

SEPTEMBER 2019

Computer Vision News

The Magazine of the Algorithm Community

We Tried for You
Data augmentation with Keras

Project by RSIP Vision
AI in Pharma: Multiplex

Women in Computer Vision
Doris Antensteiner

Project Management
Managing the Expectations of All Stakeholders

Research
Semantic Priors from Intrinsic Image Decomposition

Computer News Events

Artificial Intelligence Spotlight News




Global Leader in Computer
Vision and Deep Learning

Challenge
Irontract with Anastasia Yendiki




04



14



24



10




16



30



12



22



33

03

Editorial*by Ralph Anzarouth*

04

Research Paper ReviewIntrinsic Image Decomposition (IID)
by Amnon Geifman

10

Project by RSIP VisionAI in Pharma: Multiplex
by Dorin Yael

12

Project Management TipManaging the Expectations of All
Stakeholders *by Ron Soferman*

14

Irontract 2019 Challenge

Guest: Anastasia Yendiki

16

We Tried for YouData augmentation with Keras
by Amnon Geifman

22

AI Spotlight News

From Elsewhere on the Web

24

Women in Computer Vision

Doris Antensteiner

30

Computer Vision Events

MARSS Summer School 2019

33

More Events

Upcoming Events September-November



Computer Vision News

Editor:

Ralph Anzarouth

Engineering Editor:

Amnon Geifman

Head Designer:

Rotem Sahar

Publisher:

RSIP Vision

[Contact us](#)

[Give us feedback](#)

[Free subscription](#)

[Read previous magazines](#)

Copyright: **RSIP Vision**

All rights reserved

Unauthorized reproduction
is strictly forbidden.

Follow us:



Dear reader,

Welcome to the September 2019 issue of **Computer Vision News**. You will find in it our regular sections, as well as special articles about **AI projects at RSIP Vision**. We'd like to tell our new readers (we add hundreds every month) and remind our regular ones that RSIP Vision is much more than the editor of this magazine. It is also an international company pioneering **R&D work in computer vision and deep learning**. If you want your company to have a larger AI team and bring about more powerful AI innovations, call us and we will help you achieve your goals.

Our VP of R&D **Moshe Safran** has just relocated to lead our office **in San Jose, in the heart of the Bay Area**. Ron Soferman, Founder and CEO of RSIP Vision, says: *"The new office will allow very close collaboration from the planning phase to the practical application by the surgeon in the operating room. Moshe's leadership and the added value brought by RSIP Vision's methodology for the **medical device industry** is a major benefit for its clients"*.

*"This is a great opportunity to increase our presence in the **Silicon Valley** and collaborate with healthcare companies that need additional AI and computer vision support," Moshe says. "It's always exciting to take part in something that will disrupt the industry"*.

Enjoy the reading and, as always, **take us along on your next deep learning project!**

Ralph Anzarouth

Editor, **Computer Vision News**

Marketing Manager, **RSIP Vision**



by Amnon Geifman

Every month, Computer Vision News reviews a research paper from our field. This month we have chosen **Semantic Priors for Intrinsic Image Decomposition**. We are indebted to the authors (Saurabh Saini and P. J. Narayanan) for allowing us to use their images.

Introduction

Intrinsic Image Decomposition (IID) is the task of decomposing image into two parts; reflectance and shading. **The reflectance component** represents how the material of the object reflects light independent of viewpoint and illumination. **The shading component** captures the illumination effects caused by direct and indirect lighting in the scene. IID has many applications in computer vision and computer graphics such as contrast enhancement, image re-texturing, colorization and more. While the vision community has seen a significant advance in single image IID, it remains a challenging, highly ill-posed problem.

Since the problem is ill-posed in nature, prior assumptions must be introduced. **Reflectance sparsity** is a common prior that assumes the color details are simple, i.e. it doesn't change much between nearby pixels. Shading smoothness is a prior which assumes the shading component of the image is a smooth surface.

Optimizing these two priors separately might lead to inferior results. For this reason, the proposed method suggests a formulation to optimize these two priors in a single integrated algorithm using an iterative optimization. In order to better determine the sparse and the smooth regions, the authors incorporate semantic scene information from a DNN-based semantic segmentation. We next review the method in detail.

Method

The method composed of three building blocks: semantic feature selection, shading formulation, and reflectance formulation.

Semantic Features:

The semantic features component produces two sets of feature vectors. The first set is feature vectors generated for each 30x30 (non-overlapping) patches. Each patch is passed through Region-based Convolutional Neural network (RCNN) where the last fully connected layer is defined to be the feature vector f_b for each of the patches in the image.

The second set of features are the selective search features. These features give interesting image regions that have a high probability to contain an object. For each detection proposal, the method generates dense binary region masks and score. Each pixel is assigned with a feature vector g_i of proposal masks weighted by proposal score.

Shading Formulation:

The authors suggest using the RCNN and selective search features to define pixel neighborhoods. At each iteration of the method, the shading and reflectance are estimated by minimizing the following energy function:

$$\Psi = \lambda_g S_g + \lambda_m S_m + \lambda_l S_l$$

Here S_g , S_m , S_l are the global, mid-level and local shading priors, each weighted differently by a different coefficient. The global shading is further divided into sparse neighborhood consistency term S_c and weight propagation term S_p such that: $S_g = S_c + S_p$. We next explain the terms in the energy function:

- 1) S_c is the sparse neighborhood consistency. It is computed by summing the squared differences between each pixel and a weighted average of its neighbors' pixel (where the neighbors are defined by f_b).

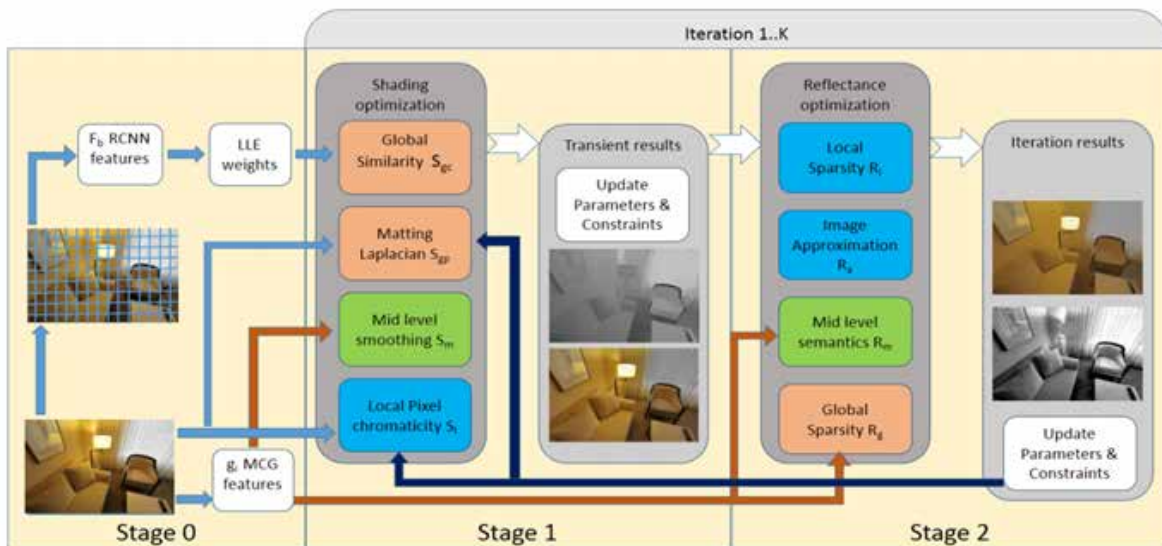
- 2) S_p is the weight propagation term. This term is computed by a weighted sum of the square of differences between each pixel shading value to its neighbor patches.
- 3) S_m is the mid-level context, generated in the same way as S_p but here the weights are defined by the selective search features g_i to penalize dissimilar pixel that belongs to the same object.
- 4) S_l is the local context, generated by a weighted sum of the log image intensity and the log shading values. This aims to enforce resemblances between the original image and the shading image.

Reflectance Formulation:

The reflectance formulation enforces color sparsity using L_1 terms. To this end, the authors use a sum of four weighted terms to define the reflectance energy:

$$\pi = \gamma_g R_g + \gamma_m R_m + \gamma_l R_l + \gamma_a R_a$$

Here again R_g , R_m , R_l , R_a stand for global mid-level local and image approximation term. Each of the first three terms is a weighted sum of the L_1 differences between the reflectance value of a pixel and its neighborhood (the difference between the three is the neighborhood and the weights). The last term R_a enforces continuity between the shading and reflectance stages by enforcing the reflectance estimate from the current stage to be similar to the intermediate reflectance solution for the previous shading iteration. These terms connect between the two stage of the optimization and make the method more robust and with fewer artifacts. Below is a figure that summarizes the method:



Results

One of the most problematic aspects of IID is the difficulty to quantify the results. It is very hard to gather direct ground truth data for this task. Previous works have collected ground truth via painting object, synthetic rendering and manual annotation, but each of these methods has a significant limitation. Here we show you some qualitative results, so that you can examine by yourself the quality of the suggested method.

Below is a figure that demonstrates the results during the optimization iterations. Each column is the reflectance image (lower) and the shading image (upper) for each iteration 1, 3, 5, 7. We can see, just like the method intended, the reflectance component of the decomposition becoming flattened/sparse while the shading becoming smoother. This demonstrates the convergence of the iterative method.



In the next figure we can see the final results of the IID task. There is a separation of shadows and illumination from the light source to shading component and color consistency in the reflectance component. These are quite good results. As we can see enforcing the suggested priors does seem to generate good quality decomposition.



Conclusion

Vision scientists have long been interested in understanding how humans separate illumination and reflectance. Many previous works have explored this task. It is quite intuitive to try to use semantic information to tackle this problem and it seems that this research manages to do it very well. The semantic information extracted using state of the art deep learning techniques enables the method to produce high-quality IID results. It seems a promising method which will hopefully lead to a breakthrough in this task.

Computer Vision Project Management



Computer Vision Project Management is a series of lectures and articles conducted by RSIP Vision's CEO Ron Soferman, many of which are published as a regular column on magazine Computer Vision News, in the project managements tips section.

Everything a project manager in computer vision should know... **at the click of a button**



**“Even the
biggest hammer
cannot replace
a screwdriver!”**

- How to solve all kind of challenges
- Team leadership and management
- How to implement deep learning
- Validation and test techniques
- What are the best practices



Did you miss an article? No worries, you can find them all in the **Project Management section of RSIP Vision's website!**



by Dorin Yael


Every month, Computer Vision News reviews a successful project. Our main purpose is to show how diverse image processing techniques contribute to solving technical challenges and real world constraints. This month we review challenges and solutions in a Pharma project by RSIP Vision: Multiplex.

Immunohistochemistry

is a powerful tool allowing the investigation of protein expression, activation and distribution within a tissue slice. Multiplex immunohistochemistry allows to repeat this process for multiple markers in the same individual slice, maximizing the amount of information acquired in it. This unique process maintains the spatial structure of the tissue and allows the identification of different cell types and subtypes, allowing in turn the investigation of

from sub-cell to tissue, and their interaction and co-localization.

The main tasks performed in the analysis of this rich information are detection, segmentation and classification, which can be performed either on single cells, in order to know the type of cell, its size, location, quantity and more; or on the whole tissue. After that, it is possible to integrate the data obtained through the staining and understand the relation between the single cells and the tissue.



Deep learning is a very powerful tool which transformed the fields of image processing and computer vision. Deep-learning architectures provide a data-driven solution, allowing a learning process which is dependent, specific and customized to the dataset. Convolutional neural networks (CNNs) are currently the state-of-the-art for automatic analysis of medical images, providing highly accurate, standardized and fast analysis integrating a large amount

of data. We have experienced much success both in the diagnostic and in the research area. Most of these studies were not possible to perform with these levels of accuracy and speed when operated manually or using traditional algorithm based automatic solutions. Today, RSIP Vision's state-of-the-art AI capabilities offer its clients all the advantages of deep learning technology on the full range of Multiplex analysis.



Courtesy of <https://themedtechconference.com>

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 Ron tells about another aspect of Project Management: Managing the Expectations of All Stakeholders. [Read the other Lectures](#)



Ron Soferman

Managing the Expectations of All Stakeholders

In a typical project, several stakeholders support the projects and contribute to designing, financing and promoting the project in all its steps.

While they all together decide on the project, the expectations of each are very different and it might be challenging to manage all these expectations along the project lifecycle. We will describe here five typical characters that you might find in your client company:

The first one is the manager who expects to get a final product, fully working and already certified within a very short lapse of time, like the time that you allocated for the Proof of Concept or for the first version. In this case, you might get much pressure at the first or second milestone in the project, to show a full-blown product.

The second character is the one that, after the development, expects the algorithm to work with 100% accuracy on all the cases that exist in the real world, with no need for a debugging process or any improvement or finetuning. In this situation, you might find that very difficult cases are expected to be solved along with the majority of the cases that your software succeeded to solve.

We may also find a risk-averting manager, who a priori wants to personally see the results before deciding to take any risk. With this kind of character, it might be very challenging to move on to the research phase, because of the unknown results expected along the way.

We might find also a head of engineering who wants the R&D team to commit to an exact a priori

at a time when the research is not completed yet. Even though such a requirement might give a direction to the requested results, it might also be expressed too early to be committed to and be too frightening for the R&D team.

The fifth character is the finance executive who expects to see the return on investment immediately after the scheduled development. This might be highly unrealistic: first, because of the many operational steps: compliance, sales force training, marketing the product and getting feedback; second, because only first adopters might be buying the product in the earlier phases.

From our experience, it is very efficient to have a weekly meeting with a large set of stakeholders, so that they understand the challenges, the pace of development and the deliverables of the research. They may prove to be precious allies in that they provide very good insights along the way for the cases that you have to cover in your algorithm, for the way to cope with difficult cases and also to decide the priorities along the development cycle.

[Read the other Lectures](#)



Photo by Feda Eid




Anastasia Yendiki

The Irontract Challenge 2019 is a collaborative effort between Harvard Medical School, the University of Rochester, and QMENTA Inc. We speak to one of the organisers, Anastasia Yendiki, who is Assistant Professor in Radiology at the Harvard Medical School.

Anastasia, what is the main expected outcome from this challenge?

The outcome will be an objective ranking of the accuracy of algorithms for tractography, a technique used to reconstruct the wiring of the brain from diffusion-weighted MRI scans. The ground truth of this wiring is not known in detail - we just know the main highways of the brain but we do not have the complete map of side streets. Although we know that diffusion MRI tractography is reasonably good at following these highways, we also know that it sometimes takes wrong turns. That is because MRI does not image brain pathways directly but instead measures how water molecules diffuse in and around these pathways. Inferring the shape of the

The image shows a small fraction (1%) of the paths produced by a typical whole-brain diffusion tractography analysis.




pathways from this diffusion is an error-prone process, therefore having access to ground truth is critical. Here we have ground truth from chemical tracer injections in the same brains that the MRI scans come from.

What have you learned to date from this challenge, that you can share with our readers?

Our preliminary studies on these data have led to some interesting findings. For example, counter to what one might have expected, we have found that there are limited gains in tractography accuracy from certain improvements in data quality that the field has been focused on recently. These include increasing the angular resolution (i.e., increasing the number of directions in which water diffusion is measured) or acquiring data at multiple levels of diffusion contrast (i.e., varying a factor called the b-value). This has been the case for the small set of tractography methods that we tested on these data, and it will be interesting to see if the conclusion holds up when we open it to all the other methods that challenge participants will deploy.

Axons, seen under darkfield microscopy, in a brain that has received a tracer injection. These can serve as ground truth for tractography.



Do you have any particular advice for the participating teams?

We are excited for the teams to use their latest and greatest algorithms. That said, it is sometimes the case that the older, "straw-man" methods don't do as poorly as one might expect. So it would be interesting to see a mix of both higher- and lower-tech approaches put to the test!

Dear reader,
If you like Computer Vision News (and also if you don't like it) we would love to hear from you:

F E E D B A C K

It will take you only 2 minutes. Please tell us and we will do our best to improve.

Thank you!



by Amnon Geifman

"Data augmentation is the common practice when training a neural network. It enlarges the data set and enables to boost the generalization capabilities of the network. Keras data generator enables to augment data on the fly without taking care of all the technical details of the transformations."

Data augmentation refers to the process of generating new data from existing one. It is well known that the generalization gap, i.e the difference in performance between the test and the train data, depends on the number of examples in the training set. As the size of the training set increases, the performance is similar on both seen and unseen data. This month we will review a few of the most common techniques to enlarge data sets using several transformations. We will use **Keras ImageDataGenerator** to create several popular augmentations and explain the rationality behind them.

The goal in data augmentation is to generate new examples without changing the class label of the data. A good example of that is shifting. The human eye can identify an object no matter where the object is located on the image. Mathematically two images with the same object in a different location are different; however, we want our network to recognize the object regardless of its location, just like humans do. **The rationale behind data augmentation** is to show the network enough examples of a shifted object, and in this way to develop invariance of the network for shifting. Of course, shifting is just one example and we will see other transformations as well.

"Training the network on the full data set may produce even better results"

Next, we will see how exactly we can create such shifted images without changing the class label and without hurting the naturalness of the image. Remember, we want to take an image and shift the main object to a different direction. In Keras we perform such transformation via the `keras.preprocessing.image.ImageDataGenerator` class.

This class generates random augmentations on the fly with a specify parameters for each augmentation. For example, to generate a shifted image we simply use:

```
1  from keras.preprocessing.image import ImageDataGenerator,
2  array_to_img, img_to_array, load_img
3  import matplotlib.pyplot as plt
4  import numpy as np
5
6  img = load_img('shoes.jpg')
7  data_generator =
8  ImageDataGenerator(height_shift_range=0.3,
9  width_shift_range=0.3)
10 X = img_to_array(img)
11 i = 0
12 FinalImage = X
13 X = X.reshape((1,) + X.shape)
14 for batch in data_generator.flow(X):
15     i += 1
16     FinalImage=concat_images(FinalImage, batch[0])
17     if i % 5 == 0:
18         break
19 plt.axis('off')
plt.imshow(array_to_img(FinalImage))
plt.show()
```

This code loads an image and iterate through the augmentations object. The resulted new images are:



Another useful and popular augmentation is a **brightness shift**. Changing the pixel intensity value is quite logic. It corresponds to two images of the same object that were taken in different illumination conditions. These two images, of course, should be classified for the same class. This is done in practice with the following change to the code:


```
1 data_generator = ImageDataGenerator(brightness_range=[0.5, 1.5])
```



There are several **additional augmentations** we can perform, such as rotations of the image, zooming, crops, flips rescaling and whitening. Using the following command, we can create multiple augmentations on the same generator call:

```
1 data_generator = ImageDataGenerator(rotation_range=90,  
2                                     width_shift_range=0.1,  
3                                     height_shift_range=0.1,  
4                                     shear_range=0.2,  
5                                     zoom_range=0.2,  
6                                     horizontal_flip=True,  
7                                     fill_mode='nearest')
```




Now, after we understood how the data generator works, we will use it to demonstrate how effective it is. The experiment that we perform is as follows: we take a VGG16 pretrained model (to avoid the training of a full model) and we use it to classify images from the dogs vs cats dataset. We first train our network without data augmentations. Then we augment data and train it again. This should boost the performance of our network. We added the code of the network with the full augmentations at the end of this article.

Note that we used the VGG16 and froze all its layers. We added a fully connected layer at the end, in order to output a binary classification. The variables `train_it` and `val_it` are iterator objects that iterate through our training and validation data. At the end, we plot the history of the training procedure.

"The goal in data augmentation is to generate new examples without changing the class label of the data"


Results

We trained our network for 50 epochs on a small portion of the dog versus cats dataset (1000 images for each class). The first training has been done on the data without the augmentations. This is done by setting the variable `datagen = ImageDataGenerator()`. We used 100 examples from each class as a test/validation set. The history of the training can be seen below:



It can be seen that we have made a nice overfitting. Our training accuracy/loss increases/decreases while the same value on the test data doesn't change. This is simply because we don't have enough training examples to generalize well. Although we trained a small number of parameters, we still need to use more than 1000 examples to train the network.

To solve the poor generalization property of the previous training, we use data augmentation. To this end, we generate data as we demonstrated at the beginning of this article. We will run a few more steps at each epoch to enable the data generator to augment enough examples. The history of the training can be seen below:



Here we can see that our training and testing error values are very close. This shows that we were able to generate some generalization by augmenting more data. Training the network on the full data set may produce even better results.

Conclusion

Data augmentation is the common practice when training a neural network. It enlarges the data set and enables to boost the generalization capabilities of the network. Keras data generator enables to augment data on the fly without taking care of all the technical details of the transformations. We implemented a network to show how it helps to generalize better when training a network. The network is quite simple and can run also on the CPU so you can try it at home.


```
1 from keras.models import Model
2 from keras.layers import Conv2D, MaxPooling2D
3 from keras.layers import Activation, Dropout, Flatten, Dense
4 import numpy as np
5 import matplotlib.pyplot as plt
6 from keras.preprocessing.image import ImageDataGenerator, array_to_img, img_to_array, load_img
7 import matplotlib.pyplot as plt
8 import tensorflow as tf
9 from keras.applications.vgg16 import VGG16
10
11 vgg = VGG16(include_top=False, weights='imagenet', input_shape= (64, 64, 3))
12
13 for layer in vgg.layers:
14     layer.trainable = False
15
16 output = vgg.layers[-1].output
17 output = Flatten()(output)
18 output = Dense(256, activation='relu')(output)
19 output = Dropout(0.5)(output)
20 output = Dense(1,activation='sigmoid')(output)
21 vgg_model = Model(vgg.input, output)
22 vgg_model.output
23 model = vgg_model
24
25 Batch_size = 16
26
27 datagen = ImageDataGenerator(rescale=1. / 255,
28                             rotation_range=30,
29                             zoom_range=0.15,
30                             width_shift_range=0.2,
31                             height_shift_range=0.2,
32                             shear_range=0.15,
33                             horizontal_flip=True,
34                             fill_mode="nearest")
35
36 # load and iterate training dataset
37 train_it = datagen.flow_from_directory('Data/Train/',
38                                       batch_size=Batch_size,
39                                       target_size=(64,64),
40                                       class_mode='binary')
41 # load and iterate validation dataset
42 val_it = datagen.flow_from_directory('Data/Validation/',
43                                     batch_size=Batch_size,
44                                     target_size=(64,64),
45                                     class_mode='binary')
46
47 model.compile(loss='binary_crossentropy',
48              optimizer='rmsprop',
49              metrics=['accuracy'])
50
51 history = model.fit_generator(
52     train_it,
53     steps_per_epoch=100,
54     epochs=50,
55     validation_data=val_it,
56     validation_steps=100 )
57
58 plt.plot(history.history['acc'])
59 plt.plot(history.history['val_acc'])
60 plt.title('model accuracy')
61 plt.ylabel('accuracy')
62 plt.xlabel('epoch')
63 plt.legend(['train', 'test'], loc='upper left')
64 plt.show()
65
66
67 plt.plot(history.history['loss'])
68 plt.plot(history.history['val_loss'])
69 plt.title('model loss')
70 plt.ylabel('loss')
71 plt.xlabel('epoch')
72 plt.legend(['train', 'test'], loc='upper left')
73 plt.show()
```


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

Five Newest Developments for Imaging in Drug Research:

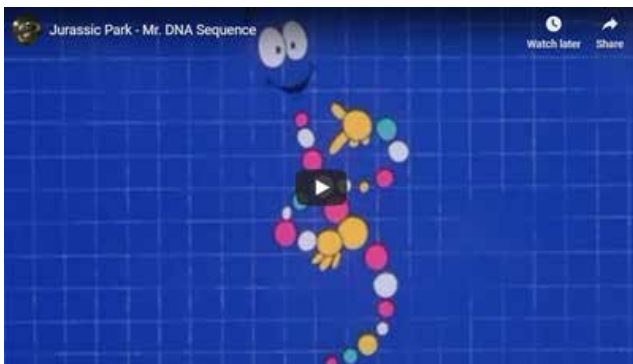
Let's start with a great pharma article: image processing techniques are used by scientists to identify therapeutic targets and improve drug efficacy. This article lists five of the latest developments for imaging methods.

[Read the article.](#)

[We list more Artificial Intelligence solution for pharma on our website.](#)



Exploring DNA with Deep Learning:



In this very nice post, FloydHub Blog explores DNA with Deep Learning. If you don't know about **population genetics** and **DNA sequencing technology**, that's where you will learn about it, with codes! They will also tell you if Jurassic Park is likely to exist during our lifetime, which is also an important question to ask.


[Read it here.](#)

Generative Adversarial Networks in 50 Lines of PyTorch Code:

"GAN is the most interesting idea in the last 10 years in machine learning!" says [Yann LeCun](#). If you always wanted to learn about GANs and you never did, this is the article to read. Using PyTorch, we can actually create a very simple GAN in under 50 lines of code. There are really only 5 components to think about. Read them in the article. [Read it here.](#)


```
class Discriminator(nn.Module):
    def __init__(self, input_size, hidden_size, output_size, f):
        super(Discriminator, self).__init__()
        self.map1 = nn.Linear(input_size, hidden_size)
        self.map2 = nn.Linear(hidden_size, hidden_size)
        self.map3 = nn.Linear(hidden_size, output_size)
        self.f = f

    def forward(self, x):
        x = self.f(self.map1(x))
        x = self.f(self.map2(x))
        return self.f(self.map3(x))
```




Artificial Intelligence in the Fight against River Blindness:

More than 21 million people in Africa are infected with the nematode *Onchocerca volvulus*, the cause of river blindness. Millions of them eventually go blind. Parasitologists at the University Hospital Bonn are looking for new, more effective weapons against the insidious parasite, using an Artificial Intelligence system that will automate the evaluation of histological



sections together with already existing deep learning models for object recognition. The objective is to reduce the time required for the otherwise very lengthy evaluation. [Read.](#)

AI Provides a Detailed Road Map for Interventional Lung Procedures:



Not only ZDNet talks about our work (see Computer Vision News of August)! Here is Medical Device and Diagnostic Industry (MD+DI) telling manufacturers of medical devices and all MedTech industry professionals what they think of [our breakthroughs in the field of pulmonary imaging](#). Our technology uses AI and deep learning techniques with the goal of pinpointing the exact location of suspicious lesions in the lungs, ensuring that surgeons can take biopsies from the correct site. [You will read about it here.](#)

Tesla's Advantage in Behaviour Prediction for Autonomous Driving

OK. 'nough said about pharma and medicine. Let's see what's up in the automotive field. This is about the importance of predicting the future behavior of road users (i.e., cars, pedestrians and cyclists) and the possible competitive advantage of Tesla, who has already over 500,000 vehicles on the road that can collect data to train neural networks to predict road user behavior. Don't miss what **Andrej Karpathy**, Tesla's Director of AI, explained about this process at Tesla Autonomy Day. [Read about it here.](#) Click on the image to follow Andrej's speech.





Doris Antensteiner

"I was hoping that I would be applying computer vision technology in the industry."

Doris Antensteiner is a scientist at the Austrian Institute of Technology (AIT), the country's largest Research and Technology Organization.

Doris, what's your work at AIT?

I work on computer vision problems for industrial applications. We focus on solving real-world problems, especially in production lines. We are working on different industrial applications, for example, solutions for inline industrial production lines. For this, we often use line scan sensors to get 3D reconstructions and analyze the surface properties, like the reflectance distribution functions from the surface... we are also working on industrial applications which are not inline, for which we are using full area scan sensors. We observe both objects in motion and objects which are not in motion.

"I really like CVPR!"

Is this what you expected to do when you started your studies as a scientist?

Pretty much, yes. I was hoping that I would be applying computer vision technology in the industry. And that's what I am doing!

"You have a lot of opportunities to meet people from your research field"

Tell me something about yourself. You are originally from Austria?

Yes, I am originally from Austria. I spent most of my life in Austria, but I did some time abroad in Denmark in Copenhagen.

Tell me about the higher scientific studies in Austria. What should we know about them?

Most of the universities in Austria are public universities so it is not usual to go to a private university. A part of the universities are also research institutes, and these research institutes are usually more directed towards applications or applied sciences. The universities are more directed towards basic and theoretical sciences.

Do you also teach?

No, I am not teaching. While I was at the university, I was a tutor for a while, but since I did my PhD, I am no longer teaching.

Do you like teaching and do you miss it now?

I think it is fun to teach, especially when you can do it and teach something you really like. You can motivate others to get into it and like it too. I guess you could miss it, but it is fine for me not to teach.

What do you like the most about your career?

When I started, it seemed that a lot of computer vision topics were still not quite solved. About 10 years ago, I started taking classes in computer vision, and with a lot of these topics, there were some solutions for it, but many problems were not in a solved state. Now, in the last 10 years, it has massively evolved to it being more and more applied, and more topics are in the state of being solved. I enjoy that a lot, to be in this area where research is finalized, creative, and pushed towards industry, and it works. I enjoy it!

Most of my readers know Vienna but certainly not as well as you do. Tell us something about Vienna or what you like about it.

Vienna is a great city to live in. It has great transportation systems, infrastructure, and good



communities. It is a good place to live.

Is there anything that you would like to bring from Denmark and put in Vienna?

There is an interesting way in Denmark to engage with people and celebrate certain events. I had the impression that people are very happy about that. That was something pretty unique. I have seen in the way they do those things in Denmark which could improve in Vienna.

Maybe our Austrian readers will read this, and they will adopt it because of you! Do you plan to work directly for an industrial company or are you still good doing research?

Right after my master's thesis, I went to an R&D department. I was working there doing computer

vision. So I was first working in industry before I joined a research center and started my PhD.

What did you like there?

It was great to work in a team, to work on actual computer vision applications, to see it running every day in the real world, and see your algorithms working out there.

Tell us more about the research institute where you work.

The Austrian Institute of Technology is the biggest research institute in Austria. It covers a lot of different topics of research areas. There are 1,300 employees.

"You can motivate others to get into it and like it too"

Which conferences is your favorite?

I really like CVPR.

What do you like at CVPR?

I like that a lot of people from the community come together, and you have a lot of opportunities to meet people from your research field and also from your research niche.

Many people attend, but only a few from the little niche you are working in. It is interesting to exchange with them.

CVPR offers a lot of opportunities, with smaller sessions to engage with people and also poster sessions where it's quite easy to talk to people. I also like when there are keynote speakers. I like the keynote speeches.

Is there anything you would like to change there so we can give a suggestion for next year?

I really like it as it is!






Women in Computer Vision

by Ralph Anzarouth



Women in Computer Vision (also called Women in Science) is a series of interviews conducted by Ralph Anzarouth. New interviews are regularly published on all RSIP Vision's publications: Computer Vision News and the Daily magazines (CVPR Daily, MICCAI Daily and many more).

Find now on the project page the direct links to almost 100 interviews... **at the click of a button** 

Competence

Talent

"The only way to succeed is to really start believing in yourself!" Michela Paganini

Leadership

Community

"Most of all, you have to believe that you can do it!" Laura Leal-Taixe

Confidence

Mentoring

"It may look like a long list of names, but behind each name there is a fascinating world in which we were let in." Ralph Anzarouth

Did you miss an interview? No worries, you can find them all on the women scientist section of RSIP Vision's website

Feel at ICCV as if you were



ICCV 2019
Seoul, Korea

October 29-30-31 and November 1



**Get the ICCV
Daily News in
your mailbox!**



**Click here and subscribe
for free to ICCV Daily!**

The Medical Augmented Reality Summer School (**MARSS 2019**) was just held in Zurich from 5 to 16 August.

We are happy to publish a few images from the event, courtesy of co-organizer Matthias Seibold.





**Feel at
MICCAI
as if you
were at
MICCAI**

October 14-15-16

**Get the
MICCAI Daily
News in your
mailbox!**

**Click
here
and
subscribe
for
free
to
MICCAI
Daily!**



COMPUTER VISION

EVENTS

BMVC - British Machine Vision Conference

Cardiff, UKd Sept 9-12

Int. Conf. on Machine Learning, Optimization...

Siena, Italyd Sept 10-13

Intelligent Health 2019

Basel, Switzerland Sept 11-12

AI for Business Summit 2019

Sydney, Australia Sept 17-19

The MedTech Conference 2019

Boston, MA **Meet Us** Sept 23-25

The AI Summit 2019

S.Francisco, CA Sept 25-26

Did we forget an important event?

Tell us: editor@ComputerVision.News

ISVC International Symposium on Visual Computing

Lake Tahoe, NV Oct 7-9

MICCAI - Medical Image Comp. and Comp...

Shenzhen, China **Meet Us** Oct 13-17

ViplIMAGE Computational Vision and Medical...

Porto, Portugal Oct 16-18

AI World

Boston, MA Oct 23-25

ICCV 2019 - International Conference on Computer V.

Seoul, Korea **Meet Us** Oct 27 - Nov 2

AI & Big Data Expo - Silicon Valley

SantaClara, CA Nov 13-14

ACPR - Asian Conference on Pattern Recognition

Auckland, New Zealand Nov 26-29

FREE SUBSCRIPTION




(click here, its free)

Dear reader,

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

You will fill the Subscription Form in less than **1 minute**. Join thousands of AI professionals and receive all issues of Computer Vision News as soon as we publish them. You can also read Computer Vision News in **PDF format** (though this online view is way better) and visit **our archive** to find new and old issues as well.


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



*IMPROVE
YOUR VISION
WITH*
**Computer Vision
News**

SUBSCRIBE
CLICK HERE, IT'S FREE

*The magazine of
the algorithm
community*



A PUBLICATION BY



Global Leader in Computer
Vision and Deep Learning