

03-20-2017 - A \$58,972.55 Machine Learning question

Hi

This is Satya Mallick from LearnOpenCV.com

Today, I am not sharing a blog post on Computer Vision and Machine Learning techniques. Instead, I am sharing a story that launched by consulting business. With this story, I hope to convey the importance of developing skills, investing in yourself and a few tricks for learning fast in this world that rewards results, not effort.

Our story begins on a Thursday morning a few years back. I received an email from a prospective client with this simple question. It would prove to be a turning point in my fledgling consulting business.

"Can you show us any of the results from the testing data?"

If I showed some nice results on their dataset, I would win the contract. There were four other vendors vying for the same contract. I had told them that their problem was "eminently solvable" using Deep Learning.

There was just one problem, though. **I had no practical experience** in Deep Learning. Zero, zip, zilch, nada. I had no results to show and they were going to make a decision the very next day!

As we shall see later in this story, **it was a \$58,972.55 question** for me!

Fortunately, I did have some theoretical understanding as I had taken Dr. Andrew Ng's and Dr. Geoff Hinton's courses on Coursera during my free time. But, my practical experience was non-existent. My initial thought was just to give up. There was not enough time to show any results. The fear of failure and rejection was paralyzing. I knew this project could pay the bills for a few

months and provide a nice boost to my consulting business.

Somehow in the middle of all the negativity, I reminded myself of my first rule for acquiring a skill quickly.

Commit to massive and relentless action until success.

Half an hour after those crippling thoughts, I was back at my desk installing Caffe on my laptop. I had stopped worrying and started working toward my goal. It was a liberating experience.

After fighting through some frustrating dependencies, I finally had Caffe installed on my machine.

Good artists copy

The second rule of learning fast is to learn by example. It just means copying what other people have done. To understand the basics of Caffe, I studied and replicated an example of digits classification on the MNIST dataset. Then I tried another example using AlexNet because I had read the paper and was familiar with the architecture. I was happy with the progress I had made so far, and I was feeling confident with these small wins. It was time to apply my working knowledge to the data I had.

Do the dirty work

Copying and replicating existing results are just the necessary first steps. They do not make you an expert. I was quickly reminded of this when I started working on my own dataset. The training error was not converging.

I knew there were knobs I could turn to make things work, but I did not understand what they did. So I spent a lot of time tweaking parameters that didn't matter. However, this step of trial and error is an extremely important one. Spending time tweaking the parameters and systematically observing their behavior builds valuable experience and expertise. I did the tedious work, read online articles, and took copious notes. A few hours later, I was still no close to debugging my problem, but I now had a very good understanding of the

debugging my problem, but I now had a very good understanding of the parameters I was trying to optimize.

Pick your tools

I am a visual learner. When I see things, I understand them better. I researched and found a great visualization tool called Nvidia DIGITS. I could now easily see the plots. The right visualization tools can bring unexpected insights especially when you are a beginner. Unfortunately, even visualizing the numbers did not help on that day.

Ask for help but come prepared (respect people's time)

I was so close to a solution and yet so far!

I needed help and there is no shame in getting help as long as you respect other people's time. Most people will help you ONCE because we all feel good about helping others; you gotta make the best of that one opportunity.

A good friend of mine is a Deep Learning expert, and he was very kind to meet me for dinner to help me out. I was very respectful of his time. I came prepared and even rehearsed what I was going to say before he showed up. After hearing my short explanation, he took one look at my error plot and without even thinking for a second said,

"Dude, your learning rate is too high."

While tweaking my parameters, I had changed the learning rate to half and had not seen any improvement. On his suggestion, I changed the learning rate to 1/10th and started training. A few minutes later, I saw the training loss falling rapidly as **the accuracy shot up to 80%**.

I felt like I was drinking a cocktail of emotions -- hope, joy, gratitude, and relief with a hint of triumph.

Great artists steal

By the time I returned home from dinner, I was tired after this long stressful day. I could go to sleep and settle for the average results I had. It was certainly sufficient to land the project. The other option was to keep working until I hit it out of the park and produced spectacular results.

Our choices define who we are. When I went to bed at 3 AM, I had stolen ideas from many different papers and created my unique solution. The next day, I showed results using a live web-based demo with a beautiful graph showing an accuracy of above 95%.

To say that the client was impressed is an understatement. None of the other vendors came even close to what I showed them. However, I was asking for a price 50% higher than what other vendors were asking.

Who do you think they chose?

This project earned me \$58,972.55 of extra income. Today, that one skill brings in hundreds of thousands of dollars in revenue every year. I hope this does not sound like I am bragging; I merely want to convey the importance of investing in yourself, gaining skills quickly and applying them to solve real world problems.

Time is the coin of your life and you gotta do whatever is needed to short-circuit the learning process. Acquire knowledge with speed and intensity even if it means paying for resources, begging for help or getting up at 4 AM!

In retrospect, the line

"Dude, your learning rate is too high."

is the best compliment I have ever received and I wish your learning rate is high as well!

I hope this long email was useful. Stay tuned for an important email next week!

Satya



LearnOpenCv.com

[Twitter](#)

[Linkedin](#)