

## Course #2 Q&A:

**Q:** What happens if two layers are connected to the same output? In the class we were told it's an error.

- Aren't there parallel paths in neural networks?

**A:** Yes there can be but this example was not setup this way. You can review the GoogleNet network to see how they split data up. There are also other ways to parallelize your data between GPUS, this is discussed in Alex Krizhevsky's One-Weird-Trick paper.

**Q:** In the lab we saw that in LeNet a RELU layer is connected to the bottom of a layer which its bottom is connected also to another layer. Why isn't it an error?

**A:** The activation function is setup to be processed by the output from that layer. The next layer will take the output from the activation function and push that into the next layer.

**Q:** Changes that we implemented in Task3

- Why did the changes improve the classification?
- What is rectified linear activation function (RELU)?

**A:** This is main point of the class, to show that with changes to the network accuracy can be improved. There are many knobs to turn. If you review the Googlenet paper, the authors describe an interesting analysis on adding layers and the impacts this had on performance. The ReLU function is an interesting activation function that I first learned about when review Alex Krizhevsky's paper. The function has a slope of one if the input is greater than zeros and zero if it is less. Some results have shown that you can train faster with this activation function. There are other activation functions such as tanh ranging from  $[-1,1]$  and sigmoid function ranging from  $[0,1]$ .

**Q:** How to customize and visualize existing model (I always had to start from the beginning with LeNet in order to get to the window with the "customize" option).

**A:** After creating a model, you can go to pretrained models. You will see all of your previously created networks listed. If you select one, the option to customize it will appear. Then you can build on your previous network configurations.

**Q:** So if my images have various different sizes, does DIGITS resize them to the intended size (e.g. 256x256), or should I do it beforehand? If it does, how are non-square images treated? (stretched/cropped/padded?)

**A:** DIGITS can handle the resizing for you when you import your dataset. It offers a variety of options for how non-square images are treated.

**Q:** What's the difference between cuDNN and Caffe ?

**A:** Caffe is a general deep learning framework for building and training DNNs. It can transparently leverage GPU acceleration through CUDA, but can also run on CPU. cuDNN is a CUDA library which accelerates specific functions important to DNN training. Caffe can optionally leverage cuDNN to further speed up training.

**Q:** How is the performance of deep learning with imbalanced data?

**A:** Comparable to any other machine learning technique - performance is better if your data is balanced. There are a number of techniques for balancing datasets across classes, e.g. data augmentation/inflation for the minority classes.

**Q:** I still don't have NVIDIA GPU available, but I do have a dataset I'd like to experiment with. Is there such a thing as DIGITS-as-a-service? Or is my best option to get an instance in Amazon AWS, and install it there?

**A:** Amazon AWS is a great option. The hands-on lab that we offer today alongside this webinar is running on AWS.

**Q:** Can I use Digits and the other frameworks without GPUs?

**A:** Yes you can. All of the frameworks support training and deploying DNNs on CPU. The lack of parallelism will greatly limit the size of network you can train in a reasonable amount of time though.

**Q:** Does Digits benefit from CUDNN?

**A:** Yes it does. DIGITS uses the Caffe framework for DNN training and Caffe can leverage cuDNN for accelerating DNN functions.

**Q:** Can I use digits on Linux?

**A:** Yes, DIGITS is officially support on Ubuntu 14.04 only. People have had success in compiling and installing on other Linux distros and Mac OSX too.

**Q:** What about redhat based releases? Build from Source?

**A:** Yes, you would build Caffe from source and then install DIGITS.

**Q:** I still don't have NVIDIA GPU available, but I do have a dataset I'd like to experiment with. Is there such a thing as DIGITS-as-a-service? Or is my best option to get an instance in Amazon AWS, and install it there?

**A:** Amazon AWS is a great option. The hands-on lab that we offer today alongside this webinar is running on AWS.

**Q:** can we use Digits to distribute computation over a network (where each computer in the network has a GPU)?

**A:** Not today, DIGITS manages multiple GPUs in a single node. That's a great suggestion for the future though.

**Q:** Any recommendations for EC2 instance suitable for installing DIGITS?

**A:** g2.2xlarge and g2.8xlarge will both work great

**Q:** Can I use Digits for other image classification tasks like face recognition?

**A:** Yes, absolutely. Any image classification task where you can create a single folder per image category will work. Note that for something like face recognition, state of the art techniques may do pre-processing of the images (“normalizing” them in one way or another). Digits does not have a specialized/optimized workflow for face recognition.

**Q:** is there support for online data augmentation or does it need to be pre-saved to the database?

**A:** There is support for cropping. Other than that you need to create augmented training images in advance.

**Q:** Hello. Nvidia is planning to bring Digits to windows any time soon?

**A:** You can use DIGITS in windows, but it is not as easy as downloading the web installer. If you want to do this, you will need to build caffe, you need to build the caffe branch (<https://github.com/NVIDIA/caffe>) and instructions for doing this can be found here - <https://initialneil.wordpress.com/2015/01/11/build-caffe-in-windows-with-visual-studio-2013-cuda-6-5-opencv-2-4-9/>.

**Q:** If I build an image classifier with digits and then want to classify a large set ~30K images producing a file with the image name and class how would I do that.

**A:** You could use our classify many, and it will post the accuracy. If I wanted to classify that many images, I would be interested in knowing the total accuracy of all of those 30K images and maybe even the accuracy of each category. In these instances I would create a script in python or C++ for more speed.

**Q:** Any plans for a Digits like front end for non-image data?

**A:** There is a lot of interest in this. We can handle other types of data if it can be put into a linear array [1x####].

**Q:** It's possible to achieve good results with a mid-end gpu using digits?

**A:** Yes. Especially if you are getting started, a mid-range GPU, or the GPUs available on Amazon cloud, will be totally fine. If you want to try to win a Kaggle competition, you probably need a high-end GPU.

**Q:** Can digits work on a cluster? I have two GPUs on different machines. If I create a cluster out of them, can digits utilise the two GPU's?  
**Q:** Can digits work on a cluster? I have two GPUs on different machines. If I create a cluster out of them, can digits utilise the two GPU's?

**A:** Yes, DIGITS can utilize two GPUs. Recall that DIGITS is built on top of 3rd party frameworks so provided those frameworks can use two GPUs, then DIGITS can also.

**Q:** I'd like to experiment with the OverFeat CNN model in DIGITS? How can I do that? Do you have any plans of maintaining a repo of popular models we can easily add to digits?

**A:** DIGITS currently has two popular networks already built in, AlexNet and LeNet. No plans of having an external repo of networks however we may choose to add more built-in networks in the future.

**Q:** When you use DIGITS, is it necessary to install GPU hardware on your machine ?

**A:** Not strictly no. DIGITS relies on the third-party frameworks for the back-end. Currently Caffe is the third-party framework used underneath DIGITS so the functionality of DIGITS directly mirrors Caffe.

**Q:** Could DIGITS work with video or sequences or images

**A:** DIGITS can work with the datatypes that the underlying frameworks can handle. Currently Caffe is the underlying framework so whatever Caffe can work with, DIGITS can work with.

**Q:** Is an amazon EC2 AWS image available that comprises DIGITS ?

**A:** NVIDIA is considering providing an official AWS image but at this time there is NOT an NVIDIA-supported image. However there are a number of other folks who have done this and published either their AMI or blog with info how to do this. It is not so difficult.

**Q:** I have an existing GTX960. For multiGPU setup Can I use a GTX980 on top of that or do I need to buy another GTX960?

**A:** You can mix and match them. They don't all need to be identical.

**Q:** Will Digits support theano in near future?

**A:** We are currently evaluating whether we will support Theano with DIGITS.

**Q:** Will Torch nn, cudnn be incorporated into Digits eventually? If so, roughly when?

**A:** We have not announced any timeline for this.

**Q:** What kind of license do you need to use Digits outside of NVIDIA?

**A:** DIGITS is open source and free for all to use. You can even contribute to its development and customize it yourself.

**Q:** Is there any benefit with SLI GPUs for Digits?

**A:** SLI is for graphics acceleration only, it doesn't have an impact for Digits. However, Digits now supports multi-GPU training, so you will get a good speedup if you run one training session over 2 GPUs versus one.

**Q:** If you are tapping your validation dataset at regular intervals, have you found it helpful to have an additional holdout dataset to judge model performance?

**A:** You are correct, it is best to have separate Test and Validation data sets. Digits supports carving out images from your main data sets for separate Test and Validation sets. I think many

folks are just skipping a final “Test” set and just using Validation, however. But that is probably not the “best practice”.

**Q:** Do you have any recommendations for GPU cloud service?

**A:** Both Amazon and Softlayer have GPU enabled instances available today. With Amazon you get their usual “by the minute” flexibility. With Softlayer, I think you have to commit to a month at a time, but they will provide more advanced GPUs on machines that are connected by Infiniband, for example...so Softlayer has a more powerful offering if you need the most powerful GPU or “HPC” style deployment. Both work great for Digits and Deep Learning.

**Q:** Can you point me to where there might be an example of how a Caffe model can be incorporated into Python and C++ code?

**A:** Yes, here are the links:

[http://caffe.berkeleyvision.org/gathered/examples/cpp\\_classification.html](http://caffe.berkeleyvision.org/gathered/examples/cpp_classification.html)

<http://caffe.berkeleyvision.org/tutorial/interfaces.html>

**Q:** What frameworks are used to set up on premise GPU platforms for DIGITS? Any documentation or best practices to accomplish this?

**A:** The easiest way to install Digits is to use the package from the NVIDIA Digits download page. This package includes the NVIDIA branch of Caffe and all the required dependencies. First step is become a “CUDA Registered Developer”. See: <https://developer.nvidia.com/digits>

**Q:** Can the input data be multi-spectral images with more than 3 channels, or imagery plus non-image meta data?

**A:** Digits does not support more than 3 color channels today. If this is valuable to you, please submit the request in the Digits discussion group. If we get a lot of requests for a feature it definitely raises its priority.

**Q:** I own a Titan X. I read somewhere that its single-precision performance (FP32) is 7 TFLOPS and double-precision performance (FP64) is only 1.3 TFLOPS. Do the frameworks discussed here all use single-precision by default? If not, how can they be configured for best performance?

**A:** By default, all the frameworks use single precision floating point.

**Q:** How is the number of GPUS set in DIGITS?

**A:** The number of GPUs to use is set on the train model page

**Q:** Is it possible to train voice datas with using NVIDIA DIGITS?

**A:** Currently DIGITS is designed for training on images, but we would like to add support for speech/voice in the future

**Q:** Although Allison did provide an initial answer that it should work. I'd like to hear a bit more about acceptable "image" sizes for digits. For example using it on one dimensional "images" or if there are certain sizes of images that work better or worse

**A:** all images will get resized and re-shaped to square during the process of creating the database. You can select the resizing size with a tradeoff between accuracy for larger resolutions vs speed for smaller resolutions

**Q:** What is advantage of gpu vs multicore amazon service?

**A:** With GPU acceleration, training will run about 8x faster than on CPU

**Q:** Can I extract cnn features using own dataset of natural scenery?

**A:** Yes, as you train, the process will extract features that are representative of the dataset that you train on.

**Q:** Mentioning about creating own network, is the own custom network restricted only image classification, or can be anything?

**A:** Currently DIGITS only supports image classification, but other data modalities are planned in future releases

**Q:** AlexNet vs LeNet, what is better?

**A:** AlexNet is designed for classifying images, while LeNet is an older, smaller network designed for handwritten digit classification

**Q:** How big of dataset can be trained with DIGITS? Because DIGITS load images into database, I am afraid DIGITS cannot deal with big dataset.

**A:** DIGITS pre-builds a database which is then passed into training in batches, whose size you can adjust for performance

**Q:** Regarding video, imagine I extracted the optical flow of the video into a texture, would it be possible to have varying sizes of the opticalflow fields still be processed properly or should I resize the smallest OF texture to match the largest?

**A:** When you create the dataset from images, you set the target size to resize to, and all the images will be resized to this value.

**Q:** I'm trying to do classification of heart beat signals. Does DIGITS support that?

**A:** If there is a way to format the signals into 2D images, then you definitely could use DIGITS to train a CNN on them

**Q:** My Titan X has 12G, I trained standard LeNet model on default MNIST dataset provided in DIGITS. It took almost all memory. Is it possible? It seems strange to me because LeNet is small and MNIST is not big.

**A:** That does seem like excessive memory usage for LeNet on MNIST

**Q:** When subsampling the the images after the convolution step, can we subsample by fractional number. E. g. 28x28 image to be converted to 20x20 image

**A:** The pooling layers use integer stride in the image sub-sampling, so fractional sub-sampling is not possible

**Q:** When DIGITS is installed is Caffe installed as well?

**A:** Yes, the NVIDIA branch of Caffe that supports multi-GPU training is installed with DIGITS when you download it from [developer.nvidia.com](http://developer.nvidia.com)

**Q:** Can I configure recurrent network and train it?

**A:** Currently only convolutional neural networks are supported by DIGITS, and not recurrent networks

**Q:** I have 1 GPU, and I am planning to buy a second GPU. Do I need to buy SLI to make DIGITS works on 2 GPUs?

**A:** You do not need SLI for multiple GPUs to be used in training. The CUDA devices are enumerated and used to train in a data parallel manner when you do multi-GPU training

**Q:** Can digits get the responses of a particular layer of a pretrained model to images?

**A:** Yes, once a network is trained, you can use DIGITS to load and classify images, and it provides a visualization of the activations at the various layers in the CNN

**Q:** Can you mix and match GeForce and Quadro ... or will driver differences get on the way?

**A:** It's not recommended to mix GeForce and Quadro due to driver incompatibilities

**Q:** What are the "stride" and "pad" parameters in convolution and pool layers?

**A:** Stride controls the size of the steps taken as the convolution moves across the image. Pad adds zeros around the edge of the image being convolved.

**Q:** Can I deploy a trained NN model on an embedded system?

**A:** Yes, we can run cuDNN, and do classification on our Jetson TK1 and TX1 platforms

**Q:** Will the model we make on digits work on Nvidia's fork of Caffe or will it work with vanilla caffe too?

**A:** It will work in the main branch of Caffe. Nvidia's fork uses the same formats and layer types.

**Q:** Is Digits for classification only? How about detection and segmentation?

**A:** Currently DIGITS is for image classification only.

**Q:** Does DIGITS include data augmentation?

**A:** DIGITS supports image cropping and resizing as pre-processing. Currently DIGITS does not support randomized data augmentation such as rotations, jitter etc. You can very easily

augment your images before creating your DIGITS database using a command line tool such as ImageMagick

**Q:** Hi. How we can use CNN for video classification?

**A:** DIGITS does not currently support classification of video. However, you could pre-process your video into frames outside of DIGITS using a tool like openCV. DIGITS could then train on individual frames and produce a model to classify them.

**Q:** What is the caffe binary that digits is looking for when the server is launched?

**A:** DIGITS looks for the caffe.bin binary in \$CAFFE\_ROOT/build/tools - this is the main Caffe command line interface for training networks.

**Q:** Can you walk us through building DIGITS on top of Windows 7 or 10?

**A:** Currently DIGITS is not supported on Windows.

**Q:** How we can use the trained model? Is it open and can be utilised in other environments? What are the common APIs for utilizing the trained models?

**A:** DIGITS produces a binary caffe\_model file and an associated deploy.prototxt file. Together you can use these within Caffe for deploying the model. This can be done through the Caffe command line, C++, Python or MATLAB interfaces.

**Q:** Can digits incrementally update a model after it is initially created?

**A:** Yes, DIGITS can fine-tune a pre-trained model. This can either be incremental update using the same data or fine-tuning for an entirely new dataset. You can also fix some layers and then modify and fine-tune other layers.

**Q:** What will happen if my model is larger than the GPU memory?

**A:** DIGITS requires that your model parameters and at least one training image can fit in GPU memory to use GPU training - otherwise, you will run out of GPU memory. Typically, this is not a problem as model parameters are in the 100s of MBs.

**Q:** Does digits have pretrained models?

**A:** DIGITS has the prototxt files for some standard architectures, e.g. AlexNet. The corresponding caffe\_model binary files can be downloaded from the Caffe Model Zoo: <https://github.com/BVLC/caffe/wiki/Model-Zoo>

**Q:** Solver sgd vs adagrad vs nag, when what to use

**A:** This is application dependent, although usually with appropriate choices of learning parameters you will achieve faster training with adagrad or nag than vanilla sgd

**Q:** Can you write code in C++ to exercise a network on a large number of images

**A:** Yes, but you would not be using DIGITS, you would use the Caffe C++ API directly



**Q:** How can we calculate the memory that is required to train a neural network?

**A:** You can find a discussion of this and simple formula here:

<https://timdettmers.wordpress.com/2014/08/14/which-gpu-for-deep-learning/>

**Q:** If I arrange my medical data as a 2D matrix ( per patient) , can I use DIGITS?

**A:** You could do that assuming you are attempting to solve a classification problem. CNNs are probably not applicable though, but MLPs might be.

**Q:** Can the DIGITS API be accessed from C++

**A:** The DIGITS API is REST, so yes you could.

**Q:** What is the minimum number of pictures representing a category in a training dataset to achieve satisfactory classification results in Digits?

**A:** This is completely application dependent. If you only have a small number of training samples there are ways to artificially increase the number of training instances to improve performance. If you use a pre-trained network and fine-tune for your new class you can sometimes get good performance with just a handful of examples.

**Q:** Is it possible to use DIGITS as an Optical character recognition (OCR) ?

**A:** Yes, that is what you are doing with the MNIST data in the lab.

**Q:** May you recommend intro link to nvidia NN? video?

**A:** <https://developer.nvidia.com/deep-learning-courses> - Class #1 - Introduction to Deep Learning

**Q:** Before using DIGITS is it necessary to install GPU hardware?

**A:** No, you can just use CPU. You will not be able to train large networks in a reasonable amount of time though.

**Q:** Imagine you have images that are hard to classify by a physician as either malignant or benign. Is it possible to classify them employing deep learning techniques? How successful is deep learning in such cases?

**A:** Deep Learning has had a number of successes in classifying medical images. For example, there was a recent Kaggle contest on Diabetic Retinopathy detection where the winning teams all used convolutional neural networks: <https://www.kaggle.com/c/diabetic-retinopathy-detection>  
Here is a link to the cell mitosis paper -<http://people.idsia.ch/~cirezan/data/miccai2013.pdf>

**Q:** Can I run Cuda and Digits in a Macbook pro, It will use GPU power ?

**A:** Yes you can. You need to compile Caffe from source before installing DIGITS. You can find instructions here: [http://caffe.berkeleyvision.org/install\\_osx.html](http://caffe.berkeleyvision.org/install_osx.html),  
<https://github.com/NVIDIA/DIGITS>

**Q:** Can you guide us on reading materials for Caffe framework and DIGITS, could not see links for Caffe! ?

**A:** In two weeks time we will have a webinar introducing Caffe and an associated hands-on lab. In the meantime, you can watch this great presentation from our GTC 2014 conference on Caffe given by one of the developers:

<http://on-demand.gputechconf.com/gtc/2014/webinar/gtc-express-deep-learning-caffee-ewan-shelhamer.mp4>

**Q:** Can I use both CPU and GPU at same time for training?

**A:** The CPU will always be used to some extent during training, but you cannot currently parallelise training a single network across both CPU and GPU. You can parallelize across multiple GPUs though.

**Q:** How does digits benefits from parallelism?

**A:** You can train multiple models in parallel by training one on each available GPU or you can parallelize the training of a single model across multiple GPUs for a speed up in training.

**Q:** if I wanted to use a model i trained with digits on a jetson, could i use main branch Caffe or I have to use NVIDIA'S fork?

**A:** You can deploy to the main branch

**Q:** Does DIGITS have an API available?

**A:** Yes, DIGITS has a REST API: <https://github.com/NVIDIA/DIGITS/blob/master/docs/API.md>

**Q:** Are the convolution filters randomly choosen and then trained?. What parameters are updated during training of DNN.

**A:** Yes, the filters are randomly initialized. The weights that connect the neurons in those filters to other layers in the network are trained.

**Q:** Is it possible to access to the codes DIGITS is using and modify the codes?

**A:** DIGITS is on github here: <https://github.com/NVIDIA/DIGITS>

**Q:** Is it possible to extract features automatically

**A:** Caffe has a tool called "extract\_features" in \$CAFFE\_ROOT/build/tools that can do that for you.

**Q:** So far, DIGITS seems mostly like a GUI tool that helps you launch computation and gather results. Is there a use for DIGITS if you have to run on non-interactive computer servers such as a cluster ?

**A:** You are correct that DIGITS is a GUI for Caffe based DNN training. It runs as a lightweight web server though with the interface being rendered in a browser, so it can be run on a remote cluster via SSH with appropriate port forwarding.

**Q:** How is multi-GPU supported? Is that in Caffe already?

**A:** DIGITS uses an NVIDIA fork of Caffe which supports multi-GPU. Multi-GPU support is not currently in the caffe master branch.

**Q:** DIGITS can recognize and run on a consumer GPU, say a GeForce GTX980?

**A:** Yes, DIGITS can run on consumer GPUs such as GeForce GTX. Many successful networks have been trained on GeForce GPUs. For deployment of DNNs we recommend Tesla GPUs.

**Q:** How do you determine the GPU capabilities number?

**A:** <https://developer.nvidia.com/cuda-gpus>

**Q:** Is DIGITS exclusively for training on images? Does it work on images with more than 3 color components?

**A:** Not at this time.

**Q:** Is DIGITS free for commercial use?

**A:** DIGITS is open source and the license details can be found here <https://github.com/NVIDIA/DIGITS/blob/master/LICENSE>

**Q:** Is DIGITS designed specifically for image classification, or can I use it with different types of data? Thanks!

**A:** We can handle non-rectangular input data. If you can format your data into a linear array it can be ingested and used with DIGITS.

**Q:** This is for classification only? Or can DIGITS do more?

**A:** Classification is supported now. We are expanding to generic inference if you are interested, check out this pull request <https://github.com/NVIDIA/DIGITS/pull/189>

**Q:** Does digits use all the caffe layers?

**A:** Yes.

**Q:** What is 7th layer of CNN features. How to give this feature as input in digit

**A:** You can add layers to the network configuration using the custom network tab.

**Q:** Hi, is it Digits suitable for 3D object classification of volumetric data, too?

**A:** You cannot do that right now.

**Q:** How flexible is Digits with creating your own neural network? For example, I want to add an RBM+CNN+DNN and then output my prediction.

**A:** DIGITS handles CNNs right now. You can download your trained network and use it with others if you want to perform inference.

**Q:** Does Digits necessarily needs a GPU? Can I install it without having a GPU?

**A:** Yes, you can using it without a GPU and use CPU. You need to build caffe with CPU\_ONLY flag.

**Q:** Can DIGITS be used for generic non-image data? (such as numerical and categorical features). Also, can it be used along with Theano?

**A:** DIGITS can be used with non-rectangular input data with a single or 3 bands. It is not working directly with Theano at this point.

**Q:** How do I know I have the needed GPUs?

**A:** You can use any GPU. If you want to use cuDNN you need a GPU with compute capabilities >3.0.

**Q:** Can a trained DNN model run on a k1 gpu e.g jetson tk1?for example a workflow of training with a tesla gpu and deploy on tegra is valid?

**A:** Yes you can classify with the Jeston TK1.

**Q:** Does DIGITS work on GPUs by other manufacturers?

**A:** No, DIGITS works with CUDA. You can run things on the CPU though.

**Q:** Is Fedora 21 instead of Ubuntu4.04 works for DIGITS?

**A:** While we haven't tested all linux distributions, installing Digits on Fedora has been reported to work.

**Q:** Does Caffe already support learning in multi-task framework and with input data being from 2 sources like audio and images, where each of them live on different lattices images on 2D and audio on time. I did not find examples on this, though I see a lot of such models in DL presentations.

**A:** Not right now

**Q:** If we run linux in a virtual machine on Windows will DIGITS benefit from the hosts GPU?

**A:** If you're virtual machine can access the GPU (and run cuda code for example), yes.

**Q:** Can you do multi label classification with digits?

**A:** Yes, in the classification tab.

**Q:** Can digits be used for non-image data?

**A:** yes, but you need to be able to enter your data as an array, like 1, by ###

**Q:** Can I install/run DIGITS on Linux openSuze ?

**A:** Yes. While we haven't tested all linux distributions, installing Digits on OpenSuse can to work.

**Q:** where can i downloaded images from? i run caffe on my local mchine

**A:** Try ImageNet or deeplearning.net.

**Q:** If we are using AWS GPUs, are there directions anywhere that could help in using the DIGITS GUI while connected to AWS?

**A:** Tunnel the port when you log in or leave the port you want to use open.

**Q:** Is there a way to do unsupervised (or semi-supervised) with DIGITS ?

**A:** Not at this time.

**Q:** Does it support a lower version of ubuntu?

**A:** Installing on previous versions of ubuntu will require you to install all the dependencies. It is possible but difficult. We strongly advise using Ubuntu 14.04.

**Q:** Can Digits combine power from both Nvidia and AMD graphics cards ? perhaps via OpenCL , etc

**A:** Unfortunately no. But you can run DIGITS in CPU mode.

**Q:** what is max number of GPUs maybe used?

**A:** The GPUs need to be in the same physical machine. Up to 8 right now.

**Q:** Can DIGITS run on older GPU arrays like S1070 without cuDNN?

**A:** Yes, usually the only gotcha is memory. You start to reach memory limits when there is less.

**Q:** Can I use DIGITS for task different than classification? What about segmentation or regressions?

**A:** Not yet. We are gathering feedback right now and will expand DIGITS functionalities in the future.

**Q:** Are there any problems with using Digits in a virtualized environment?

**A:** Digits work well in a virtualized environment. The GPU has to be accessible to your virtual environment for GPU acceleration to work.

**Q:** Can Digits train on a new dataset which is not included in the package?

**A:** Yes. The dataset has to be set up in the dataset tab.

**Q:** What are the main advantages of DIGITS over using directly Caffe?

**A:** DIGITS provides an easy to use interface for creating datasets and configuring and training models. It also provides visualization of the network, training progress, and classification.

**Q:** How easy to modify a network architecture in DIGITS and Caffe? Where do I need to modify in Caffe or DIGITS

**A:** You can modify the network architecture directly in DIGITS by editing the text in the customize window in the model page

**Q:** Would digits would with alex krizhevsky's network?

**A:** Yes. Alexnet is available in DIGITS.

**Q:** Do we have to use images? I have NLP data set I want to use.

**A:** Not yet.

**Q:** Can I use caffe + DIGITS with 2 GPUs installed on 2 different nodes of a cluster?

**A:** You will have to have to separate installs of DIGITS (one on each node).

**Q:** Does training epochs = 30 mean to process all training data 30 times?

**A:** Yes

**Q:** Why would we use something as theano if we have this high level tool? In other words, does theano do something we cannot do with DIGITS?

**A:** Theano is a more general symbolic math library. So it offers great expressivity, but it is more complex to use than DIGITS for CNN training on images. In four weeks time we will be having a webinar and hands-on lab all about Theano.

**Q:** Does Digits support X-forwarding over SSH so that it could be installed on a Linux server but interacted with over ssh on a Windows machine?

**A:** Yes. But it is often not needed as DIGITS is provided as a web server (using nginx), once running, you can connect from any web browser (on any device, including your phone ;)).

**Q:** Is DIGITS just an interface for Caffe (as user interface), or it works independent of Caffe?

**A:** DIGITS requires the NVIDIA fork of Caffe.

**Q:** Is there a single installer for DIGITS which will install it with all the dependencies? As per the instructions in GIT, it seems I'd hv to install all the dependencies separately. What does the web installer (mentioned in the slides) do?

**A:** If you are using Ubuntu 14.04, the web installer will install the dependencies for you. For other Linux distros you will need to build Caffe from source, so the web installer is not appropriate.

**Q:** Once I've trained a DNN, is it possible to run it on an Nvidia Jetson?

**A:** Trained models can be downloaded as caffe model files. That means you can deploy them to Caffe installations elsewhere, such as on the Jetson TK1

**Q:** Can DIGITS take a video input? like i could write numbers on a white board and point a camera at them?

**A:** Not today, but that is a great suggestion for future functionality

**Q:** Is it possible to import a trained caffe model into digit (e.g. for troubleshooting purposes)

**A:** Yes the classification tab offers a pretrained model option.

**Q:** Will "Nvidia Caffe" interfere with "original Caffe" when both are installed?

**A:** You can have both installed simultaneously.

**Q:** Can I use DIGITS in my code and avoid GUI based training and testing

**A:** Caffe can easily be incorporated into C++, Python and MATLAB code.

**Q:** Where can I find the syntax specs for defining custom layers and parameters?

**A:** <http://caffe.berkeleyvision.org/tutorial/layers.html>

**Q:** Is this regular Caffe syntax? What are Nvidia-specific additions (such as for parallelism)?

**A:** Yes, the syntax in the prototxt files is the same as Caffe. NVIDIA additions are for additional capability such as multi-GPU parallelism.

**Q:** Could I connect a data layer or any other layer to two or more layers simultaneously?

**A:** Yes you can, you can define those kind of network connections in the model prototxt file

**Q:** Do you get a factor of 2 improvement by having 2 GPUs?

**A:** Not quite as there is a communication overhead introduced when you parallelize across multiple GPUs. You can find more info here and see some performance results:

<https://developer.nvidia.com/digits>

**Q:** Can digits be used to fine-tune a pretrained network, for example replacing the number of outputs in the output layer to adapt it to a different dataset?

**A:** Yes it can. You can import a pre-trained caffe model and fine-tune.

**Q:** Is DIGITS usable through command prompt?

**A:** For command line use, you would typically use Caffe directly.