


CAS Ratemaking & Product Management Seminar
Overview and Practical Application of Machine Learning
Methods in Pricing – Part 1

Wednesday March 21, 2018

Claudine Modlin, Graham Wright


© 2018 Willis Towers Watson. All rights reserved.

Willis Towers Watson 



Agenda

Agenda	<p>Objective: to give you a working knowledge of some machine learning methods that may be used to improve GLM results and/or offer valuable insights in their own right in the field of P&C insurance pricing</p> <p><i>Please note that the on-site presentation will also include example results from particular methods that will not be included in this printed version; consequently, page numbers will differ.</i></p>
Context of machine learning in pricing	
Session 1:	
Decision trees Random forests Gradient boosting machines	
Session 2:	
“Earth” Neural networks Penalized regression Generalized additive models	
Conclusions	
Q&A	

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

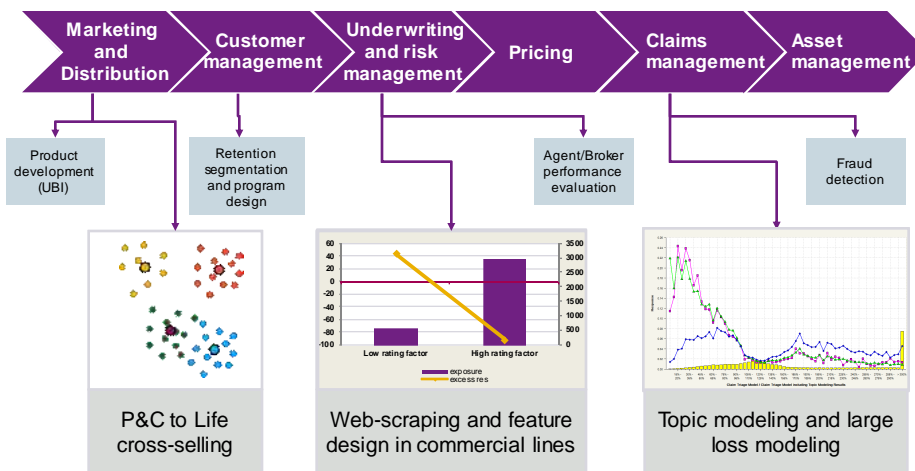
Willis Towers Watson  2

Who's interested in what?

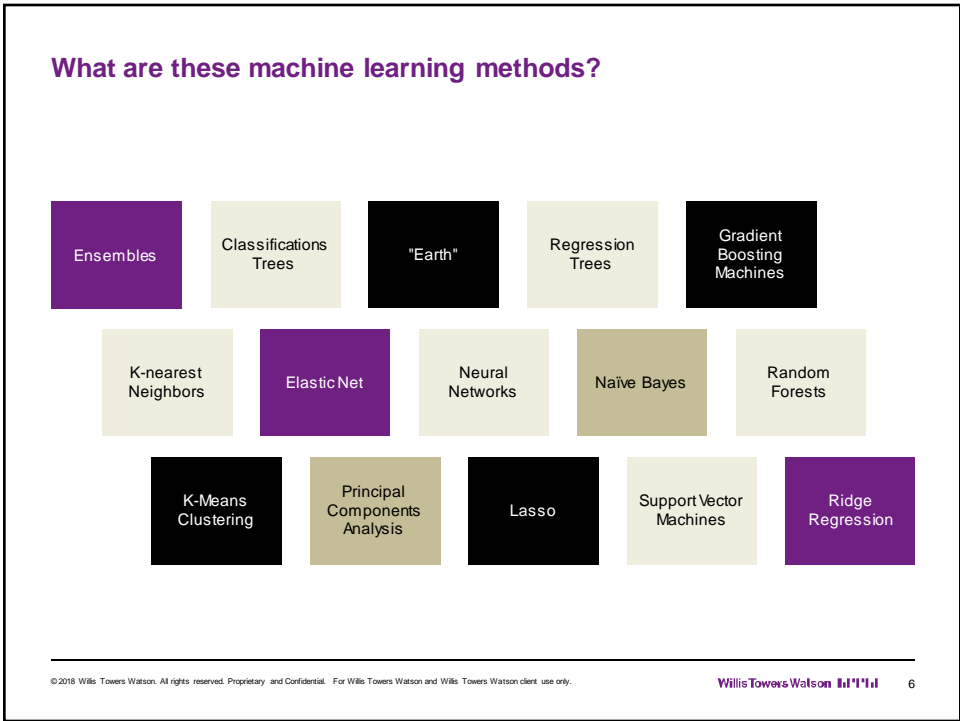
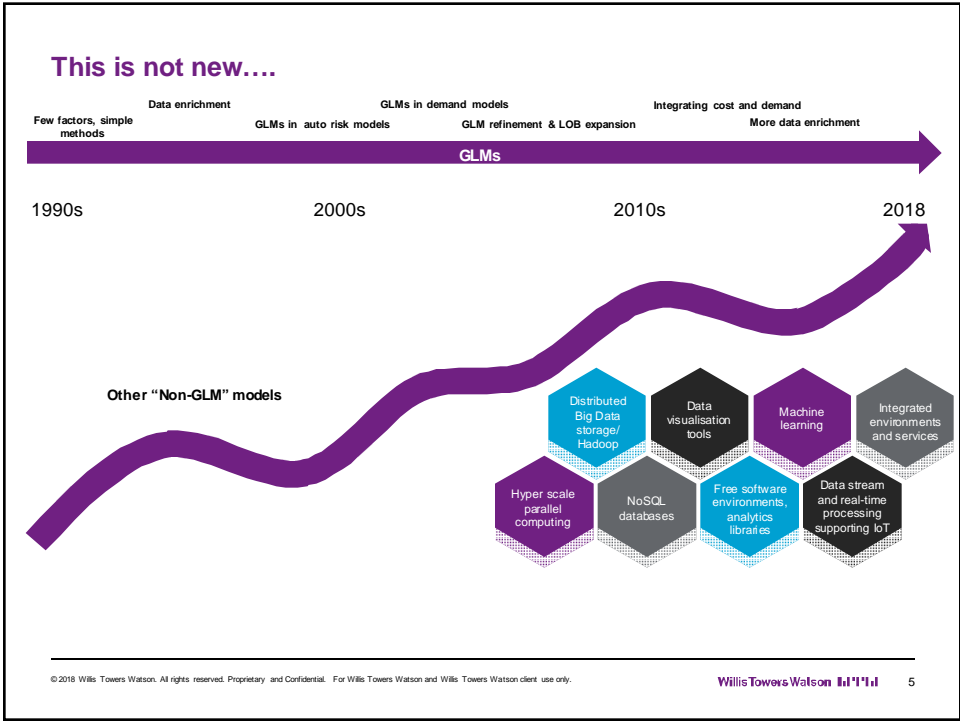


© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Applications of machine learning in the insurance sector



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.




Kaggle

The image shows two screenshots of the Kaggle website. The left screenshot is the homepage, featuring a navigation bar with 'Host', 'Competitions', 'Datasets', 'Scripts', 'Jobs', and 'Community'. Below the navigation bar, there are three main sections: 'Download', 'Build', and 'Submit'. The 'Active Competitions' section lists several ongoing challenges, including 'State Farm Distracted Driver Detection', 'Santander Customer Satisfaction', 'Home Depot Product Search Relevance', 'BNP Paribas Cardif Claims Management', '2016 US Election', '2013 American Community Survey', and 'World Development Indicators'. The right screenshot shows the 'Kaggle Rankings' page, which displays a grid of user profiles with their names, avatars, and performance statistics across various competitions.


© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 7

Is it really all about the method?



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 8

Is it really all about the method?



Is it really all about the method?

Data




Physical facticity
E.g., height, length, weight



Mechanical nature
E.g., engine size, fuel type




Qualitative descriptors
E.g., body type, model range



Performance
E.g., maximum speed, torque, BHP

Is it really all about the method?

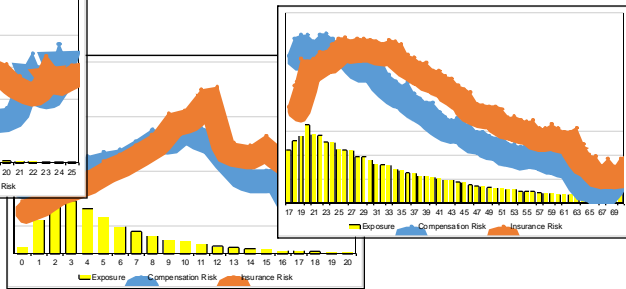





© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 11

Is it really all about the method?


Response selection



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 12


Is it really all about the method?



How do you know if a method works?



How do you measure value?



- Rank hold out observations by their **fitted values** (high to low)
- **Plot cumulative response** by cumulative exposure
- A **better model** will explain a **higher proportion of the response** with a **lower proportion of exposure**
- ...and will give a **higher Gini coefficient** (yellow area)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.


Willis Towers Watson  15

But...

- Think of a model...
 - Multiply it by 123
 - Square it
 - Add 74½ billion
-
- ...and you get the same Gini coefficient!

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  16



Financial value estimate

- Errors in insurance pricing are not symmetrical
- Financial benefit can be estimated



Example results redacted from printed version

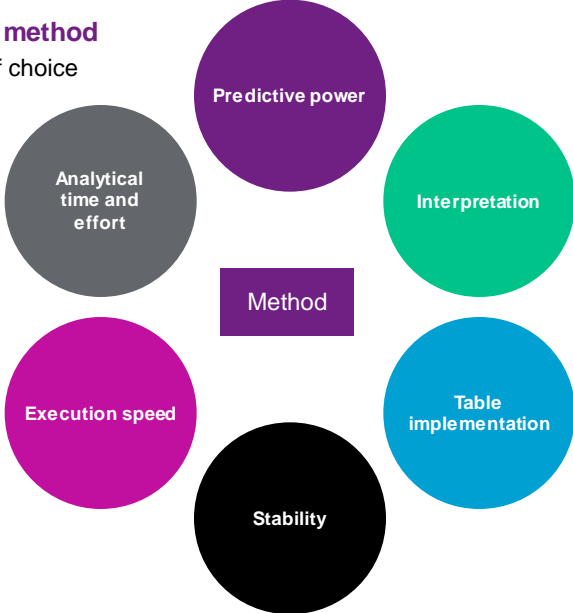
Is there more to it...?



Predictive power

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. **Willis Towers Watson** 19

Choosing a method
Dimensions of choice



Analytical time and effort

Predictive power

Interpretation

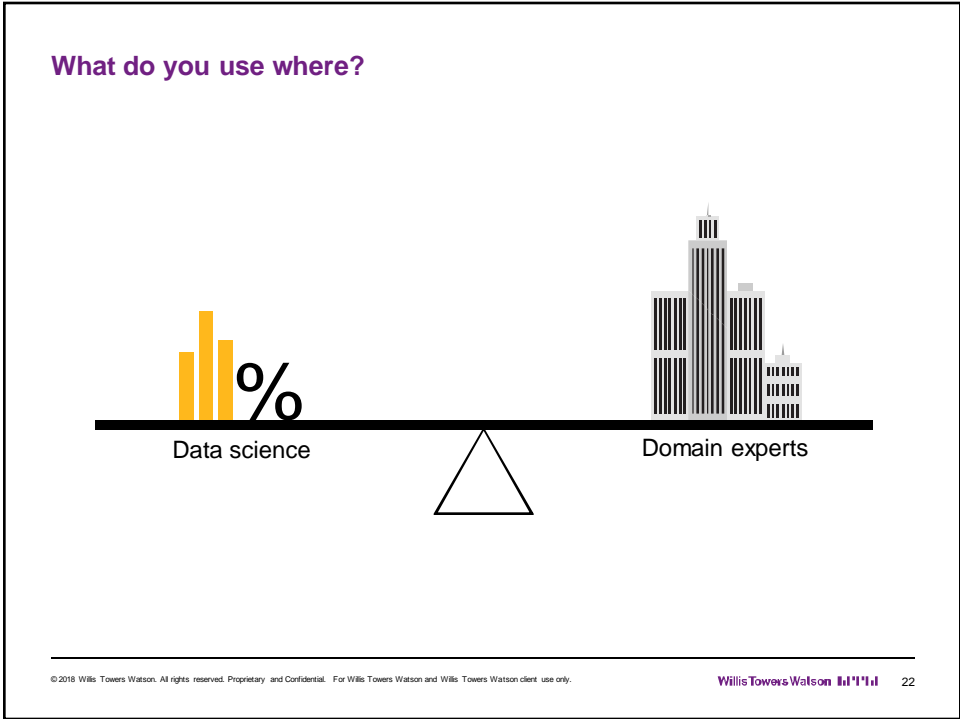
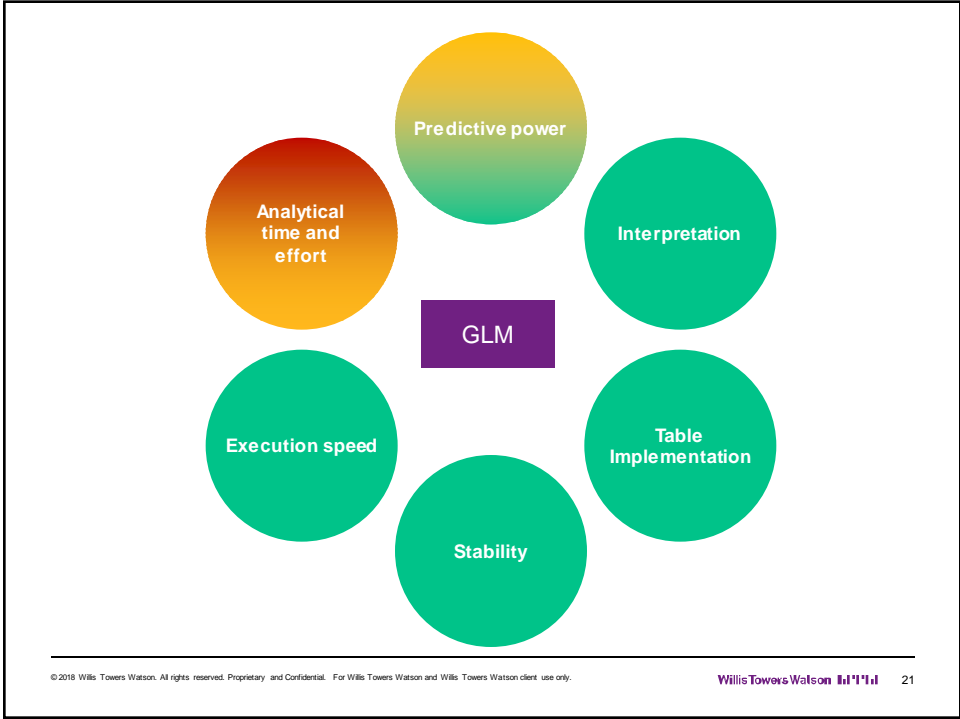
Method

Execution speed


Table implementation

Stability

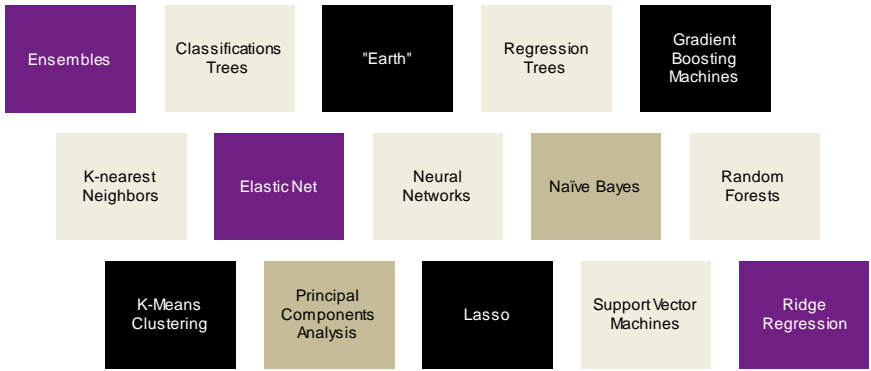
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. **Willis Towers Watson** 20



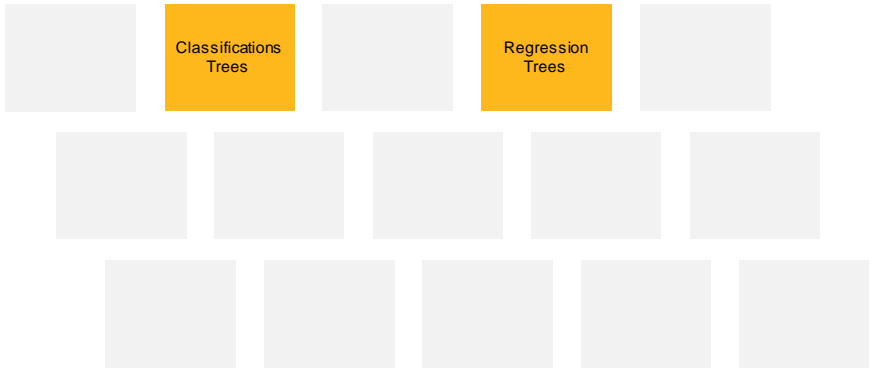
It's domain expertise that helps decide



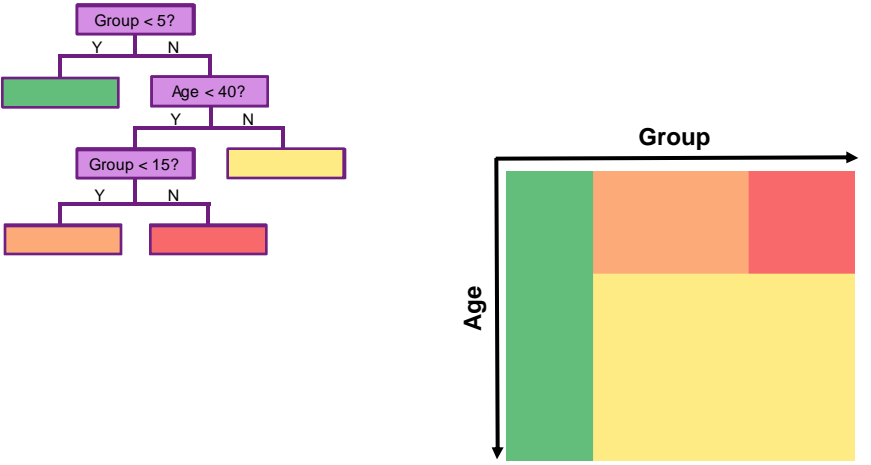
Some machine learning methods




Focus on Trees



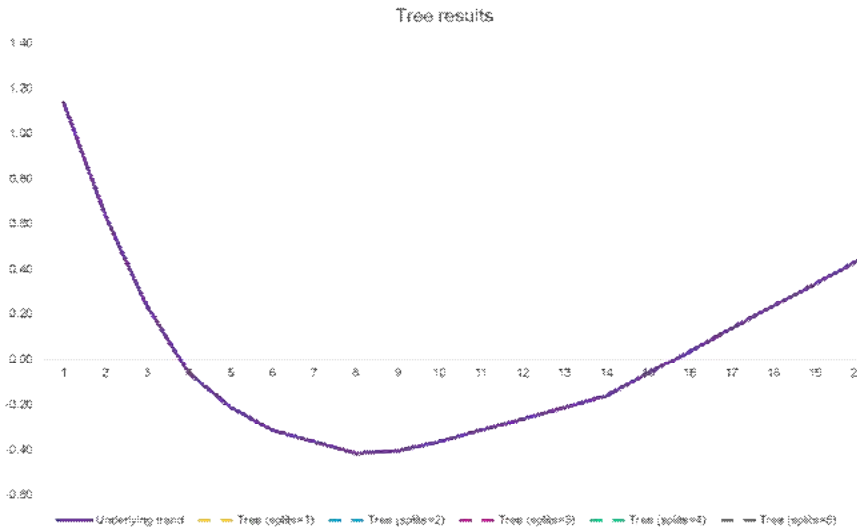
Decision Trees





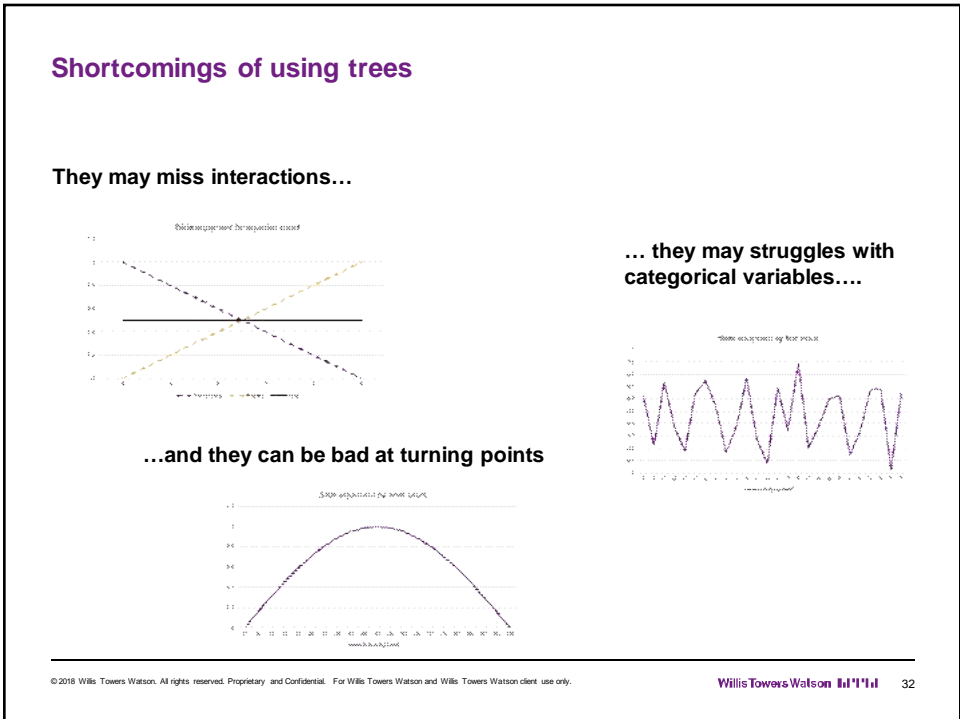
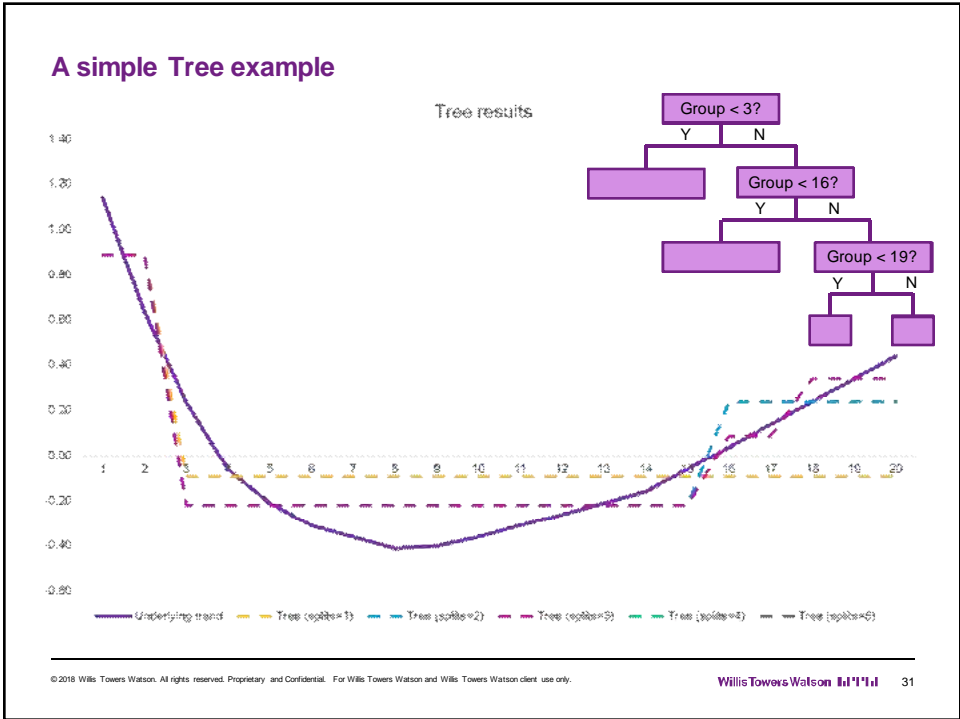
Decision Trees

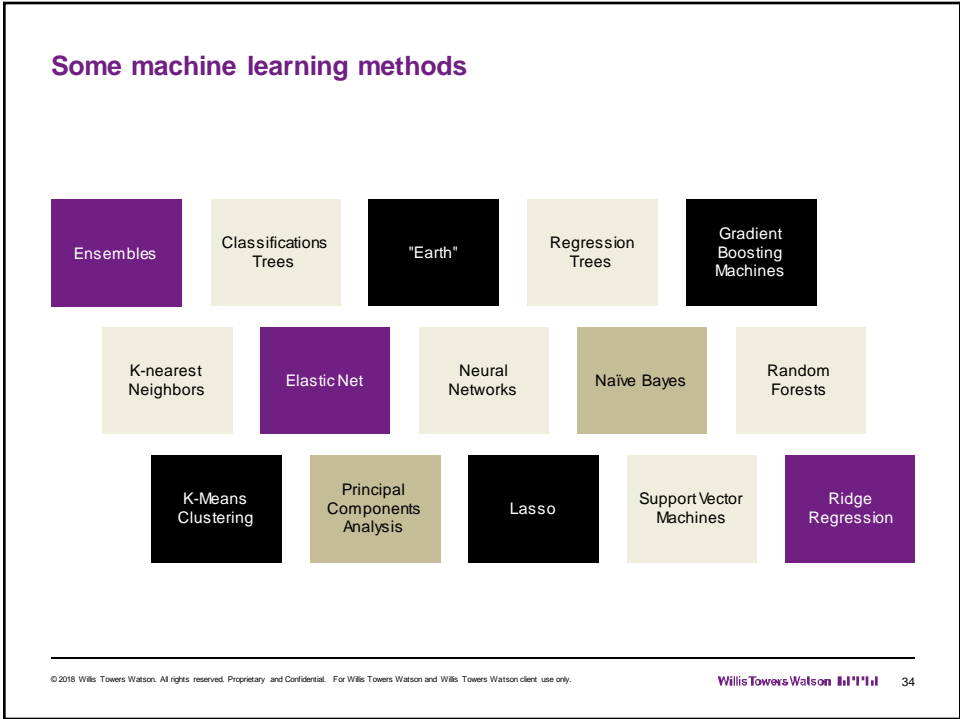
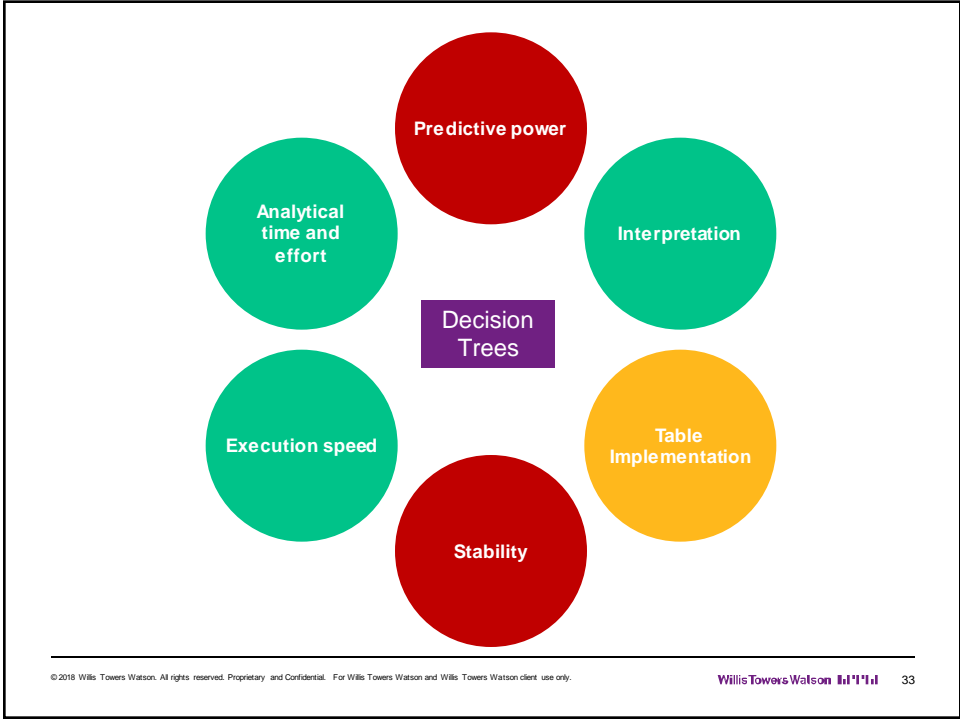


A simple Tree example










Focus on Random Forests



Random Forests

- Tree 1: Prediction 1 = Signal 1 + Noise 1
- Tree 2: Prediction 2 = Signal 2 + Noise 2
- Tree 3: Prediction 3 = Signal 3 + Noise 3
- ...
- Tree 1000: Prediction 1000 = Signal 1000 + Noise 1000


Random Forest:

$$\text{Prediction} = \text{AVERAGE}(\text{Tree Predictions})$$


$$= \text{AVERAGE}(\text{Tree Signal}) + \text{AVERAGE}(\text{Tree Noise})$$

- Average Noise → 0 if the trees are independent
- Independence of trees achieved by fitting each tree to:
 - Random subset of data (bootstrap sample)
 - Random subset of factors
- Average Signal → Underlying trend, provided trees are complex enough to represent it
- This is **bagging** (bootstrap aggregation) – fit lots of independent models and take an average


A simple Random Forest example



A simple Random Forest example




A simple Random Forest example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

A simple Random Forest example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

A simple Random Forest example



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

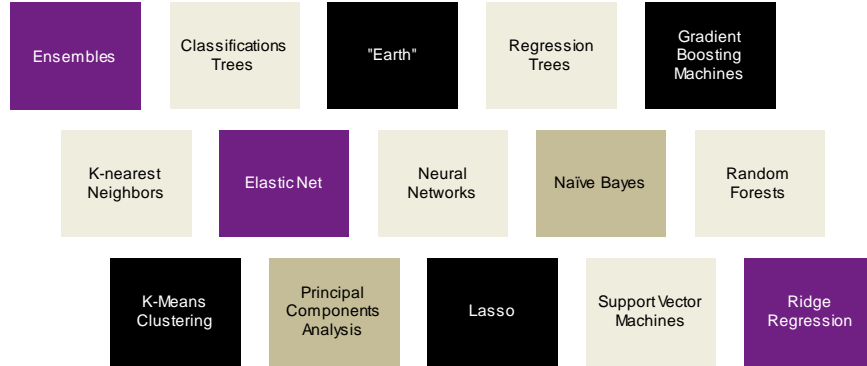
Willis Towers Watson 41




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 42

Some machine learning methods




Focus on Gradient Boosting Machines




Gradient Boosted Machine or “GBM”

A tree
 $f_i(x)$



A GBM


$$f(x) = \lambda \sum_{n=1}^N f_n(x)$$





© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Four main assumptions

- **λ Learning rate / “shrinkage”**
 - Amount by which the old model predictions are varied for the next model iteration
 - New model = Old + (Prediction x Learning rate)
- **Interaction depth**
 - Number of splits allowed on each tree (or the number of terminal nodes – 1)
- **N Number of trees** (iterations) allowed
- **Bag fraction**
 - Trees are fitted to a subset of the data (the bag fraction) on a randomized basis
 - Additional noise-reduction can be achieved by using a random subset of the available factors at each iteration




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.




A simple GBM example



GBM results at iteration 0




A simple GBM example

GBM results at iteration 1







A simple GBM example



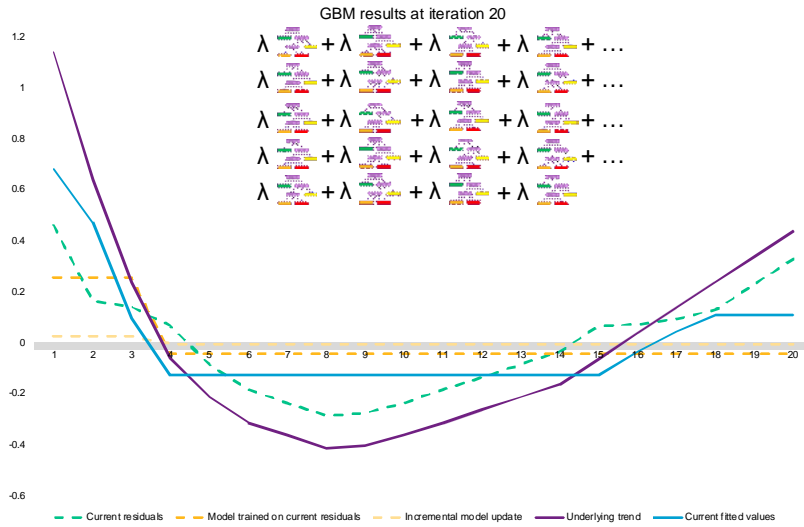
A simple GBM example




A simple GBM example




A simple GBM example



A simple GBM example




A simple GBM example



A simple GBM example

GBM results at iteration 50




© 2016 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 65

A simple GBM example

GBM results at iteration 100




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 66

A simple GBM example

GBM results at iteration 200




© 2016 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 67

A simple GBM example


GBM results at iteration 300



© 2016 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.


WillisTowersWatson 68

A simple GBM example




Calibrating the assumptions

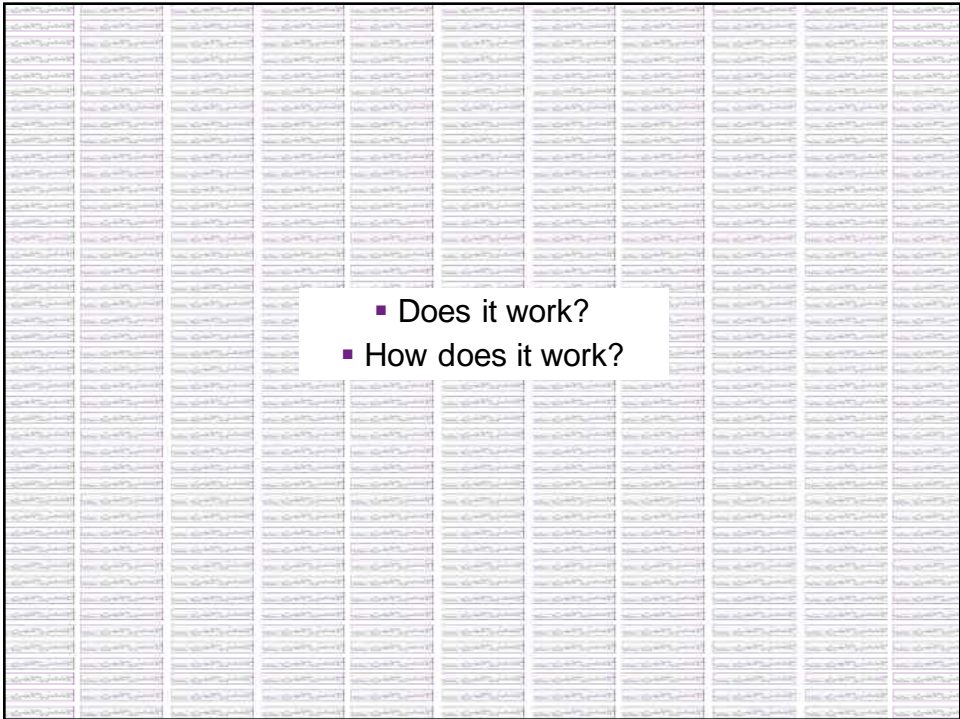
- *n*-fold cross validation used to develop the interaction depth and learning rate assumptions
 - Eg for 3-fold validation, split into 3, fit on purple, test on blue parts, take average



- Resulting plots can be used to determine the optimal assumption choice
 - Including how many trees to run


What does a GBM look like?


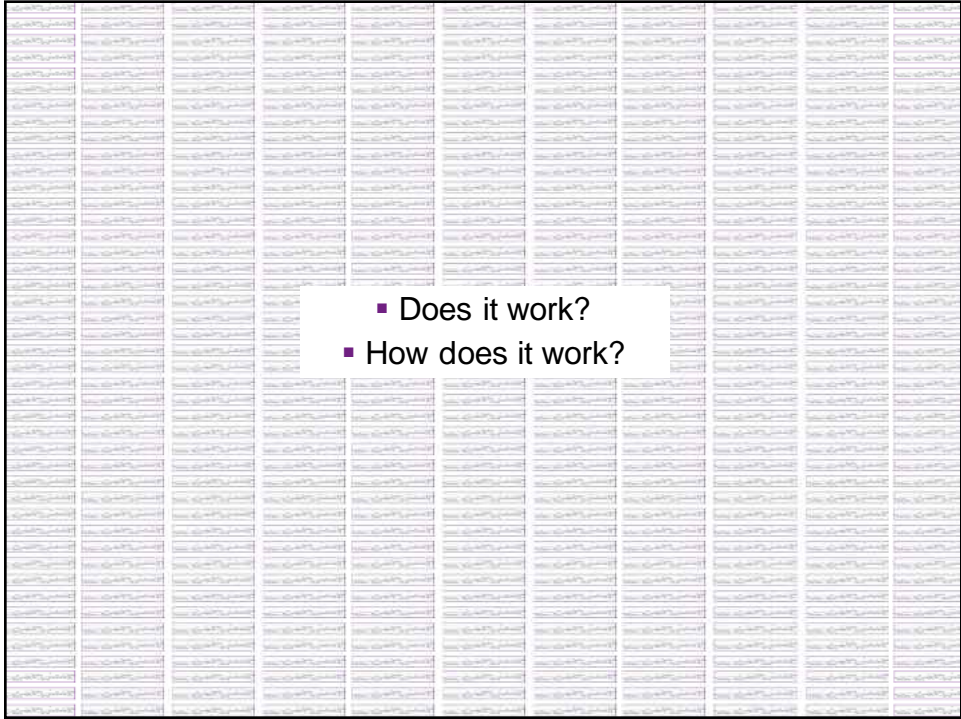




- Does it work?
- How does it work?

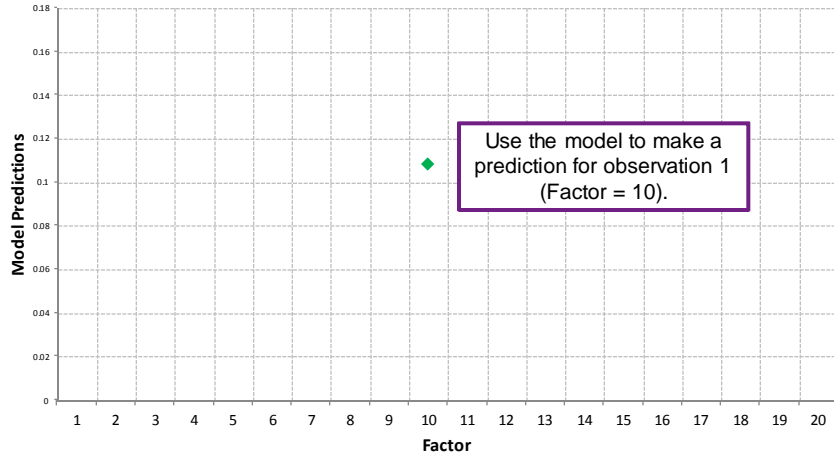
What about an automated GLM?





Partial dependency plots


Example



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Partial dependency plots


Example



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 81

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 82

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 83

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 84

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 85

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 86

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 87

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 88

Partial dependency plots

Example




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 89

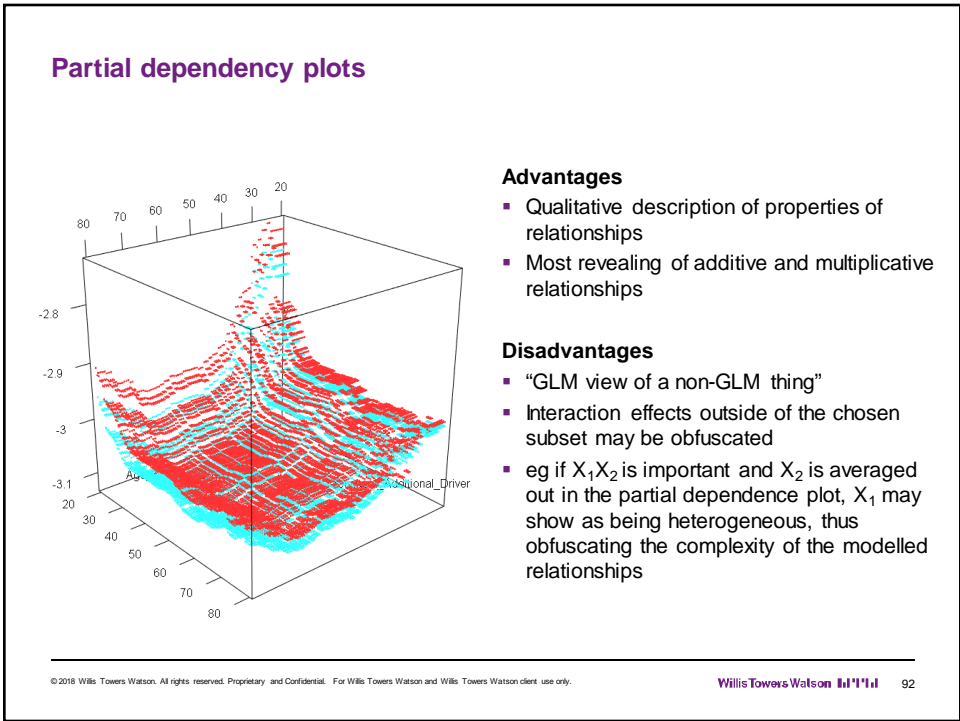
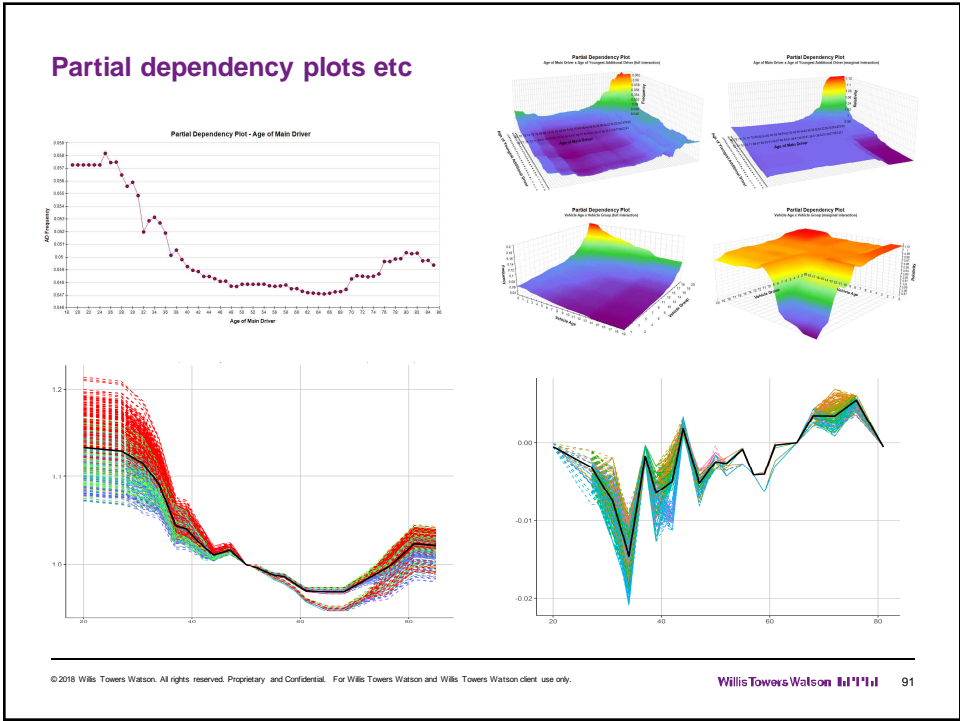
Partial dependency plots

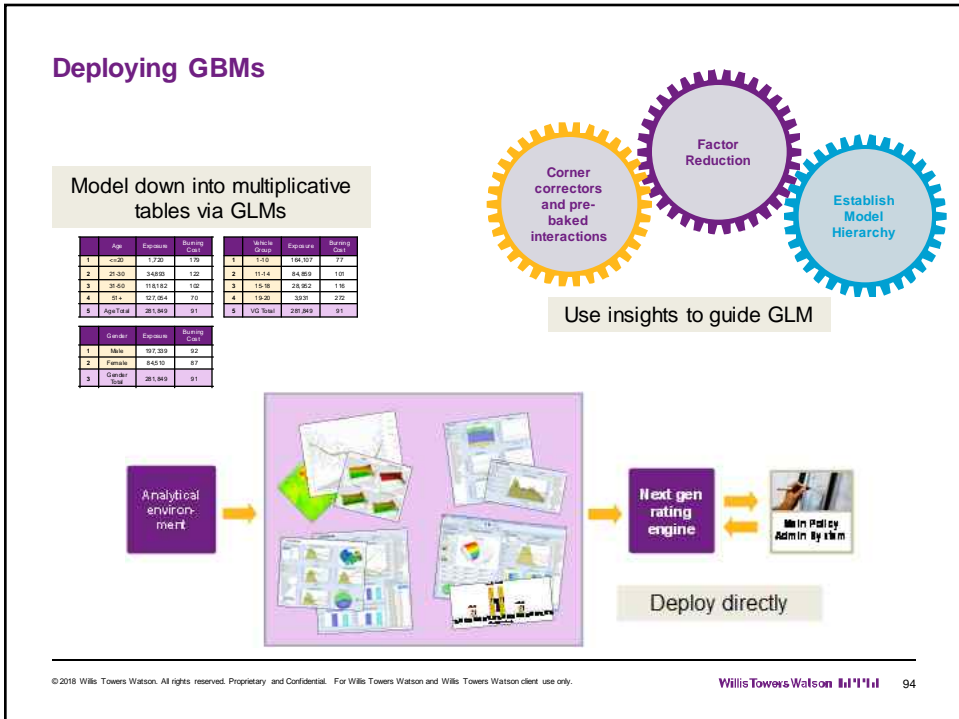
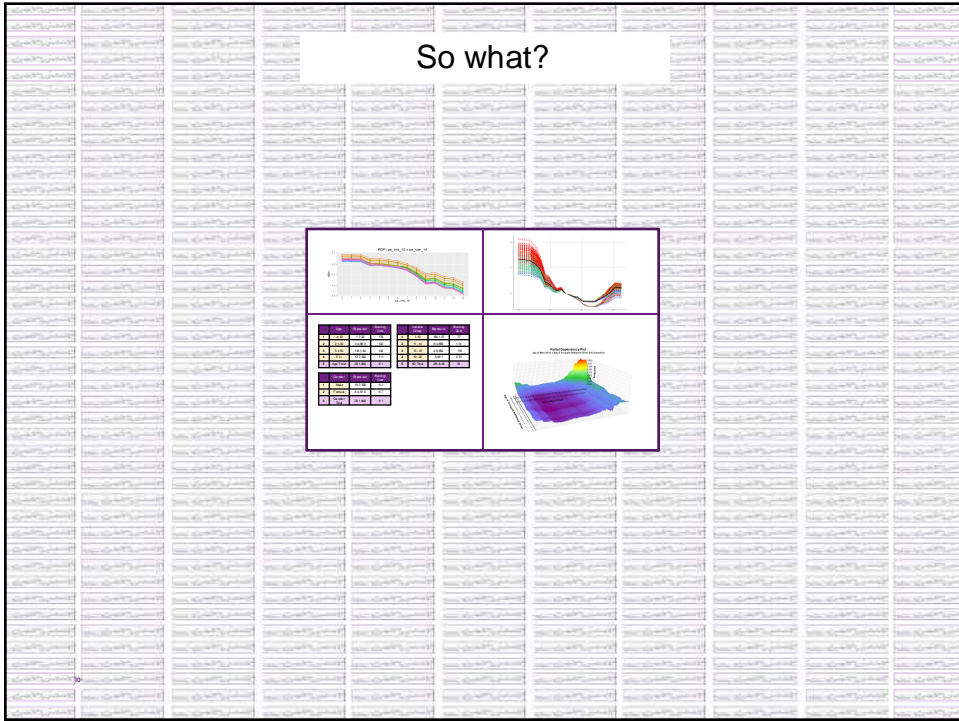
Example

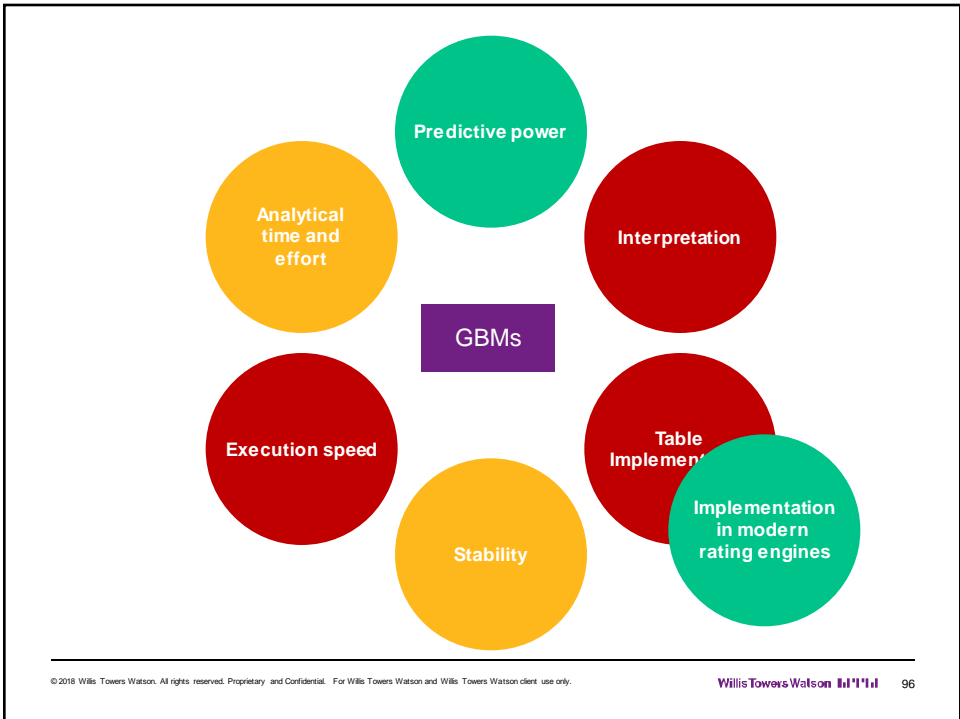
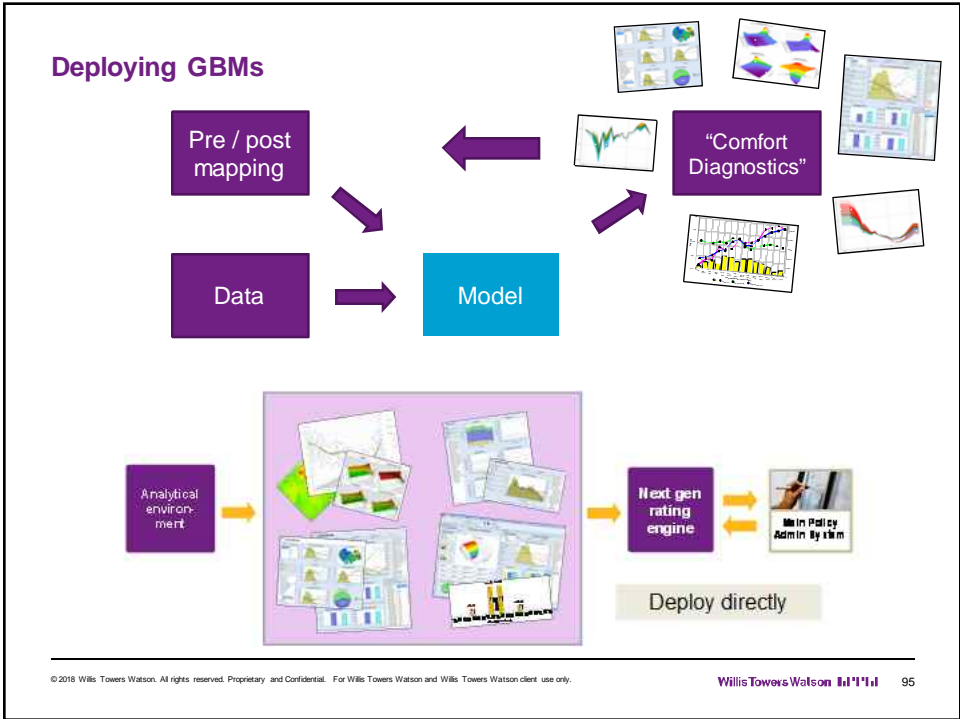


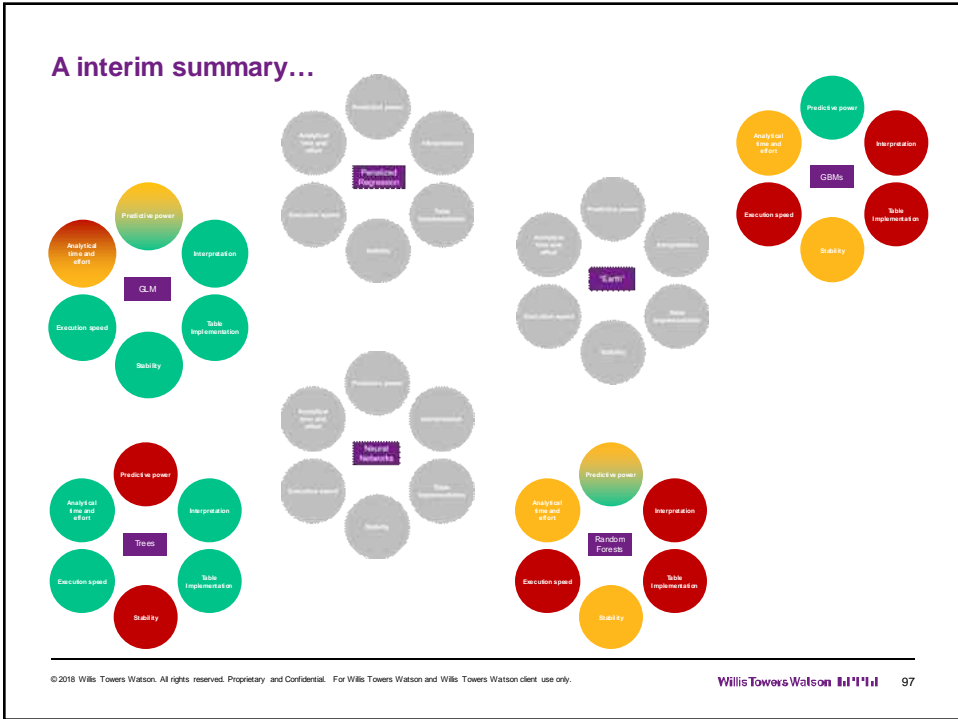
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 90









Machine Learning in Pricing

Conclusions (Part 1)

- There are many forms of ML models
- New data and feature/response engineering generally add more value than new methods BUT we need to continuously explore which methods work on which problems
- Traditional measures of prediction value may not reflect applications in insurance
- And it's not all about predictive power anyway – other criteria are important
- GBMs can provide predictive lift benefits by capturing higher order effects ... BUT
 - Can you cope with not seeing the model and instead use broad diagnostics
 - Effort is required to expose/understand higher order effects in an expeditious manner
 - How will business leaders and regulators respond to this method?
 - Do you have the software and hardware to fit to large dataset
 - Do you have a rating engine that can implement a GBM
- More methods and insights to follow in Part 2...

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 98

What's coming in Session 2?

Agenda

Context of machine learning in pricing

Session 1:

- Decision trees
- Random forests
- Gradient boosting machines

Session 2:

- “Earth”
- Neural networks
- Penalized regression
- Generalized additive models

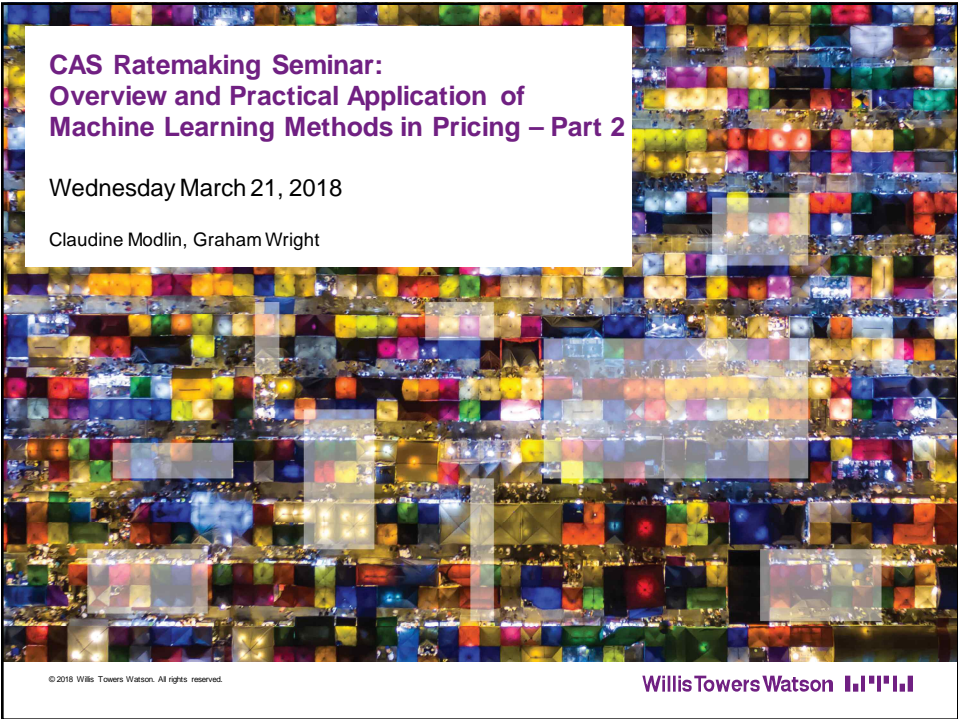
Conclusions

Q&A

Objective: to give you a working knowledge of some machine learning methods that may be used to improve GLM results and/or offer valuable insights in their own right in the field of P&C insurance pricing

Questions






**CAS Ratemaking Seminar:
Overview and Practical Application of
Machine Learning Methods in Pricing – Part 2**

Wednesday March 21, 2018

Claudine Modlin, Graham Wright


© 2018 Willis Towers Watson. All rights reserved.

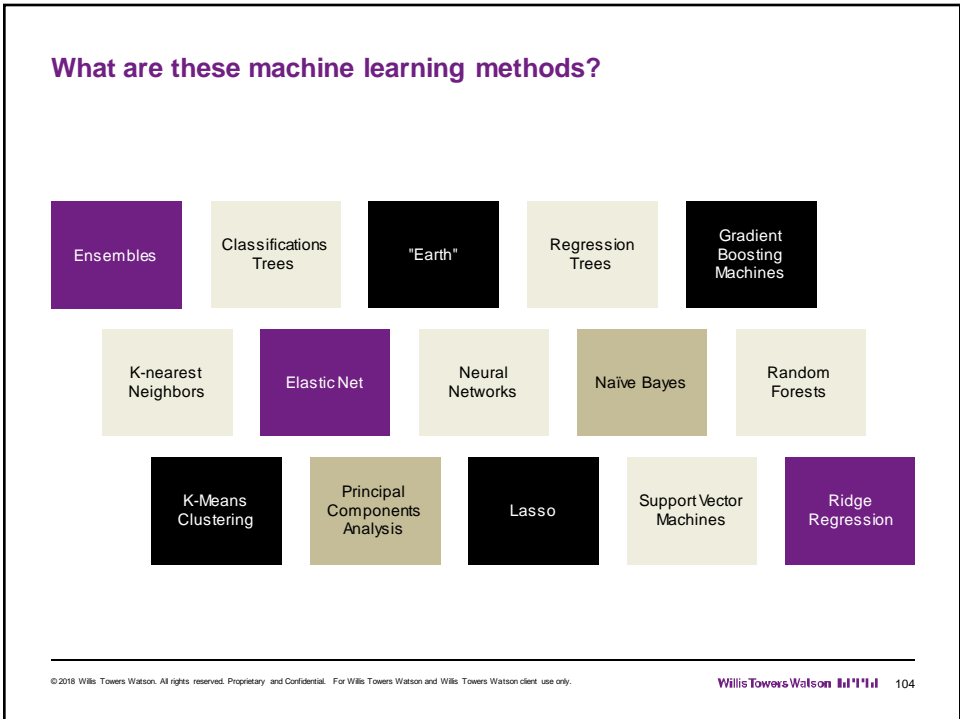

Willis Towers Watson 

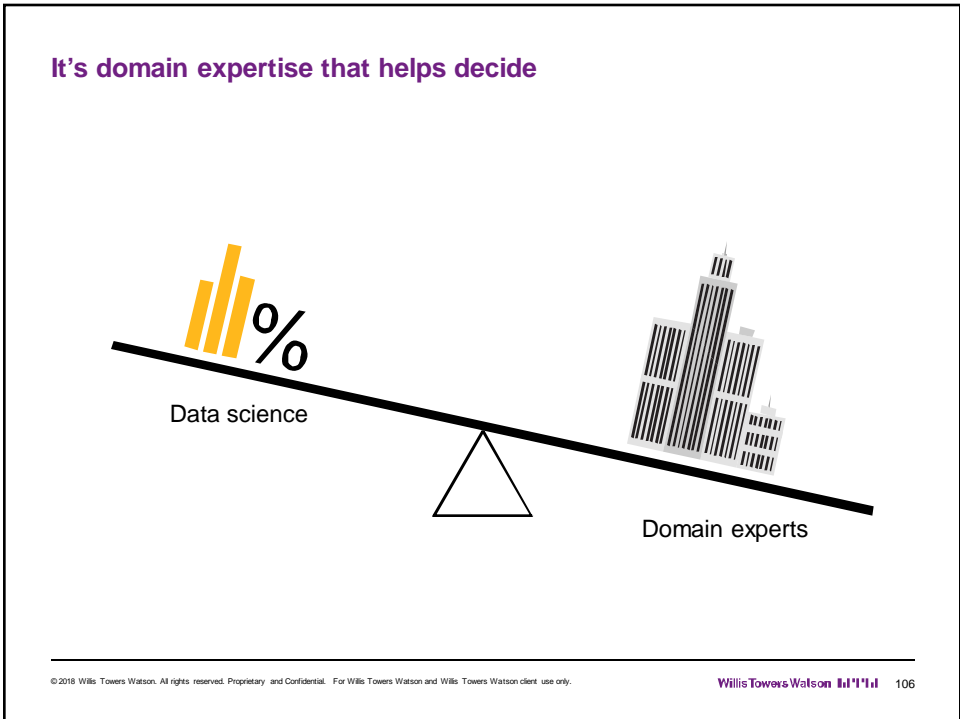
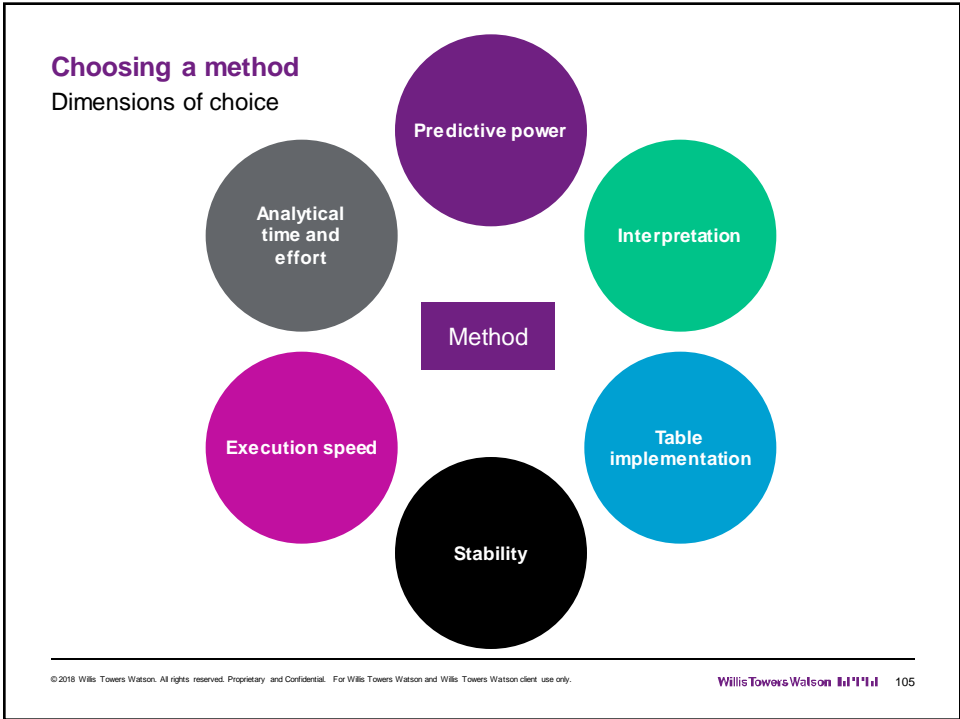
Agenda

Agenda	<p>Objective: to give you a working knowledge of some machine learning methods that may be used to improve GLM results and/or offer valuable insights in their own right in the field of P&C insurance pricing</p> <p><i>Please note that the on-site presentation will also include example results from particular methods that will not be included in this printed version; consequently, page numbers will differ.</i></p>
Context of machine learning in pricing	
Session 1:	
Decision trees Random forests Gradient boosting machines	
Session 2:	
“Earth” Neural networks Penalized regression Generalized additive models	
Conclusions	
Q&A	

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  102





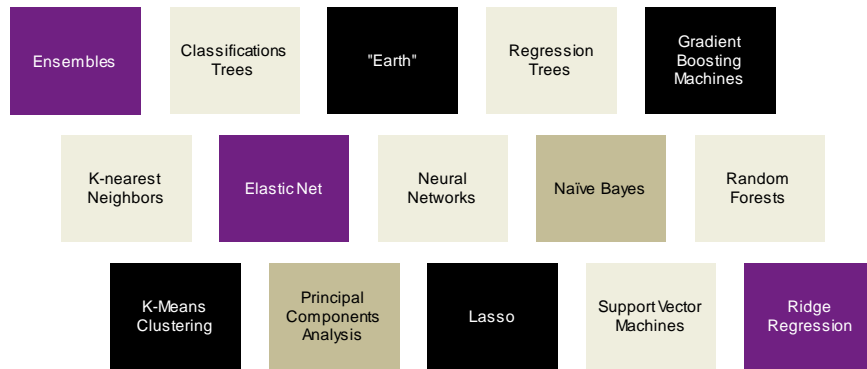
Financial value estimate

- Errors in insurance pricing are not symmetrical
- Financial benefit can be estimated

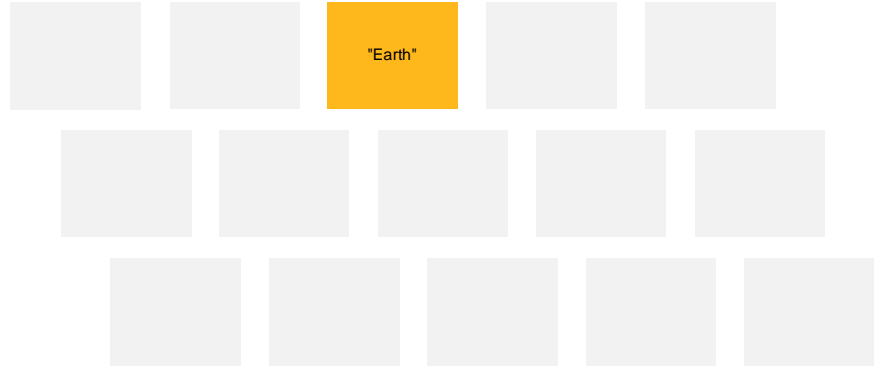


Example redacted from printed version

Some machine learning methods




Focus on "Earth"



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 109

Multivariate adaptive regression splines ("Earth")




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 110

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 111 111

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 111 111

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 113

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 114

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 115

Multivariate adaptive regression splines (“Earth”)

Categorical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 116

Multivariate adaptive regression splines (“Earth”)

Numerical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 117

Multivariate adaptive regression splines (“Earth”)

Numerical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 118

Multivariate adaptive regression splines (“Earth”)

Numerical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 119

Multivariate adaptive regression splines (“Earth”)

Numerical factors




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 120

Multivariate adaptive regression splines (“Earth”)

Interactions




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 121

Multivariate adaptive regression splines (“Earth”)

Interactions




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 122

Multivariate adaptive regression splines (“Earth”)

Interactions



Multivariate adaptive regression splines (“Earth”)

Advantages

- Minimum manual setup required
- Fast run time
- Highly interpretable results

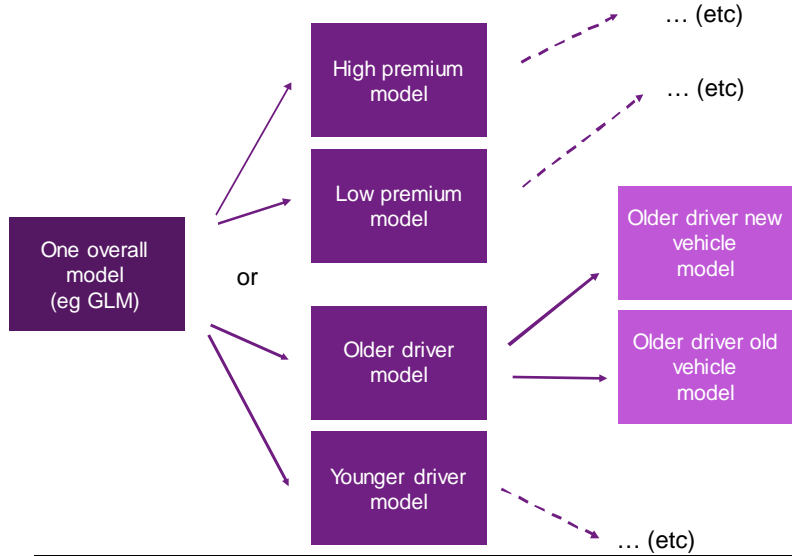
Disadvantages

- Model will contain discontinuities around knot points
- Hand-crafting likely to improve results

How might “Earth” be applied?

- Historically pricing models have been fit by coverage and/or peril – are these still the most suitable splits?
- When should models be split/combined? (e.g., homeowners and landlords policies or fire and lightning perils)
- How many models should we build and what should they predict?
- Increasing use of machine learning to answer these structural/strategic questions

Case study - model hierarchy



Case study - model hierarchy

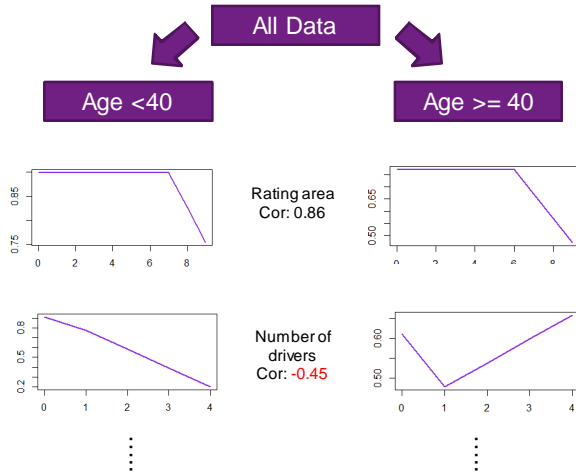
Automated evaluation of model structures

Split Points to Consider

- Policyholder age
- Vehicle age
- Premium size
- Payment method
- ...

Test Factors used for Evaluation

- Source
- Vehicle owned months
- Youngest additional driver age
- Days from cover start
- Vehicle kept overnight
- Class of use
- Claims free years
- Voluntary deductible
-



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

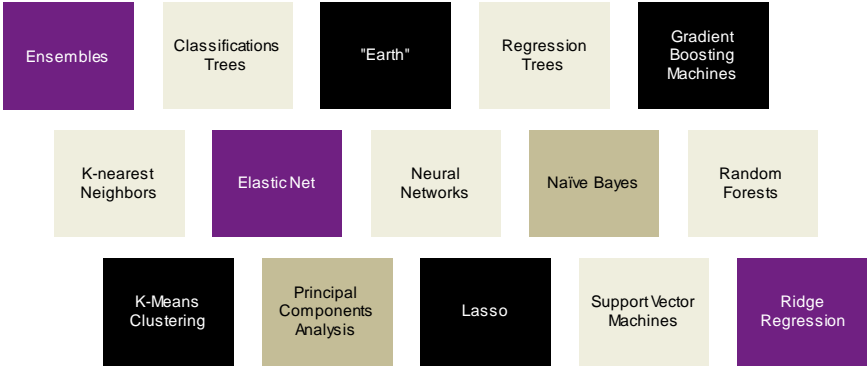
Willis Towers Watson 127



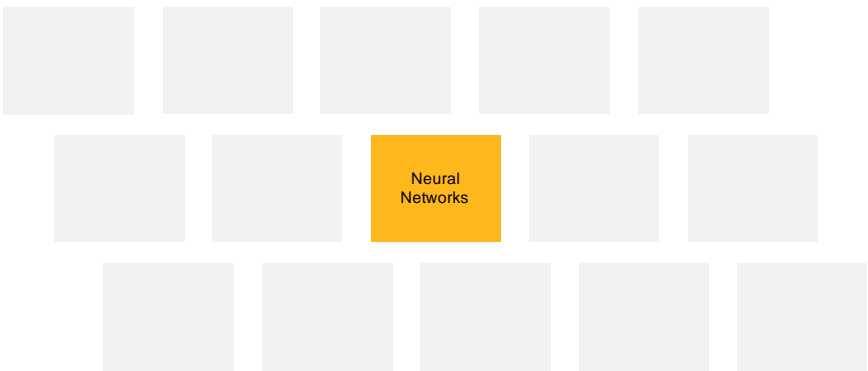
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 128

Some machine learning methods




Focus on Neural Networks



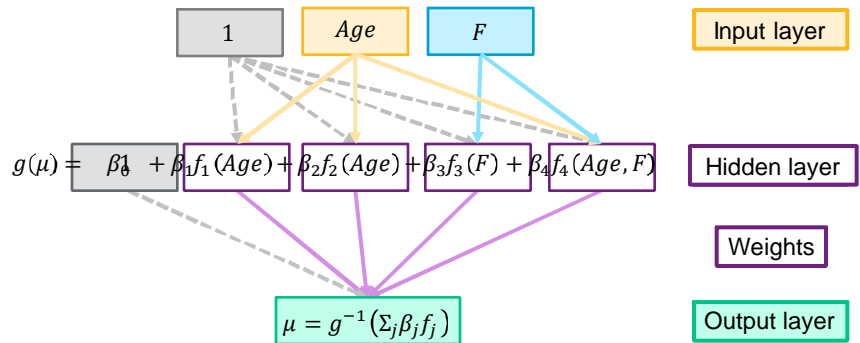
Start with a simple GLM...

- Log link function, g
- Age (piecewise-linear variates)
- F (indicator of Gender = Female)
- Age x Gender interaction

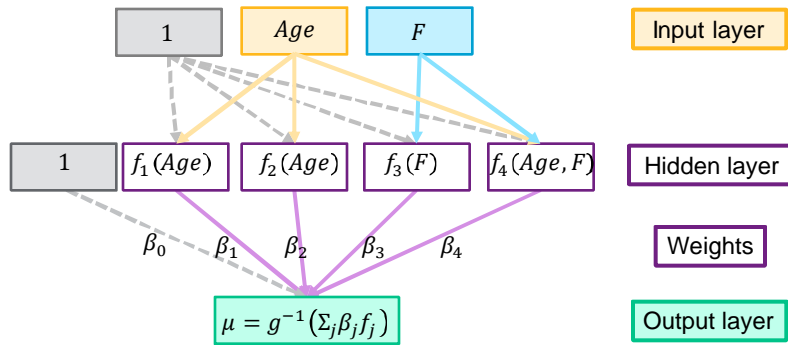
$$g(\mu) = \beta_0 + \beta_1 f_1(Age) + \beta_2 f_2(Age) + \beta_3 f_3(F) + \beta_4 f_4(Age, F)$$



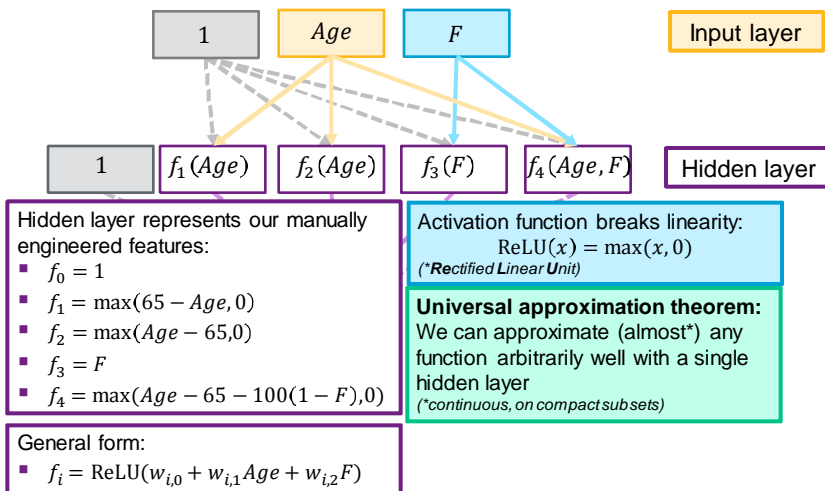
We can represent GLMs as a network...



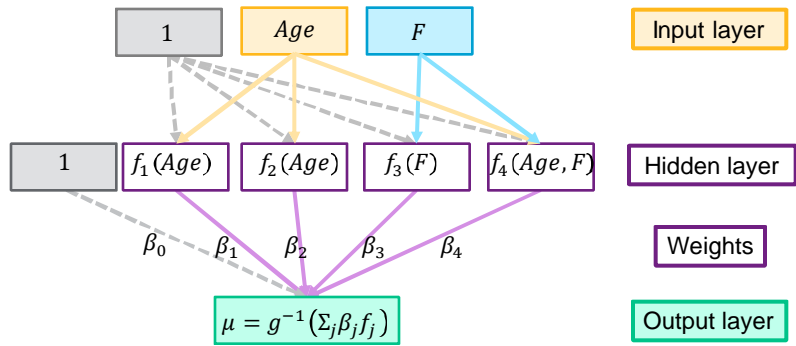
We can represent GLMs as a network...



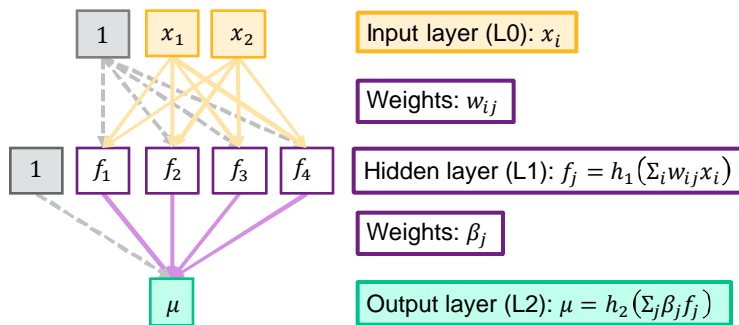
We can represent GLMs as a network...



We can represent GLMs as a network...




Generalizing to neural networks



Generalizing to neural networks


Model structure decisions



- Input features
- Number of hidden layers
- Size of each hidden layer
- Activation functions
 - Typically specified by layer
 - ReLU is most commonly used
- Connectivity of layers and weight sharing
 - Typically **fully connected** with **unique weights**
 - Many variants exist, eg: **Convolutional Neural Networks** for image classification connect nearby blocks of pixels and apply the same shared weights across each block

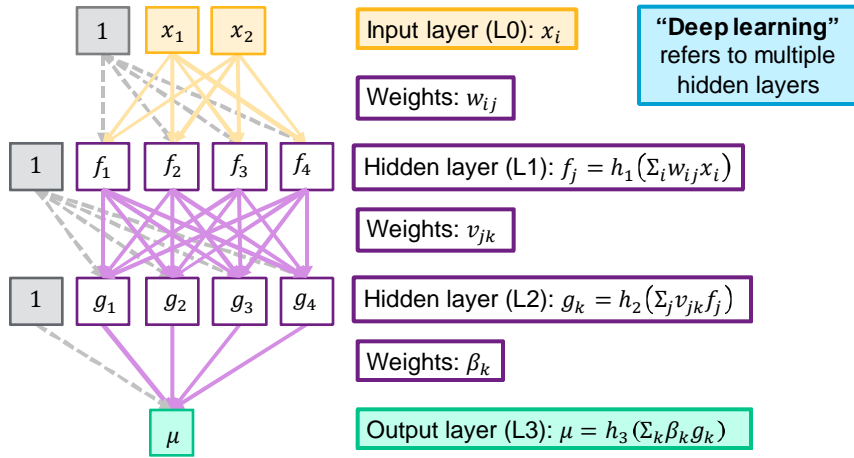
Generalizing to neural networks

Key model fitting decisions




- **Optimization algorithm**
 - Typically variants of **Back-Propagation**
- **Loss function** – to be minimized
- **Batch size** – number of rows to consider in each iteration
- **Epochs** – number of passes through full data
- **Initial weights**
- **Regularization parameters**, eg:
 - L1 / L2 penalties
 - Learning rate and decay
 - Dropout

Generalizing to neural networks




Where is the value?




Which policyholder is more likely to make a claim?



Where is the value?



Which picture is more likely to be of a cat?



willistowerswatson.com
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  141

Where is the value?



Which picture is more likely to be of a cat?




willistowerswatson.com
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  142

Neural networks

Evolution or revolution?



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  143

Neural networks

Case study – market models

Context

- UK aggregator sites provide some historic quote data
- We wanted a model of “Average top 5 premium” for auto quotes to understand the market’s pricing structure
- One month of data (~1m quotes)
- Limited subset of factors (no data enrichment beyond simple rating area & vehicle group)

Approach

- 60/40 split for training and holdout data
- Modelled as Log-Normal (ie $\ln(\text{Premium}) \sim N(\mu, \sigma^2)$) as Normal distributions well supported across packages
- Compare Neural Network performance to GLM (using existing model parameterizations) and GBM with RMSE of log-Premium on holdout data

willistowerswatson.com
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  144

Neural networks
Require some work!

Input layer

Dropout

Optimization algorithm

Output layer

Hidden layers

Epochs

Batch size


Learning rate

Regularization

Initial weights

Activation functions

willistowerswatson.com
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  145

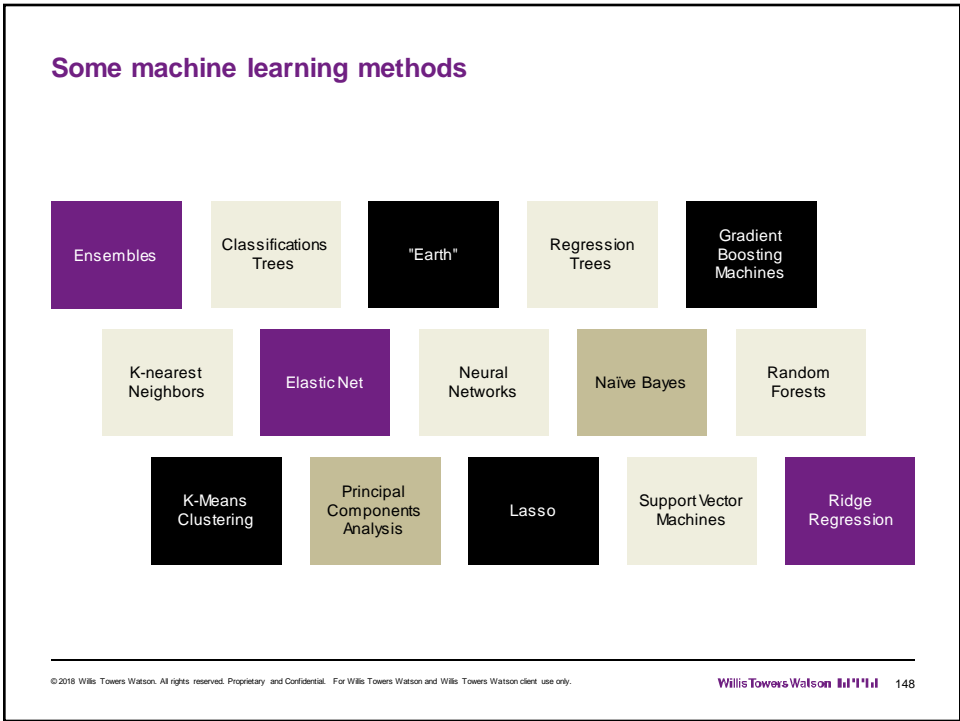
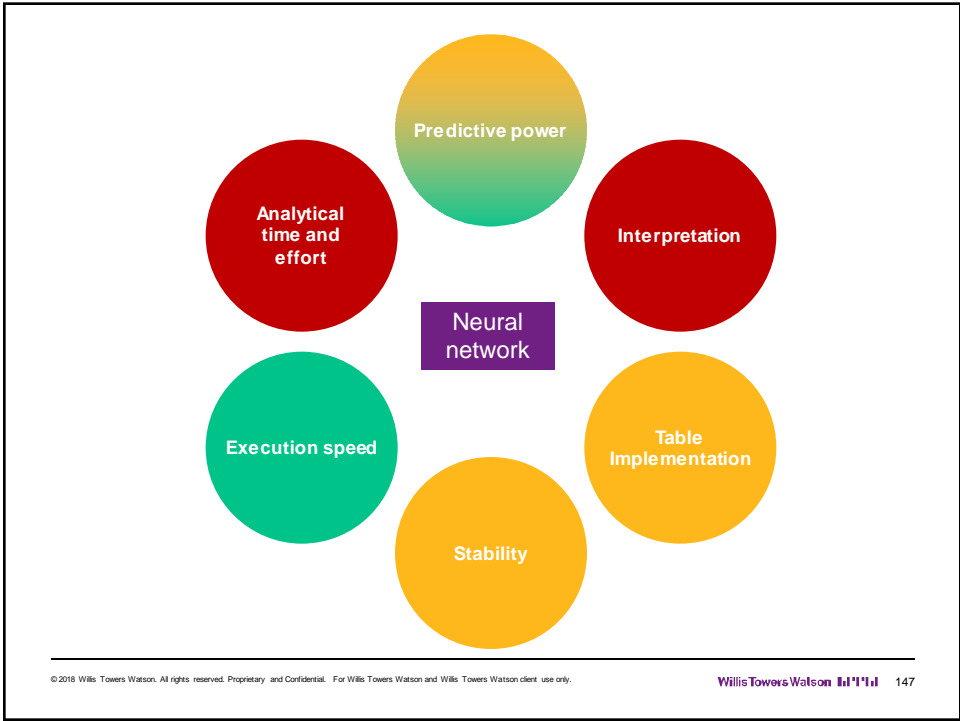
Neural networks



A photograph showing a hammer striking a nut, with a large white question mark overlaid on the scene. The nut is broken into pieces on the surface below.

willistowerswatson.com
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  146



Focus on Penalized Regression

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 149

Penalized Regression

Overview

GLMs

- Predictions are given by $f(\underline{x}) = g^{-1}(\mathbf{X}\cdot\underline{\beta})$
- $\underline{\beta}$ is estimated by minimizing a loss function $L(\underline{\beta}|\mathbf{X},\underline{y})$ (\mathbf{X} is data & model, \underline{y} the response)

Penalized regression

- The same, except the objective function becomes $L(\underline{\beta}|\mathbf{X},\underline{y}) + \lambda$. "Penalty on $\underline{\beta}$ "

Elastic Net

Minimize: $L(\beta|X,y) + \lambda_1 \sum_i |\beta_i| + \lambda_2 \sum_i \beta_i^2$

Lasso - just the blue part

- Penalty reduces insignificant parameter values to zero – useful for variable selection

Ridge - just the purple part regression models


- Penalty heavily penalize extreme parameters, but do not reduce parameters to zero

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 150

Penalized Regression

GLM

$f(x) = g^{-1}(X.\beta)$ where β estimated by minimizing $L(\beta|X, y)$



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 151


Penalized Regression

GLM

$f(x) = g^{-1}(X.\beta)$ where β estimated by minimizing

$$L(\beta|X, y) + \lambda_1 \sum_i |\beta_i| + \lambda_2 \sum_i \beta_i^2$$

Elastic Net




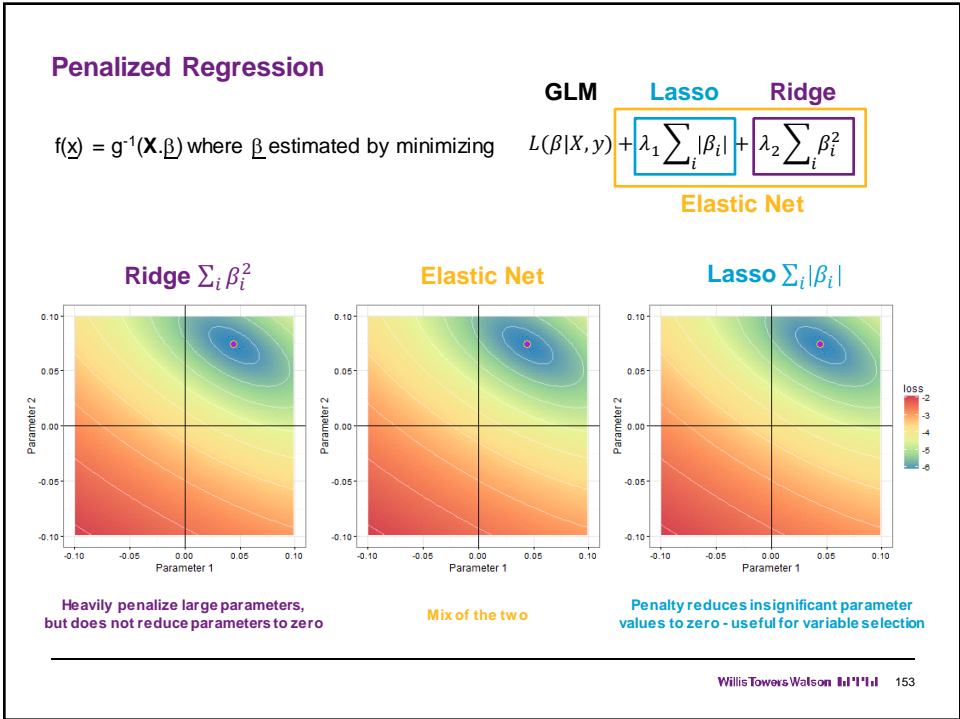
Heavily penalize large parameters, but does not reduce parameters to zero



Mix of the two

Penalty reduces insignificant parameter values to zero - useful for variable selection

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 152





Penalized Regression

Parameter selection

There are costs to allowing too many factors in our models

- Computational cost of processing more data / fitting more parameters
- Time cost of analysts needing to consider more potential effects
- Reduced comprehensibility of interplay of many different correlated effects in our models
- Financial cost of licensing and maintaining many different data sources, and hosting/updating tables to use them in rating
- Performance cost as increased number of tests makes it more likely that we will find false-positives and overfit to noise in our data

Penalized Regression

Case study – vehicle classification



Physical facticity

E.g., height, length, weight




Mechanical nature

E.g., engine size, fuel type



Qualitative descriptors

E.g., body type, model range



Performance

E.g., maximum speed, torque, BHP


Penalized Regression

Vehicle classification – categorical factors

Exposure	# Claims	Policy Factors	Ext Code	Vehicle Make	Engine Size
1	0	...	000001	Ford	1400
1	1	...	000002	Porsche	3000
0.5	0	...	000001	Ford	1400
1	0	...	000001	Ford	1400
0.5	1	...	000003	Honda	1300
1	0	...	000002	Porsche	3000
1	0	...	000001	Ford	1400
0.5	0	...	000003	Honda	1300
0.3	0	...	000003	Honda	1300
1	1	...	000002	Porsche	3000
1	0	...	000001	Ford	1400
...

Make = Ford	Make = Honda	Make = Porsche
1	0	0
0	0	1
1	0	0
1	0	0
0	1	0
0	0	1
1	0	0
0	1	0
0	1	0
0	0	1
1	0	0
...

- One 0-1 column per level (excluding base)
- Equivalent to adding a “simple factor” to a GLM




Penalized Regression

Vehicle classification – numerical factors

Exposure	# Claims	Policy Factors	Ext Code	Vehicle Make	Engine Size
1	0	...	000001	Ford	1400
1	1	...	000002	Porsche	3000
0.5	0	...	000001	Ford	1400
1	0	...	000001	Ford	1400
0.5	1	...	000003	Honda	1300
1	0	...	000002	Porsche	3000
1	0	...	000001	Ford	1400
0.5	0	...	000003	Honda	1300
0.3	0	...	000003	Honda	1300
1	1	...	000002	Porsche	3000
1	0	...	000001	Ford	1400
...

Engine Size = 1300	Engine Size = 3000
0	0
0	1
0	0
0	0
1	0
0	1
0	0
0	0
1	0
1	0
0	1
0	0
...	...

- Adding one 0-1 column per value/band allows full flexibility, but loses knowledge of ordering




Penalized Regression

Vehicle classification – numerical factors

Exposure	# Claims	Policy Factors	Ext Code	Vehicle Make	Engine Size
1	0	...	0000001	Ford	1400
1	1	...	0000002	Porsche	3000
0.5	0	...	0000001	Ford	1400
1	0	...	0000001	Ford	1400
0.5	1	...	0000003	Honda	1300
1	0	...	0000002	Porsche	3000
1	0	...	0000001	Ford	1400
0.5	0	...	0000003	Honda	1300
0.3	0	...	0000003	Honda	1300
1	1	...	0000002	Porsche	3000
1	0	...	0000001	Ford	1400
...

Engine Size	(Engine Size)^2	...	(Engine Size)^5
1400	1960000	...	5.38E+15
3000	9000000	...	2.43E+17
1400	1960000	...	5.38E+15
1400	1960000	...	5.38E+15
1300	1690000	...	3.71E+15
3000	9000000	...	2.43E+17
1400	1960000	...	5.38E+15
1300	1690000	...	3.71E+15
1300	1690000	...	3.71E+15
3000	9000000	...	2.43E+17
1400	1960000	...	5.38E+15
...

- Adding one 0-1 column per value/band allows full flexibility, but loses knowledge of ordering
- Adding variates retains ordering, but limits flexibility
 - Model fit also impacted by scale of x-values as parameters are scaled, affecting the penalty size
 - Orthogonal variates/splines can help with scaling and convergence



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 161


Penalized Regression

Vehicle classification

Exposure	# Claims	Policy Factors	Ext Code	Vehicle Make	Engine Size
1	0	...	0000001	Ford	1400
1	1	...	0000002	Porsche	3000
0.5	0	...	0000001	Ford	1400
1	0	...	0000001	Ford	1400
0.5	1	...	0000003	Honda	1300
1	0	...	0000002	Porsche	3000
1	0	...	0000001	Ford	1400
0.5	0	...	0000003	Honda	1300
0.3	0	...	0000003	Honda	1300
1	1	...	0000002	Porsche	3000
1	0	...	0000001	Ford	1400
...

Engine Size <= 1300	...	Engine Size <= 3000
0	...	1
0	...	1
0	...	1
0	...	1
1	...	1
0	...	1
0	...	1
1	...	1
1	...	1
0	...	1
0	...	1
...

- Adding one 0-1 column per value/band allows full flexibility, but loses knowledge of ordering
- Adding variates retains ordering, but limits flexibility
 - Model fit also impacted by scale of x-values as parameters are scaled, affecting the penalty size
 - Orthogonal variates/splines can help with scaling and convergence
- Adding a series of "less than or equal" indicators retains as much flexibility as a column per band, and also retains knowledge of ordering



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

WillisTowersWatson 162

Deploying Penalized Regression

Same as GLMs!

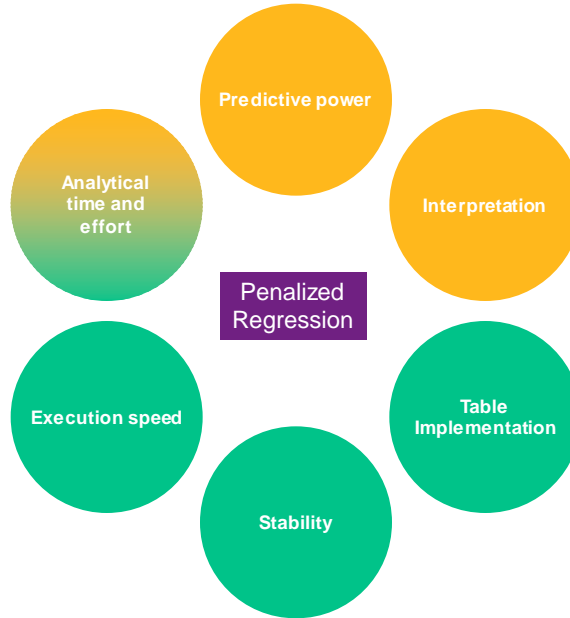
	Age	Exposure	Loss Cost
1	<=20	1,720	179
2	21-30	34,893	122
3	31-50	118,182	102
4	51+	127,054	70
5	Age Total	281,849	91

	Vehicle Group	Exposure	Loss Cost
1	1-10	164,107	77
2	11-14	84,859	101
3	15-18	28,952	116
4	19-20	3,931	272
5	VG Total	281,849	91

	Gender	Exposure	Loss Cost
1	Male	197,339	92
2	Female	84,510	87
3	Gender Total	281,849	91

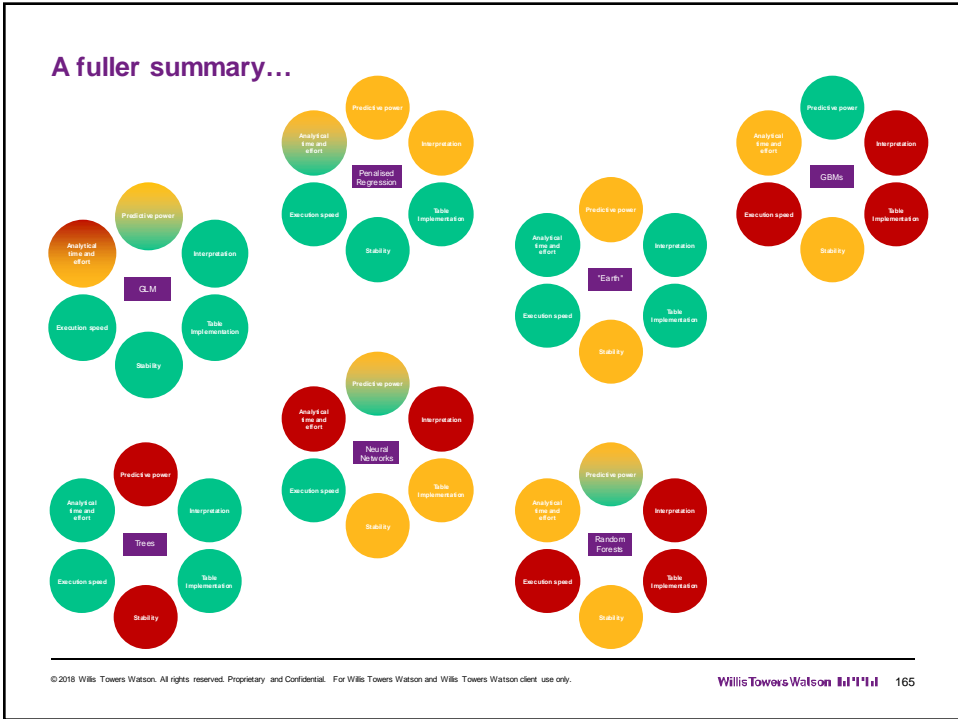
© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 163




© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 164



Machine learning in pricing

Conclusions (Part 2)



- Machine learning brings a proliferation of new methods
- Improving models is more than just finding the best method. Consider:
 - What data are available and how can data be transformed to give insight
 - What is the optimal model structure and target variable?
 - How can information be transferred between models?
- Earth is a fast, interpretable method that can improve overall lift by informing when/where to segment models
- Neural networks are complex and require numerous input decisions; analyzing unstructured data (e.g., imagery) is an intuitive application for this method ... but where else may it be helpful?
- Penalized regression can aid in factor selection decisions and may in fact be a good method in its own right – particularly when the modeler has less of a “feel” for the data
- Machine learning in pricing is not all about improving predictive power. Consider:
 - Fast investigation of new data
 - Quick assessment and response of emerging experience

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only. Willis Towers Watson 166


So what? How is the US market doing with machine learning
Some critical success factors

Component	Rating	Directional trend
Data availability		Static
Appetite to try new approaches		
Installing tools and platforms		
Internal skills sets		
Managing vendors		
Application		

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 167

What are the three biggest challenges preventing your company from becoming more data driven? (Q.21)



Base: U.S. respondents (n = 51)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 168


Top-growing new data sources for insurers

	Now	Two years
Personal lines		
Smart home/smart building data	0%	52%
Usage-based insurance/telematics	26%	70%
Social media	26%	52%
Unstructured internal claim information	39%	61%
Unstructured internal underwriting information	30%	52%
Images	13%	35%
Commercial lines		
Unstructured internal claim information	46%	92%
Other unstructured customer information	11%	54%
Unstructured internal underwriting information	25%	39%
Usage-based insurance/telematics	11%	47%
Web/clickstream/phone/email customer interactions	11%	36%
Images	3%	39%

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 169

Approximately what percentage of the external data that you collect is proprietary as opposed to open source? (Q.24)



Base: U.S. respondents (n = 51)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 170


So what? How is the US market doing with machine learning
Some critical success factors

Component	Rating	Directional trend
Data availability		Stable
Appetite to try new methods		Slowly upward
Modeling tool capabilities		
Internal skills sets		
Management beliefs		
Application		

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  171

So what? How is the US market doing with machine learning
Methods used



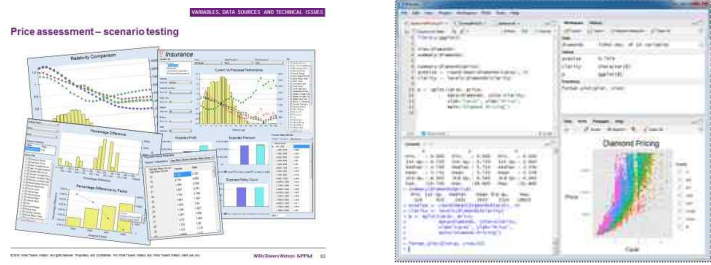
Base: U.S. respondents using advanced analytics for underwriting/pricing (n = 49), claims (n = 37) and/or marketing (n = 33)

Willis Towers Watson  172

So what? How is the US market doing with machine learning

Some critical success factors

Component	Rating	Directional trend
Top management		Stable
Appetite to try new approaches		
Modeling tools and platforms		Slowly upward
Internal skills sets		
External vendors		
Application		



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 173

Cloud-based environments and Hadoop

Regardless of size, insurers are actively exploring technology to manage big data

	Large		Medium		Small	
	Now	Exploring	Now	Exploring	Now	Exploring
Cloud-based (Amazon Web Services, Azure)	19%	48%	7%	50%	0%	40%
Hadoop	19%	37%	7%	14%	0%	20%

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 174


So what? How is the US market doing with machine learning
Some critical success factors

Component	Rating	Directional trend
Internal skill sets	?	Slowly upward
Application		
Costs		
Appetite to try new approaches		
Customer and business outcomes		
Business opportunity		

"We're also seeing an influx of quantitative talent to the insurance industry. In addition to actuaries, insurers are hiring statisticians, data scientists, marketing scientists and behavioral scientists. The industry is challenging these professionals to solve a wider range of problems across the customer value chain"

- Recent article by Claudine Modlin and Graham Wright

What are the three biggest challenges preventing your company from becoming more data driven? (Q.21)



Base: U.S. respondents (n = 51)


So what? How is the US market doing with machine learning
Some critical success factors

Component	Rating	Directional trend
Data availability		Static
Appetite to try new approaches		
Installing tools and platforms		
Internal skills sets		
Measuring value		Static
Application		

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 177

How do you determine the value of your advanced analytic models? (Q.11)
How well understood are your advanced analytic models by those who need to use them, outside of the modeling team? (Q.12)



Base: U.S. respondents using advanced analytics to evaluate fraud potential (n = 41)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 178


So what? How is the US market doing with machine learning
Some critical success factors

Component	Rating	Directional trend
Non-predictive		Stable
Appetite to try new approaches		
Installing tools and platforms		
Internal skills sets		
Measuring value		Slowly upward
Application	?	Slowly upward

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 179

For which aspects of underwriting/pricing does your company group currently use or plan to use advanced analytics? (Q.2)




Base: U.S. respondents using or planning to use advanced analytics for underwriting/pricing (n = 51)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 180

For which aspects of claims does your company group currently use or plan to use advanced analytics? (Q.4)




Base: U.S. respondents using or planning to use advanced analytics for claims (n = 39)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 181

Beyond underwriting/pricing and claims, in which other areas does your company group currently use, or plan to use, advanced analytics? (Q.9)



Base: Total U.S. respondents (n = 51)

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson 182

How AI and machine learning will streamline processes

Top applications insurers plan to use two years from now for artificial intelligence (AI) and machine learning

	Now	Two years
Reduce time spent by humans	8%	49%
Identify high-risk cases	10%	45%
Build risk models for better decision making	8%	45%
Help humans identify appropriate risk attributes	6%	43%
Better understand risk drivers	20%	41%
Identify patterns of fraudulent claims	6%	39%
Augment human-performed underwriting	6%	37%



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

So what? How is the US market doing with machine learning

Some critical success factors


Component	Rating	Directional trend
Data availability		Static
Appetite to try new approaches		Slowly upward
Modeling tools		Slowly upward
Internal skills sets	?	Slowly upward
Measuring value		Static
Application	?	Slowly upward

© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

The data and analytics “iceberg”

Evolution or cyclical? Innovation and standardization in the analytical approach

- | | | | |
|------------------|-------------------|-------------------------|-----------------------|
| Early GLM | Mature GLM | Early ML/Current | Mature ML/Next |
| ▪ Early 2000s | ▪ 2010 to 2015 | ▪ 2015 to 2017? | ▪ 2017 onwards? |



- | | | | |
|--|---|---|--|
| <ul style="list-style-type: none"> ▪ Data not organized for analytics ▪ Statistical / coding skills ▪ Business leaders inexperienced ▪ Regulators unfamiliar | <ul style="list-style-type: none"> ▪ Data organized ▪ Speed of thought analytics ▪ Business & regulators experienced ▪ Workflow, governance, deployment | <ul style="list-style-type: none"> ▪ Data not yet organized for ML ▪ “Coding” type skills ▪ Numerous tools ▪ Business leaders interested, concerned ▪ Regulators - ? | <ul style="list-style-type: none"> ▪ Data organized for ML ▪ Integrated tools with slick interface ▪ Speed of thought analytics ▪ Business comfortable ▪ Regulators - ? |
|--|---|---|--|

Machine learning beyond pricing



- Carriers are experimenting with ML, it is becoming established within insurance analytics
- It opens up a broader set of problems to analytics, and offers a broader tool set for familiar problems
- New (wider) data beats new methods – think UBI!
- Factor definition, problem specification and method selection are critical for success
- There’s opportunity to reveal actionable, first-order insights in applications to which analytics have not been deployed previously
- With this broad new opportunity, spotting strong initial use cases is important

Questions



© 2018 Willis Towers Watson. All rights reserved. Proprietary and Confidential. For Willis Towers Watson and Willis Towers Watson client use only.

Willis Towers Watson  187