₅₀	AP ₇₅
<i>RoIPool</i> [12]			max	26.9	48.8	26.4
<i>RoIWarp</i> [10]		✓	max	27.2	49.2	27.1
		✓	ave	27.1	48.9	27.1
<i>RoIAlign</i>	✓	✓	max	30.2	51.0	31.8
	✓	✓	ave	30.3	51.2	31.5

The table above details the results of an evaluation of the proposed RoIAlign layer, compared with RoIPool and RoIWarp. RoIWarp, similar to RoIPool, also quantizes the RoI, losing alignment. ResNet-50-C4 backbone, which has stride 16, was used for the experiment. RoIAlign improves AP by about 3 points over RoIPool, with much of the gain coming at high IoU (AP75). This highlights the importance of proper alignment for predicting pixel-accurate masks.


	backbone	AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅	AP ^{bb} _S	AP ^{bb} _M	AP ^{bb} _L
Faster R-CNN+++ [19]	ResNet-101-C4	34.9	55.7	37.4	15.6	38.7	50.9
Faster R-CNN w FPN [27]	ResNet-101-FPN	36.2	59.1	39.0	18.2	39.0	48.2
Faster R-CNN by G-RMI [21]	Inception-ResNet-v2 [37]	34.7	55.5	36.7	13.5	38.1	52.0
Faster R-CNN w TDM [36]	Inception-ResNet-v2-TDM	36.8	57.7	39.2	16.2	39.8	52.1
Faster R-CNN, RoIAlign	ResNet-101-FPN	37.3	59.6	40.3	19.8	40.2	48.8
Mask R-CNN	ResNet-101-FPN	38.2	60.3	41.7	20.1	41.1	50.2
Mask R-CNN	ResNeXt-101-FPN	39.8	62.3	43.4	22.1	43.2	51.2

The table above compares Mask R-CNN to COCO 2016 Object Detection Challenge bounding-box object detection top performers. For this result, even though the full Mask R-CNN model is trained, only the classification and box outputs are used (the mask output is ignored). Mask R-CNN using ResNet-101-FPN outperforms the base variants of all previous state-of-the-art models, including the single-model variant of G-RMI, winner of COCO 2016. With ResNeXt-101-FPN as backbone, Mask R-CNN further improves its results, beating the best previous single model entry (using Inception-ResNet-v2-TDM) by 3 points.



Results on human pose estimation are shown above, demonstrating the flexibility of Mask R-CNN. Keypoint detection results on COCO test using Mask R-CNN (ResNet-50-FPN used a backbone), with Mask R-CNN predicting K masks, one for each of K keypoint types (e.g., left shoulder, right elbow). This model has a keypoint AP of 63.1 and runs at 5 fps.

More results of Mask R-CNN on COCO test images:



Conclusion:

Without bells and whistles, Mask R-CNN surpasses all previous state-of-the-art single-model results on the COCO instance segmentation task, including the heavily engineered entries from the 2016 competition winner. The proposed method also outperforms at the COCO object detection task. The authors ran ablation experiments, evaluating basic instantiations, demonstrating the method's robustness and allowing analysis of the effects of core factors.

Arterys



ARTERYS

An automated, intelligent medical imaging analytics cloud platform

Arterys provides an automated, intelligent medical imaging analytics cloud platform for physicians. Based in Silicon Valley with offices in Calgary, Canada and Europe, the company has about 40 employees working together to provide this unique medical imaging platform.

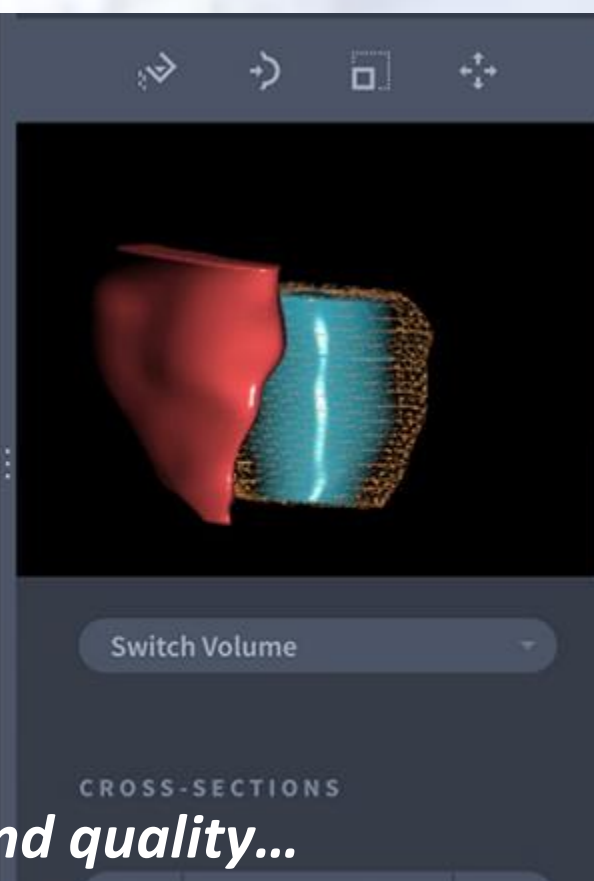

CTO and Founder John Axerio-Cilies says Arterys has created a way to host multiple applications that are ultimately served up to physicians. They focus on automation and bringing machine learning to medical imaging in order to provide a level of intelligence that does not exist today in medical clinics.

Arterys received approval earlier this year for its first FDA-cleared

application. Their first case involved cardiac MRI providing an application for cardiologists and radiologists to automate much of the tedious process of annotating heart muscles, or the myocardium. Today it can take up to 30 minutes or longer to annotate every phase of the cardiac cycle. Radiologists don't have the time to spend drawing circles when they could be helping patients instead. Arterys leverages deep learning in order to automate that entire process in seconds.

Physicians also have a problem with accuracy and consistency in diagnosing medical imaging. Ordinarily, up to 30% of medical imaging diagnoses have some sort of error. That happens, not only between hospitals, but also between readers in an institution. A different radiologist will diagnose the images differently.

This occurs on a regular basis. **Deep**



Deep learning brings consistency and quality...

learning brings consistency and quality, says Axerio-Cilies.

In medical imaging, many companies provide algorithms or some sort of machine intelligence. Axerio-Cilies thinks this accounts for about 5% of the work required to create a clinical solution. Anyone can create some sort of algorithm to detect, track, segment, or automate any other kind of medical imaging process. That's the easy part. The hard part is actually bringing that to the clinic, to deliver that sort of intelligence straight to where the physicians actually look at the images.

What they have been able to do, in collaboration with several channel partners, is to set up a distribution channel so that the radiologists can buy the software from major vendors and have access to it. Then it can be integrated into their existing workflow. Integration means leveraging a zero-footprint viewer. As a cloud based solution, Arterys allows any physician to interact with and report on their data from anywhere in the world. That's one unique problem that they solve. They serve up the data and have deep learning run inference in seconds while still maintaining a user-experience which is as perfect as possible.

“To build an entire global infrastructure so that the machines could communicate from anywhere in the world”

In order to accomplish this, they had to build an entire global infrastructure so that the machines could communicate from anywhere in the world. That was a major challenge that they faced. It was more of a data organization system that was difficult to set up, rather than

algorithmic machine learning or deep learning work. Arterys is completely link-based with zero-footprint, which means it does not require any installation. **Users can easily access the technology right from a URL.**

From a technical perspective, optimizing the algorithm is technically challenging, but in terms of time, it's really not the bottleneck. It became an issue of solving more of the high-level commercial problems.

Other companies that provide some sort of niche application typically offer their product as a CD-ROM, USB, downloadable software, or a client-server based architecture stored on their local intranet. However, the lack of a central repository means that these systems cannot receive daily auto-updates, and therefore, do not get smarter with time.

Arterys faced the challenge of providing cutting-edge technology that could quickly adapt in a rapidly developing industry. They turned to the cloud to



CTO and Founder John Axerio-Cilies

solve this problem. Axerio-Cilies had seen the cloud completely overtake local, static software in many industries. After all, moving to the cloud has many advantages: it allows for more effective distribution, as well as control and management of everything each user has access to from a central position. It enables to upsell and provide additional applications with virtually no installation, configuration, setup, hardware or infrastructure. In addition, the cloud has the unique advantage to scale and provide massive distribution through word of mouth.


software takes months. You need to go through security, the provisioning of the hardware, and the installation. A single click in a cloud-based ecosystem will handle all of that in seconds.

Hackers have gotten into hospitals and held the data hostage

Arterys uses a SaaS-based model, which means that every user pays a license per year. They went with this approach because it's very standard even for local workstation-based software. Going forward into the future, many markets could potentially be ready for a bundling based on a certain volume or even on a per case basis where they would report back their bill. However, he has noticed that a lot of countries and costs are not quite ready for that jump yet.

In the past, there has been skepticism about the cloud for privacy or security reasons. The industry is shifting now because of high profile hijackings in which hackers have gotten into hospitals and held the data hostage. Now people realize that **a local, closed environment that stores all of the data on the premises is not any safer than cloud-based storage**. They're seeing more and more people adopt the cloud on their system.

The team at Arterys has created a unique software that actually de-identifies the data and keeps patient health information within the hospital network. Arterys only keeps random pixel data outside of the hospital, but when a user logs in, they see all of the information. This *"magic happening behind the scenes"* has helped Arterys to leverage the cloud, not only in the United States, but also in Europe. They sell their product in markets such as Italy, France, and Germany where it is



Axerio-Cilies also believes in the importance of leveraging existing channels. For example, GE and Siemens are building ecosystems, what he calls "app stores". These app stores would host cloud-based applications within a central environment so that a user could quickly download and get access to all of that information, having it at their fingertips in seconds. Installing some sort of workstation based

difficult to receive clearance. They've been able to do that particularly because of their software infrastructure, and also because no other company is doing this today.

Looking ahead, they have very strong ambitions to grow their team, both from an FTE perspective and also from a commercial-sales perspective. It took them several years to build a platform, and they just became commercial this year. Currently, they are leveraging GE and Siemens to distribute their product. They have very strong ambitions to grow their sales organization, not only in the United States and Europe, but in other countries as well.

"It's really a revolution that's about to happen!"

Axerio-Cilies shared a unique story with us regarding the development of Arterys:


*"I think there is a lot of skepticism around **deep learning**, or around **machine learning**, as to what it can do. Cardiologists are renowned for being particular and maybe don't believe in every hype, to their credit. When we started this test and created this algorithm, we had no idea how accurate it was. In order to get **FDA clearance**, we had to test the algorithms and run it through evaluation and verification.*

*We were going into this clinical study completely blind where we had physicians manually annotate or manually contour images. They said that machine learning could never be as accurate as the way they annotate. When we compared our results with all of these expert annotators', what was amazing was the **massive inter-reader***

variability of these expert annotators. For example, you need to measure the ejection fraction of the heart. One person would say that it was 65%. Another person would say that it was 46%. These are experts in the field. It was crazy! We did this clinical study and found out that we were actually right in the center of the average of all of these expert physicians.

*That was remarkable for us because we didn't know what our algorithm could do. By proving that we were right down the center, we could say that our algorithm was just as accurate as expert physicians. It also showed that we can solve this inter-reader variability. This range is so broad. I really do think that once automation comes into this picture, the level of variability will come down and it will be more consistent across users and across sites. That's a world where we want to move to. **It's a world of data-driven medicine.***

*Right now, physicians do not have access to historic data in a very simple way. What we want to move to is combining the data from all different sites and to leverage **a central intelligence** to help improve the patient care path. Right now we're just scratching the surface of this. There are several different companies in this space doing the same thing. **It's really a revolution that's about to happen!"***



parfor in MATLAB

Image processing algorithm development is often done in Matlab. The trick of this month will review a way to improve run-time efficiency, parallel execution using a `parfor` loop.

A `parfor` loop is just like a regular loop: a sequence of instructions iterated for a range of values. The uniqueness of `parfor` is that the loop's iterations aren't executed serially, running instead in parallel as much as it is made possible by the number of CPUs available.

It's easy -- some would say, too easy -- to convert `for` loops into `parfor` loops. In fact all we have to do is add the prefix `par-`. This change was made deliberately simple by MathWorks (which is great!). On the other hand, the seemingly trivial change hides significant work taking place.

It's important to understand that every parallel iteration is sent to a worker. A worker is an entity responsible for executing a single iteration of the loop. A number of workers can run parallel on the same computer, each with its own dedicated memory space. Thus, the data needed to perform its tasks is sent to the worker by the client, and the worker's product is sent back to the client. The great advantage of using `parfor` is that data transfer is performed transparently for the user. Matlab automatically analyzes the content of the loop's instructions and decides what data needs to be sent to which worker, and what data (i.e. which variables) gets sent back to the client. Understanding this automatic process is the key to understanding what constraints limit the use of the `parfor` loop and why -- what can and what cannot be performed in parallel by a `parfor` loop.

The first thing to consider when working with `parfor` is how Matlab treats each of the variables in the loop.

Matlab distinguishes between 5 types of variables:

- Loop Variables -- Defines the loop index value for each iteration.
- Sliced Variables -- Arrays whose segments are operated on by different iterations of the loop.
- Broadcast Variables -- Variables defined before the loop, whose value is required inside the loop, but never assigned inside the loop.
- Reduction Variables -- Variables that accumulate a value across iterations of the loop, regardless of iteration order.
- Temporary Variables -- Variables created inside the loop, and not accessed outside the loop.

The following code template shows all five types of variables in context:


```
TempVar1 = 0;
BroadcastVar = 50;
ReduceVar = 0;
sliceInput = rand(1,100);
parfor loopVar = 1:100
    TempVar1 = loopVar;
    ReduceVar = ReduceVar * loopVar;
    sliceOutput(loopVar) = sliceInput(loopVar);

    if (loopVar < BroadcastVar)
        TempVar2 = TempVar1 + 5;
    end
end
```

Let's understand each variable:

The loop variable: this is the loop index value for each iteration. It is set in the first line of a `parfor` statement.

Temporary variable: variables created inside the loop, and not accessed outside the loop.

It's important to understand the properties of temporary variables in `parfor`: (1) `parfor` clears temporary variables before each iteration; (2) they must be set inside the `parfor`-loop, so that their values are defined separately for each iteration; (3) they cannot be passed from one iteration of the loop to another; (4) they do not get sent back to the client; (5) they have no effect on other variables with the same name that exists outside that iteration of the loop. These properties / behaviors are different from ordinary `for`-loops.

Let's take a look at an example:

```
T = true;
parfor ii = 1:100

    if T && foo(ii)
        T = false;
    end
end
```

T is a temporary variable because it is assigned in the body of the loop. It is cleared ahead of each iteration, so that its use in the `if` condition is uninitialized. If you change `parfor` to `for`, the value of T assumes sequential execution of the loop.

Broadcast Variables: a *broadcast variable* is any variable that (1) is not a loop variable or (2) a sliced variable and (3) does not change inside the loop; (4) it is declared before the `parfor` loop.

You should note that broadcast variables get sent to all workers. So large broadcast variables can weigh down the loop with significant communication. Therefore, it's sometimes more efficient to use temporary variables in their place.

Reduction Variables: a *reduction variable* accumulates value based on all iterations, independent of iteration order. All assignments for a reduction variable must have the same operation.

As in 2D SIFT, here too, interpolation is conducted between the histogram bins to mitigate the effects of quantization. Specifically, at the selected orientation, interpolation between the vertices of the triangular tile, as presented in the following illustration:

For example:

The code-snippet on the left is not valid because the reduction has two different assignments ($A=A+4$) and ($A=[A,4]$). The `parfor` loop on the right is valid.

```
%%
n = 100;
A = [];
parfor i = 1:n
    if (i>n/2)
        A = A + 4;
    else
        A = [A, 4];
    end
end
```

The variable 'A' is used like a PARFOR reduction variable, but it has additional (invalid) uses. [Details](#)

```
%%
n = 100;
A = [];
parfor i = 1:n
    if (i>n/2)
        A = A + 4;
    else
        A = A + 4 * i;
    end
end
```

Sliced variable: a *sliced variable* is a variable that is broken up into segments/slices; those slices are operated on separately by different workers. Each worker is sent only the slices it needs.

With `ii` as the loop variable, the `Mat` variables shown on the left are not sliced, while the `Mat` variables on the right are sliced:

invalid for slicing

```
% foo(k) invalid for slicing
mat(ii+foo(kk),jj,:,30)

% 10:20 not scalar
mat(ii,10:20,end);

% struct.field1 not simple broadcast var
mat(ii,:,struct.field)
```

valid for slicing

```
%valid for slic
mat(ii+k,jj,:,3
mat(ii,:,end)
```

Another issue in `parfor` loops are iterations defined in terms of the previous iteration: in this case, the dependency problem is reported.

```
XX = 1:10;
parfor k = 1:10
    XX(k) = XX(k+1);
end
```

The PARFOR loop cannot run due to the way variable 'XX' is used. [Details](#)

Yet another example:

The left code-snippet has multiple assignments to different parts of the variable `Y` in `parfor`, therefore causing an error as variable `Y` cannot be classified. The code-snippet on the right works, because the value of `x` is completely overwritten. Now, `parfor` can unambiguously determine that `x` is a temporary variable.

```
out1 = zeros(1,50);
parfor ii = 1:50
    Y(1) = 10;
    Y(2) = 20;
    out1(ii) = avg(Y);
end
```

The PARFOR loop cannot run due to the way variable 'Y' is used. [Details](#)

```
out1 = zeros(1,50);
parfor ii = 1:50
    Y = [10, 20];
    out1(ii) = avg(Y);
end
```

Using struct in parfor loop:

Structures cannot be assigned in a `parfor` loop using dot notation. In the code-snippet on the left, both lines in the body of the loop generate classification errors. In the code-snippet on the right, as a workaround, you can create the structure in the loop or in the first field by using the `struct` function.

```
parfor ii = 1:40
    myStruct.structFiled = 7;
end
```

The PARFOR loop cannot run due to the way variable 'myStruct' is used. [Details](#)

```
parfor ii = 1:40
    myStruct = struct();
    myStruct.structFiled = 7;
end
```

Variable classification problems in `parfor` loops can usually be solved by converting the body of the `parfor` loop into a function.

In the code-snippet on the right, the problem with the `struct` variable on the left is solved by converting the body of the `parfor` loop into a function.

```
d = rand(5,5);
A = zeros(1,5);
parfor ii = 1:5
    myStruct.max = max( d(:,ii) );
    A(ii) = myStruct.max;
end
```

The PARFOR loop cannot run due to the way variable 'myStruct' is used. [Details](#)

```
d = rand(5,5);
A = zeros(1,5);
parfor ii = 1:5
    A(ii) = myMax( d(:,ii) );
end

function res = myMax( data )
    myStruct.max = max(data);
    res = myStruct.max;
end
```

Let's see now what happens with a nested `for` loop:

The range of a `for` loop converted to a `parfor` loop must be properly defined as constant numbers or variables.

The code-snippet on the left causes an error because the upper limit of the `for` loop is defined by a function call. The code-snippet on the right provides a workaround by first defining a broadcast or constant variable outside the `parfor` loop body.

```
Mat = zeros(10, 10);
parfor ii = 1:size(Mat, 1)
    for jj = 1:size(Mat, 2)
        Mat(ii, jj) = ii * jj;
    end
end
```

The PARFOR loop cannot run due to the way variable 'Mat' is used. [Details](#)

```
Mat = zeros(10, 10);
inAryIdxSize = 10;
parfor ii = 1:size(Mat, 1)
    for jj = 1:inAryIdxSize
        Mat(ii, jj) = ii * jj;
    end
end
```

In the following example, the code-snippet on the left causes an error because the upper limit of the `for` loop is defined by a function call. The code-snippet on the right provides a solution to this error by setting the broadcast or constant variable outside the `parfor` loop:

```
Mat = zeros(10, 10);
inAryIdxSize = 10;
parfor ii = 1:size(Mat, 1)
    for jj = 1:inAryIdxSize
        Mat(ii, jj+1) = ii * jj;
    end
end
```

The PARFOR loop cannot run due to the way variable 'Mat' is used. [Details](#)

```
Mat = zeros(10, 10);
inAryIdxSize = 10;
parfor ii = 1:size(Mat, 1)
    for jj = 1:inAryIdxSize
        Mat(ii, jj) = ii * jj + 1;
    end
end
```

Nested `parfor`: A regular `for` loop, nested within a `parfor` loop, can only use a single index, as seen on the right. The code-snippet on the left causes an error because `jj` and `kk` loop over different values.

```
Mat = zeros(10, 10);
parfor ii = 1:10
    for jj = 1:7
        Mat(ii, jj) = ii * jj;
    end
    for kk = 8:10
        Mat(ii, kk) = 10;
    end
end
```

The PARFOR loop cannot run due to the way variable 'Mat' is used. [Details](#)

```
Mat = zeros(10, 10);
parfor ii = 1:10
    for jj = 1:7
        Mat(ii, jj) = ii * jj;
    end
    for jj = 8:10
        Mat(ii, jj) = 10;
    end
end
```


Important issues to consider:

(1) When using a `parfor` loop it is important to remember that any variables defined as *broadcast variables* will be sent to all workers, therefore when defining the data structure `<?!>` used by the `parfor` loop you should consider and plan the minimal set of data that will be required within the loop

(2) It is also important to consider the number of workers that you want to define (which can be set by the `parpool` command). By default, Matlab initializes workers according to the number of logical CPU cores available. However, cases in which hyperthreading doesn't improve (and sometimes even degrades) performance have been reported. In case you suspect this is happening to you, it is recommended to check if setting the number of workers to the number of physical (rather than logical) cores improves performance. [Click for additional reading.](#)

(3) You should note that the way to achieve true multi-threading in Matlab is via MEX, Java or .Net, or by spawning external standalone processes. [Read more about this option.](#)



The cover of Computer Vision News magazine features a collage of images related to computer vision and image processing. At the top, the title "Computer Vision News" is prominently displayed in large white letters, with the subtitle "The Magazine Of The Algorithm Community" below it. A central white box with the word "Subscribe" and the text "(click here, it's free)" is overlaid on the cover. The background includes various images: a 3D visualization of a face, a street scene with cars, a medical scan of a brain, and a person in a lab coat. Text on the cover includes "The only magazine covering all the fields of the computer vision and image processing industry" and "A publication by RSIP VISION".

Computer Vision News

The Magazine Of The Algorithm Community

The only magazine covering all the fields of the computer vision and image processing industry

Subscribe

(click here, it's free)

A publication by

RSIP VISION


Kaggle - Can you improve lung cancer detection?

Breathe in the Future

The 3rd annual Data Science Bowl aligns with Vice President Joe Biden's Cancer Moonshot™ to achieve one of the key strategic goals: unleashing the power of data against lung cancer. Together, we can pit machine learning and artificial intelligence against cancer, advancing the state of the art in future screening, care, and prevention. Let's **breathe out the past** and end the disease as we know it.

Vital Energy

Lungs fuel us with the oxygen we need in our bloodstream to provide our cells with energy, while removing the lethal waste of carbon dioxide.




Lung Cancer is the most common type of cancer with...

225,000

new cases in the U.S. in 2016¹


\$12 billion

were accounted for in healthcare costs in the U.S. every year²



Malignant Tumors

Tumors form when gene changes in the DNA of the cells mutate and promote unnatural growth. Uncontrolled growth can spread to surrounding areas or metastasize to other organs if not treated early.



Low-Dose CT scans help assess if a person is at risk of lung cancer or other pulmonary disease. Scientific research reports...

20%

of lung cancer deaths can be reduced with early detection³

However, the image assessments in use today are identifying lung lesions as potentially cancerous that later turn out to not be cancer.

High false positive rates

lead to unnecessary patient anxiety, additional follow-up imaging and interventional treatments^{3,4}

SOURCES:

¹ Siegel RL, Miller KD, Jemal A. "Cancer Statistics," 2016. CA: A Cancer Journal for Clinicians. 2016; 66:7-30.

² National Institutes of Health, "Cancer costs projected to reach at least \$158 billion in 2020," <https://www.nih.gov/news-events/news-releases/cancer-costs-projected-reach-least-158-billion-2020>, (January 12, 2011).

³ Aberle DR, Adams AM, Berg CD, et al. "Reduced lung-cancer mortality with low-dose computed tomographic screening." N Engl J Med. 2011;365:395-409.

⁴ Low-Dose CT has historically resulted in high false positive rates of around 25% (Aberle, et. al., New England J Med, 2011, 365:395-409).

Breathe in the future. Breathe out the past.



#DATASCIOWL

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 have chosen to review the **Data Science Bowl 2017: Turning Machine Intelligence Against Lung Cancer**, organized by **Kaggle** and concluded only a few weeks ago. The websites of the challenge are [here](#) and [here](#).



Background

Lung cancer is today one of the main causes of death in the world among both men and women, with an impressive rate of about five million deadly cases per year. Early detection is critical to surviving lung cancer: it enables treatment options not available when cancer is detected at a more advanced stage.

Early detection can reduce lung cancer deaths by up to 20% thanks to Low-dose Computed Tomography (CT), the potential of which is very promising, but unfortunately still inaccurate, generating a large number of false positives. [Read here about a very efficient technique to eliminate false positives in the findings.](#)

Motivation

The purpose of this 3rd Data Science Bowl is to ask the global community of data scientists to help radically improve doctors' abilities to accurately detect lung cancer at its first stages. This will enable more people to get the life-saving treatments they need to prevail over this deadly disease.

With 1 million USD in prize money provided by the Laura and John Arnold Foundation, participants developed artificial intelligence algorithms to accurately determine when lesions in

the lungs are cancerous, and reduce the excessive false positive rate (about 25%), due to suspicious lesions initially identified as high risk of cancer (and which subsequent follow-up tests proved to be non-cancerous).


Results

Winning team grt123 shared both their [code](#) and their [final document](#). Their solution is a unified framework of 3 modules: pre-processing, nodule detection and cancer classification based on 3D Convolutional Neural Networks.

Their model was implemented using Pytorch, appreciated for its flexible and efficient deep learning framework.

Below is the architecture of the nodule detection network developed by the winning team. Patch-based (rather than image-based) inputs were used for training due to GPU memory constraint when working on fine-level lung scans.

Nodule Net (detection)



Augmented Reality - medical and non-medical

Every month, Computer Vision News reviews a successful project. Our main purpose is to show how diverse image processing applications can be and how the different techniques contribute to solving technical challenges and physical difficulties. This month we review **RSIP Vision's** algorithms and techniques for the use of **Augmented Reality in both medical and non-medical applications**. RSIP Vision's engineers can assist you in countless application fields. [Contact our consultants now!](#)



Why an always increasing number of projects include Augmented Reality? Let's take an example from the medical world. Augmented Reality can use the pre-operative planning and transform it into a guide for the operation itself: it shows the surgeon what needs to be done in real-time and in a comfortable way.

How does Augmented Reality do it?

Augmented Reality generates an overlay over the area which needs treatment and it visually displays the next steps to the doctor, following the pre-op plan.

What are its main fields of application, besides medical? Augmented Reality is widely used in other industries, like in training, gaming and in a very wide range of applications. A good example

is Augmented Reality for bicycle riders: glasses equipped with Augmented Reality can show the cyclist where to drive, as well as useful statistics, which are much easier to consult this way than on a smartphone; in addition, their direct communication with the track enables better safety.

Bringing AR to your eyes

What are the latest development of Augmented Reality affecting our industry? Once, Augmented Reality was restricted to television or computer video. Now it is used in smartphones too and, during the last few years, much work has been done on Augmented Reality glasses, like Meta and [Microsoft's HoloLens](#). These glasses can literally bring Augmented Reality to your eyes, without the limits imposed by the mediation of a television screen. The environment as we see it perfectly integrates the virtual objects and we see it exactly as if they were really in that space, opening limitless possibilities in many application fields.

What is **RSIP Vision's** experience in Augmented Reality and which of its projects involved it? RSIP Vision's main use of Augmented Reality has been to date in the field of [ophthalmic imaging projects](#) like the development of software for intra-ocular retinal surgeries and [cataract interventions](#).

Again, the purpose is to offer guidance to the surgeon: for instance, by keeping track of the areas in which the cuts have been done or by superposing an **Augmented Reality planning overlay** on the eye itself. The map stays on the eye even when it moves, as an accurate navigation system should do.

Reduce recovery time and risks

Navigations allows the surgeon to go through exactly the same cut, reducing recovery time and risks.

How is this done? Before the surgery, the doctor generates an operational plan for it. This plan is positioned in real-time on the retina of the patient, to see exactly what area deserves treatment and how. This procedure is carried out in a very accurate way.

In what other application fields did RSIP Vision use Augmented Reality? We have also made a large use of Augmented Reality in **orthopedic surgery**. The most significant challenge in Augmented Reality is to make a [registration of the plan with the real world](#). Since we are talking about bones



which are found within the body, what we need to do is a registration of the bones in the space.

Registration is the most basic and important phase in Augmented Reality work, enabling to display the virtual layer (the pre-op plan) in the real space. It is done by finding the alignment of points in the two coordinates systems which need to be registered. The procedure to do so is a rigid transformation, which is composed of Rotation and translation $[R, t]$.

Two possibilities exist: either the pairing of points among the two datasets is known; or it is not. The optimal configuration is when pairing is known, so that its solution is in closed form (analytical). When pairing between point is unknown, an iterative solution needs to be found, which requires initialization and only guarantees convergence to local optimum.

The analytical procedure to find the optimal solution of known pairing uses a least-square technique as follows:

$$P_a = \{p_{a,1}, p_{a,2}, \dots, p_{a,n}\} \quad \text{and} \quad P_b = \{p_{b,1}, p_{b,2}, \dots, p_{b,n}\}$$

are two sets of corresponding points, one from image a and one from image b.

Given two points $p_{a,i}, p_{b,i}$ and the transformation $[R, t]$, the residual pairwise distance error is:

$$e_i = p_{b,i} - Rp_{a,i} - t$$

We seek a rotation R and a translation vector t that minimize the sum of squared errors:

$$\sum_{i=1}^n \|e_i\|^2$$

We show the points around the center of mass of each set of points:

$$p'_{a,i} = p_{a,i} - \mu_a \quad p'_{b,i} = p_{b,i} - \mu_b$$

The error term will be shown as:

$$e_i = p'_{b,i} - Rp'_{a,i} - t'$$

in which:

$$t' = t - \mu_b + R\mu_a$$

$$\begin{aligned} \sum_{i=1}^n \|e_i\|^2 &= \sum_{i=1}^n \|p'_{b,i} - Rp'_{a,i} - t'\|^2 \\ &= \sum_{i=1}^n \|p'_{b,i} - Rp'_{a,i}\|^2 - 2t' \sum_{i=1}^n [p'_{b,i} - Rp'_{a,i}] + n\|t'\|^2 \end{aligned}$$

Term 1 should be minimized with respect to the rotation operator

Term 2 is nil, since it relates to the center of mass

Term 3 is non-negative and depends only on t'

$$\sum_{i=1}^n \|p'_{b,i} - Rp'_{a,i}\|^2 = \sum_{i=1}^n \|p'_{b,i}\|^2 - 2 \sum_{i=1}^n p'_{b,i} \cdot Rp'_{a,i} + \sum_{i=1}^n \|Rp'_{a,i}\|^2$$

Term 1 is composed of the following

Constant

Maximizing this term minimizes the error

Constant

RSIP Vision has much experience in both 2D and 3D registration. [Consult our engineers](#) to learn how Augmented Reality can enhance your results!

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!

SOMA: Predicted 3D Human Shape & Motion Test

This is a human-aware Artificial Intelligence platform, trained using body data and computer vision using just a smartphone camera to capture 3D human motion and shape. It comes directly from [Michael Black's Body Labs!](#) [Watch](#)



What do 50 million drawings look like?

Google Creative Lab has open sourced the world's largest doodling data set to help developers train new neural networks. It includes 50 million drawings by over 15 million players, including ours. In a couple of minutes it will include also yours! [Play the Game](#) or [Get the Data](#)



Google Lens offers a clear view of the company's future

Google's efforts in computer vision are paying off: it has moved from organizing the Web to basically organizing everything, inside and outside of screens. With computer vision reaching an "inflection point" where error rates are lower than a human's, with better cameras and with more capable devices and networks, Google Lens is set to turn Augmented Reality into a computing tool which can make sense of the world without any user input. Like allowing Android users to analyze their surroundings via the Google Assistant: just point the phone at a restaurant and get ratings information, or aim the camera at a flower to identify the species. [Read&Watch...](#), [Read](#) and [Read](#)



Top 20 AI experts you should follow on Twitter

It's not us saying that and it's not our list. But we pretty much agree with most of it. You don't miss by following these fine folks and you might even learn a lot! Before you ask, the picture is Angelica Lim's at [@petitegeek](#). And before you ask, women make less than 30% of the list. [Read...](#)



The most accurate digital model of a human face

Computer scientists used 3D morphable model (3DMM) which represents an average face, but also contains information on common patterns of deviation from that average. James Booth of Imperial College London (ICL) and friends have developed a new method that fully automates the construction of 3DMMs, avoiding the need for tedious labeling all of their features. [Read...](#)



Great map: [The Top 263 Companies Racing Toward Autonomous Cars](#)

Deepti Pachauri - 3M

We continue our series of interviews with **women in computer vision**. This new section, which we started with the [CVPR Daily at CVPR 2016](#), hopes to help mitigate the severe gender imbalance in the computer vision community by getting to know better some **remarkable female scientists** and their career paths. This month, we interview **Deepti Pachauri**, a research scientist in Machine Learning and Computer Vision at **3M**.

Deepti, you spent the first part of your life in India. Whereabouts?

I grew up in Haridwar, a very small town in the state of Uttaranchal India. I stayed there until my undergrad. In the same town there was a small college called CD College and I did my Bachelors of Science there.

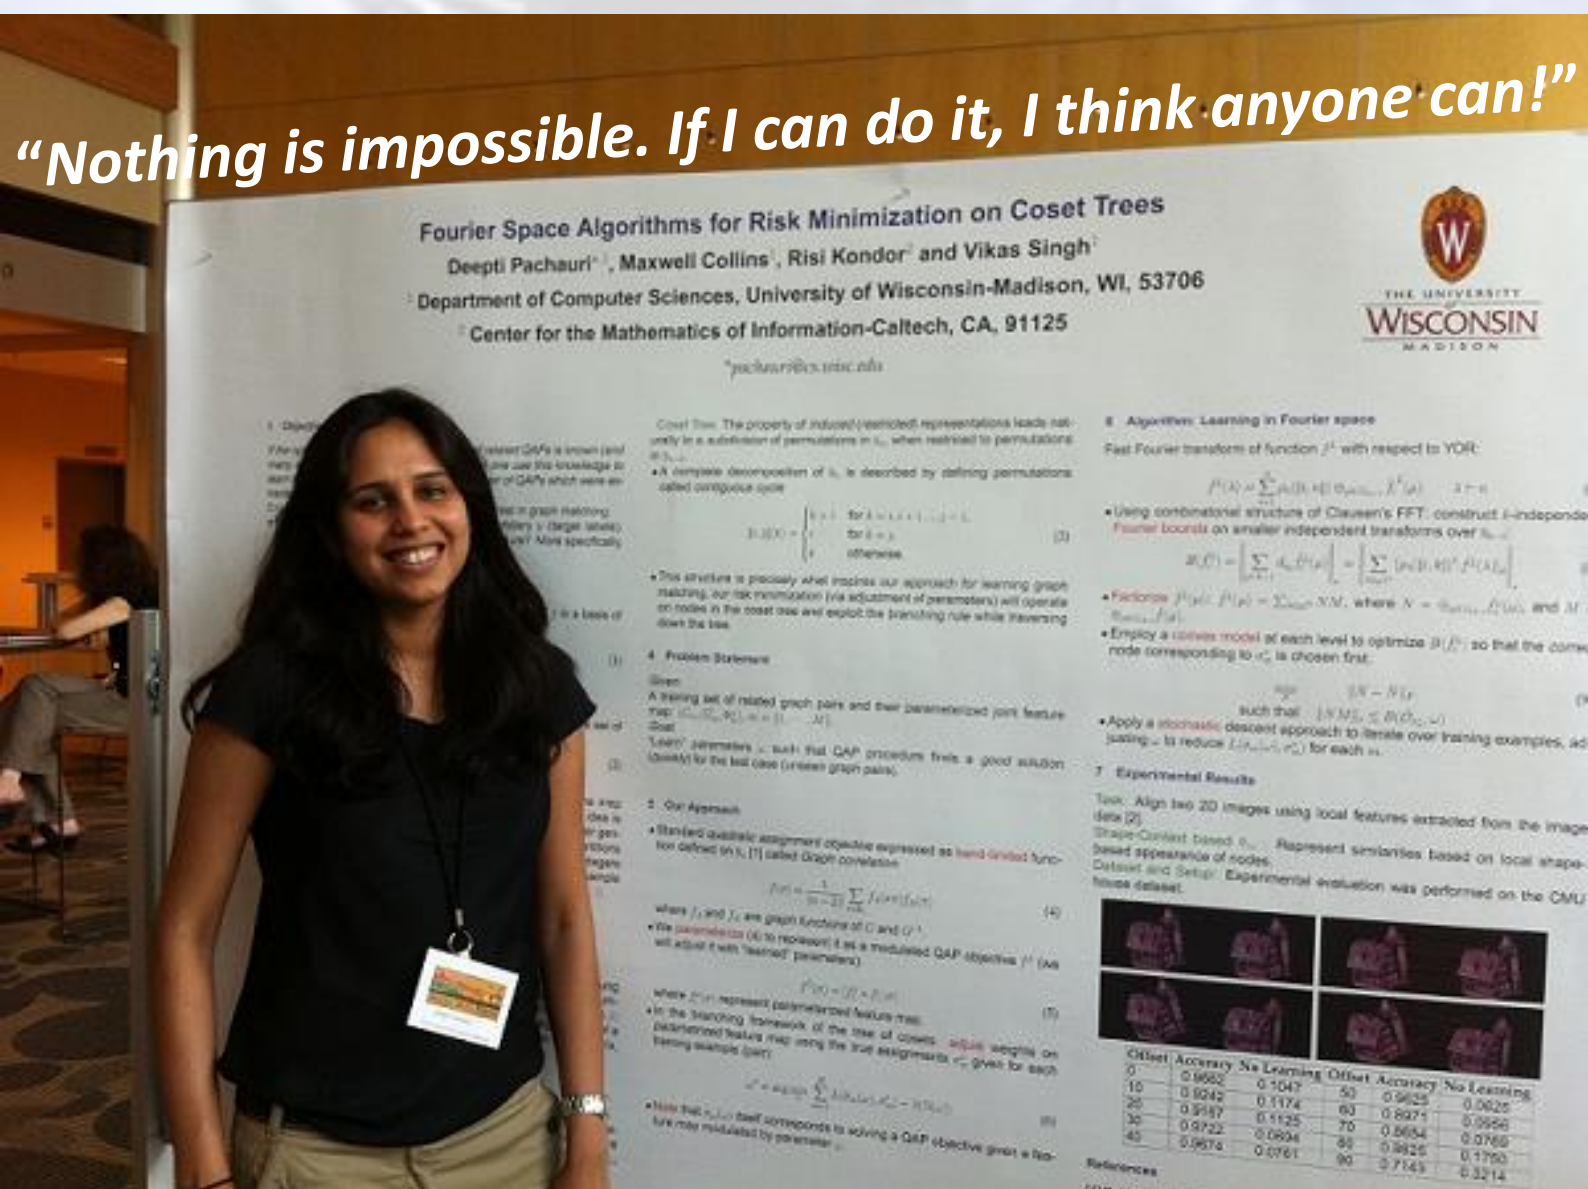
Science was an early choice in your life.

Yes, I can say that: in the environment around me, everyone was expected to become either an engineer or a doctor.

It's only when you are not good at your studies, that you study something else.

I wouldn't say that it was a very conscious decision. It was more because I was a good student and my parents didn't know any better. They

"Nothing is impossible. If I can do it, I think anyone can!"



At the machine learning summer school (MLSS) held in 2011 at Purdue University, Indiana

just relied on the ecosystem: everyone who is good at their studies should study science.

Were you among the best in your class?

Yes, I always was in the top 3 students... but very early on I realized that class competitions didn't matter as much as being good at higher entrance exams. In India (still today), if you're not getting good enough grades on entrance exams, then your class grades can help you very little. I didn't try hard to be at the top of this group. I always enjoyed science, math, physics, and chemistry, but I wouldn't study to be at the very top. It's not a competition in my mind.

What made you think differently?

I think it was my parents. My mom is a housemaker. She only studied until 12th grade and didn't get a chance to go and work. I have an elder sister, and my mom made a point that she wanted us to have enough education under the belt so that we'll be independent.

In another family, your destiny would have been different.

Yes - I totally think that would have been the case.

Do you feel lucky that your path was different from other girls your age?

Yes - I would say that my parents did a very good job in mentoring us as a good human being and giving us focus. They didn't know what needed to be done exactly, but my dad used to say that whatever you do, be excellent at it. My parents were good at mentoring and helping us decide what needed to be done. They tried to provide the best that they could in the financial situation that we had.

Later on you found yourself in university.

It's a long story! [laughs] I really wanted to study in an IIT, Indian Institutes of Technology. To study there, you have to pass a very competitive exam. Until my 12th, I was planning only for it. I also knew that the curriculum that was followed in my school was not meant to prepare me for this entrance exam. I had this agreement with my dad that if I got certain grades in my 12th, he would put some money together to send me to coaching for one year, so that I could write my IIT entrance exam.

Did you get in?

No - this is the crazy part, actually. At that time, my dad had some financial issues, and even though I got much better grades than we agreed upon, he couldn't send me. That was a step back. I got admissions to other engineering colleges, but I didn't get to the IIT and I didn't have the money to even get the coaching for it. At that time I told my dad that I didn't want to just do engineering from any random college: I wanted to go to IIT and I would try again for my Masters. That's when I took admission in the science college just to do my Bachelor's in physics. I think it was a brave decision because everyone around us was skeptical, but my parents supported me in that decision.

What is the motivation that brings a very young Asian girl to leave her family and find herself alone in a very foreign culture and environment? I ask you that because there are so many of you (from China, India, Pakistan, Turkey or else), all very talented and brave.

I can speak for myself because everyone has a different journey and motivation. I prepared for my Master's entrance exam for IIT and I was able to get through. Whilst there, I was debating if I wanted to join the workforce

“If you are really committed to something, you may have to do hard work, but it is possible!”



after my Master's or if I wanted to do something else. I decided to do a PhD and that's something no one in my family had done. I wanted to do it and I wanted to do it from the best possible place. That was my only reason to apply for a US university.

Did you do it for science?

Yes - I wanted to do a good PhD.

What kind of career did you dream of when you made that decision?

Mostly I was thinking to join academia afterwards. At the same time, I always wanted to solve real-world problems. Although my PhD is very theoretical and I'm very proud of it. I wanted to learn new things and, at the same time, I wanted to develop solutions which could practically be used in a real problem.

I hadn't decided yet what topic I wanted to do for my PhD because I didn't know any better. I made sure that I put myself in the place where I could learn from others and decide what to do. I came to the United States to join a PhD in physics. After spending a year over there, I was fascinated by computer science. I never had any practical training in software development before that. The first time I used a computer, it was MATLAB in 2006 or 2007. You need to learn computation to actually get some realistic results. After a couple of months of learning MATLAB by myself, I realized that I wanted to do something where I could leverage my physics and math background, but use the computation to do something intelligent. To be honest, I had no idea about artificial intelligence at that time.

But you chose it anyway?

Yes. The idea of taking this knowledge of math and physics and apply it in a different domain... that really fascinated me.

"I don't have regrets because I explored every possible option"

There is a pattern that comes up all of the time when you talk. You have an explanation for every decision that you make. Is this rationalization a part of your character, that you want everything to be explained?

Yes - I think I have the habit of analyzing my options and then deciding on things. I need to analyze because I don't like to regret. Once I have analyzed all of my options and made a decision, even if it doesn't work out, I don't have regrets because I explored every possible option at that time.

If you could change anything that you regret, what would you change?

To be honest, I wouldn't. I thought a lot about my journey. I don't even regret my struggles of not being able to attend very good coaching. I really believe to my core that whatever I am today is just because of the way things happened.



With her father S. P. Pachauri at graduation walk 2015: "My dad used to say that whatever you do, be excellent at it!"

Did you ever feel any racism against you because of color, origin or accent?

I wasn't aware of the term racism when I studied in the US. But as a teaching assistant, I instructed many undergrads and I had accent issues. They were having difficulties and made some comments. I didn't have the right accent back then. I still don't, but I think I'm still fine. I did understand a bit about racism. One of the things that I have learned growing up is that if something isn't working out, figure out what you can do to fix it. I joined the English as a second language courses that the university had at the international student services. I went to evening classes to work on my English. Maybe the students were really struggling, and that's why they were nagging me. The first two years as a teaching assistant really helped me because I could see how to interact with young students.

It was not something personal with you...

At that time, I didn't take it personal. I didn't know how it was. I could see their frustration. At the same time, it was a 101 course so that those who do not want to do physics have to do that 101 level course. They didn't want to be in that course because they were not going to study physics. I think it was a combination of a lot of things that probably came out.

"They just don't know how to tell a female that she is wrong"

Were there any obstacles in your path due to the fact that you were born female?

I didn't feel that actually. It's about being aware. For the first few years, I wasn't even aware of these issues

because growing up in my family, I didn't feel it. I was a very headstrong girl. No one ever told me what I can do or cannot do, although I grew up in a small town in India. During my PhD, when I started to attend some of the conferences such as Grace Hopper Conference, I realized that it may not have happened to me directly, but I learned from others that it happens frequently. Interestingly, in my PhD group with my advisor, I was the only female at that time. Once I started attending these conferences, I started talking about it with my male colleagues. It's so interesting that they, themselves, were aware that something was wrong when they had to discuss anything technical with females. They wanted to avoid it because they were not comfortable. Once you start talking about it openly, they admit to these issues. It's not that they are sexist. For example: they just don't know how to tell a female that she is wrong. Either they were doing a very bad job at it or they will just not say it because they don't know how the female will take it. Those conversations really helped me understand. The problem is not just that some people



*in Edinburgh, Scotland
for ICML 2012*

are really bad, and they intentionally do such things. A lot of it is very subconscious.

What is the best way to tell you that you are wrong?

Actually, I would say, just come up and say it. Then be ready for a reverse question where I may ask you to explain what is wrong. You have to explain to me how I can correct it or what you have to counter it. An open discussion is important.



Again, your answer is not emotional, but very rational: “Don’t treat me the way you think I should be treated. Just tell me the real reason for it.”

I can’t just accept “because you said so”.

Do you feel like a very rational girl?

I am very logical.

Do you want to pursue your future career in industry or academia, or both?

I think I really like both. The reason I took an industry job is because I never experienced industry first hand. I stayed in academia for so long that it was one thing that I never experienced... and I

don’t like not knowing something. I preferred to go into industry. 3M is a wonderful place for people in computer vision and machine learning because it is a unique place which is so diverse in terms of product range. It has like 50,000 products!

Can you tell us what you do there?

I am a computer vision scientist in their computational intelligence group which is a part of the corporate research laboratory. My responsibility is to generate intellectual property, and at the same time, work with the businesses to either make their products smarter or identify gaps so that new products can be launched

You have chosen a career which might be difficult for a female. How do you prepare yourself to overcome obstacles?

One of the really helpful things is having some senior female mentors. Since they may have experienced some of the issues themselves, they sometimes can guide you too. Having a supportive partner really helped in navigating through some of these issues because you don’t have to keep compromising your career for family. My partner is very respectful of my career.

Is he a scientist too?

Yes, he is also a scientist. There comes a mutual respect for each other. I really feel like if something is important for you... you have to figure out a way. I always follow this example in my life. I moved from physics to computer vision for a PhD. I had no idea how it would work out. One thing that experience has taught me is that if you are really committed to something, you may have to do hard work, but it is possible. Nothing is impossible. If I can do it, I think anyone can!

Communication in Project Management



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 learn about proper handling of **Communication in Project Management**.

You need great communication skills, when there are many stakeholders in the company!

As promised in our May article, we dedicate the lecture of the month of June to securing proper communication though the project. Before we dive into the specific case of computer vision projects, let's describe the general case of communication in project management

A typical project is initiated by top management, after the acquisition of the specifications from sales and/or marketing. A project manager is assigned to the project and a kick-off meeting is held, which gathers all participants. Project should be held following a top-down scheme, from the major activities to the minor tasks, and it is recommend not to go beyond a weekly planning.

All the activities are grouped with milestones, each having its own delivery and testing procedures. The project managers acts as a coordinator, from the center of the project, using his communication skills. For this purpose, he or she schedules regular meetings with participants. Among these meetings: a weekly report to top management of project highlights, risks and potential alternatives (with their costs, delivery date, schedule and the like).

A decision should be taken about the setting of meetings with the other participants: should meetings with project leaders (technical, mechanical, R&D and so on) include their teams or not? The larger the project, the most difficult it is to involve everybody in these meetings, which are designed to solve special events: deadlocks, engineering problems and the like.

Responsible for managing the milestones, their content and the deliveries

Back to the project cycle, the project manager is responsible for managing the milestones, their content and the deliveries: once the scheduled time for a milestone is reached, its deliverables are generated and tested. In this way, the project manager can assess the status of the progress in the project, according to the initial scheme defined at the beginning. If integration is involved, it can start at the initial phases and feedback can be supplied online.

Once the project gets to its final step, when the last milestone is reached with its specific deliverables and it satisfies the predefined test (i.e. alpha site and beta site), the project is closed and the maintenance schedule starts.

Computer Vision projects:

Let's now analyze this typical structure when the project at hand is in computer

vision. First of all, when getting the top management's blessing to start the project, it is important to make sure that all data and definitions are in place: in many cases, data is developing through the project and it must be taken into account. Maybe the engineering of the hardware will take longer than expected and it has to be scheduled properly in order to gather some intermediate data to work on.

Requirements can be tricky too: in many cases, the results given by the system do not work for all cases. The project manager must define which are the cases on which software must work and design ways to deal with false alarms and misdetections. When there is a good plan to handle those cases, risk management will be much easier during the project and product phases. Unmet management expectations could jeopardize the whole activity.

Another role which must be fulfilled by the project manager since the beginning is that of having a tight control



and thorough understanding of the development stages and results during the R&D process. In fact, R&D involves risks: there is no guarantee that research will bear the wished results; thus, manager must be aware of the situation and manage it in all its details, in order to be in control of all risks and of the ways to mitigate them.

Controlling the risks involved in R&D

In complex projects, involving advanced technologies, this is a real challenge: on one hand, you want to count on people and let them the initiative and the freedom to develop by themselves, giving them responsibility for the result; on the other, you need great communication skills when there are many stakeholders in the company.

First, he or she must ascertain that developers communicate very sharply and inclusively about the state of the project, so that the project leader is aware at all times of current status, results to date, algorithms being used. Gaps cannot be completely avoided and all the team must be aware of

them and be ready to be part of the effort to overcome any challenge.

We advise to keep a very tight control of the database, because in the first phase there are too many exceptions, due to the initial engineering status of the system. Nobody wants to be stuck with those images, but at the same time it is key to know how to handle that kind of images in the future.

During delivery, it is also important to keep track of the version and the data being worked on, by running unit tests on all cases within reach. It is also advisable to perform continuous integration, which means that at each stage, when running any part of the algorithm, it is tested permanently to make sure that results are improving and the system is not stuck on some bugs or neglected cases.

To summarize, management of computer vision projects needs to focus on controlling the risks involved in research and development, exposing in a transparent way all the challenges to all stakeholders and to all the creative engineers that can solve them.

*Did you read
Boston Vision yet?
Full reports on
9 great computer
vision companies
from the Boston
area and their
engineering work*

A publication by

Boston Vision

The vibrant Boston algorithm community
A special edition of **Computer Vision News**

Vision Systems:
Intelligent, automatic and robust computer vision technologies

Realeyes:
Measure people's reactions to something that they see

StopLift:
Video analytics and computer vision for POS revenue assurance

Quanterix:
Detect and measure molecules at very low concentrations

Teledyne DALSA:
Software & hardware for machine vision applications

MERL:
Electronics & Communications,

Engagement

Brow Raise	2%
Valence	100%
Smile	100%
Brow Furrow	0%
Lip depressor	0%

Affectiva:
Bringing emotional intelligence into the digital world

Computer Vision in Boston:
20 pages - click here to read!



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.

IEEE Intelligent Vehicles Symposium 2017

Redondo Beach CA, USA Jun 11-14 [Website and Registration](#)

SCIA 2017 Scandinavian Conference On Image Analysis

Tromsø, Norway Jun 12-14 [Website and Registration](#)

CARS 2017 Computer Assisted Radiology and Surgery

Barcelona, Spain **MEET US!** Jun 20-24 [Website and Registration](#)

RE•WORK Machine Intelligence Summit

Amsterdam, Netherlands Jun 28-29 [Website and Registration](#)

ICIAR - International Conf. on Image Analysis and Recognition

Montreal, Canada July 5-7 [Website and Registration](#)

EEVCV - 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 July 25-27 [Website and Registration](#)

ECVP - European Conference on Visual Perception

Berlin, Germany August 27-31 [Website and Registration](#)

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

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!

MICCAI 2017 - Preview with Simon Duchesne

Computer Vision News points its spotlights on specific events which deserve the attention of our readers. **MICCAI 2017** will be held on Sept 10-14 in **Quebec City, Canada**. Notifications for accepted authors have just been sent and all satellite events are online and accepting submissions. Eight months after [our interview at MICCAI 2016](#) and with only a few months to go before the MICCAI 2017, the heat is on and it is time to ask **Simon Duchesne, General Chair** of the upcoming 20th edition, to update us and our readers about the preparations to the event.

***“Our main keynote speakers are
Yoshua Bengio, Polina Golland and Curt Rice”***

Simon, what’s on at MICCAI 2017?

We think we had an amazing response in a 2-stage process: first, more than 1,100 intentions to submit, a couple of weeks before the actual submission. And then **800 actual submissions** which were reviewed by **600 reviewers** and the 53 area chairs, members of the **Program Committee**. Each papers received at least 3 reviews: there was a first ranking, then a ranking by every area chair and finally by the second area chair. The Program Committee got

together in **Quebec City**: the thirty attendees had to come on their own expense (travel and lodging), which shows how committed they are to the event. It has always been the case at MICCAI conferences. Half of the time was dedicated to reviewing papers and the other half to a mini-conference, in which 19 of them gave a presentation.

Out of that process, 255 papers were accepted, which is about 32%, in line with usual acceptance rate at MICCAI. We are very happy for those who were



“Quebec is a nice place!”



“All the conditions are there to make the event memorable...”

accepted and the others should not despair, since there are about 40 different challenges and workshops which are still open. Each sets its own deadline, most of them being around June 10-15: website MICCAI2017.org links to all of them. We encourage to submit papers there, also because the 2 days of workshop at MICCAI attract about the same number of participants as the 3 days of the conference.

How many people did already register?

We have already 215 registrations and those people have access to our partnering platform on Jujama, which I am sure you will like by the way.

Why?

Once you are in, you can see the profiles of everybody, set up meetings and plan your whole week ahead of time!

That's great! I could write all my articles even before MICCAI! *[we laugh]*
What are the next MICCAI deadlines?

Before the workshops, on the first week of June, is the time for the camera-ready papers and all authors are notified to send theirs. In July people will know if they were accepted to the workshop.

Any message to first-time attendees?

This edition is going to be cheaper than previous editions, so it's a good time to come, even if you do not present. I think we are going to have the most

workshops, we expect a lot of people and Quebec is a nice place! So we are looking forward to a great event, in the sense that all the conditions are there to make the event memorable...

Do you think there will be a [MICCAI Daily magazine](#) also this year?

Of course! *[we both laugh]* We are going to have a MICCAI Daily which will be distributed via the Jujama partnering platform to every single attendee! When the organizers get the magazine, they can send an Email to all attendees, who will be able to read it on the device of their choice. *Voilà*, that's the plan!

I would like to add that, with regards to entering the country, the visa situation is much better than in the United States. If you are used to a Schengen-style system, Canada and the United States have a separate visa policy, and the Canadian requirements are less demanding. Just [contact us](#) if you want a letter of invitation. There are tons of hotels and tons of Airbnb. The keynotes will have a strong emphasis on machine learning: it's not taking the field by storm, we are *in the storm*! Our main keynote speaker is **Yoshua Bengio**, one of the leaders of machine learning, and who is from Montreal.

We'll have a session dedicated to the past, present and future of MICCAI with [Polina Golland](#); and another with **Curt Rice** on gender equality and the presence of women in natural science, followed by the [Women Networking Lunch](#). Finally, the student events, 17 exhibitors and 2 sport events: football and Canadian curling!

“About 40 different challenges and workshops are still open”

Improve your vision with

Computer Vision News

The Magazine Of The Algorithm Community

The only magazine covering all the fields of
the computer vision and image processing industry

Subscribe

(click here, it's free)






REWORK

```
% invoke the matlab debugger
function STOP_HERE()
    [ST,~] = dbstack;
    file_name = ST(2).file; fline = ST(2).line;
    stop_str = ['dbstop in ' file_name ' at ' num2str(fline+1)];
    eval(stop_str)
```



Gauss Surgical



Mapillary

A publication by

