MSBAPM Newsletter

November 2017 Edition

Fall 2017 Capstone ended last week successfully for all four Classes. One of the teams with a snap with Prof John Wilson

THIS MONTH

in the MS BAPM Newsletter..

- Image Classification - Brief Notes
- Interview with the Faculty - Ramesh Shankar
- Prayer - A humble perspective by Vinay Srivastava
- Lessons in photography by Saurav Gupta

Winter is here!
See you in December!

Image Classification – Brief Notes

- **Karthik Ramachandra**

Consider an image of a T-shirt.

Ok, no. Let us take a simpler image, a greyscale, black and white image of a T-shirt.

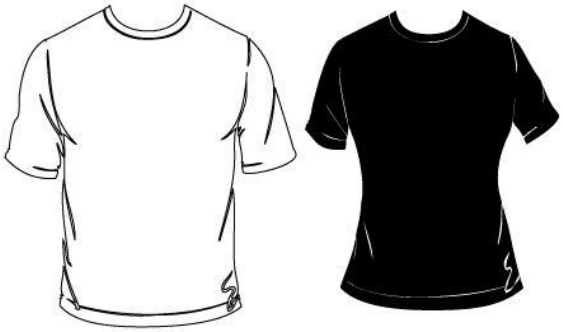

Good, so two t-shirts, black and white in nature. But, our human minds already have this memory to know these are "T-Shirts" right? Our brains are "trained" already and hence "classifying" them into a category is easy.

But what about machines? How do we train algorithms to recognize images and then classify them into respective labels? Let us see what we can learn in simpler terms.

Image classification is also called contextual image processing. The whole process of teaching or training an algorithm to learn an image is entirely based on the pixels digit data of the image. This is "contextual" in nature. As in, based on the neighborhood intensity, the algorithm estimates various facets of the image and dissects it among the 255 pixels to understand the true nature of the image.

Now that we have considered t-shirts before let us look at a data set, popularly called as fashion MNIST for our reference here. MNIST is Modified National Institute for Standards and Technology and is a database of digit based pixel image data which consists of various datasets that have been extensively used for decades as a starting point to learn image classification.

Zolando, a fashion retail e-commerce store based in Europe released this dataset called the Fashion MNIST, around last year with an intention to make it the standard for all image classification problems. Zolando thinks that basic MNIST dataset has lost the aura it once had as the dataset itself has been used extensively world-wide and can be accurately classified by simpler algorithms.

A lot of research and literature is already available on the internet on Fashion MNIST already. People have extensively employed deep neural networks to extract higher accuracies. The standard go-to algorithm had been Convolutional Neural Networks (at this moment referred to as CNN). Like any other deep learning algorithm, any other neural network, a CNN is also a Black box. Although we know the various layers and theoretically understand what is happening at each layer, it is difficult to comprehend and costlier to implement.

You can read more about CNN here: http://cv-tricks.com/tensorflow-tutorial/training-convolutional-neural-network-for-image-classification/

Let us see if we can apply the same dataset and train a simpler classification algorithm to get a similar or (being ambitious) higher accuracy than a CNN.

A quick look at the dataset:
Training set has pixel information of 60,000 images of clothing.
Test set has pixel information of 10,000 images of clothing.
Each image is a 28 * 28, greyscale pixel data.

One thing to be noted here is, the dataset consists of pixel data and not images. If you have picked a dataset with images, you need to convert the data first into pixels and then train an algorithm.
Each of the clothing image data had to be categorized into any of the ten labeled classes as given.

Label	Description
0	T-shirt/top
1	Trouser
2	Pullover
3	Dress
4	Coat
5	Sandal
6	Shirt
7	Sneaker
8	Bag
9	Ankle boot

To start with, the first algorithm to be applied was K-Nearest Neighbor(KNN). KNN plots all the data points and separates them into classes based on the labels. It smoothens the boundaries between labels using what is called a majority vote of specified K nearest neighbors to each point. However, here there is a drawback. The algorithm, in reality, does not learn the data but does a comparison when it comes to prediction. So, unlike other algorithms which learn the training data and then predict on the test data, KNN compares each observation in the test data with every single observation in the training data to see if the two observations being compared are close to each other or not. This requires high computational power to predict and increases as we increase the value of K.

This is undesirable in any real-world scenario and has very few business applications when used for image processing.

Though KNN achieved good accuracies, considering the cumbersome nature of its classification method, a random forest is now applied for classification. Here, three different methods are used: a) A random forest without tuning hyper parameters b) A random forest after tuning hyper parameters using Grid Search and c) Segmented the pixel data to obtain different classes of similar intensity and used it as a feature and then tuned the parameters. All these models gave us pretty robust results. The accuracies of all the models, including KNN, are as below:

Models	**Training**	**Test**
K- Nearest Neighbors (K = 5)	90	85.2
Random Forest - Untuned	98	88
Random Forest - Tuned	88	85
Random Forest – Segmented & Tuned	**90**	**87**

The random forest algorithm with segments and tuned to hyper parameters provided the best results with accuracies of 90% on training and 87% and test datasets.

Since most CNN used as an industry standard achieved an accuracy of 89% on the test data, maybe this can be a good measure to understand that simpler algorithms do work well on this dataset. Using these algorithms save a lot of computational power and costs of implementation. But, the tradeoff comes in terms of the accuracy and what is the business need. Can the business afford a loss of accuracy on the scale of 2% when millions are at stake? That truly is a million-dollar question, isn't it?

Note: All work and results are a product of an academic project of the author and may not be reused/reproduced without permission

----

Interview with the Faculty: Ramesh Shankar

Briefly, give us an introduction about yourself, please.

I did my Ph.D. at NYU Stern and joined UConn OPIM department in 2004. I have taught in the undergraduate MIS program, MS-BAPM, as well as Executive MBA programs. In MS-BAPM I teach Big Data Analytics and Data Science with Python. I also do research using large datasets, using a variety of econometric and data mining techniques.

How did you land your first teaching job?

UConn is my first and only teaching job. I interviewed in Winter of 2004, in a visit that lasted two days and involved meeting perhaps some 15 to 20 faculty members and administrators. I enjoyed

meeting everyone and loved UConn. Fortunately, they liked me too and made me a job offer.

What excited you to be part of the BAPM Faculty?

I find the BAPM students to be enthusiastic and eager to learn. And that makes teaching in BAPM fun. Moreover, I love data analytics. It is synergistic with my research, so I get to teach topics which I can relate to in my research as well.

What would you like to improve at BAPM?

Meghan, Kathy, Akilah, Tashika, Lynsi, Jose and others are already doing a fantastic job; it is hard to think of ways to improve BAPM.

It will be nice if we can have a systematic way of inviting back alumni for frequent tech talks. It will help renew connections with alumni and help current students with their job search and career guidance. A student group should take charge of inviting alumni, scheduling talks, and promoting the events. Perhaps we can get some funding to fly in alumni from other states like California.

What advice would you like to give to graduating BAPM students?

Your BAPM education is only the start of your business analytics journey. Never stop learning and improving yourself in your professional and personal life. And once you graduate and move on to your new job, don't forget to help subsequent batches of students – make connections for interviews or be available for career advice via email or skype.

What are your hobbies and interests apart from making BAPM great?

I like to listen to music and read books on current topics in business, technology, and society.

Prayer – A humble perspective

Vinay Srivastava

We stand in front of Almighty and beg for his blessings. We ask HIM to grant us this wish or that wish.

If there is someone or something who is "Almighty" then isn't that Almighty omniscient?

If HE is omniscient then why do we need to ask HIM for something? Doesn't HE already know what we want?

With this food for thought, I would like to humbly state that instead of asking we must be thanking HIM every now and then.

Many a times we complain that our prayer was unanswered. We didn't get what we wanted. The truth is that our prayer failed at that exact moment when we begged.

By seeking blessings (wishes/desires), we are doing nothing but complaining. Complain implies that we are unhappy with our current state. It also means that what we want is not what we have **and** what we currently have is not what we wanted. In other

words, we consider ourselves far wiser than the Almighty. What HE is doing is not right. HE must seek our advice and then act.

True prayer is absolute love. It has no demands and no qualms. True love has always been zero-demand and unconditional.

The day we learn this, that day would be the first day when we shall have everything in life. Life will be far more beautiful than it is. It would be blissful.

It shall be the pinnacle of our existence.

Lessons in Photography

-Saurav Gupta

"How to take better pictures with your smartphone?"

I live by the quote *"A good photo is a good photo, no matter what device it's taken with"*

Today, in this era of Machine Learning and Artificial intelligence the cameras in mobile phones are at par with conventional SLR's. One doesn't need to carry a bulky camera anymore. For the last decade smartphones have gotten thinner and faster and thinner and faster, and well, you get the picture. Nowadays, pictures from smartphones will catch your curiosity and attention. So, today I will be sharing some tips for taking better pictures with your smartphone.

1. **Clean the lens:**
   Your phone spends a lot of time in your pocket, a bag or in your hand, and as a result, it will get dirty. Forget about rugged and expensive cases; they won't help. Dust, dirt, and grease have a big

effect on the quality of your photos. Dirt blocks light from entering the camera sensor and will leave blur, smudges and dust spots on your images. On the contrary, a clean lens will ensure you get sharp, clear images from your smartphone. This is a very basic thing, and most of the people don't focus on this. So, next time make sure you clean the lens before you take out your phone to snap a picture.

2. **Never use the zoom:**
   Zooming in to get closer to your subject? Well, don't do it.
   Our phones are equipped with digital zoom, not optical zoom. In essence what happens with a digital zoom is that the image gets cropped and you lose image quality. If you want to take the photo of a faraway subject, walk closer to the subject and take a picture without zoom. Alternatively, you can crop the picture itself in the editing process to bring the view closer as desired. Cropping afterward will give you more control over how much of the image you want to remove.

3. **Use Rule of Thirds for Composition:**
   Without good composition, your photo isn't likely to be very eye-catching. Don't get intimidated by the term Rule of thirds; it is a jargon in the photography world. Let me lay it for you; so, while taking up a picture, you mentally divide up your image using two horizontal lines and two vertical lines such that your position important elements/subject along those lines, or at the point they meet.

Check these examples:

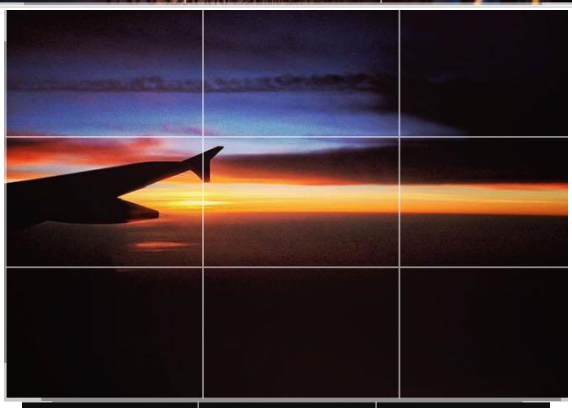

While taking up pictures, we usually keep our subject in the middle that is not always interesting. The idea behind the rule of thirds is that off-center composition makes for a more interesting shot.
On iPhone, you can actually turn on grid over the screen.
Go to Settings -> Camera, then turn the grid on. My android friends have to figure it out! Google it!

In all of these pictures the pathway, the road, the railings lead your eyes to the end of the pathway. It draws attention and creates a depth effect. Mostly leading line start in the foreground and move off to the background which creates the most dynamic effects.

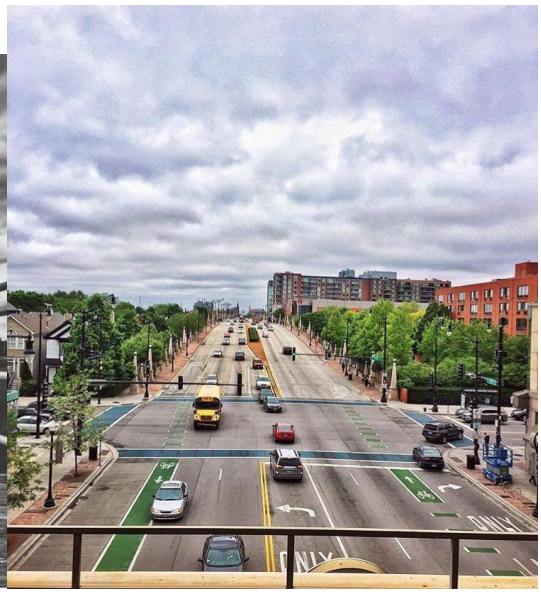

4. **Use Leading Lines:**
   Leading lines are another very useful and powerful composition tool. They help to focus the viewer's eye on the main subject and lead the eye deeper into the image. It's a simple technique that involves using vertical, horizontal or converging lines to focus attention on the subject.

Leading lines happen everywhere in nature so next time lookout for these interesting frames.

5.Perspective | Angles

Usually, we are in the habit of shooting everything from eye level, isn't it?

Most common example, when we are walking and if something catches our eye, we take a picture right from where we are standing. We usually don't explore another perspective. If you seriously want to make an impact in your photography, you need to get out of an eye-level rut. You need a change in perspective.

Get ready for some bending, turning, walking and climbing, start working with perspective in photography and your pictures will thank you.
See below pictures for examples

Get up High
See this picture above, how well I used Rule of thirds for composition and I was standing on a chair to photograph her innocence. Getting well above your normal line of sight will certainly give you better perspective.

Get Low
Get your camera down towards ground level, and see how it Is impacting the perspective. I did the same in this picture. My subject appears more imposing.

Look up high:
I had to lay down on the floor to take up this picture. See how the perspective has changed.

So, do not fall into the trap of shooting everything you see an eye-level, just as you see it. Take some time to explore your subject and consider changing your perspective.
Keep Shooting!

PS: All the pictures in the post are shot by me and are subject to copyright