Neural Networks
Could we actually file one?

Christopher Cooksey, FCAS, MAAA, CSPA
Head Actuary, Data & Analytics

20 March 2018

GUIDEWIRE

Agenda

- Brainstorming – Of course we can't
- A Regulator's Perspective
- How Neural Networks Work
- The "Black Box"
- Real Issues, Problems and Solutions

Brainstorming – Of course we can't

Conventional Wisdom Says…

We could never use a Neural Network for a rating algorithm because…

How Neural Networks Work

Neural Networks

Artificial Neural Networks come in many varieties…

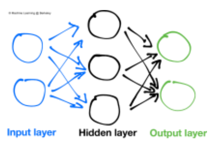

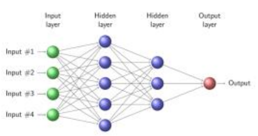

Neural Networks

Artificial Neural Networks come in many varieties…

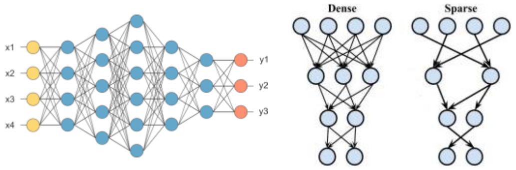

Neural Networks

Artificial Neural Networks come in many varieties…

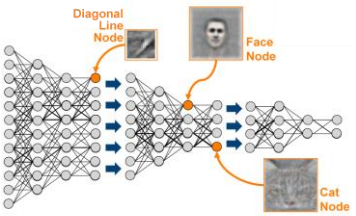

Neural Networks

…so let's simplify our discussion.

**A single layer,
feedforward,
backpropagation
neural network
with one output node.**

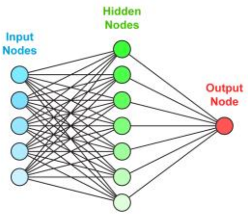

Neural Networks

A single layer, feedforward, backpropgation neural network with one output node.

Input nodes are the predictor fields. Think design matrix:

- **Each input node for a categorical predictor is one level, and the values are 0 or 1.**
- **Each input node for an ordinal predictor is the normalized value of that predictor.**

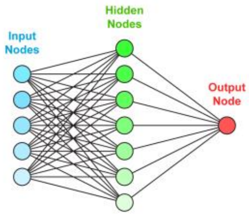

Neural Networks

A single layer, feedforward, backpropgation neural network with one output node.

Each hidden node is a linear predictor, literally just a weighted combination of the input (i) for a given hidden node (j).

$$\mu_j = \sum_i (Input_i) w_{ij}$$

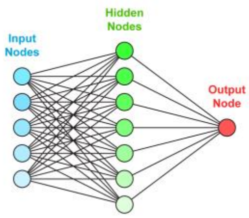

Neural Networks

A single layer, feedforward, backpropgation neural network with one output node.

Except…the output from the linear predictor, μ_j, is put through an "activation" function, such as a logistic function…

$$f(x) = \frac{1}{1 + e^{-x}}$$

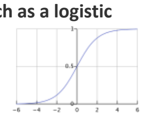

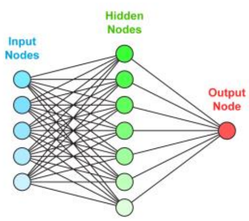

Neural Networks

A single layer, feedforward, backpropgation neural network with one output node.

Finally, the transformed output from each hidden node is weighted together into the output.

Just like GLMs, we're trying to find optimal weights.

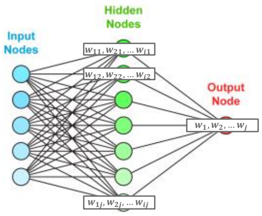

Neural Networks

A single layer, feedforward, backpropgation neural network with one output node.

GLM optimal betas are found through minimizing a function on all the data.

Neural Net optimal weights are adjusted one record at a time until the error can't be reduced.

Each hidden node has different starting weights as well!

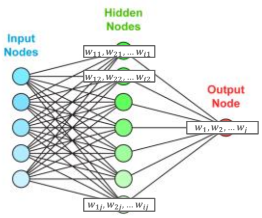

Neural Networks

The activation function is also a key element of Neural Networks.

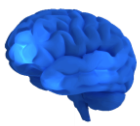

Just like the brain has different neurons fire for different tasks, each hidden node is sensitive to only certain inputs.

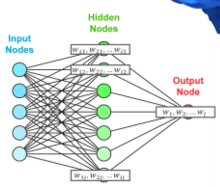

Neural Networks

Finally, for purposes of this discussion, we will not focus on getting a good performing Neural Network.

We will assume we've created one for the purposes of rating and that it works well.

The question in this presentation is, how would we file it?

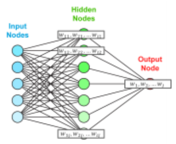

The "Black Box"

The "Black Box"

What do people mean when they say that Neural Networks are Black Boxes?

- You don't get to know what is happening in the hidden layer.
- Others?

The “Black Box”

What makes neural networks harder to explain is that we can give no simple meaning to the weights.

However, the following statements are not necessarily true.

- **Neural networks are complex.**

Neural networks, just like GLMs or ensembles or nearly any other technique, will be as simple or complex as the modeler chooses to make them.

This picture to the right would have 42 separate coefficients estimated.

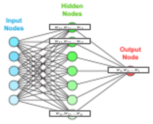

The “Black Box”

What makes neural networks harder to explain is that we can give no simple meaning to the weights.

However, the following statements are not necessarily true.

- **Neural networks are complex.**
- **Neural networks are not transparent in their operation.**

Neural networks, just like GLMs, are a fairly straight-forward formula. How they operate is completely transparent. It is only the meaning of the individual components that is difficult.

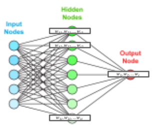

The “Black Box”

What makes neural networks harder to explain is that we can give no simple meaning to the weights.

However, the following statements are not necessarily true.

- **Neural networks are complex.**
- **Neural networks are not transparent in their operation.**
- **Nothing can be said about individual predictors in a neural network.**

Not true! There are different ways to tackle this – relative importance of predictors, reason codes, observed versus modeled (OvM) charts, etc.

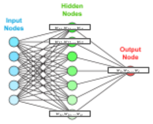

Real Issues, Problems and Solutions

Does the model perform well?

Does the model differentiate between customers (lift)?
Does the model predict hold-out data?

It is easy enough to look at lift charts and OvM (observed versus modeled) charts to establish model performance.

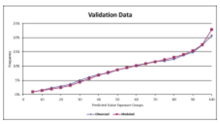

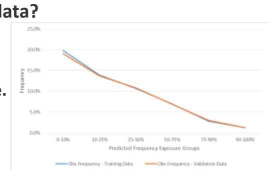

Is the model stable over time?

Are the predictions relatively stable year over year?

This is an important question. Lift curves found individually for each year of experience can answer it.

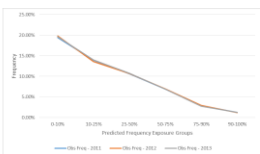

Can we specify a rating plan? Calculate a rate?

Neural networks are a formula with weights. Not too hard to express; not too hard to calculate.

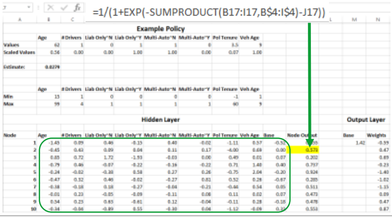

Can we specify a rating plan? Calculate a rate?

Neural networks are a formula with weights. Not too hard to express; not too hard to calculate.

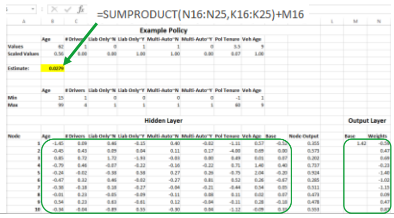

Again, the “rating plan”?

The rating plan is now a collection of input characteristics and specified weights. There is no Age curve in the same sense as the GLM.

Let’s look at the example Rating Manual Handout.

Thoughts?

Defining minimums and maximums may be an issue.

Need to account for expenses too.

Which fields are more important?

Once the model is specified, it is a straight-forward task to calculate various metrics around variable importance.

- **Replace all values of a field with its mean; refit the network. The loss of accuracy for each field quantifies its importance.**
- **For a single-layer network, you can take the absolute value of all the weights.**

This question has been kicked around a lot. You can get an answer.

How do different fields behave in the model?

Observed versus Modeled (OvM) charts can give insight into the model. Just let the x-axis be the levels of the field instead of some sort by model prediction.

Don't forget that reality doesn't change just because we pick a different modeling technique.

If the severity increases as limit increases, then an OvM from any model should show this.

Getting more practical – Rate Swings?

Rate swings are always an issue. Calculating them is easy. Controlling the size of them? Not so easy.

Is this different than a GLM though?

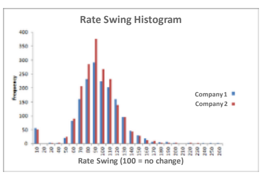

Getting more practical – Reversals?

How about reversals? No guarantees that there won't be reversals across certain fields.

Is this different than a GLM though?

Um, yeah. GLMs are easier for managing reversals…
…except when using interactions and fields show up in multiple places. Then GLMs can be hard to manage reversals too.

More practicalities – FCRA?

Compliance with the Fair Credit Reporting Act is one of those things that we need to provide for even if not included in a filing.

Rates can be calculated with the old and new credit. If the rates increase, well, then you know.

More practicalities – Customer complaints?

Whenever rates change, there can be questions. Typically we need to provide reasons for a rate.

Age	BI	PD	MP
15	3.463	3.020	3.092
16	3.298	2.876	2.945
17	3.141	2.739	2.805
18	2.991	2.609	2.671
19	2.849	2.485	2.544
20	2.713	2.366	2.423
21	2.584	2.254	2.308
22	2.461	2.146	2.198
23	2.344	2.044	2.093
24	2.232	1.947	1.993
25	2.126	1.854	1.898
26	2.025	1.766	1.808
27	1.928	1.682	1.722
28	1.836	1.602	1.640
29	1.749	1.525	1.562
30	1.666	1.453	1.488
31	1.586	1.383	1.417
32	1.511	1.318	1.349
33	1.439	1.255	1.285
34	1.370	1.195	1.224
35	1.305	1.138	1.166
36	1.243	1.084	1.110

multi-unit indicator	discount
Y	0.850
N	1.000

Model Year	Factor
2018	1.100
2017	1.050
2016	1.000
2015	0.960
2014	0.920
2013	0.880
2012	0.840
2011	0.800
2010	0.760
2009	0.720
2008	0.680
2007	0.640
2006	0.600
2005	0.560
2004	0.520
2003	0.480

For a GLM, just look at the relativities that the policy gets.

Age makes this policy's rate higher…

Model Year makes this policy's rate lower…

Neural Network weights don't work this way!

More practicalities – Customer complaints?

Whenever rates change, there can be questions. Typically we need to provide reasons for a rate.

Pulling a few relativities from a GLM is a simplification of what is happening, *especially* if there are interactions.

So…simplify the Neural Network into a GLM that captures the main patterns.

This GLM would use the NN predictions as the target objective.

Getting creative

It may be that this is all too much. The simplicity of the GLM approach in how it goes on to be implemented may be worth it.

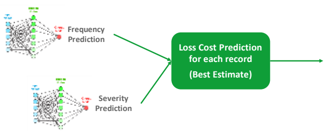

Simple Regression targeting the NN-predicted Loss Costs

What set of relativities best approximates the Neural Network predictions?

You *will* lose something in the translation of the NN-based result to a GLM structure.

Summary

- **Unlike GLMs, the weights for neural networks do not have any straight-forward meaning.**
- **Nevertheless, like GLMs their operation is a well-defined formula.**
- **It is no harder to prove that a neural network is predicting well than any other model.**
- **It is still possible to get insights into how the model is working in relation to its predictors.**
- **Practical issues like reversals, rate swings and customer complaints are real issues that need to be, but can be, accounted for.**

Questions?

Christopher Cooksey, FCAS, MAAA, CSPA
Head Actuary, Data and Analytics

Guidewire Software

ccooksey@guidewire.com