Deep Neural Networks

Actuarial Applications for emerging AI
CAS RPM
Prashant De

Agenda

1) Neural Networks - Introduction
2) What does “Deep” mean in Neural Network terminology
    a) Stochastic Gradient Descent
    b) Backpropagation is the key
    c) What drives the network growth
3) Architecture is everything: How Deep Learning architectures solve problems
    a) Convolutional Neural Networks with Images + Demo
    b) RNN/LSTM with Text + Demo
    c) Autoencoders with unsupervised learning + Demo(if time)
4) Actuarial Applications
    a) Applications for deep neural nets in insurance
    b) Challenges and risks : Questions actuaries should ask during neural net discussions

1) Perceptrons the building blocks of neural nets

Perceptron: basis of a neural network

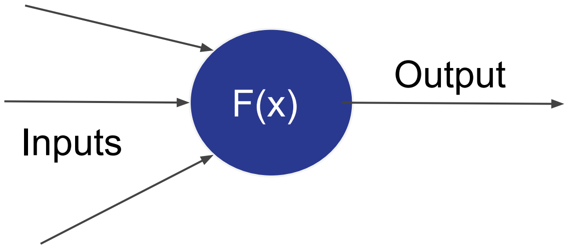

F(X) is a definable function

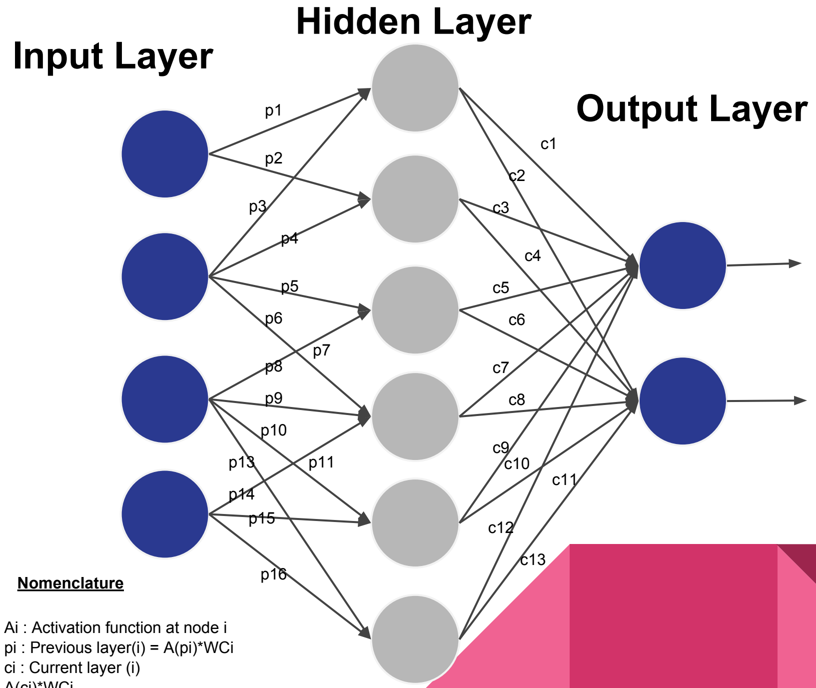

Nomenclature

Ai : Activation function at node i
pi : Previous layer(i) = A(pi)*WCi
ci : Current layer (i)
A(ci)*WCi

Terminology before we dive in

- **“Pre-processing”:** Data transforms readying data for input to a neural net. Key decisions include
- **“Feed-Forward”:** Neural Networks where signal from functions travel in direction from input to output
- **“Convolutional”:** The merging of multiple, if not comprehensive, results from functions estimating an objective function
- **“Recurrent”:** Networks where outputs are re-used as inputs at the next time step
- **“Long Short Term Memory”:** Commonly used in a recurrent network, a static layer recording the previous output with functions to control input, memory and output to the next step
- **“Pooling”: TBD**
- **“Dropout”:** A method to randomly drop connections to reduce over-fitting
- **“Fully Connected”:** All neurons receive an input signal and deliver an output signal
- **“Backpropagation”:** Adjusting the network weights by first minimizing the loss function and working back towards the input, adjusting connected weights at each step
- **“Stochastic Gradient Descent”:** A step towards a smaller loss function using partial derivatives
- **“Hidden Layer”:** Series of functions that process input signal and output to another layer. Hidden since input and output layers are transparent
- **“Epoch”:** A single run through the neural network from input to output (iteration)

Sigmoid Function: (1)/(1+exp(-x))

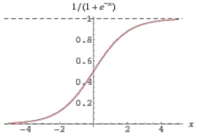

Purpose: Similar to logistic, binomial classification

Tanh Function: TBD

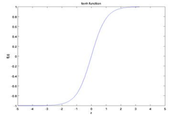

Purpose: To transform a -ve 1 scale to +ve 1

ReLU function: ~ max(0 ,ln(1+exp(x))) x:[0:1]

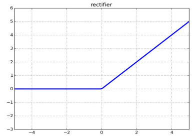

Purpose: Activate as +ve multiplier at gradient 1 to output

SoftMax Function: Prob (yi | input) = exp(xi*wi)/(Sum: exp(xi*wi: for 1 to D)

Purpose: Transform multinomial scalar outputs to posterior log normalized vector squashed between 0 and 1.Multinomial.

2) The “Deep” in Deep Neural Nets

GoogleNet for Image Classification

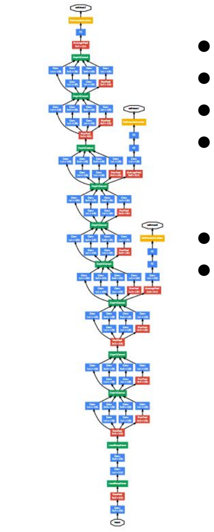

Figure 3: GoogLeNet network with all the bells and whistles.

- 22 Layers deep
- Multiple entry points and merges
- Image classification focus
- Example of a feed-forward network where the outputs are not re-used as inputs
- Purpose is to classify images
- A tangential purpose is to support internet memes (referenced in paper)

https://www.cs.unc.edu/~wliu/papers/GoogLeNet.pdf : Szegedy et al.

Recurrent Neural Networks for Text Processing

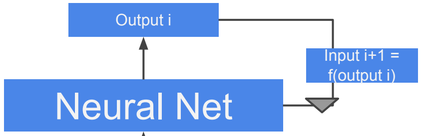

Unpacking the network

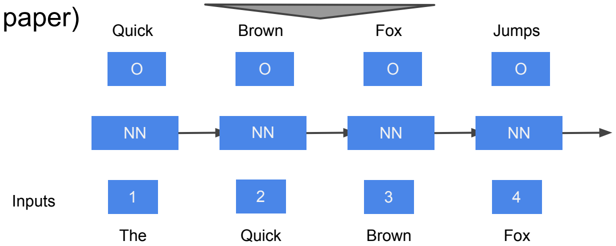

Notes to consider:

1. **The layers of abstraction are akin to “feature engineering”**. By abstracting the data into “hypotheses” to test that a human might not arrive at to develop as a hypothesis for the model
2. **Works well for specific problems** such as images, and text in some cases where there is complex data problem that needs a general solution
3. **This is important to Actuaries for three specific reasons:**
   a. Abstraction layers are not easy to explain, nor are results
   b. Network architecture are important but complex; mistakes can be made
   c. Overfitting is an issue(use dropouts)

2a) Stochastic Gradient Descent

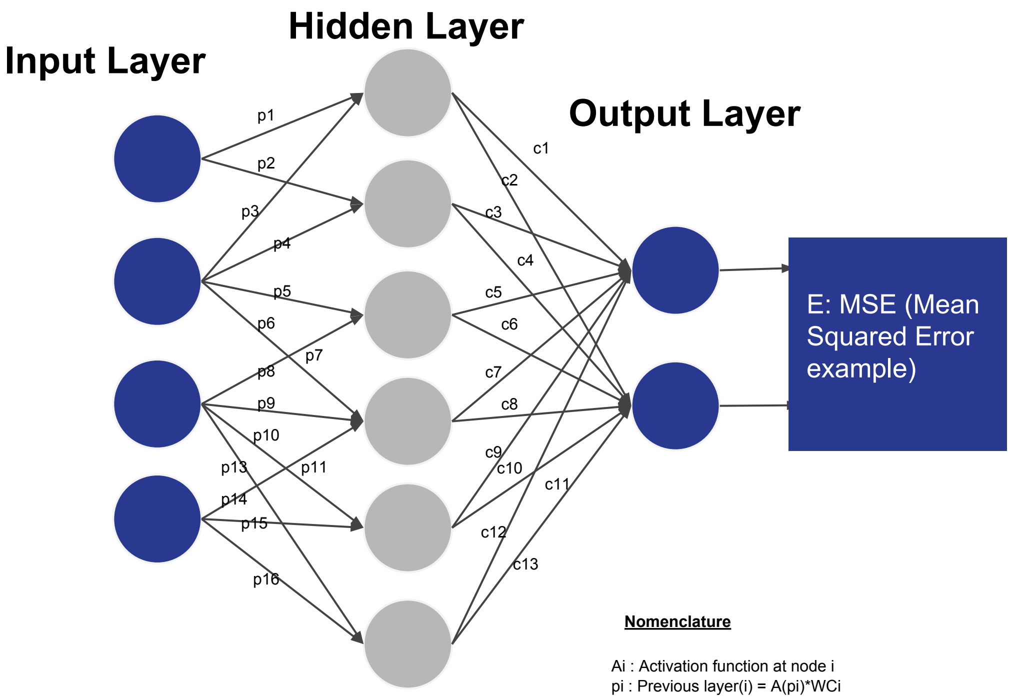

Nomenclature

Ai : Activation function at node i
pi : Previous layer(i) = A(pi)*WCi
ci : Current layer (i)
A(ci)*WCi

Stochastic Gradient Descent is a popular method to reduce the error function and fit the model closer to the data

- Efficient because it assesses one connection at a time (unlike batch gradient descent which tries the whole or a large sample of the network)
- Series of partial derivatives. Let's take one example:
  - Let's assume an error function that measures the deviance between observed(o) and expected(e)
  - Squared difference between o and e is loss function L
  - Derivative with respect to single weight wi is ∆L(wi)
  - A single step learning parameter lamda(lam) is introduced
  - The weight is re-adjusted to wi=w0 - lam*∆L(wi)
  - Why negative gradient? We want to reduce the error!
- This is done repetitively until a stopping criteria is reached
- Some issues
  - Saddle points explored in Bengio
  - Overfitting

2b) Backpropagation is key to fitting models

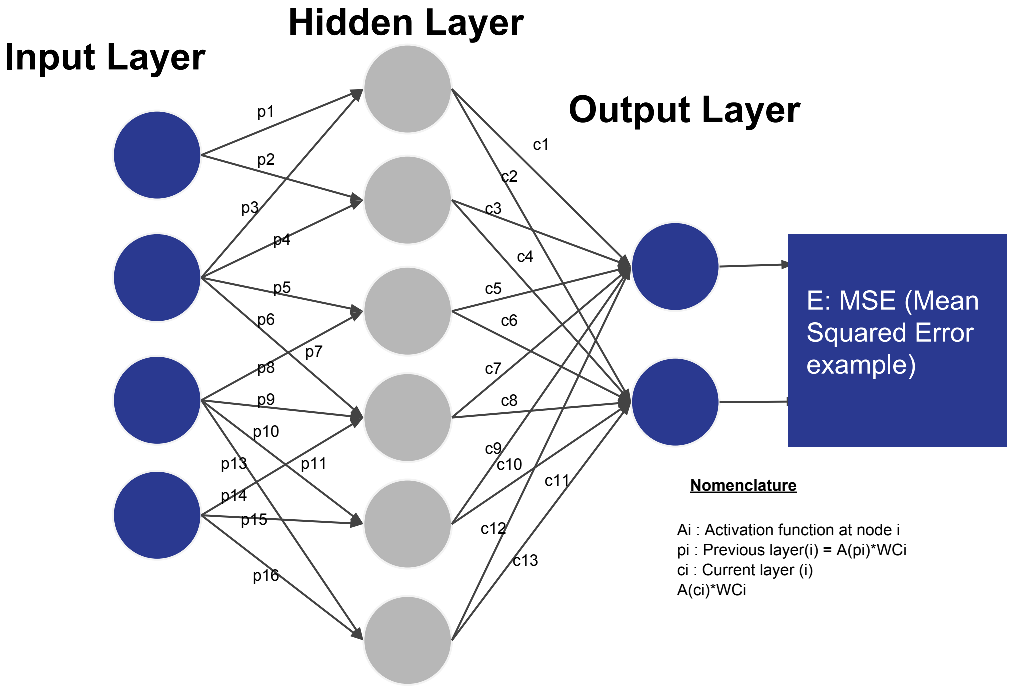

Nomenclature

Ai : Activation function at node i
pi : Previous layer(i) = A(pi)*WCi
ci : Current layer (i)
A(ci)*WCi

Backpropagation Steps:

1) Initialize the weights at each pi and ci
2) Calculate the Error Value
3) Take a random connection ci
4) Peturb the weight, ci*wi, value by a small amount δ(ci)wi
5) Relate back to connected pi by derivative of activation function (δpi = A’(pi) Z wij δcj)
6) Re-calculate the Error Value
7) Repeat until a stopping criteria is activated

3a) Convolutional Neural Net with images in R

Demo on MNIST data

3b) RNN/LSTM with Text

Demo on Text and Context

4a) Deep Learning Actuarial Applications 1/2

Actuaries have been exploring neural networks for some time!

Some examples from Actuarial Lit or "Actuarial Neural Network" history:

1996 General Insurance Convention

Neural Networks v. GLMs in pricing general insurance

Workshop to be presented by:

Julian Lowe (Chairman)
Louise Pryor

417

COMPARISON OF INCURRED BUT NOT REPORTED (IBNR) METHODS

Sponsored by
Society of Actuaries Health Section

Prepared by
Cabe Chadick, FSA, MAAA
Wes Campbell, ASA, MAAA
Finn Knox-Seith, ASA, MAAA

October 2009

L&E Actuaries & Consultants

© 2009 Society of Actuaries, All Rights Reserved

Comparison of IBNR Methodologies

Appendix II
Advance Methods: Neural Network

Neural Networks Demystified

Louise Francis, FCAS, MAAA

253

Major changes in the industry since these papers

- **Development** of open source software to build neural networks at scale
  - The Apache software foundation creating Spark, Storm and Cassandra and Hadoop
  - Python and R emerging as open source statistical and data munging programs with a vibrant community of developers
  - Data Science Community open sourcing code: TensorFlow from Google, Caffe, Keras, Torch, Theano
- **Hardware** and specifically GPUs that support parallel processing
- **Web services** such as Amazon AWS, Google, Microsoft Azure and Rackspace offering managed services
- **Data** especially **unstructured data** such as text holding value with an objective in mind

4a) Deep Learning Actuarial Applications 2/2

	Observed in Industry	**Value add with DL in value chain**
Submission	Logistic models for underwriters for a final **(Yes/No) using thresholds**	Finer tuned and multinomial models detailing Yes/No/Potential to write with change as an example of multiple outcomes with actions attached
Pricing	Frequency-Severity or Pure Premium **parametric or non-parametric** (typically shallow ML) on historical data	Finer tuned pricing - company owns the architecture of the model
Claims	Shallow Machine Learning **Triage** models using some n-gram text models to predict Severity	Deep NN architecture development that can combine text, sound, image and regular structured data - with valuable additional capabilities to absorb highly dimensional data
Reserving	Chain Ladder,BF and Cape Cod at the AY-Dev period grain. Stochastic reserving using parametric methods	Claims grain with fine tuned cohort prediction across time, including RNN architectures that can estimate future states
Fraud	Unsupervised clustering, network and graph analyses and shallow ML where target data available	Ability to combine data types and develop finely tuned suspicion models

1) You own the model architecture

2) Can use increasingly dimensional and complex data being created

3) Finer tuning of model to objective (dangers of overfitting)

4b) Challenges and risks: Questions to ask

- What is the objective of the model and how does this solve the business problem?
- What pre-processing steps have you taken?
- Is there a resource constraint? (In terms of people, systems and time)
- Why does the architecture actually work? Have you tested other architectures?
- Are you collecting human decision data? How are you adjusting for bias?
- Is the model overfitting? How have you controlled for this?
- Does the model need to be explainable to a business?
- How much upkeep is needed?
- Has ensembling with another model improved results?