A publication by RSIP VISION

Computer Vision News

The magazine of the algorithm community

October 2016

Guest
Prof. Jürgen Schmidhuber

Image: FAZ/Bieber

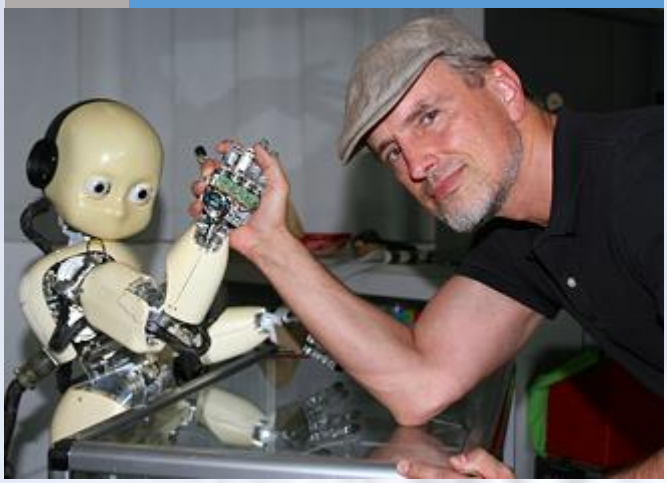

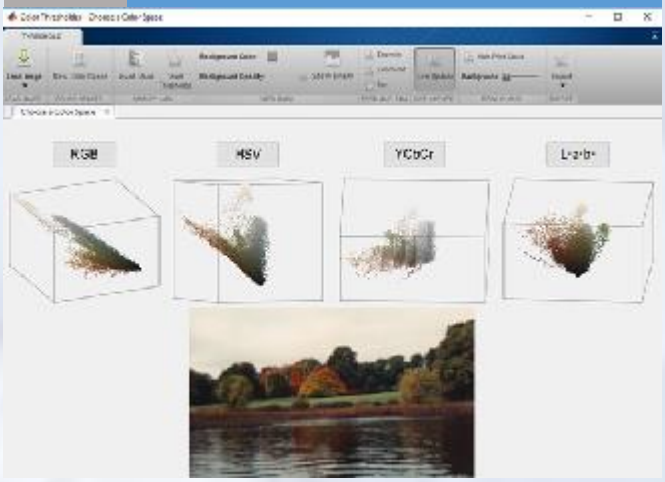

```
// Query Image
imRead{
    id : queryImageLoader
    visible : false
    file : queryImage
}
FastFeatureDetector{
    visible : false
    id : queryFeatureDetect
    input : queryImageLoader.output
}
BriefDescriptorExtractor{
    visible : false
    id : queryFeatureExtract
    keypoints : queryFeatureDetect.keypoints
}
// Matching
FlannBasedMatcher{
    id : descriptorMatcher
    queryDescriptors : queryFeatureExtract.descriptors
    knn : 2
    params : {'indexParams' : 'lsh'}
}
DescriptorMatchFilter{
    id: descriptorMatchFilter
    matchesIto2: descriptorMatcher.matches
    minDistanceCoeff: 2.5
    nndrRatio: 0.6
}
DrawMatches{
    keypoints1 : queryFeatureDetect.keypoints
    keypoints2 : trainFeatureDetect.keypoints
    matchesIto2 : descriptorMatchFilter.matchesItoOut
    matchIndex : 0
```


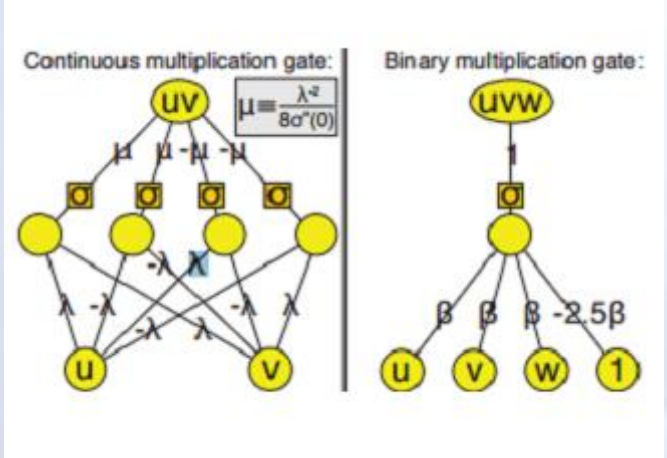

Computer Vision News
Editor: **Ralph Anzarouth**
Publisher: **RSIP Vision**

Click to contact us

Click to give us feedback

Click for free subscription

Dear reader,

The month of October will be full of events. We will no doubt meet many of our readers at ECCV 2016 in Amsterdam, on October 11-14. The European Conference on Computer Vision has a special meaning for us this year, since we are partnering with ECCV to publish the daily magazine of the conference: **ECCV Daily.** It's the first time that the attendees of ECCV will enjoy a community tool of this kind: its mission will be to enhance communication between participants, promote their work and give them additional exposure within academia and the industry at large.

This daily magazine is not a new idea: we did it earlier this year at CVPR 2016, where we published the now famous CVPR Daily. We shall do it again later this month at MICCAI 2016, which will be held in Athens on Oct. 17-21: the International Conference on Medical Image Computing and Computer Assisted Intervention partners with us for the publication of a brand new **MICCAI Daily**.

Of course, we are very proud of being called on for these delightful projects, and we are grateful to all the conference chairs (and to the readers too).

Let's now turn our attention to this Computer Vision News of October. All the regular sections feature interesting and stimulating stories, starting at page 4 with the exclusive interview of our guest, **Professor Jürgen Schmidhuber**, one of the **pioneers of Deep Learning Neural Networks** since 1991. You will love this magazine: please keep sharing it with friends and colleagues.

Enjoy the reading!

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

Professor Jürgen Schmidhuber

Prof. Jürgen Schmidhuber is Scientific Director of the Swiss AI Lab IDSIA and Professor of Artificial Intelligence at the University of Lugano (USI). He has pioneered Deep Learning Neural Networks since 1991 and his research group is a serial winner of international contests. His **Recurrent Neural Networks (RNNs)** are now available to billions of users through Google, Apple and others. For example, since 2015, Google's speech recognition on smartphones is based on his team's "Long Short-Term Memory (LSTM)" RNNs. Google DeepMind is also heavily influenced by his former students. **Computer Vision News** was very lucky to have a fascinating discussion with Prof. Schmidhuber: his views on the future of the human kind cannot leave anybody indifferent…

Computer Vision News: *The field of Deep Learning has seen during the last 25 years (as your impressive research page testifies) a long list of successful technology breakthroughs. Isn't there anything which disappointed you, like a finding or an innovation which you expected but never arrived?*

Professor Schmidhuber: That's an interesting question. I often thought that some breakthrough would come earlier than it actually did and I was sometimes surprised by a development arriving faster than I thought. But I cannot say that anything I thought was going to happen, suddenly does not have a chance to happen anymore. It is a question of **predicting the speed of technological evolution**. Details of my predictions have been wrong, but the predictions were generally in line with what has happened during recent decades.

"I am living at a time when the universe wants to make its next step towards greater complexity"

Image: Tomás Donoso on Vimeo

CVN: *What area of study would you suggest to a young engineer entering now the field?*

Prof. Schmidhuber: Machine Learning and Artificial Neural Networks.

CVN: *Would it be the same answer if that engineer wanted to turn those studies into a business?*

Prof. Schmidhuber: Sure. That's the technology used by many big and small companies to make people happier and healthier and even more addicted to their smartphones.

CVN: *You have used techniques of computer vision to solve problems in so many fields: scientific, financial and even in the artistic domain. Where would you see its main contribution to improving our society at large?*

Prof. Schmidhuber: First, let me point out that computer vision is just one special area where artificial Neural

Networks are applicable. In particular, recurrent neural networks, such as LSTM, are now widely applied not only to vision but also to speech recognition, machine translation, automatic email answering, natural language processing, etc.

Generally speaking, such networks learn by changing the strengths of connections between their artificial neurons to solve practical tasks.

"For the first time intelligent machines may be better general problem solvers than humans"

Such a task could be pretty much anything that you can imagine. One of the most important applications is in health care, which can be greatly improved by artificial Neural Networks: they can imitate doctors, for example, by learning to recognize cancer in images of breast tissue. In 2012, our deep neural nets were the first to win a contest of that kind, having learned from examples how to distinguish between harmless conditions and dangerous pathologies.

In the future it should be possible to provide reasonable diagnostics to many people who today cannot afford it, e.g. in developing countries where decent health care is not available. For example, take an image of a dark brown spot on your skin with a smartphone camera and send it to an AI medical doctor, who may decide whether that needs to be examined further by a human expert.

CVN: *In conclusion, what would you like to tell our readers?*

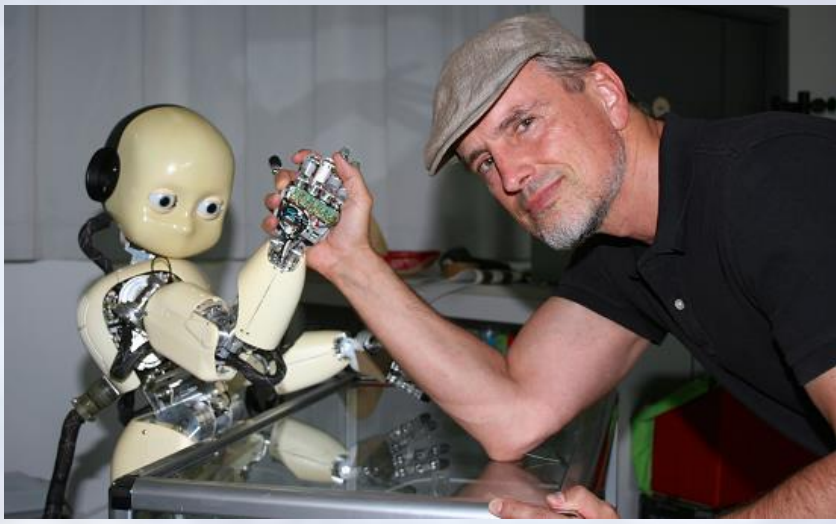

Image: FAZ/Bieber

Prof. Schmidhuber: We are living in extremely exciting times, when great things are happening: in not so many decades, for the first time **intelligent machines may be better general problem solvers than humans**.

Soon, the most important decision-makers will not be humans anymore: every aspect of civilization will be transformed by that. It will not stop there, since these smart AI beings will realize that most of the resources of the solar system are not within the thin film of biosphere but out there in space. They will colonize the solar system and the entire Milky Way and beyond, in a way which is completely impossible for humans.

CVN: *That sounds fascinating. Do you look forward to that civilization?*

Prof. Schmidhuber: I do. Since the 1970s it has seemed clear to me that I am living at a time when the universe wants to make its **next step towards greater complexity**, a step comparable to the invention of life itself over 3 billion years ago. This is more than just another industrial revolution. **This is something new that transcends humankind and even biology**. It is a privilege to witness its beginnings, and contribute something to it.

New Matlab Version - R2016b

Every month, Computer Vision News shows you a nice trick that will make your work easier. This time, our engineers will tell you about the new features in the **new version of Matlab 2016 (R2016b)**. We will briefly highlight the new features from the computer vision toolbox and the image processing toolbox. For the complete list of new features we invite the reader to check the MathWorks website.

(1) Image Browser App.

The first new feature we'll cover is the new and improved image browser app, which now supports quick viewing of multiple images and using basic computer vision operations on the images. The image browser app lets you quickly upload a set of images and perform some basic image processing on those images.

The Matlab image browser looks as follows:

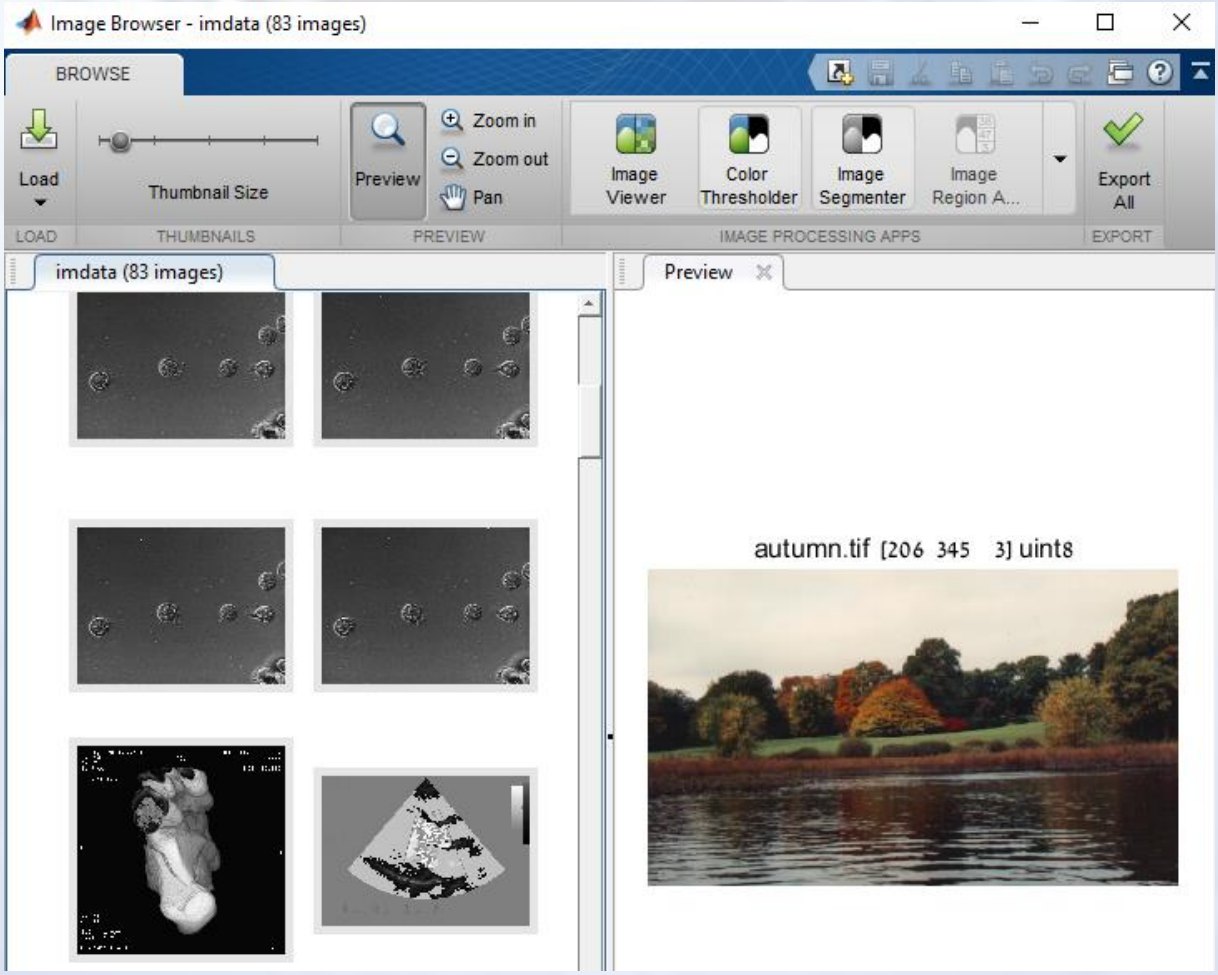

Now, for instance, if you select a color image with the image browser app - by double clicking on the image - you can use the new features of the color threshold app on it, which we shall demonstrate next.

(2) Color Thresholder app: View color data as point cloud for segmentation

The color threshold app has also been upgraded, with the inclusion of a new option which lets you view image color data as a point cloud in four different color spaces: RGB, HSV, YCbCr, L*a*b*. You can use these point clouds to evaluate which color space provides the clearest separation of color that makes it easy to select particular colors for segmentation. You use the point clouds in conjunction with the existing controls offered for each color space. For example:

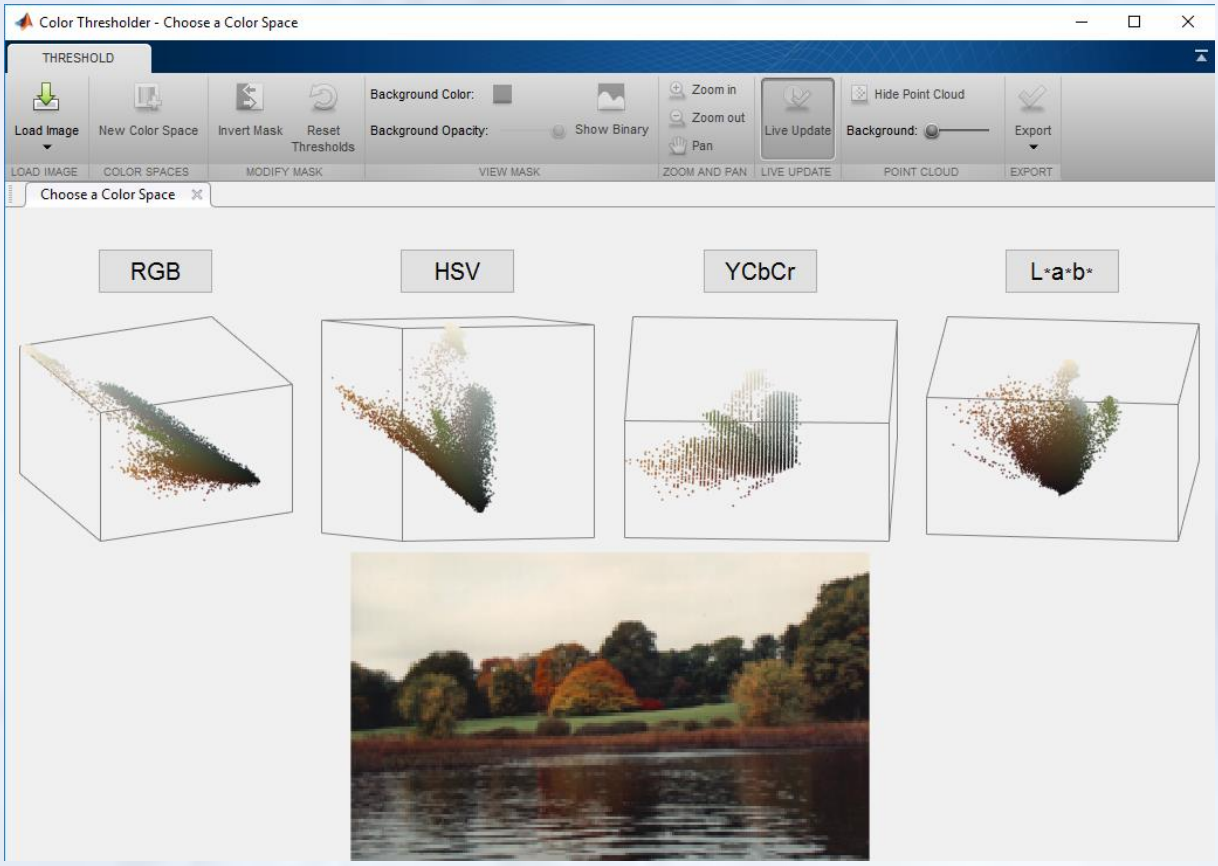

FEEDBACK

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!

(3) 3-D superpixels:

Use simple linear iterative clustering (SLIC) with volumetric image: remember the SLIC superpixel algorithm we have reviewed in our July issue? Now Matlab comes with built-in support for 3D superpixels, enabling you to preprocess 3D images (like CT or MRI scans) with the superpixel oversegmentation, as can be seen in the example on the right. Though we only include a 2D slice here, all the computation is done in 3D, making the superpixels more precise. Indeed, you can see the main organs in the CT are recognized. This result could be the used as input of a more advanced segmentation algorithm and improve the latter's precision due to the superpixel prep work.

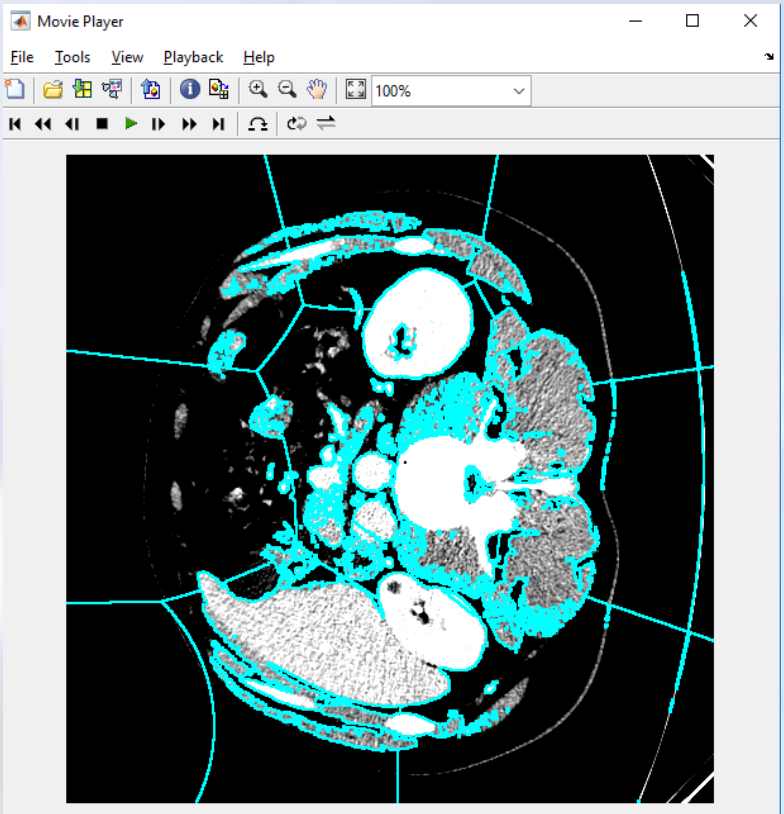

Axial slice from a CT scan, where the main organs are segmented with the superpixel

(4) imregdemons Supports 3-D Images on GPU:

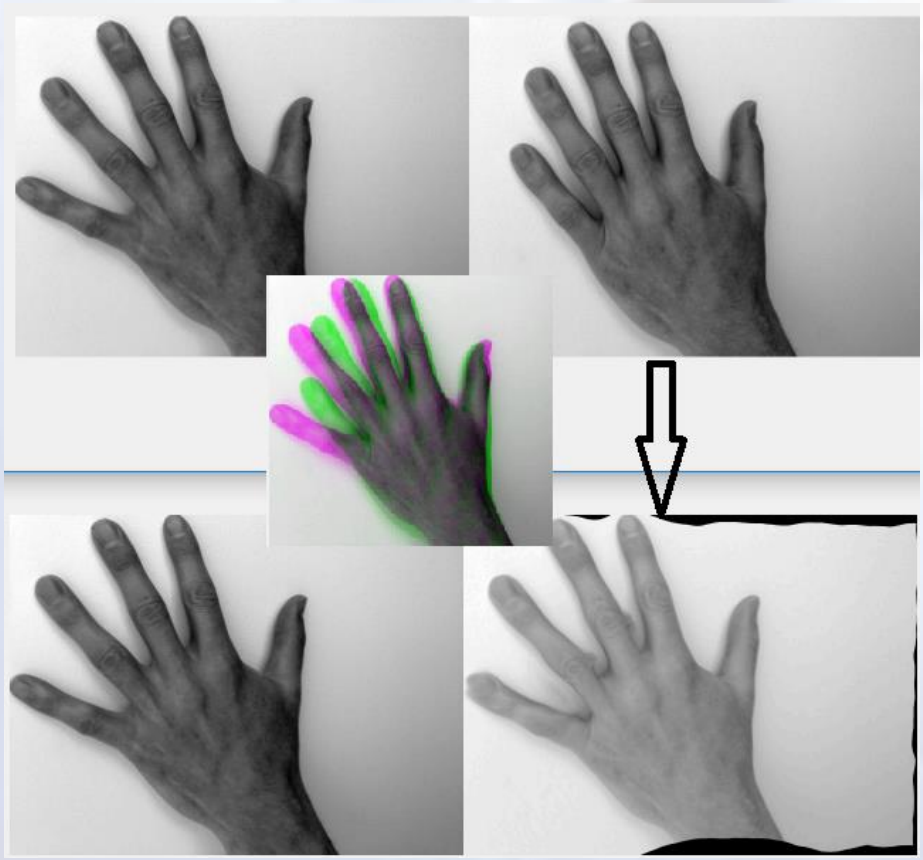

The imregdemons function now comes with a built-in fast implementation on GPU. The capabilities of the function can be seen in the example on the left: two hand images in different positions are the input, while the demons algorithm registers them one on top of the other and aligns them together.

New in Computer Vision Toolbox:

(1) Deep Learning for Object Detection: detect objects using region-based convolution neural networks (R-CNN)

Matlab keeps adding additional deep learning features: this time, the R-CNN capability. The following example demonstrates the R-CNN capability to identify a stop sign. After a short learning period on an average laptop computer like mine (a Lenovo Y50 with a simple GPU), this is the result:

(2) Structure from Motion:

Another new feature is built-in support for estimating the essential matrix and computing camera pose from 3-D to 2-D point correspondences and Project 3-D world points into image.

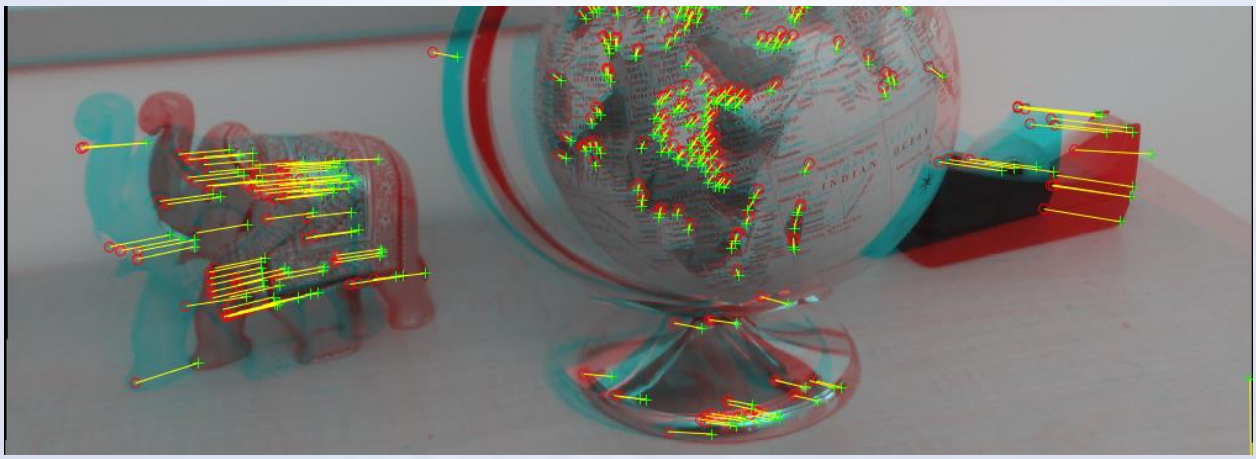

LiveCV

LiveCV (developed by Dinu SV) is a computer vision coding tool enabling quick prototyping of computer vision algorithms while seeing results update live as soon as the code changes. The aim of LiveCV is to **simplify the way people learn, configure and interact with computer vision algorithms**. LiveCV runs on both Linux and Windows. Portable stand-alone (i.e. no installation is required) and source code versions are available. For download, installation and instructions see here.

At its core, LiveCV uses three main technologies: OpenCV for the algorithms, QML for the programming language and QT for the display and the user interface (UI).

The Qt Meta Language (QML)

The QML (Qt Meta Language) is a user interface markup language. It is a JSON-like declarative language aiming at designing UI applications. QML represents a tree of elements which can be integrated to create UI ranging from simple buttons and sliders, to complete Internet-enabled programs. QML uses the JavaScript runtime V4 engine and Qt Quick for the scene graph-based UI framework. In QML, objects are specified by their type, followed by a pair of brackets. Object types always begin with a capital letter.

Programming in LiveCV is viewed as a pipeline:
'Read-image' -> 'processing it'…. -> 'observing the results'. Let's see a first example, the code snippets below read an image and blur it with a Gaussian filter.

```
import lcvcore 1.0
import lcvimgproc 1.0

Row{
    ImRead{
      id : src
      file : '/.image.jpg'
    }

  GaussianBlur {
        Input: src.output
        ksize:   "21x21"
        sigmaX : 5
        sigmaY: 5
  }
}
```

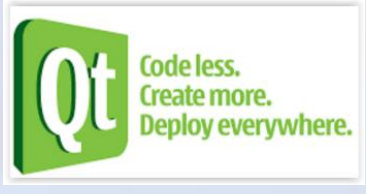

In the code on the left, there are three objects: a Row and its two children, an ImRead and GaussianBlur. Each object could have a special unique property called an id. Id enables the object to be referred to by other objects. Between braces, information and properties about the object can be specified; properties are specified as pairs <property>:<value>. For example, ImRead has a property id and a file specifying the file name to be read.

In addition, objects can be linked between themselves through their inputs and outputs. This is done by the property binding. The 'property binding' specifies the value of a property in a declarative way. The property value is automatically updated if the other properties or data values change. For example, GaussianBlur has a property 'input' which is bound to the ImRead. As mentioned, LiveCV wrappers OpenCV function. Thus, most of the documentation about the function, their usages and parameters can be found in the original OpenCV documentation.

Here is how the GaussianBlur OpenCV API looks:

> **GaussianBlur**
>
> Blurs an image using a Gaussian filter.
>
> **C++:** void **GaussianBlur**(InputArray **src**, OutputArray **dst**, Size **ksize**, double **sigmaX**, double **sigmaY**=0, int **borderType**=BORDER_DEFAULT)

It is evident that the parameters ksize, sigmaX and sigmaY are the same (boarderType is also available in LiveCV).

LiveCV interface

We will demonstrate how this code looks in the LiveCV GUI. We will start with a simpler example which only reads the image.

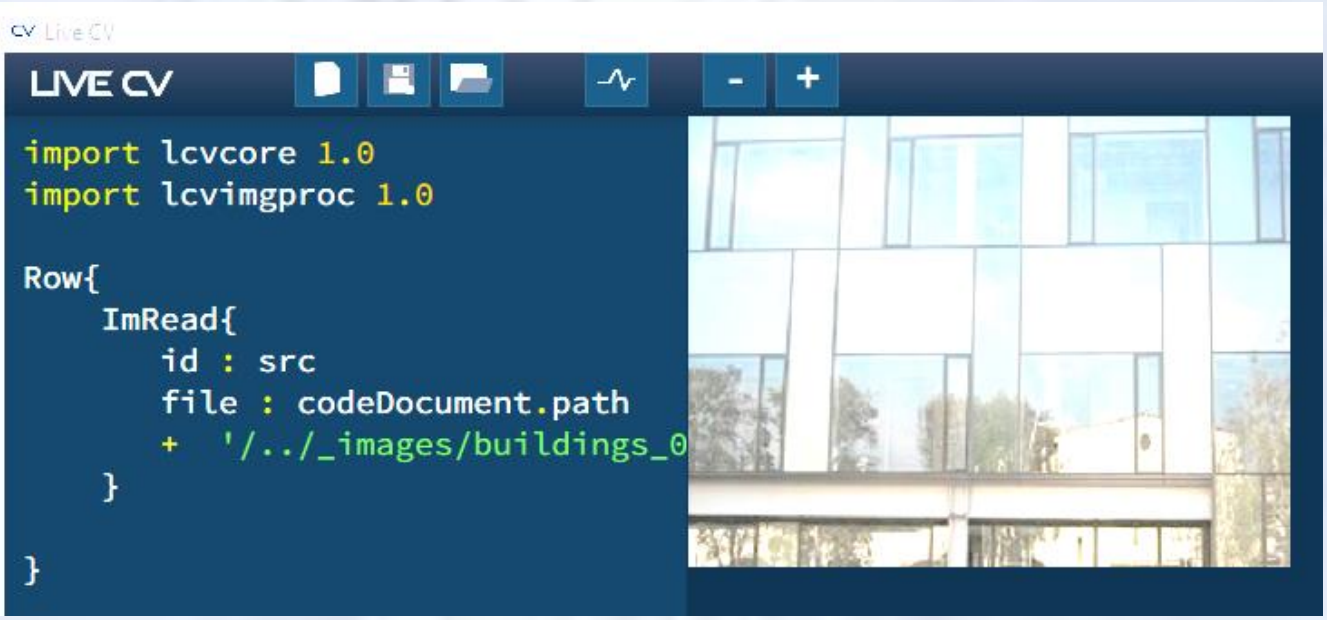

The LiveCV interface: on the top, control buttons for saving, loading, 'log' - for debugging purposes and '-/+' increase/decrease font size; on the left, a panel for writing the code (the QML code); on the right, the results are displayed.

Now we will add the Gaussian blur filter:

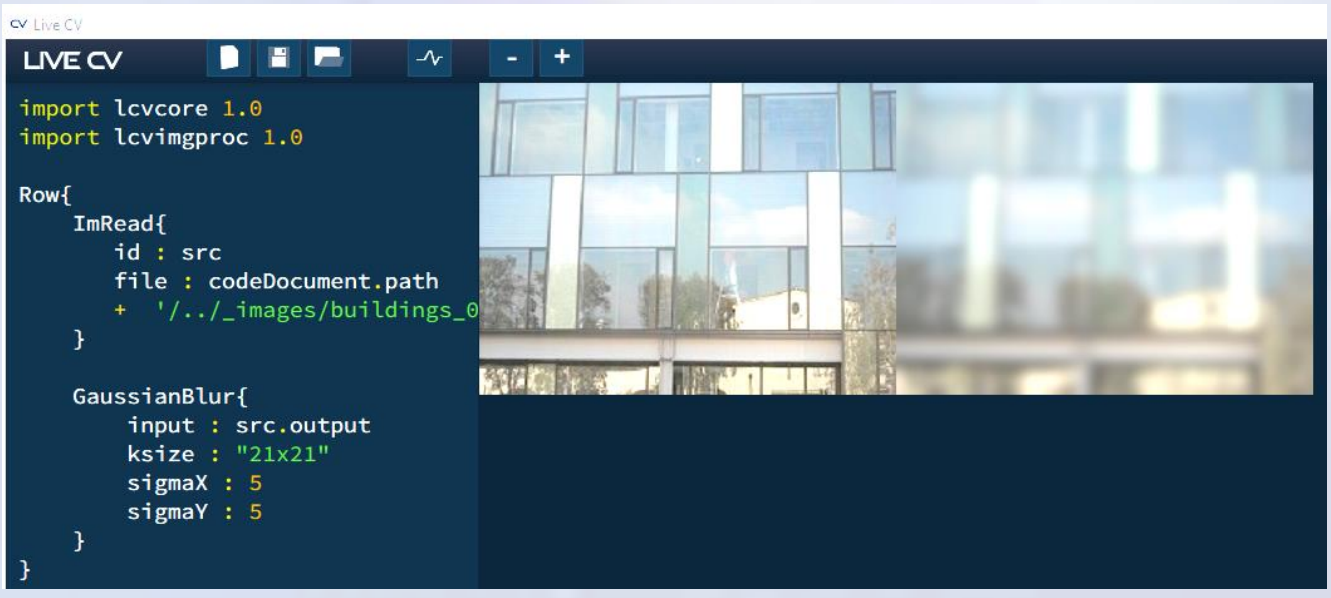

The image and the blur image are displayed side-by-side since we define those two under a “Row” object.

Now comes the fun part :) including any change that you’ve made to the code, either the sigmaX, sigmaY or the ksize, or even specifying a different file to load. The effect of those changes will automatically appear on screen to let you quickly learn, configure and interact with the algorithms and their parameters.

Let’s see another quick example: the canny edge detector. As in the GaussianFilter, the parameters are the same as in the OpenCV. The input to the Canny is bound to the ImRead. Note that here we define the filename as a properties string and not inside the ImRead section.

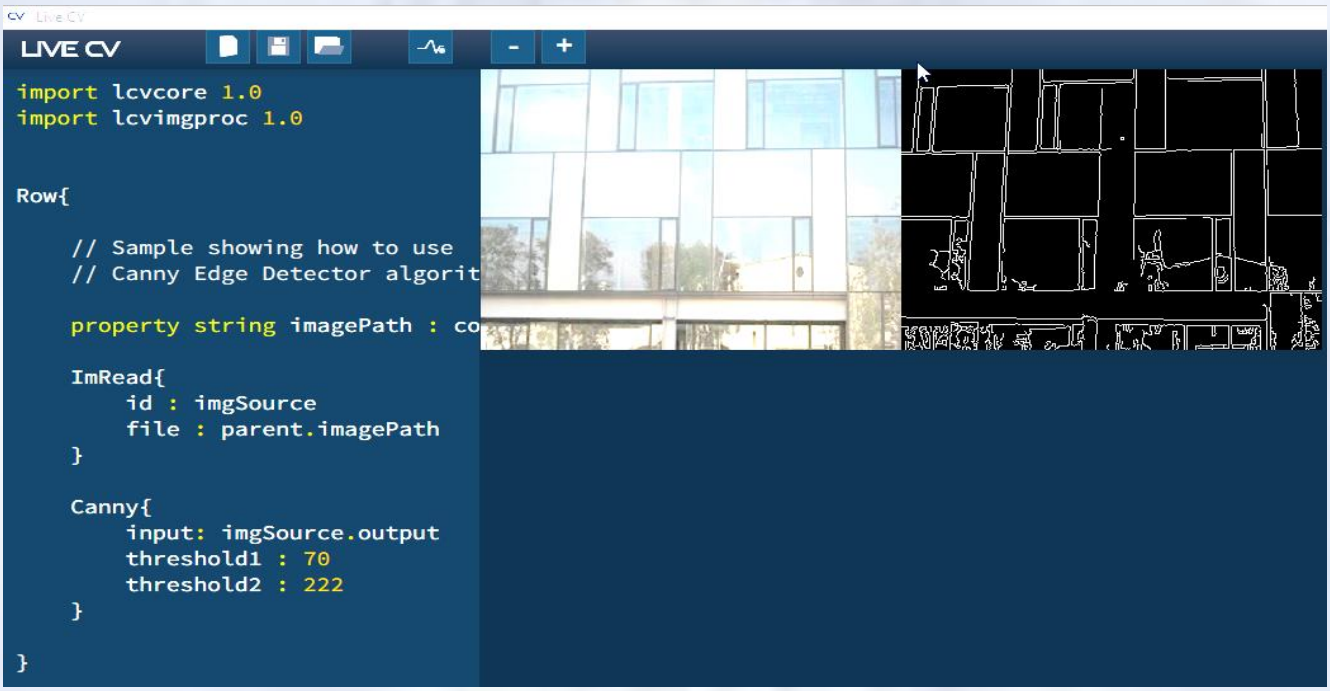

“Code less.
Create more.
Deploy everywhere.”

As before changing any of the Canny parameters (i.e. thresholds), the result will immediately be reflected on screen. This lets you quickly evaluate whether edges can be extracted from the image, which edges will be easy to extract, which are subtler and so on.

More elaborate examples include object recognition through feature detection. We won’t go into all the details but will mention the most important ones. In general, Image feature matching is a process that includes (1) finding features in the image, (2) describing those features, (3) comparing those descriptors between images to find similar images or to detect objects in images etc.

Each step in this pipeline may involve many variants of algorithms, each with its own parameters. Any of those can be evaluated with LiveCV and, as before, the result is visible on screen as soon as the change is made.

```
        // Query Image
        ImRead{
            id : queryImageLoader
            visible : false
            file : queryImage
        }
        FastFeatureDetector{
             visible : false
            id : queryFeatureDetect
            input : queryImageLoader.output
        }
        BriefDescriptorExtractor{
            visible : false
            id : queryFeatureExtract
            keypoints : queryFeatureDetect.keypoints
        }
    }
    // Matching
    FlannBasedMatcher{
        id : descriptorMatcher
        queryDescriptors : queryFeatureExtract.descriptors
        knn : 2
        params : {'indexParams' : 'Lsh'}
    }
    DescriptorMatchFilter{
        id: descriptorMatchFilter
        matches1to2: descriptorMatcher.matches
        minDistanceCoeff: 2.5
        nndrRatio: 0.8
    }
    DrawMatches{
        keypoints1 : queryFeatureDetect.keypoints
        keypoints2 : trainFeatureDetect.keypoints
        matches1to2 : descriptorMatchFilter.matches1to2Out
        matchIndex : 0
    }
}
```

In the above example, the first code section (//Query Image) starts by reading the image (ImRead) and detecting the feature with the Fast method (FastFeatureDetector). Other methods such as STAR, BRISK (and more) can be quickly changed and evaluated simply by renaming the object name from “FastFeatureDetector” to “ORBFeatureDetector” etc. Follows the description of those detected features performed with BriefDescriptorExtractor. Other alternative includes: BRISK, ORB and more.

The second block of code (// Query Image) starts by the FannBasedMatcher (Fast Approximate Nearest Neighbor Search Library), which finds matching points. The next block DescriptorMatchFilter filters out some matches by removing non-related ones, such as far distance matching. The last block DrawMatching (implanted by QT plugin) is responsible for displaying the matching points of the two images (the query and the training image).

Interactive interface

“It works quickly and efficiently”

In the last example we will show you an interactive interface for investigating image features. Here the image is shown on the left and on top of the image features are marks with circles. In addition, the feature points are listed with their values on the right. The feature points were detected and described by one of the point detection methods (i.e. FastFeatureDetector and BriefDescriptorExtractor). Clicking on one of the points to the left marks it with a big circle on top of the image: its values are displayed below, along with its histogram values.

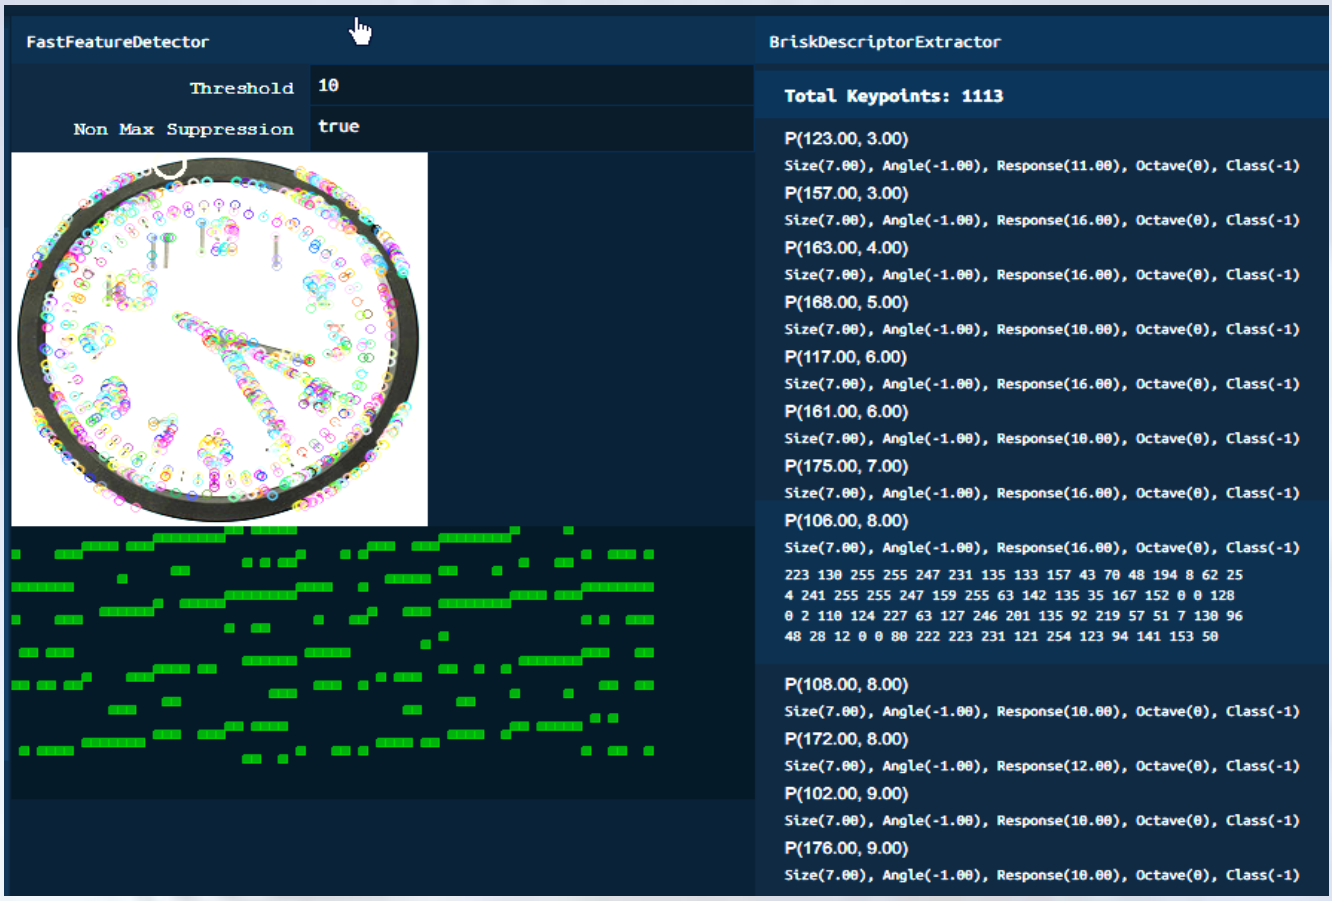

The interested reader can find the code in the LiveCV package itself (in the sample section). The LiveCV package comes with abundant additional samples, including interactive components which let you tweak parameters from the UI -- we strongly recommend you try them out.

Lastly, it's worth mentioning that you can write your own wrapper in case there is an OpenCV function that do not already exist in LiveCV. It is quite intuitive and the instructions on how to do that can be found on the LiveCV website.

To summarize:

LiveCV is a tool which helps you learn and understand computer vision algorithms, with the abilities to combine algorithms and progressively achieve the desired results. It has interactive components used to tweak parameters and also includes support for cross-platform, executable scripts, libraries to link and extend existing C++ algorithms

It works quickly and efficiently. As a portable stand-alone, it doesn't require any installation and can be launched even from a USB drive, which is great for demonstration / teaching.

Note that LiveCV is not reliable for replacing the developing platform (Visual studio for example): its purpose is for quick prototype familiarization of basic computer vision elements. It could be used for performing demonstrations in a teaching environment and getting a quick sense/feeling of various algorithms on your dataset of images.

Fully Automated Surveillance with PTZ Cameras

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 can help to solve technical challenges and physical difficulties. This month we review software for **fully automated surveillance with PTZ cameras**, developed for a client by **RSIP Vision engineers**. Do you have a project in computer vision and image processing? Contact our consultants.

Computer vision technologies contributed **major improvements** to the field of surveillance. From fully manual systems, in which videos were monitored by humans, the **first step** towards modern surveillance methods happened when software was able to **monitor the video feeds of cameras**: using predefined rules, alarms were automatically activated when intruders were detected entering specific areas.

The **second step** occurred recently when **big data and deep learning** were introduced into this field. Big data allowed to process the **huge amount of information** coming from cameras and sensors and integrate it into some usable database. Deep learning enabled the **recognition of patterns** (behavior and others) which could not be learned before, nor set by simple rules.

Thanks to these breakthroughs, we can use computers to track events which previously could not be brought to our attention. **RSIP Vision** has participated to this progress along both said steps. Regarding the first step, we have developed algorithms and complex mathematic transformations to allow tracking intruders using a **PTZ camera**: this is a camera providing pan, tilt and zoom functions, so that humans or objects can always be in the center of the frame. From a fully human intervention based process, RSIP Vision algorithms were able to automatically track the target intruder or event (unattended baggage, suspicious behavior), starting from the moment it was detected. On any movement of the target, the algorithm detecting its new location inside the frame adjusts the camera to keep it in the center. The algorithm can keep track of the target also in case of partial occlusion or temporary complete disappearance. Thanks to our work, security personnel can continue to observe the target as long as needed.

Regarding the second step, the breakthrough is only beginning: technologies to **process big amounts of information** already exist and engineers are gaining expertise in the use of deep learning to make further progress in the surveillance area. We shall learn more about these new developments in next month's Computer Vision News.

Image: DVTEL ioIMAGE on YouTube

Imama Noor

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. We think that some of what they did might serve as an example for other young females who wish to enter this field. This month, we interview **Dr. Imama Noor**, who earned her PhD from the University of Memphis.

Computer Vision News: *Where do you come from, Imama?*

Imama Noor: I was born in northern Pakistan, in Islamabad. I came to the United States in 2009.

CVN: *What are you currently doing?*

Imama: In my work, I collaborate with people who are trying to solve different problems which can leverage machine learning and computer vision. The project I'm working on right now is developing a social robot: it will be used as a telepresence robot or an appliance to monitor elderly people.

It involves a lot of tasks where you have to know the position of the person, their daily activities, recognize the person, and develop something which people can quickly adapt to and not be afraid of.

CVN: *Who is the group behind this project?*

Imama: I work with **E-Motion Robotics**. They are a startup based in Houston, Texas.

CVN: *Do you think your role is in the industry?*

Imama: I'm working in industry right now because industry has more resources when it comes to developing applications for consumers. Academia is more focused on theoretical problems and things which do not have immediate practical implementation.

I graduated in 2013 with a PhD and I joined industry right away. I didn't go for academia. Though, in my long term plan, at some point I might steer my career towards academia, when I'd like to slow down.

CVN: *Did you ever consider teaching or mentoring?*

Imama: I currently mentor high school kids who are interested in **robotics and machine learning**. They try to prototype or follow a certain blog for implementing a phase detector. They don't have in-depth knowledge of what's going on, but they can follow a certain guide to implement a particular application.

CVN: *Do you enjoy it?*

Imama: It's interesting to compare their experience with mine, when I was in high school in Pakistan. Back then we didn't have a lot of exposure to advanced technology. It's interesting that these kids have that kind of exposure, and that they are already building applications so that they might be able to do more things going forward.

CVN: *Did you ever experience any difficulty in being a woman scientist?*

Imama: Scientists in general are very rare in Pakistan because it's mostly an industrial-based economy. There are very few opportunities for research. In engineering, woman were always fewer than the male counterparts. In the US, I think it's more open. You do get acceptance if you are given that opportunity, so I think it's easier than in Pakistan.

There is a stereotype that women struggle to get accepted in the community. Males have the advantage of being perceived as engineers and scientists. If you are a female, you have to prove your abilities in a concrete line of work. Other fellows were able to get accepted without very significant contributions, but the women have to put forth concrete contributions to get accepted and move up in the field.

Women could provide **a valuable resource in computer vision or machine learning**. They should be treated equally. This may take some time, but I think the trend is positive. It should happen in our lifetime, in this generation or the next.

"Into the future"

CVN: *If you hadn't become a scientist, what would you have done?*

Imama: If I hadn't become a scientist, perhaps I would have become **an airplane pilot**. I wanted to be a pilot when I was young. I think it's a different kind of challenge from what I do now. The representation of women is fewer there as well, but now it's getting more automated and access will become easier. I would think there are similarities between the two, but doing what I do now I feel much more involved compared to just operating a flying device.

CVN: *What is your dream?*

Imama: I am very intrigued by **the concepts of space and time**, and how there is uncertainty in both. I would like to find out what is the truth behind space and time, and **understand the complexities of the universe**. I don't think that I can achieve it, but this is my dream.

CVN: What would be needed in order to achieve it?

Imama: I keep on following updates in that field. We are still very far from understanding how the universe was created and where we are going. If I had this information, **I would time travel**.

"It could be scary, but I think that shouldn't prevent you from trying"

CVN: *What time would you like to visit?*

Imama: If I could, I'd probably go into the future.

CVN: *Isn't that scary?*

Imama: I think it would be useful. It's something which is beyond our imagination. It could be scary, but I think **that shouldn't prevent you from trying**.

CVN: *You have worked a lot with space and time. With which of them you like to work the most?*

Imama: I think that time has interesting properties, but my expertise is mostly in space. I would like to explore time more, but I had more opportunity to explore space. I never had the opportunity to explore time, but **time looks more interesting to me**.

CVN: *Which teacher was particularly meaningful to you?*

Imama: My supervisor, Dr. Eddie Jacobs from the University of Memphis, was very helpful during my dissertation. He basically gave me the freedom to take any problem and figure out any improvements or other findings we can demonstrate.

CVN: *Do you plan to go back to Pakistan one day?*

Imama: After working more in the industry in the United States, I may go back to Pakistan. The work is more advanced in the United States, and there are a lot of problems in Pakistan which can be solved using this kind of technology. When I have the opportunity to spend time working on those problems then I will, but not right now.

"Understand how the brain functions, be able to model that kind of functionality, and help people with disorders to recover"

CVN: *What are you passionate about?*

Imama: One thing I'm really passionate about in technology is understanding how the brain functions, to be able to model that kind of functionality, and help people with different disorders to recover. I think that would be really helpful.

CVN: *What if this technology ends up in the wrong hands?*

Imama: Everything has positive and negative forces. I think it's the role of the law to make a policy to prevent this from happening.

Onfido

Traditionally, the identity of a person has always been verified face to face: when you start to work somewhere, on the first day you would often show your passport or ID to your new employer. Then, they would visually check if the photograph on your passport is the same as your face. These too may become old time memories thanks to **Onfido, a British startup** that has developed a **technology using machine learning** to bring all this process online and make it automated. In addition, it also prevents counterfeiters and fraudsters from creating sophisticated fakes.

Basically, the Onfido platform comes with many features that provide identity verification online remotely. For example, it provides identity verification if someone wants to sign up for an online bank, become an Uber driver, clean houses, and other things which are done with no face to face interaction.

The basis of the machine learning is in document verification. If you hold up a passport or driver's license, the first step the app does is verifying the document: for example, it checks if it seems to be a genuine passport or a fraudulent one.

"Comparing images to verify the person and the document"

Onfido does much more than comparing user data with an existing database: it actually compares images to verify the person and the document. As a user, you enter your name, date of birth, and address, then you hold up your driver's license or passport, the app takes an image of that and verifies that it is genuine. Then you record your face in a new image and **Onfido's platform compares the photo of your face with the photograph on the identity document**, in order to tell if your face matches the face on the documents.

We asked Husayn Kassai, Onfido's very young CEO and co-founder, about the application's approach: it primarily uses deep learning so that, based on all of the historic documents already checked, the app knows if there are any anomalies. If any fraudulent document is found during the verification, the system learns all of its

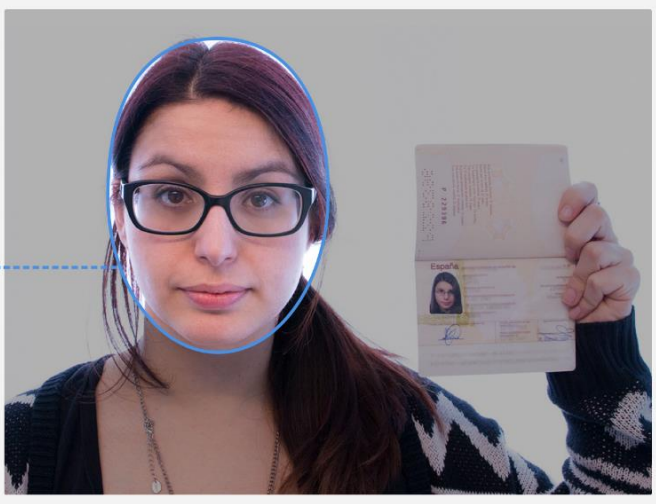

patterns and traits. In that way, if any similar patterns emerge on future documents in the future, the app will give an additional alert to their clients. Then clients can double check and focus more resources on the individuals that pose a higher risk.

Kassai told us that usually other identity verification providers work manually. Each time an image comes on their computer screen for the data processors, they compare these templates manually. Instead, Onfido verifies these documents using an **automated process with six layers of verification**.

The first layer is **Optical Character Recognition**. When a passport comes through the system, the latter uses **OCR** to extract all of the characters and verify whether they are correct.

A good thing with passports and many other identity documents is that you can do a parity check. Usually passports contain a standardized machine readable zone at the bottom which includes a parity check, so you can tell if the extracted passport number expiration date is consistent with the passport machine readable zone. That is like an internal consistency accuracy check.

Whether or not the app checks with authority databases or with an autonomous system depends very much on the country. There are two types of databases you can check. The first type of database is one which allows you to check whether the government issued that document or not. In India for example, there is a Unique Identification Authority (UIDAI) and Onfido checks the Indian identity card (called Aadhaar) by

Husayn Kassai, Onfido's CEO and co-founder

extracting its number and sending it to the government's API database. Then, they receive a confirmation as to whether the name, date of birth, and Aadhaar card number was issued by them or not.

The process is the same in the 15 states across the US which allow to go through it via the Department of Motor Vehicles (DMV). They also do it in the UK and in Brazil (with the CPF numbers). On the other hand, Onfido does not need to check passports with anyone elsewhere in the world, because consistency and parity checks are built in the machine readable zone. This is one type of database research which would validate whether or not the issuing body had issued that document.

The other databases that they search are police databases, which depend on each different country where Onfido establishes relationships with law enforcement authorities. Every time police raids a warehouse with thousands of fake passports, they extract the numbers and communicate them to Onfido. When the latter finds

fake documents, they report it to police as well. This doesn't happen in the USA however, since they don't share this information, which according to Kassai is a big shame. Such exchanges are much better accepted in Europe.

Of course, there are different levels of fake documents: poor quality fakes which are regularly caught; moderate quality fakes which are sometimes caught; and sophisticated fakes which the client reports to Onfido, giving them precious feedback for the future. Essentially, the more checks they do, the more robust the app becomes, which is why technology and machine learning works better than with humans. It is so much more robust.

The **specific algorithms** that help solve these challenges have come about after four years of research, during which they tried different techniques and found the ones that work the best. The technology uses an ensemble of techniques to get the needed results, ranging from what would be considered more **classical machine learning and computer vision techniques** (SVM's, feature extraction and the like) to the use of **deep learning** and a combination of these methods. The nature of their work, Kassai explains, often means that they are dealing with a considerable amount of uncertainty and they are often balancing this risk and reward with the techniques that they are using.

Their goal in doing that is to create an end-to-end trainable system. However, to achieve that goal they had to break the problem into a number of smaller problems which they are iterating over and then pulling these solutions together to create the end-to-end approach. This consolidation process may sometimes be quite painful: the team may throw away weeks or months of work and fine tuning to replace it with a more generalised model or approach, however Onfido feels that it is all part of the process, when the system is constantly evolving and growing. Onfido makes use of a huge variation of techniques, including **convolutional networks, LSTMs, SVMs** and many more.

Machine learning engineers at Onfido

Privacy is a very big concern with all identity verification systems. When asked about what happens if an individual is concerned with exposing his or her own personal documents, Kassai replies:

*"There are two things that happen. One is that we only check all persons with their consent. When you sign up with an online bank, a remittance platform, or with any client of ours, **we first take your consent** before we are able to verify you.*

"The second thing is that most people understand that in order to open a bank account, you have to show up to the bank, queue up and show your details. It takes a long time. Whereas, if you want to do it online, then there needs to be a robust way to do it. Otherwise there would be so much

fraud that it wouldn't happen.

"Equally, if you have a tutor who comes to your house to tutor your child, you want to be sure that the tutor has been checked. Likewise, you want to check the person with whom your daughter wants to share a car on BlaBlaCar.

"Most people recognize that in order to be a part of the new economy, there is ***a need for verification****. However, what we do is we always ensure that security is of the utmost importance. Penetration testing, ISO 27001, we adhere to all these standards to ensure security; in addition, everyone gives their full consent before they are checked; finally, data is deleted after the check and it does not enter any database: we only keep a record of fraudulent data and documents. That is for fraud prevention."*

"As soon as you are able to robustly verify people online, it opens up solutions to many problems"

Kassai admits that the system is exposed to **false positives**, but this happen nearly always because the quality of the document isn't good enough. Instead of giving a definitely positive or negative judgement on a document, they use a score. Usually when the quality of the image isn't good enough, it may well be a false positive. In this case, the client asks the person checked to submit another form of document, to re-upload or resend documents or go to the branch for face to face verification. This essentially means that only 2 or 3 out of every 100 people need to go through a more rigorous manual process as opposed to all of them, which saves a lot of time.

When they suspect that it may be a compromised document, they proceed with double checking: 90 to 95% of documents that come through pass the verification successfully. Out of the other five, two may be problematic and three may be questionable, so they only need to double check these ones.

Kassai remarks that in the primary elections in the US, people needed to stand in line for 6 to 9 hours in order to vote, while all this could now be done online. As soon as you are able to **robustly verify people online**, it opens up solutions to many problems, Kassai explains. This includes online voting, the gun control problem, or even the Syrian refugee crisis: the project is to work with governments so that they can verify someone in Jordan as a refugee. When they land in the US, they can verify that it is the same person. It alleviates all the concerns that on the transition they may steal someone else's identity.

Onfido gives much importance also to the work they do with banks around the world. There are 2 billion people, representing almost 1/3 of the world's population, that don't have a bank account or any access to finance and to the global economy. The reason for that is ridiculous: they are not able to verify their identities, so no one will give them credit or debit cards, cell phones, etc. But there is now a new wave of companies such as PayJoy or Pockit that are extending this kind of services. Pockit, for example, will give anyone in the world a prepaid debit card. Onfido can **capture somebody's face and passport** in an emerging country at a library or anywhere with a computer. That's all is needed to verify location and identity in order to sign up

and use this service.

When migrants come to the US or the UK, they have to wait 6 months before they are on credit reference agencies, can be verified, get a house, and start their lives because there is no history of them in that country. Whereas, Onfido can instantly verify them so they can start working for Uber or whoever else much faster.

"The only provider that uses machine learning and computer vision for document verification"

Onfido claims to be the only provider that combines all the elements of this check to make it robust and the only provider that uses machine learning and computer vision for document verification. It's the only way to make it robust so it can't be cheated.

Onfido is managed by a young co-founding team of three people, who left university and started the company, four years ago. At the time, when they tried to sell to their first customers, Kassai was only 22 years old. It is said that one of those prospects wouldn't even believe that he was 18: "***How can we trust you and buy anything from you, if we're not even sure that you are of legal age!***" Joking or not, they later became Onfido clients.

Onfido is constantly growing and looking for people interested in joining a research project or having something to contribute (in Lisbon, London, and San Francisco).

The Onfido Dashboard. Onfido has also been supported with 1 M€ by the Eurostars 2020 program of the European Commission

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!

The Extraordinary Link Between DNN and the Nature of Universe

How do you explain that **Deep Neural Networks** are better than humans at tasks such as face recognition and object recognition? There seem to be no mathematical reason why networks arranged in layers should be so good at it. Or so it seemed until now: **Henry Lin (Harvard)** and **Max Tegmark (MIT)** found out that the answer comes from **physics, the nature of the universe and the hierarchy of its structure**. [Read this fascinating work here!]

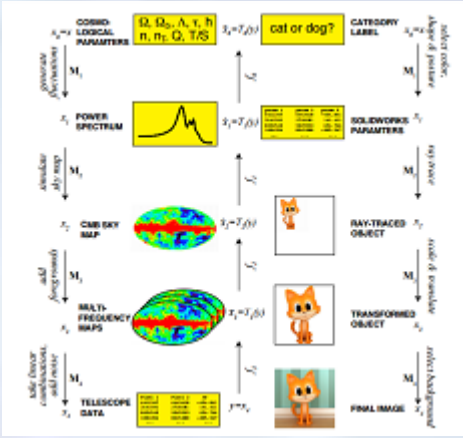

Engineers Teach Machines to Recognize Tree Species

Did you know that data from satellite and street-level images (like Google Maps') is used to automatically create an inventory of street trees? The **California Institute of Technology** can help municipalities manage the urban stock of trees. Other practical application of computer vision and image processing in forestry exist. **Read...**

3 OCR Libraries for Java Compared: Which is the Winner?

Gábor Vecsei did a great work to compare **Tesseract**, **Asprise** and **Google Could Vision**, three **Optical Character Recognition libraries for Java**. Want to know which library is the winner? **Read it here** and **find the codes here**. You might be tempted to read other articles in this blog...

Salesforce Offers CRM Artificial Intelligence for Business

Marc Benioff, chairman and CEO of Salesforce, thinks that the availability of **Einstein** - which embeds AI capabilities across every Salesforce Cloud - **will make Salesforce the world's smartest CRM**. Salesforce Einstein leverages all customer, activity, social and IoT data to to train predictive models powered by advanced machine learning, deep learning, predictive analytics, natural language processing and smart data discovery. **Read...**

What's in the Photo? Google's Caption Tool now Open-source

Show and Tell is **Google's automated captioning system**, which learns to identify and describe photos. Now it is available for open-source use with **TensorFlow**, Google's open machine-learning framework. **Read...**

See also this **lovely video from MIT CSAIL**, consult this **excellent list of TensorFlow resources** and don't miss this **report about NASA's 3D mapping of the Sun's edge**.

DSB - Transforming How We Diagnose Heart Disease

Every month, Computer Vision News reviews a challenge related to our field. If you can't find time to read challenges, but are interested in the new methods proposed by the scientific community to solve them, this section is for you. This month we have chosen to review the **Second Annual Data Science Bowl**, intended to catalyze a change in cardiac diagnostics: **Transforming How We Diagnose Heart Disease**. The website of the challenge, with all its related resources, is **here**; the Kaggle page is **here**.

Background

The **Second Annual Data Science Bowl**, created and sponsored by **Booz Allen Hamilton** with **Kaggle**, was designed to take action to transform how we **diagnose heart disease**. Thousands of people are diagnosed every day with heart failure, a life-threatening event. Data science applied to cardiology can help physicians save more lives.

Declining cardiac function is a key indicator of heart disease: **assessing the heart's squeezing ability** can therefore give clues about the heart conditions, enable an early diagnose and improve the effectiveness of heart disease treatment.

Two are the properties which need to be measured: end-systolic and end-diastolic volumes (i.e., the size of one chamber of the heart at the beginning and middle of each heartbeat). From these two measures, it is possible to calculate the **ejection fraction** (EF), which is the percentage of blood ejected from the left ventricle at each heartbeat.

To learn more about the medical considerations behind these measures, see the study by **RSIP Vision** on cardiac left ventricle segmentation.

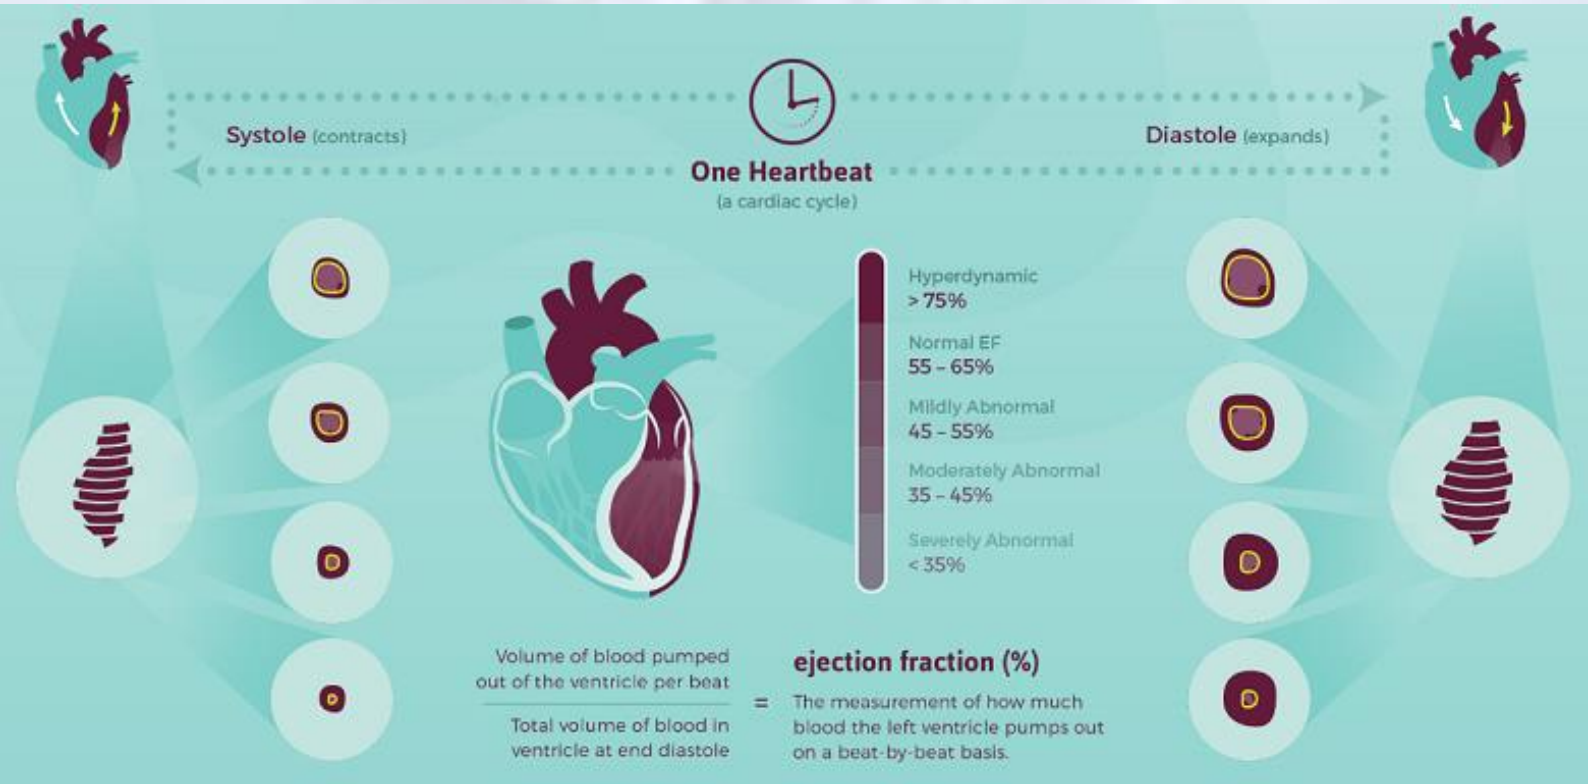

Challenge

Magnetic Resonance Imaging (MRI) is the reference standard used to assess those measures, also due to the reproducibility of making these important measurements.

The challenge is to help replace the slow manual process used by the doctor to derive the ejection fraction with an automatized procedure, which would arrive to the same conclusion in a more efficient way. Participant were requested to create an algorithm which would automatically measure end-systolic and end-diastolic volumes in MRIs, examining cardiac images from more than 1,000 patients.

"The leading models reviewed keep the diagnosis categories quite tightly grouped together"

Results

A relatively short development time was allotted for the challenge, nonetheless the top teams obtained excellent results. This is particularly notable for those who hand-labeled patient data without being trained physicians. You can read here further medical perspectives on the results of this challenge by Andrew Arai, MD (NIH, Bethesda).

Coming to the leading teams submission analysis, an EF under 35% is a dire emergency, around 60% is normal, and above 73% is considered hyperactive. The good news is that the leading models reviewed keep the diagnosis categories quite tightly grouped together: even though they are not always perfectly precise, there

Data science applied to cardiology can help physicians save more lives

is a very low probability that an emergency EF will be incorrectly categorized in the mild or the average categories. The normal to mild diagnoses are very likely to stay within their domain of the matrix.

Michael Hansen, co-PI (Principal Investigator) for this Data Science Bowl, noted that the best models can fail, but they should "*fail loudly*": in other words, they should be able to predict their own level of accuracy, so that when the system flags the confidence in the prediction as insufficient, it is possible to consider retaking a measurement. Results show low correlations between the model error and "confidence" distributions, suggesting that further improvement could be done in assessing the prediction value.

Deformable Part Models are CNNs

Every month, Computer Vision News reviews a research from our field. This month we have chosen to review **Deformable Part Models are Convolutional Neural Networks**, a research paper showing a synthesis of these two widely used tools for visual recognition. In fact the paper, presented at **CVPR 2015**, shows that a **deformable part model** (which is a graphical model) can be formulated as a **convolutional neural networks**. We are indebted to the authors (**Ross Girshick**, **Forrest Iandola**, **Trevor Darrell** and **Jitendra Malik**) for allowing us to use their images to illustrate this review. The full paper is **here** and the source code is **here**.

Background:

Deformable Part Models (DPMs) and **Convolutional Neural Networks** (CNNs) are two common tools for computer visual recognition. They are typically viewed as distinct approaches: DPMs are graphical models (Markov random fields), while CNNs are "black-box" nonlinear classifiers.

DPMs typically operate on a scale-space pyramid of gradient orientation feature maps (HOG). Nowadays, the state of the art object detection is done with deep convolutional networks. The authors propose a new method DeepPyramid-DPM in which the HOG features can be replaced with features learned by a fully-convolutional network.

Challenge:

The challenge is to replace the HOG feature with its CNN counterpart and finding the optimal way mapping each step in the DPM inference algorithm to an equivalent CNN layer.

Novelty:

The implementation details of CNN are time-consuming and challenging to setup correctly. As a result, HOG-based detectors are still used in a wide range of systems, especially where region-based methods (i.e. poselets) are involved. The DeepPyramid-DPM should therefore be of broad practical interest to the visual recognition community. In addition, DPMs are usually thought of as flat models, but DeepPyramid-DPM suggest that DPMs actually have a second, implicit convolutional layer.

Method:

A DeepPyramid DPM is a convolutional network that takes as input a color image pyramid and produces as output a pyramid of object detection scores. We will start by defining input and output for both the training and the testing as well as the dataset used.

	Training	Testing
Input	Set of images Trained CNN network (SuperVision. AlexNet)	Color image pyramid
Output	Optimized hyperparameter: distance transform (DT), geometry filter	DPM score map for the input pyramid level
Dataset	PASCAL VOC 2007	PASCAL VOC 2010-2012

The DeepPyramid DPM model overview (figure below): (1) an image pyramid is built from a color input image; (2) each pyramid level is fed through a truncated SuperVision/AlexNet that ends at convolutional layer 5 ($conv_5$); (3) the pyramid of $conv_5$ feature maps depth is 256 as in SuperVision/AlexNet; (4) each $conv_5$ level is fed into a DPM-CNN (more details in the next paragraph), which (5) produces a pyramid of DPM detection scores.

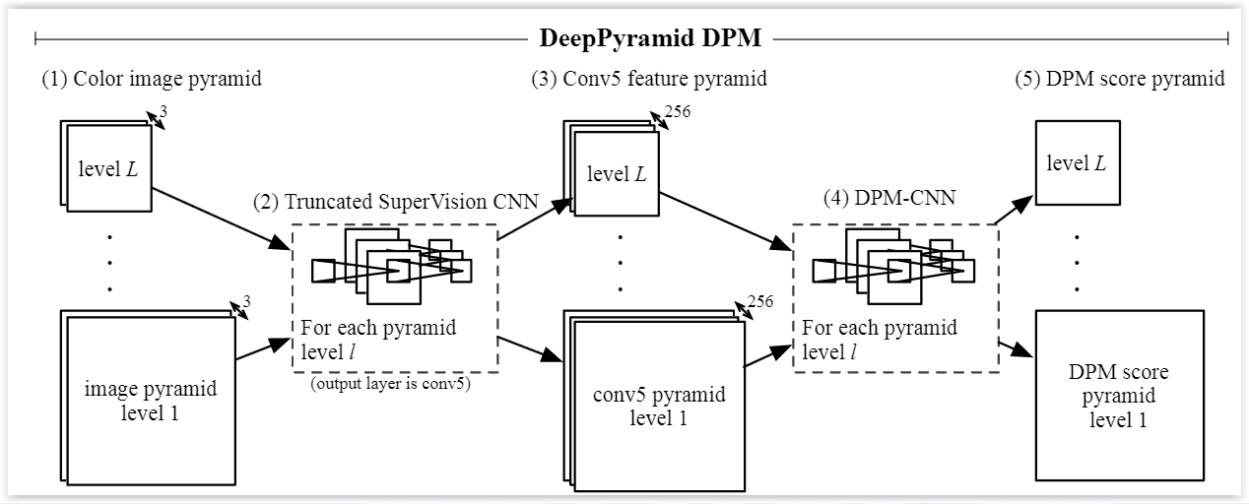

A single-component DPM-CNN operates on a feature pyramid level: (1) the pyramid level is convolved with the root filter and the P filter parts, generating P+1 convolution maps; (2) those are processed with a distance transform (DT) pooling layer and (3) a sparse object geometry filter. The output is a single-channel score map for the DPM component.

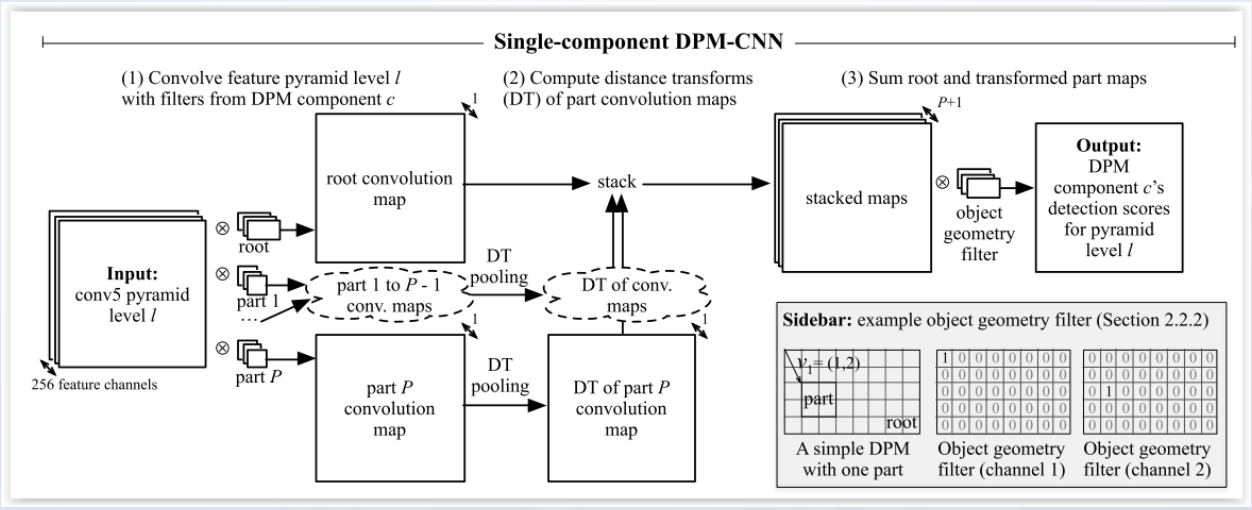

The **distance transform** (DT) pooling is a generalization of the familiar max-pooling operation, the **object geometry** layer that encodes the relative offsets of DPM parts. More details about that can be found in the paper.

Lastly, a multi-component DPM-CNN composed of one DPM-CNN per component and a maxout layer that takes a max over component outputs at each location.

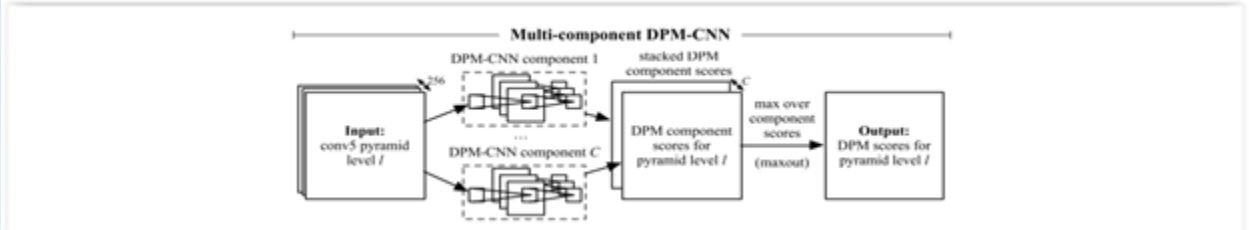

The model is trained in two stages: first, the front-end CNN is fitted; second, the DPM-CNN is trained using latent SVM while keeping the front-end CNN fixed.

Results:

The authors start by comparing HOG with the $conv_5$ features on different pyramid levels. Next, we evaluate different DPM-CNN settings and compare them to R-CNN.

HOG vs $conv_5$

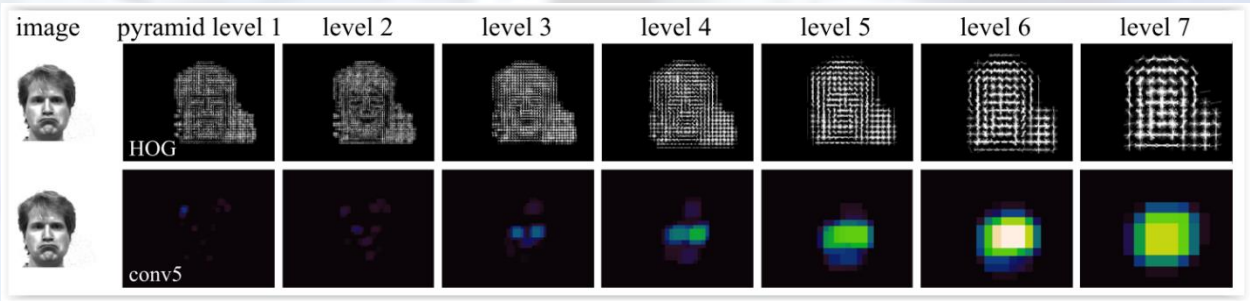

The input is the same two face image (left). The first row shows a HOG feature pyramid, while the second shows the "face channel" of a $conv_5$ pyramid. In HOG pyramid level, an appropriately sized template could detect the face. In contrast, the $conv_5$ face channel is scale selective: it is almost entirely zero (black) in the first pyramid level and it peaks in level 6.

"Can any DPM be formulated as an equivalent CNN?"

Evaluated DPM-CNN

Evaluated different DPM-CNN settings:

Line 1-3: The results start with a "root-only" model (i.e., no parts) and then show results after adding 4 or 8 parts to each component. With 4 parts, mAP increases by 0.9 percentage points.

Line 4: The hypothesis is that the convolution filters in conv5 already act as a set of shared "parts" on top of the conv5 features, this idea is implemented by applying a 3×3 max filter to conv5 and then training a root-only DeepPyramid DPM with three components achieving 45.2. Those are the best results for this section supporting the hypothesis. In addition, the hypothesis is also evidenced in the face heat maps above (HOG vs conv5) in that it selects specific visual structures at their locations and scales (face at conv5 level5).

Line 5-6: Computed to HOG-DPM baseline using 6 components and 8 parts, removing parts decreases the performance to 25.2%. Both do not perform as well compared with DeepPyramid -DPM max5.

Line 7: Compared to the recently proposed R-CNN [Goodfellow, ICML 2013]. The R-CNN compared here has pool5 features without fine-tuning.

Line 8-9: Fine-tuned R-CNN [Goodfellow, ICML 2013] clearly outperform DeepPyramid-CNN. However, DeepPyramid-CNN runs at about 20x.

The R-CNN results (lines 7-9) suggest that the gains from fine-tuning come in through the non-linear classifier (layers fc6 and fc7) applied to pool5 features. This suggests that similar levels of performance might be achievable with DeepPyramid DPM through the use of a more powerful nonlinear classifier than the SVM.

	Number of components	Number of parts	Train	Test
			mAP	
DP-DPM $conv_5$	3	0	43.3	N/A
DP-DPM $conv_5$	3	4	44.2	
DP-DPM $conv_5$	3	8	44.4	
DP-DPM max_5	3	0	45.2	42.0
HOG-DPM	6	0	25.2	N/A
HOG-DPM	6	8	33.7	33.4
R-CNN $pool_5$ [Goodfellow, ICML 2013]	N/A	N/A	44.2	N/A
R-CNN FT fc_7 [Goodfellow, ICML 2013]	N/A	N/A	54.2	50.2
R-CNN FT fc_7 BB	N/A	N/A	85.5	53.7

Sum up:

For decades, visual recognition models have made a wide use of part-based representation techniques, such as deformable part models (DPMs). In recent years, good performance on image classification, object detection and a wide variety of vision tasks made CNNs extremely popular among researchers. DPMs and CNNs are generally viewed as distinct approaches to visual recognition, the former being graphical models and the latter being non-linear classifiers. The question asked by this paper is whether these models are actually distinct or whether it can be shown that any DPM can in fact be formulated as an equivalent CNN. In order to do so, the authors introduce the notion of distance transform pooling, and the object geometry layer. The authors find that DeepPyramid DPMs significantly outperform DPMs based on histograms of oriented gradients features (HOG) and slightly outperform a comparable version of the R-CNN detection system, while running significantly faster.

EVIC - IEEE Latin-American Summer School on Computational Intelligence

The **School of Engineering and Applied Sciences of the Universidad de Los Andes** (Santiago de Chile), jointly with the **Chilean Chapter of the IEEE Computational Intelligence Society**, is organizing **EVIC, the Latin-American Summer School on Computational Intelligence** (December 14-16).

Thanks to **José Delpiano**, we have been able to ask **Pablo Zegers, General Chair of the Summer School**, to tell us more.

CVN: *How many students do you expect at the Summer School and what percentage of international students?*

Pablo Zegers: We expect more than 100 people to attend EVIC (the Spanish acronym for the meeting - Escuela de Verano en Inteligencia Computacional). Most of them will be students, a few attendants coming from industry, and also faculty members. Most of the students are Chilean, both from Santiago and regions; the usual **ratio of international students is around 10%**.

CVN: *What is in this program that students will not find anywhere else?*

Zegers: This event is unique in Latin-America. Every version of the Summer School has attracted world-renowned researchers in the field of computational intelligence as keynote and tutorial speakers. Chile is an appealing travel destination for researchers around the world, as the record-breaking attendance of last **International Conference on Computer Vision (ICCV)** showed to the organizers.

CVN: *This is already the 12th edition of this program: what has changed since the first edition and what has remained the same?*

Zegers: Since the first edition, the world-class level of keynote and tutorial speakers has remained. We continue to aim at covering current research and applications of **neural networks, fuzzy logic, and evolutionary computation**. Nevertheless, the selection of relevant topics has changed with the interests of computational intelligence researchers, providing a level of novelty to each edition. This year, the program focus will be on **analytics and deep learning**.

Event

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.

ECCV - European Conference on Computer Vision
Amsterdam, Netherlands Oct. 8-16 Website and Registration

RE•WORK Women in Machine Intelligence & Healthcare
London, UK Oct. 12 Website and Registration

MICCAI - Medical Image Computing and Computer Assisted Intervention
Athens, Greece Oct. 17-21 Website and Registration

RE•WORK Deep Learning Summit
Singapore Oct. 20-21 Website and Registration

International Conference on Computer Vision and Image Processing
Malaga, Spain Oct. 20-21 Website and Registration

ACIVS - Advanced Concepts for Intelligent Vision Systems
Lecce, Italy Oct. 24-27 Website and Registration

HCOMP 2016 - Human Computation
Austin TX, USA Oct. 30-Nov. 3 Website and Registration

RE•WORK Machine Intelligence Summit
New York NY, USA Nov. 2-3 Website and Registration

International Conference on Computer Vision and Image Processing
Paris, France Nov. 21-22 Website and Registration

ACCV'16: The 13th Asian Conference on Computer Vision
Taipei, Taiwan Nov. 20-24 Website and Registration

International Conference on Pattern Recognition and Computer Vision
London, UK Nov. 24-25 Website and Registration

DICTA 2016: Digital Image Computing Techniques and Applications
Gold Coast, Australia Nov. 30-Dec. 2 Website and Registration

ICPR 2016: International Conference on Pattern Recognition
Cancun, Mexico Dec. 4-8 Website and Registration

ISVC 2016: International Symposium on Visual Computing
Las Vegas NV, USA Dec. 12-14 Website and Registration

IPTA 2016: Image Processing Theory, Tools and Applications
Oulu, Finland Dec. 12-15 Website and Registration

ICCVISP 2016 : Computer Vision, Image and Signal Processing
London, UK Dec. 15-16 Website and Registration

ICVGIP 2016: Computer Vision, Graphics and Image Processing
Guwahati, India Dec. 18-22 Website and Registration

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

Full-kit of Requirements

Our CEO **Ron Soferman** has launched a series of lectures to provide a robust yet simple overview of how to ensure that computer vision projects respect goals, budget and deadlines. This month, we will learn about the **Full kit of Requirements**.

"Without a full kit including all the detailed requirements, it might take several development cycles until the product is complete"

One of the principles in project and software development is the need for a full kit of requirements before starting the actual development.

It may happen that the image processing team is given many tasks with tight deadlines, in order to support the roadmap of product development. This often leads the process into a bottleneck.

If the project is not launched in a systematic way, it might take several development cycles until the desired product is complete. In many cases, the reason is that the full specifications of the product are not well defined since the beginning: the development team is eager to start its work, managers are very determined to keep up with schedule and everyone wants to start as soon as possible. But without a full kit including all the detailed requirements which are needed for the development, deadlines can't be met.

"Include complete and realistic data in sufficient quantity"

Full kit for research:

Sometimes, project managers sense that full requirements cannot be defined upfront because of technical and algorithmic unknowns. General product requirements are probably known, but they may be fluctuating until the software development step.

Nevertheless, I recommend that the full kit in the research stage include all that is required for research: complete and realistic data in sufficient quantity, without which the developed algorithm might behave differently when confronted with real data. The research dataset should include all scenarios, especially in the medical field where an algorithm tested on data belonging to healthy people might be useless when used on patients with pathologies. Collecting relevant data and analysing its possible appearances is a difficult but necessary task.

The reason for the need of a full kit in research is that algorithm must take all cases into account. With only limited examples, it will be a naïf algorithm that will not hold when used on the full variety.