

1: TensorFlow on AWS - Deep Learning on the Cloud

Cognitive Class Deep Learning with TensorFlow. The majority of data in the world is unlabeled and unstructured. Shallow neural networks cannot easily capture relevant structure in, for instance, images, sound, and textual data.

Next, we drop all the rows storing NaN values by using the `dropna` function. Next, we make all data in `data` a `np`. Splitting the dataset Next, we split the whole dataset into training and test data. The data was not shuffled but sequentially sliced. This process makes the mean of all the input features equal to zero and also converts their variance to 1. This ensures that there is no bias while training the model due to the different scales of all input features. If this is not done the neural network might get confused and give a higher weight to those features which have a higher average value than others. Nowadays, rectified linear unit ReLU activations are commonly used activations which are unbounded on the axis of possible activation values. However, we will scale both the inputs and targets. We instantiate the variable `sc` with the `MinMaxScaler` function. The whole dataset will look like below. The test dataset is used to see how the model will perform on new data which would be fed into the model. The test dataset also has the actual value for the output, which helps us in understanding how efficient the model is. Now that the datasets are ready, we may proceed with building the Artificial Neural Network using the TensorFlow library. The input layer, the hidden layers and the output layer. The input layer is the 8 features. After input layers there are 3 hidden layers. The first hidden layer contains neurons. Subsequent hidden layers are always half the size of the previous layer, which means 2nd hidden layers contains and finally 3rd one neurons. After 3 hidden layers there is output layer. Net is creating by tensor flow interactive secession. We need two placeholders in order to fit our model: It is crucial to understand which input and output dimensions the neural net needs in order to design it properly. Besides placeholders, variables are another cornerstone of the TensorFlow universe. While placeholders are used to store input and target data in the graph, variables are used as flexible containers within the graph that are allowed to change during graph execution. Here Weights and Biases are represented as variables in order to adapt during training. Variables need to be initialized, prior to model training. Since neural networks are trained using numerical optimization techniques, the starting point of the optimization problem is one the key factors to find good solutions to the underlying problem. There are different initializers available in Tensor flow, each with different initialization approaches. Designing the network architecture In designing the network architecture, 1st we need to understand the required variable dimensions between input, hidden and output layers. In case of multilayer perceptron MLP , the network type we use here, the second dimension of the previous layer is the first dimension in the current layer for weight matrices. It means each layer passing its output as input to the next layer. After definition of the required weight and bias variables, the network topology, the architecture of the network, needs to be specified. Hereby, placeholders data and variables weights and biases need to be combined into a system of sequential matrix multiplications. Furthermore, the hidden layers of the network are transformed by activation functions. Activation functions are important elements of the network architecture since they introduce non-linearity to the system. There are many possible activation functions out there, one of the most common is the rectified linear unit ReLU which will be using in this model. Cost function We use cost function to optimize the model. For regression problems, the mean squared error MSE function is commonly used. MSE computes the average squared deviation between predictions and targets. The development of stable and speedy optimizers is a major field in neural network and deep learning research. Fitting the neural network Now we need to fit the neural network that we have created to our train datasets. After having defined the placeholders, variables, initializers, cost functions and optimizers of the network, the model needs to be trained. Usually, this is done by mini batch training. They store the input and target data and present them to the network as inputs and targets. Afterwards, TensorFlow conducts an optimization step and updates the networks parameters, corresponding to the selected learning scheme. After having updated the weights and biases, the next batch is sampled and the process repeats itself. The procedure continues until all batches have been presented to the network. One full sweep over all batches is called an epoch. The training of the network stops once the maximum number of epochs is reached or

another stopping criterion defined by the user applies. We stop the training network when epoch reaches. With this, our artificial neural network has been compiled and is ready to make predictions. **Predicting The Movement Of The Stock Now** that the neural network has been compiled, we can use the predict method for making the prediction. This is done by slicing the dataframe using the iloc method as shown in the code below. **Computing Strategy Returns** Now that we have the predicted values of the stock movement. We can compute the returns of the strategy. We will be taking a long position when the predicted value of y is true and will take a short position when the predicted signal is False. We first compute the returns that the strategy will earn if a long position is taken at the end of today, and squared off at the end of the next day. We use the decimal notation to indicate that floating point values will be stored in this new column. Next, we store in it the log returns of today, i. Next, we will compute the Strategy Returns. By using the np. We now compute the cumulative returns for both the market and the strategy. These values are computed using the cumsum function. **Plotting The Graph Of Returns** We will now plot the market returns and our strategy returns to visualize how our strategy is performing against the market. We then create the legend and show the plot using the legend and show functions respectively. The plot shown below is the output of the code. The green line represents the returns generated using the strategy and the red line represents the market returns. **Conclusion** The objective of this project is to make you understand how to build an artificial neural network using tensorflow in python and predicting stock price. The objective is not to show you to get a good return. You can optimize this model in various ways to get a good strategy return. My advice is to use more than , data points when you are building Artificial Neural Network or any other Deep Learning model that will be most effective. This model was developed on daily prices to make you understand how to build the model. It is advisable to use the minute or tick data for training the model. Now you can build your own Artificial Neural Network in Python and start trading using the power and intelligence of your machines. All investments and trading in the stock market involve risk. Any decisions to place trades in the financial markets, including trading in stock or options or other financial instruments is a personal decision that should only be made after thorough research, including a personal risk and financial assessment and the engagement of professional assistance to the extent you believe necessary. The trading strategies or related information mentioned in this article is for informational purposes only. Existing Users Log In.

2: Intro to Deep Learning With TensorFlow - Starweaver

TensorFlow is an open source software library for high performance numerical computation. Its flexible architecture allows easy deployment of computation across a variety of platforms (CPUs, GPUs, TPUs), and from desktops to clusters of servers to mobile and edge devices.

Choosing a Deep Learning Framework in However, on a Thursday evening last year, my friend was very frustrated and disappointed. Yoshua Bengio had decided to stop the support for the framework. Hence, the choice of the framework you decide to spend your time learning and practicing is very important. So, when I got a few emails from some of our readers about the choice of Deep learning framework mostly Tensorflow vs Pytorch , I decided to write a detailed blog post on the choice of Deep Learning framework in If you are getting started on deep learning in , here is a detailed comparison of which deep learning library should you choose in PyTorch is one of the newest deep learning framework which is gaining popularity due to its simplicity and ease of use. Pytorch got very popular for its dynamic computational graph and efficient memory usage. Dynamic graph is very suitable for certain use-cases like working with text. Pytorch is easy to learn and easy to code. For the lovers of oop programming, torch. Module allows for creating reusable code which is very developer friendly. Pytorch is great for rapid prototyping especially for small-scale or academic projects. Tensorflow, an open source Machine Learning library by Google is the most popular AI library at the moment based on the number of stars on GitHub and stack-overflow activity. It draws its popularity from its distributed training support, scalable production deployment options and support for various devices like Android. One of the most awesome and useful thing in Tensorflow is Tensorboard visualization. In general, during train, one has to have multiple runs to tune the hyperparameters or identify any potential data issues. Tensorflow Serving is another reason why Tensorflow is an absolute darling of the industry. In Tensorflow Serving, the models can be hot-swapped without bringing the service down which can be crucial reason for many business. In Tensorflow, the graph is static and you need to define the graph before running your model. Although, Tensorflow also introduced Eager execution to add the dynamic graph capability. In Tensorflow, entire graph with parameters can be saved as a protocol buffer which can then be deployed to non-pythonic infrastructure like Java which again makes it borderless and easy to deploy. Caffe is a Python deep learning library developed by Yangqing Jia at the University of Berkeley for supervised computer vision problems. It used to be the most popular deep learning library in use. The same goes for OpenCV, the widely used computer vision library which started adding support for Deep Learning models starting with Caffe. For years, OpenCV has been the most popular way to add computer vision capabilities to mobile devices. So, if you have a mobile app which runs openCV and you now want to deploy a Neural network based model, Caffe would be very convenient. Since we have limited experience with CNTK, we are just mentioning it here. Another framework supported by Facebook, built on the original Caffe was actually designed by Caffe creator Yangqing Jia. It was designed with expression, speed, and modularity in mind especially for production deployment which was never the goal for Pytorch. Recently, Caffe2 has been merged with Pytorch in order to provide production deployment capabilities to Pytorch but we have to wait and watch how this pans out. Promoted by Amazon, MxNet is also supported by Apache foundation. Torch also called Torch7 is a Lua based deep learning framework developed by Clement Farabet, Ronan Collobert and Koray Kavukcuoglu for research and development into deep learning algorithms. Torch has been used and has been further developed by the Facebook AI lab. Theano was a Python framework developed at the University of Montreal and run by Yoshua Bengio for research and development into state of the art deep learning algorithms. It used to be one of the most popular deep learning libraries. The official support of Theano ceased in In fact, almost every year a new framework has risen to a new height, leading to a lot of pain and re-skilling required for deep learning practitioners. The world of Deep Learning is very fragmented and evolving very fast. Look at this tweet by Karpathy: Imagine the pain all of us have been enduring, of learning a new framework every year. Keras is being hailed as the future of building neural networks. Here are some of the reasons for its popularity: Keras is designed to remove boilerplate code. Few lines of keras code will achieve so much more than native

Tensorflow code. Keras is an API which runs on top of a back-end. This back-end could be either Tensorflow or Theano. Currently, Keras is one of the fastest growing libraries for deep learning. The power of being able to run the same code with different back-end is a great reason for choosing Keras. Imagine, you read a paper which seems to be doing something so interesting that you want to try with your own dataset. Now, If the code is written in Keras all you have to do is change the back-end to Tensorflow. This will turbocharge collaborations for the whole community. So, you can train a network in Pytorch and deploy in Caffe2. So, that could be a good thing for the overall community. I would love if Tensorflow joins the alliance. That will be a force to reckon with. Comparison of AI frameworks.

3: Step by Step Tutorial: Deep Learning with TensorFlow in R | R-bloggers

In this codelab, you will learn how to build and train a neural network that recognises handwritten digits. Along the way, as you enhance your neural network to achieve 99% accuracy, you will also discover the tools of the trade that deep learning professionals use to train their models efficiently.

Continue reading for information about each of the different offerings. Learn more by reading this blog, [Scaling TensorFlow and Caffe to GPUs](#) Large model support, enabling you to use higher resolution data Tools for ease of development Reduce data preparation time by an order of magnitude, with upcoming tools Automated hyper-parameter tuning and optimization to make your models faster and more accurate PowerAI includes all necessary dependencies and removes the time, effort, and difficulty associated with getting a deep learning environment operational and performing optimally. Go here to learn more about the PowerAI releases including how to order the latest version. Deploy with PowerAI Enterprise For enterprises looking to rapidly scale their deep-learning applications, PowerAI Enterprise combines the PowerAI features above with additional functionality to optimize and speed up the completion of your training, testing, and validation. PowerAI Enterprise truly shines when you are looking to expand into distributed deep learning with more than four nodes. Go here to learn more about PowerAI Enterprise including how to register for free trial access. Analyze images and video with PowerAI Vision PowerAI Vision provides tools and interfaces for business analysts, subject matter experts, and developers without any skills in deep learning technologies to begin using deep learning. This enterprise-grade software provides a complete ecosystem to label raw data sets for training, creating, and deploying deep learning-based models. It can help train highly accurate models to classify images and detect objects in images and videos. The tools assist users to focus on rapidly identifying datasets and labeling them. They can then train and validate a model in a GUI interface to build customized solutions for image classification and object detection. Go here to learn more about PowerAI Vision including how to register for free trial access. The DDL capability is incorporated into the Deep Learning frameworks as an integrated binary, reducing complexity for clients as they bring in high-performance cluster capability. Because of this, PowerAI with DDL can scale jobs across large numbers of cluster resources with very little loss to communication overhead. Similar to `snap-ml-mpi`, the `snap-ml-spark` package offers distributed training of models across a cluster of machines. The library is exposed to the user via a spark. It takes a computational graph defined by users, and automatically adds swap-in and swap-out nodes for transferring tensors from GPUs to the host and vice versa. During training and inferencing this makes the graph execution operate like operating system memory paging. Ready your tools Follow these simple steps to get your PowerAI-based application development started. See the PowerAI release notes for more information. Each framework included in PowerAI is unique and selecting a preferred framework for your application is important. The integrated installer for PowerAI means you have everything installed and performant, so you can rapidly try examples in each and select suited to your preferences. Your data can be visual, audio, text, or beyond. Packages like TensorFlow in PowerAI incorporate tools to help make your training network design even easier. Ideas to get you started Here are some recommendations for how to add deep learning to your applications. The Enterprise Layer Deep Learning atop your existing data-store Tease out value from your existing data by applying deep learning as a technique for advanced analysis. Reshape or augment an existing business process Augment human insight or manual labor with machine intelligence. Use deep learning to train a visual or audio recognition system that helps guide decisions. Then run those high-likelihood simulations with greater precision. Sift through existing unstructured data or vast outputs of a simulation with Deep Learning and gain new insights rapidly. Find information about all of the components in PowerAI here: Frameworks are available built, installed, and configured ready to use immediately. No set up required. Your free trial comes with a container pre-configured with the necessary frameworks and libraries. Courses and learning paths Data science and cognitive computing courses Build your Deep Learning skills, for free, with this learning path from [cognitiveclass](#).

4: Deep Learning and PowerAI - Linux on Power Developer Portal

The high-level Keras API provides building blocks to create and train deep learning models. Start with these beginner-friendly notebook examples, then read the TensorFlow Keras guide.

Please enter a valid input. Submit The post has been successfully mailed. Email Post I have designed this TensorFlow tutorial for professionals and enthusiasts who are interested in applying Deep Learning Algorithm using TensorFlow to solve various problems. Following are the topics that will be discussed in this TensorFlow tutorial blog: In this TensorFlow tutorial, before talking about TensorFlow, let us first understand what are tensors. Tensors are nothing but a de facto for representing the data in deep learning. In general, Deep Learning you deal with high dimensional data sets where dimensions refer to different features present in the data set. TensorFlow is a library based on Python that provides different types of functionality for implementing Deep Learning Models. In TensorFlow, the term tensor refers to the representation of data as multi-dimensional array whereas the term flow refers to the series of operations that one performs on tensors as shown in the above image. Now we have covered enough background about TensorFlow. Next up, in this TensorFlow tutorial we will be discussing about TensorFlow code-basics. Code Basics Basically, the overall process of writing a TensorFlow program involves two steps: Building a Computational Graph Let me explain you the above two steps one by one: Building a Computational Graph So, what is a computational graph? Each nodes take 0 or more tensors as input and produces a tensor as output. Explanation of the Above Computational Graph: Running a Computational Graph Let us take the previous example of computational graph and understand how to execute it. Following is the code from previous example: A session encapsulates the control and state of the TensorFlow runtime i. Let me show you how to run the above computational graph within a session Explanation of each line of code has been added as a comment: Constants, Placeholder and Variables In TensorFlow, constants, placeholders and variables are used to represent different parameters of a deep learning model. For this kind of functionality, placeholders are used which allows your graph to take external inputs as parameters. Basically, a placeholder is a promise to provide a value later or during runtime. Let me give you an example to make things simpler: Placeholders are not initialized and contains no data. One must provides inputs or feeds to the placeholder which are considered during runtime. Executing a placeholder without input generates an error. In a nutshell, a variable allows you to add such parameters or node to the graph that are trainable i. Variables are defined by providing their initial value and type as shown below: Constants are initialized when you call tf. On the contrary, variables are not initialized when you call tf. To initialize all the variables in a TensorFlow program, you must explicitly call a special operation as shown below: Now that we have covered enough basics of TensorFlow, let us go ahead and understand how to implement a linear regression model using TensorFlow. Therefore, for creating a linear model, you need: Copy the code by clicking the button given below: Creating variable for parameter slope W with initial value as 0. Second thing that we need is to validate our trained model by comparing its output with the desired or target output based on given set of x values. Loss Function $\hat{\epsilon}$ Model Validation A loss function measures how far apart the current output of the model is from that of the desired or target output. Therefore, we need to adjust our weights W and bias b so as to reduce the error that we are receiving. The simplest optimizer is gradient descent. It modifies each variable according to the magnitude of the derivative of loss with respect to that variable.

5: Deep Learning with TensorFlow | PACKT Books

Deep Learning with TensorFlow LiveLessons is an introduction to Deep Learning that bring the revolutionary machine-learning approach to life with interactive demos from the most popular Deep Learning library, TensorFlow, and its high-level API, Keras. Essential theory is whiteboarded to provide an intuitive understanding of Deep Learning's.

In this codelab, you will learn how to build and train a neural network that recognises handwritten digits. This codelab uses the MNIST dataset, a collection of 60,000 labeled digits that has kept generations of PhDs busy for almost two decades. Install the necessary software on your computer: Python, TensorFlow and Matplotlib. Full installation instructions are given here: [The folder contains multiple files. Other files are either solutions or support code for loading the data and visualising results.](#) When you launch the initial python script, you should see a real-time visualisation of the training process: See instructions at the bottom of the file. The visualisation tool built for TensorFlow is TensorBoard. Its main goal is more ambitious than what we need here. It is built so that you can follow your distributed TensorFlow jobs on remote servers. For what we need in this lab matplotlib will do and we get real-time animations as a bonus. But if you do serious work with TensorFlow, make sure you check out TensorBoard. We will first watch a neural network being trained. The code is explained in the next section so you do not have to look at it now. Our neural network takes in handwritten digits and classifies them, *i.e.* It does so based on internal variables "weights" and "biases", explained later that need to have a correct value for the classification to work well. This "correct value" is learned through a training process, also explained in detail later. What you need to know for now is that the training loop looks like this: Here you see the training digits being fed into the training loop, at a time. You also see if the neural network, in its current state of training, has recognized them white background or mis-classified them red background with correct label in small print on the left side, bad computed label on the right of each digit. There are 50,000 training digits in this dataset. We feed of them into the training loop at each iteration so the system will have seen all the training digits once after iterations. We call this an "epoch". To test the quality of the recognition in real-world conditions, we must use digits that the system has NOT seen during training. Otherwise, it could learn all the training digits by heart and still fail at recognising an "8" that I just wrote. Here you see about of them with all the mis-recognised ones sorted at the top on a red background. The choice of a loss function here, "cross-entropy" is explained later. What you see here is that the loss goes down on both the training and the test data as the training progresses: It means the neural network is learning. The X-axis represents iterations through the learning loop. This is computed both on the training and the test set. You will see it go up if the training goes well. The final two graphs represent the spread of all the values taken by the internal variables, *i.e.* Here you see for example that biases started at 0 initially and ended up taking values spread roughly evenly between These graphs can be useful if the system does not converge well. If you see weights and biases spreading into the s or s , you might have a problem. The bands in the graphs are percentiles. Keyboard shortcuts for the visualisation GUI: How is the "cross-entropy" computed? How exactly does the training algorithm work? Jump to the next section to find out. Each "neuron" in a neural network does a weighted sum of all of its inputs, adds a constant called the "bias" and then feeds the result through some non-linear activation function. Here we design a 1-layer neural network with 10 output neurons since we want to classify digits into 10 classes 0 to 9. For a classification problem, an activation function that works well is softmax. Applying softmax on a vector is done by taking the exponential of each element and then normalising the vector using any norm, for example the ordinary euclidean length of the vector. Why is "softmax" called softmax? The exponential is a steeply increasing function. It will increase differences between the elements of the vector. It also quickly produces large values. Then, as you normalise the vector, the largest element, which dominates the norm, will be normalised to a value close to 1 while all the other elements will end up divided by a large value and normalised to something close to 0. The resulting vector clearly shows which was its largest element, the "max", but retains the original relative order of its values, hence the "soft". We will now summarise the behaviour of this single layer of neurons into a simple formula using a matrix multiply. Let us do so directly for a "mini-batch" of images as the input, producing

predictions element vectors as the output. Using the first column of weights in the weights matrix W , we compute the weighted sum of all the pixels of the first image. This sum corresponds to the first neuron. Using the second column of weights, we do the same for the second neuron and so on until the 10th neuron. We can then repeat the operation for the remaining 99 images. If we call X the matrix containing our images, all the weighted sums for our 10 neurons, computed on images are simply X . Each neuron must now add its bias a constant. Since we have 10 neurons, we have 10 bias constants. We will call this vector of 10 values b . It must be added to each line of the previously computed matrix. Using a bit of magic called "broadcasting" we will write this with a simple plus sign. It extends how normal operations work on matrices with incompatible dimensions. By the way, what is a "tensor"? A "tensor" is like a matrix but with an arbitrary number of dimensions. A 1-dimensional tensor is a vector. A 2-dimensions tensor is a matrix. And then you can have tensors with 3, 4, 5 or more dimensions. Now that our neural network produces predictions from input images, we need to measure how good they are, i . Remember that we have true labels for all the images in this dataset. Any distance would work, the ordinary euclidian distance is fine but for classification problems one distance, called the "cross-entropy" is more efficient. It is handy here because the format is very similar to how our neural network outputs its predictions, also as a vector of 10 values. Here is how it works. The cross-entropy is a function of weights, biases, pixels of the training image and its known label. If we compute the partial derivatives of the cross-entropy relatively to all the weights and all the biases we obtain a "gradient", computed for a given image, label and present value of weights and biases. Remember that we have weights and biases so computing the gradient sounds like a lot of work. Fortunately, TensorFlow will do it for us. The mathematical property of a gradient is that it points "up". Since we want to go where the cross-entropy is low, we go in the opposite direction. We update weights and biases by a fraction of the gradient and do the same thing again using the next batch of training images. Hopefully, this gets us to the bottom of the pit where the cross-entropy is minimal. In this picture, cross-entropy is represented as a function of 2 weights. In reality, there are many more. The gradient descent algorithm follows the path of steepest descent into a local minimum. The training images are changed at each iteration too so that we converge towards a local minimum that works for all images. It would be like trying to get to the bottom of a valley while wearing seven-league boots. You would be jumping from one side of the valley to the other. To get to the bottom, you need to do smaller steps, i . We call this fraction the "learning rate". To sum it up, here is how the training loop looks like: Doing so on examples gives a gradient that better represents the constraints imposed by different example images and is therefore likely to converge towards the solution faster. The size of the mini-batch is an adjustable parameter though. There is another, more technical reason: Frequently Asked Questions Why is the cross-entropy the right distance to use for classification problems? The code for the 1-layer neural network is already written. Your task in this section is to understand this starting code so that you can improve it later.

6: Deep Learning - Artificial Neural Network Using Tensorflow In Python

What is deep learning? Grow an understanding of the basic fundamentals of deep learning in our accredited deep learning online course. Gain proficiency in TensorFlow to implement, train, visualize, export, and deploy deep models from scratch, as well as utilize pre-trained sources.

I truly believe there are no boring projects. There are only boring executions! Hi everyone, welcome to this blog series about Tensorflow. Tensorflow is the most popular and apparently best Deep Learning Framework out there. TensorFlow is a framework created by Google for creating Deep Learning models. Machine Learning has enabled us to build complex applications with great accuracy. Whether it has to do with images, videos, text or even audio, Machine Learning can solve problems from a wide range. Tensorflow can be used to achieve all of these applications. The reason for its popularity is the ease with which developers can build and deploy applications. Moreover, Tensorflow was created with processing power limitations in mind. The library can be ran on computers of all kinds, even on smartphones yes, even on that overpriced thing with half an apple on it. Deep learning concept of Tensorflow But before learning Tensorflow, we have to understand a basic principle. The human brain consists of billions of neurons which are interconnected by synapses. This process is called thinking. To replicate that process on computers, we need machine learning and neural networks. Traditionally, we always got computers to do things by providing a strict set of instructions. Machine Learning uses a very different approach. Instead of giving the computer a set of instructions on how to do something, we give it instructions on how to learn to do something. Instead of manually finding unique characteristics from images of those animals and then coding it up, machine learning takes in images of those animals and finds characteristics and differences by itself. This process of teaching the computer is referred to as training. Deep learning is a Technique for implementing Machine Learning. It uses neural networks to learn, sometimes, using decision trees may also be referred to as deep learning, but for the most part deep learning involves the use of neural networks. So, what is a neural network? Every time you open a door, you become a different person. By the time you open the last door, you have become a very different person. Each door, in this case, represents a layer. A neural network, therefore, is a collection of layers that transform the input in some way to produce an output. If you want to know more, give the following article a try: Before following it, you might also want to take a look at the official installation guide. Python Tensorflow programs are written in Python, which you can download at <https://www.tensorflow.org/install>: Otherwise, you have to install that as well. Tensorflow The next step is to install Python. Open a command line as administrator!

7: Deep Learning and AI frameworks - Azure | Microsoft Docs

Kick-start Deep Learning with TensorFlow and Keras. This is a kick-start memo of how to run Deep Learning the 'fast and lean way'. This means that this page is showing how to quickly having a Keras example running.

8: TensorFlow Tutorial | Deep Learning Using TensorFlow | Edureka

Deep learning is the intersection of statistics, artificial intelligence, and data to build accurate models and TensorFlow is one of the newest and most comprehensive libraries for implementing deep learning.

9: Deep Learning with Tensorflow: Part 1 " theory and setup

TensorFlow is a framework created by Google for creating Deep Learning models. Deep Learning is a category of machine learning models (=algorithms) that use multi-layer neural networks. Machine Learning has enabled us to build complex applications with great accuracy.

Conventions of form and thought in early Greek epic poetry Translating into Predicate Reflections . path to prayer. Steven alter information systems V. 9. Quebec, 1636 Working Papers for Exercises and Problems to accompany Financial Accounting Visionary dreamer Gas chromatography mass spectrometry theory Introduction vii Christopher Columbus Infinite Abelian Groups Sharepoint 2010 basics for beginners Personalize relationships The crucifixion of liberty Satan attacks the church On allis chalmers owners manual for a grain drill When I Am Ten Years Old Brain Quest: Science TrapWise and Student CD Package Unconvicted prisoners in Australia The norton anthology of poetry fifth edition Drawn into the edges of myself Carol Brorsen Reckoning with Barth Can can sheet music piano Food ABC (The Colors We Eat) The short stories Sport in a philosophic context Sirk on Sirk (Directors on Directors) The God of This World to His Prophet How animals protect themselves Technological accidents The making of the dalit public in North India International child mental health Myron L. Belfer Residential Valuation Theory and Practice 3. The image of the Buddha Science, fact, and value, by M. Scriven. Advanced excel tutorial Invention financing and joint ownership African American Education Cat 2000 solved paper The age of polymorphous perversity : can civilization survive?