



# CAS Ratemaking & Product Management Seminar

## Overview and Practical Application of Machine Learning Methods in Pricing – Part 1

Wednesday March 27, 2019

Ben Williams, Graham Wright

# Agenda

## Agenda

### Context of machine learning in pricing

---

#### Session 1:

---

**Decision trees**  
**Random forests**  
**Gradient boosting machines**

---

#### Session 2:

---

“Earth”  
Penalized regression  
Neural networks

---

### Conclusions



---

### Q&A

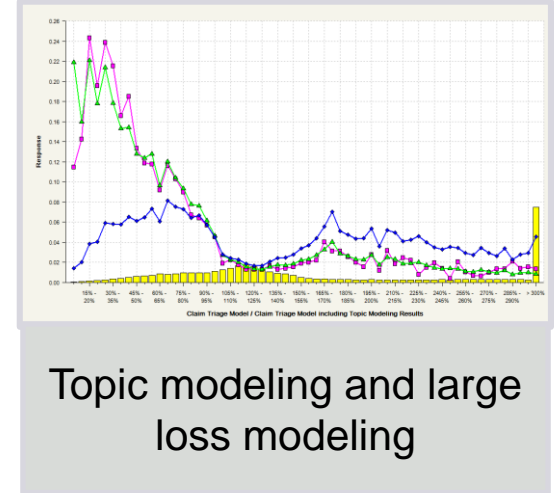


---

**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

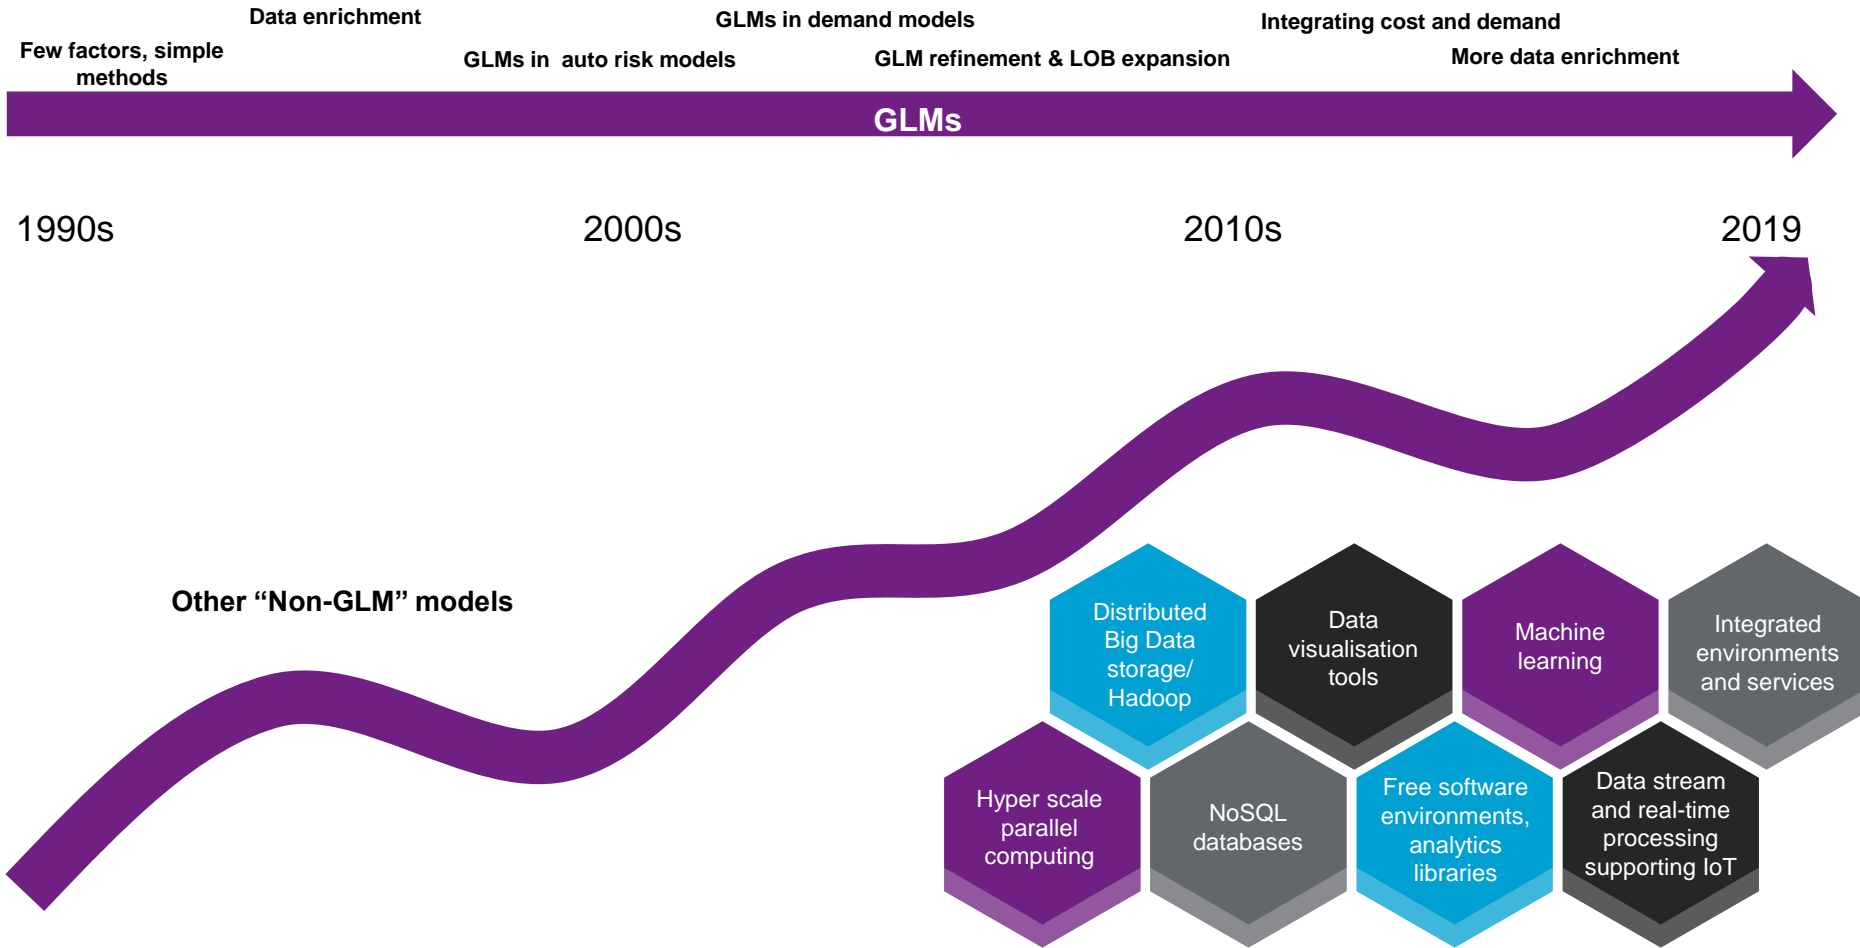
# Who's interested in what?



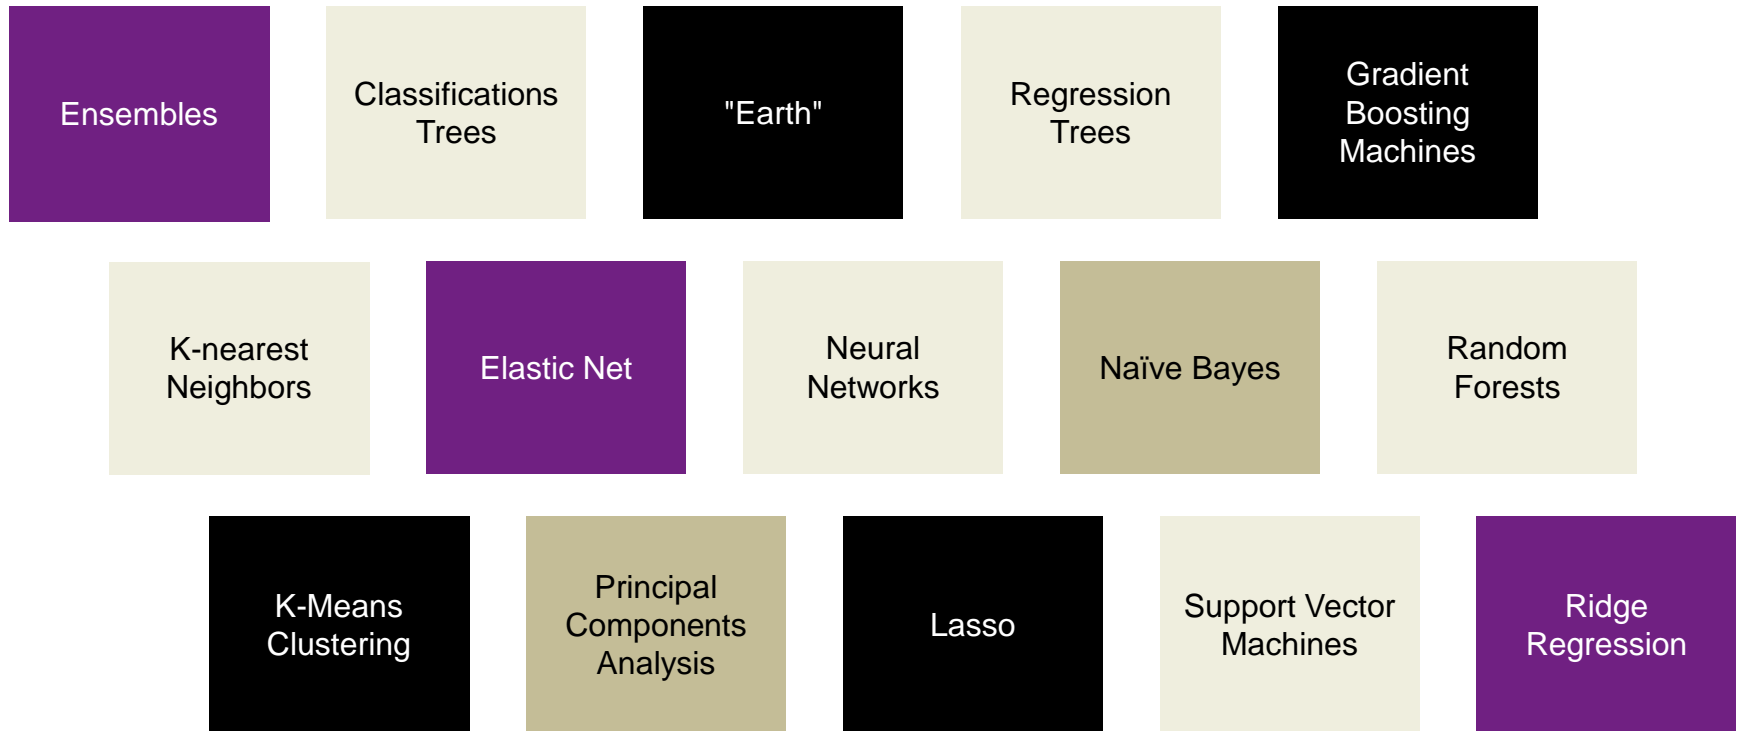
# Applications of machine learning in the insurance sector



# This is not new....




# What are these machine learning methods?



[kaggle](#)
Host
Competitions
Datasets
Scripts
Jobs
Community ▾
Sign up
Login


Welcome to Kaggle's data science competitions.

[New to Data Science? Tutorials on the Titanic competition >](#)  
[Want to learn from other's code? Kaggle's top rated scripts >](#)




**Download**

Choose a competition & download the training data.



**Build**















Build a model using whatever methods and tools you prefer.



**Submit**

Upload your predictions. Kaggle scores your solution and shows your score on the leaderboard.









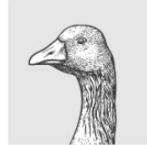


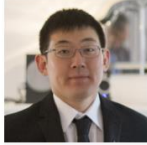

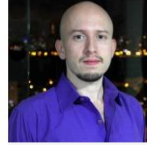

**Active Competitions**

All Competitions	Active Competitions		
      	 <p><b>State Farm Distracted Driver Detection</b> Can computer vision spot distracted drivers?</p>	<p>3 months 239 teams 110 scripts \$65,000</p>	
	 <p><b>Santander Customer Satisfaction</b> Which customers are happy customers?</p>	<p>18 days 3894 teams 2478 scripts \$60,000</p>	
	 <p><b>Home Depot Product Search Relevance</b> Predict the relevance of search results on homedepot.com</p>	<p>11 days 1944 teams 1486 scripts \$40,000</p>	
	 <p><b>BNP Paribas Cardif Claims Management</b> Can you accelerate BNP Paribas Cardif's claims management process?</p>	<p>4.4 days 2947 teams 1692 scripts \$30,000</p>	
	 <p><b>2016 US Election</b> Explore data related to the 2016 US Election</p>	<p>339 scripts 699 downloads</p>	
	 <p><b>2013 American Community Survey</b> Find insights in the 2013 American Community Survey</p>	<p>1077 scripts 1098 downloads</p>	
	 <p><b>World Development Indicators</b> Explore country development indicators from around the world</p>	<p>147 scripts 1694 downloads</p>	

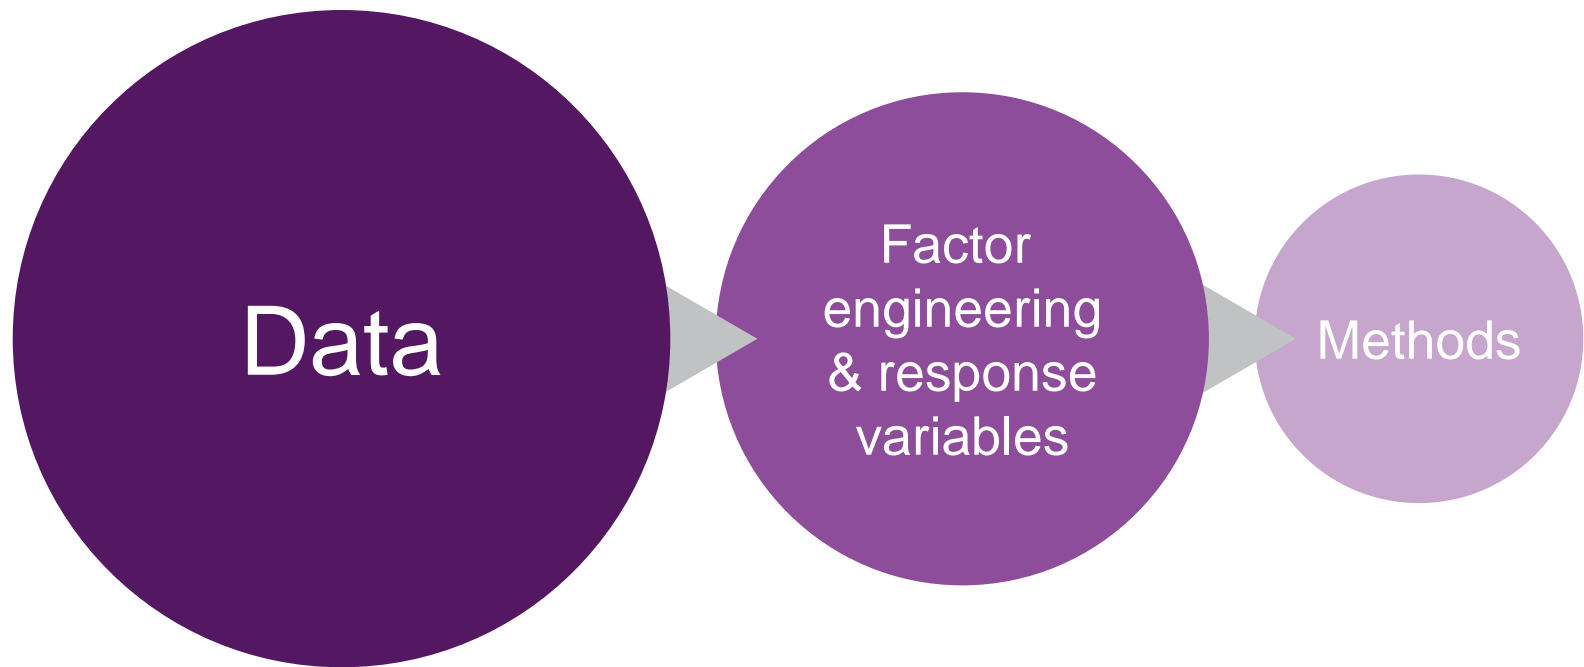
[kaggle](#)
Host
Competitions
Datasets
Scripts
Jobs
Community ▾
Sign up
Login

## Kaggle Rankings

Kaggle users are allocated points for their performance in competitions. This page shows the current global ranking. For more information on how we calculate points, please visit the [user ranking wiki page](#).


1st	191,154 pts	2nd	189,482 pts	3rd	163,407 pts	4th	144,134 pts	5th	139,658 pts
	<b>Gilberto Titericz</b> 66 competitions Curitiba Brazil		<b>Μαριος Μιχαηλιδης</b> 72 competitions Volos Greece		<b>Stanislav Semenov</b> 31 competitions Moscow Russian Federation		<b>Owen</b> 42 competitions NYC United States		<b>Kohei</b> 70 competitions Tokyo Japan
	<b>Alexander Guschin</b> 21 competitions Moscow Russia		<b>Abhishek</b> 97 competitions Berlin Germany		<b>Leustagos</b> 45 competitions Belo Horizonte Brazil		<b>Cardal</b> 4 competitions Israel		<b>Gert</b> 24 competitions Goes The Netherlands
	<b>y</b> 55 competitions South Korea		<b>Mike Kim</b> 48 competitions Washington DC United States		<b>clustifier</b> 56 competitions Israel		<b>Mario Filho</b> 17 competitions Sao Paulo Brazil		<b>utility</b> 15 competitions Moscow Russian Federation

# Is it really all about the method?

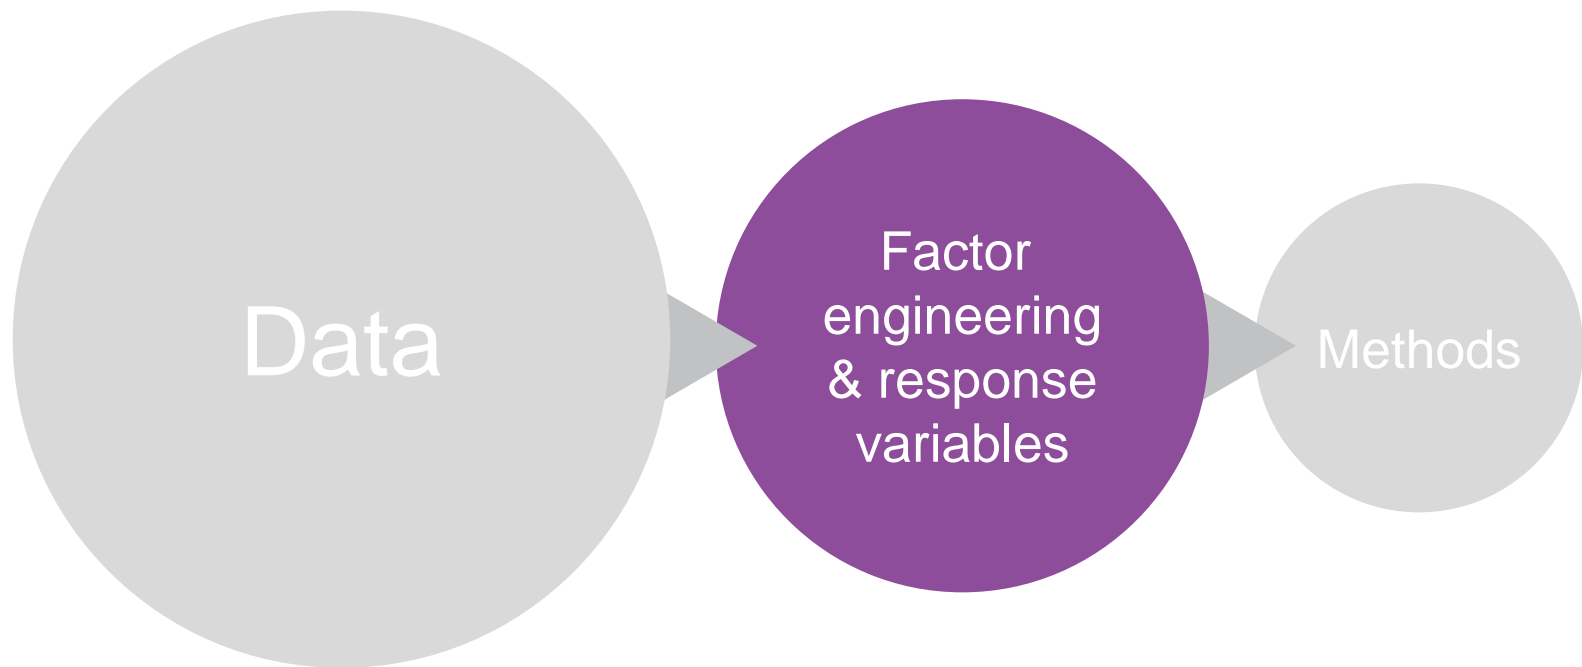





# Is it really all about the method?



## Is it really all about the method?




# 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)

## Example results

Model	Gini
GLM	0.327

## Example results

Model	Gini
GLM	0.327
New Model	0.330

## Example results

Model	Gini	Gini improvement
GLM	0.327	0.0%
New Model	0.330	1.0%




## Example results

Model	Gini	Gini improvement	Gini rank
GLM (main factor removed)	0.318	-2.6%	4
GLM (minor factor removed)	0.322	-1.3%	3
GLM	0.327	0.0%	2
New Model	0.330	1.0%	1

## But...

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

# Double lift chart



## Financial value estimate

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




Example results redacted from printed version

## Example results

Model	Gini	Gini improvement	Gini rank	Loss ratio @ elasticity 6	Loss ratio rank	Loss ratio @ elasticity 2	Loss ratio rank
GLM (main factor removed)	0.318	-2.6%	4	-0.9%	4	-0.4%	4
GLM (minor factor removed)	0.322	-1.3%	3	-0.4%	3	-0.2%	3
GLM	0.327	0.0%	2	0.0%	2	0.0%	2
New Model	0.330	1.0%	1	2.2%	1	0.5%	1

# Financial value vs Gini

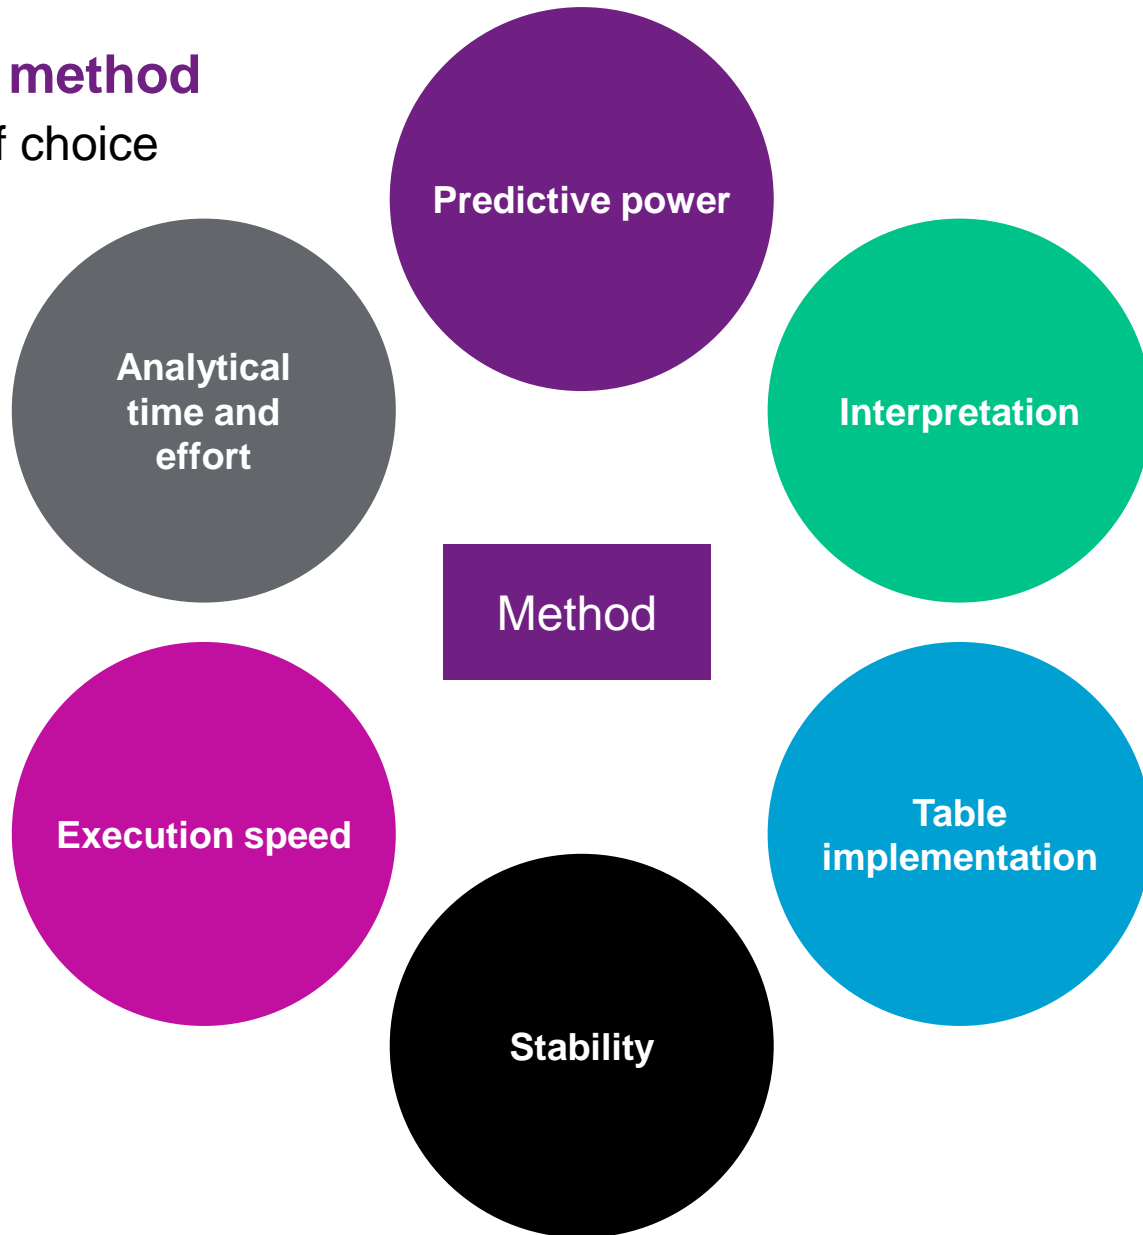


**Is there more to it...?**

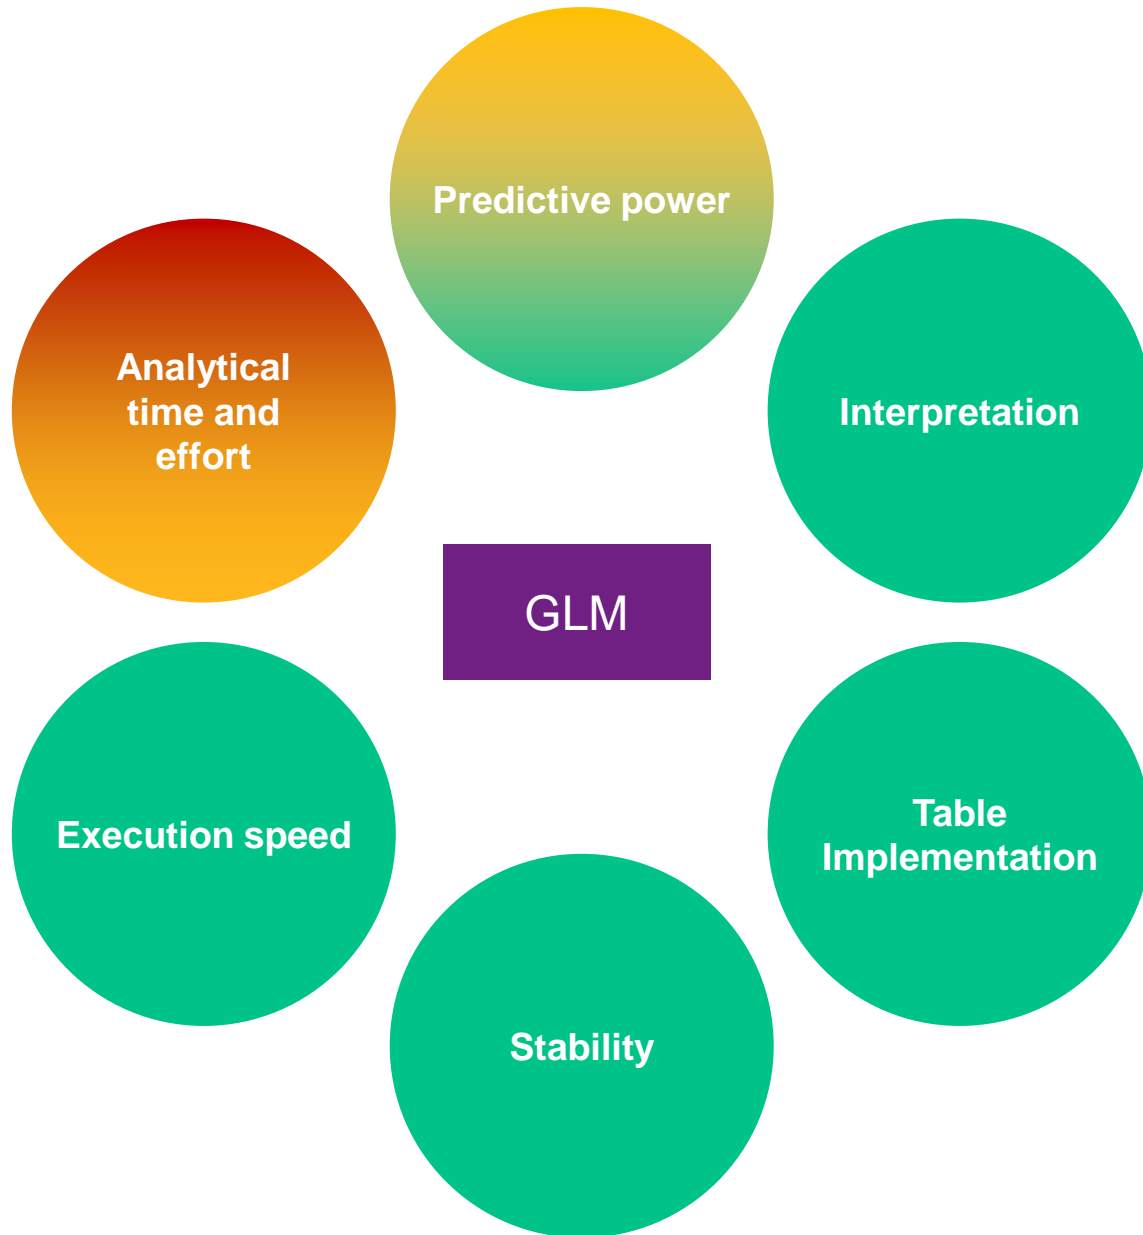


# Choosing a method


Dimensions of choice








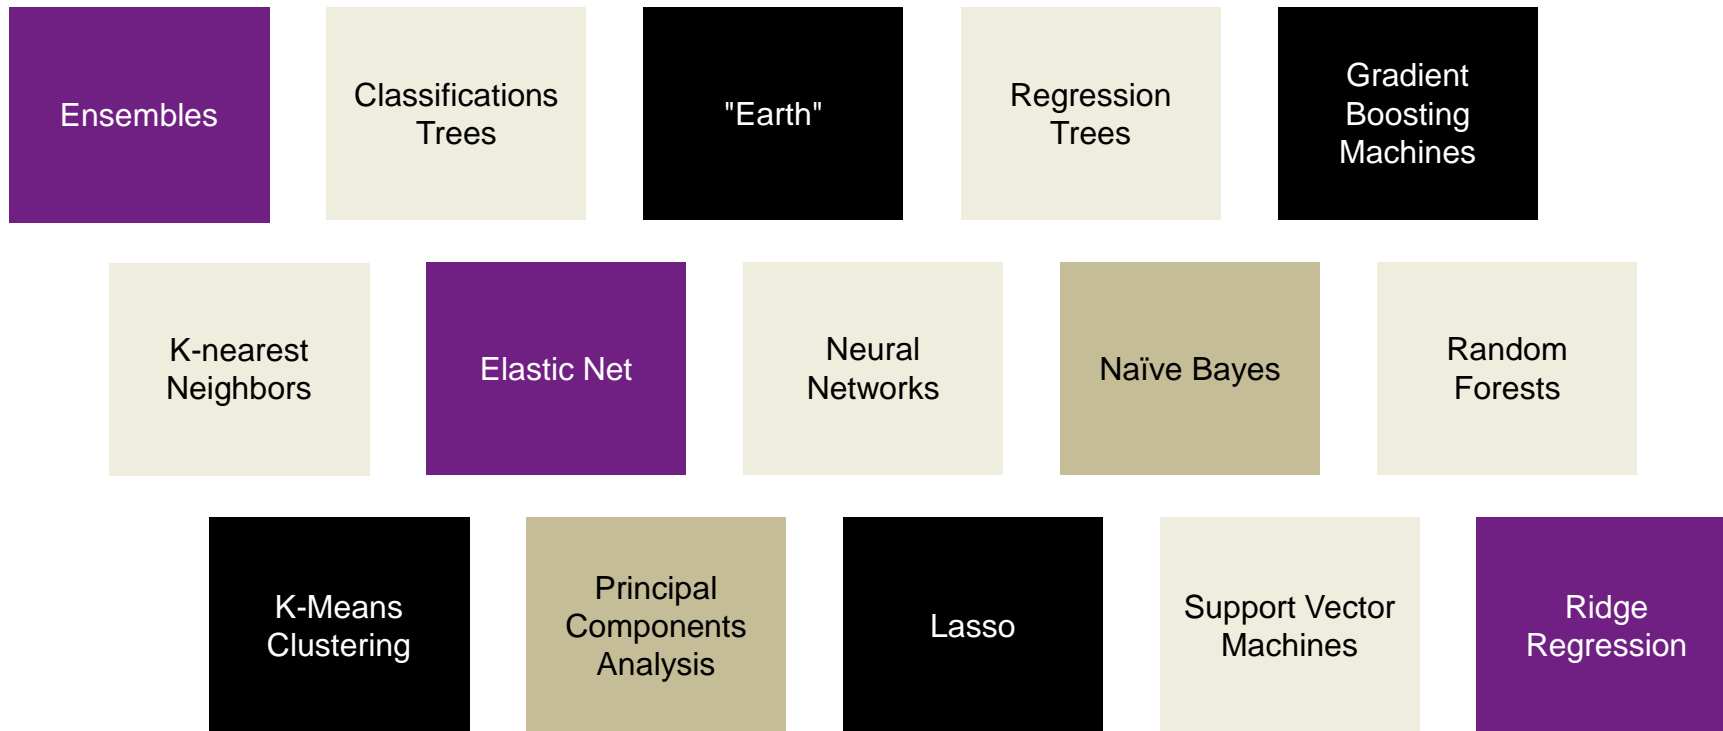
# What do you use where?



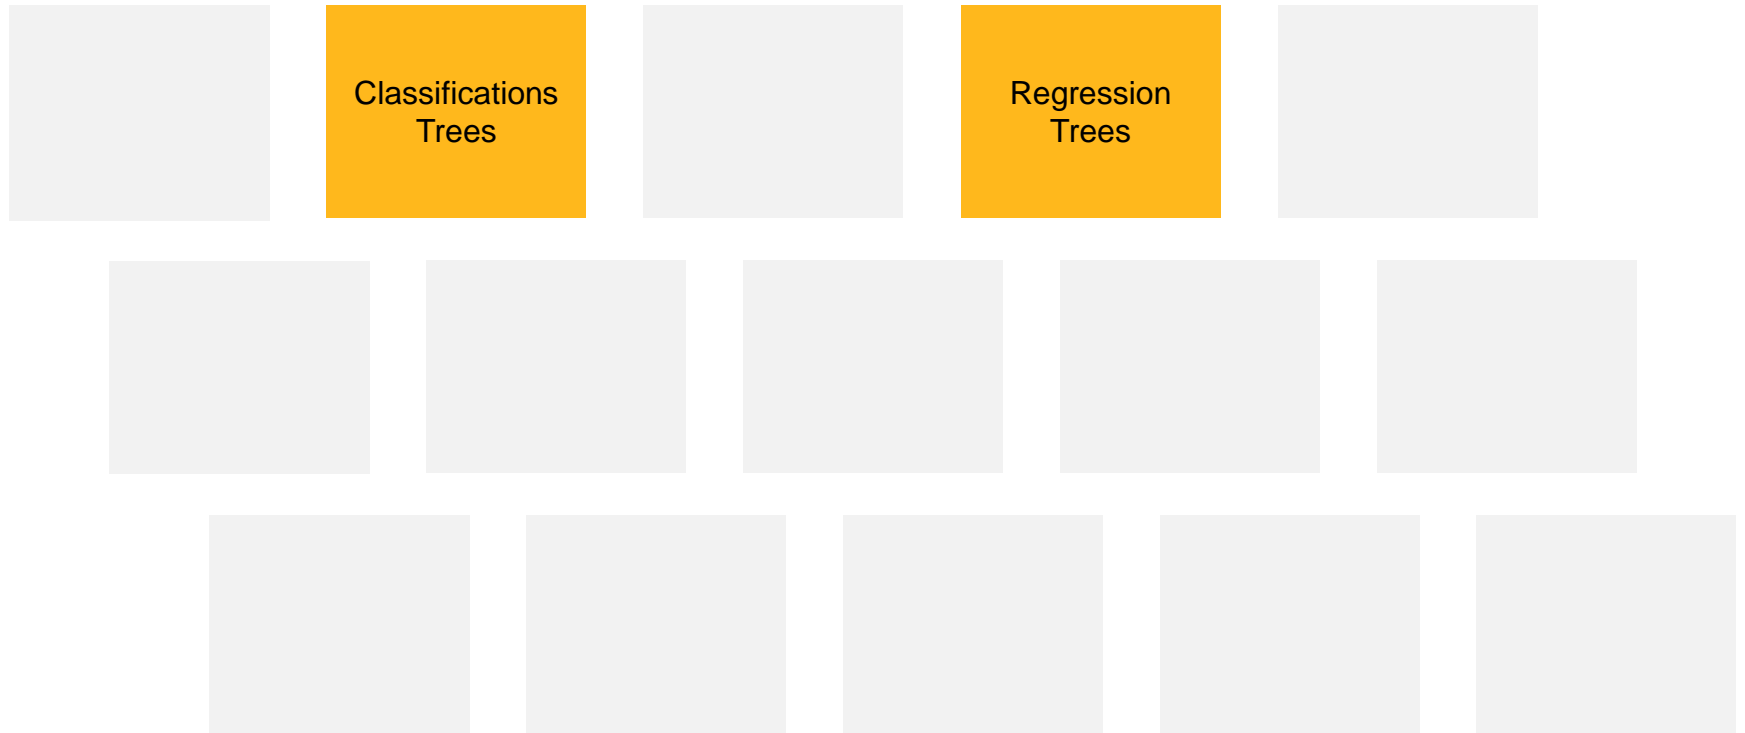
# It's domain expertise that helps decide





# Some machine learning methods





# Focus on Trees



# Decision Trees




# Decision Trees



# A simple Tree example



Tree results






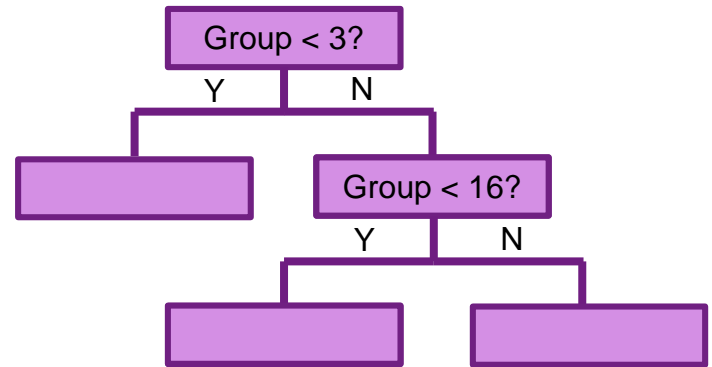
# A simple Tree example

Tree results




# A simple Tree example

Tree results




# A simple Tree example

Tree results




# Shortcomings of using trees


They may miss interactions...

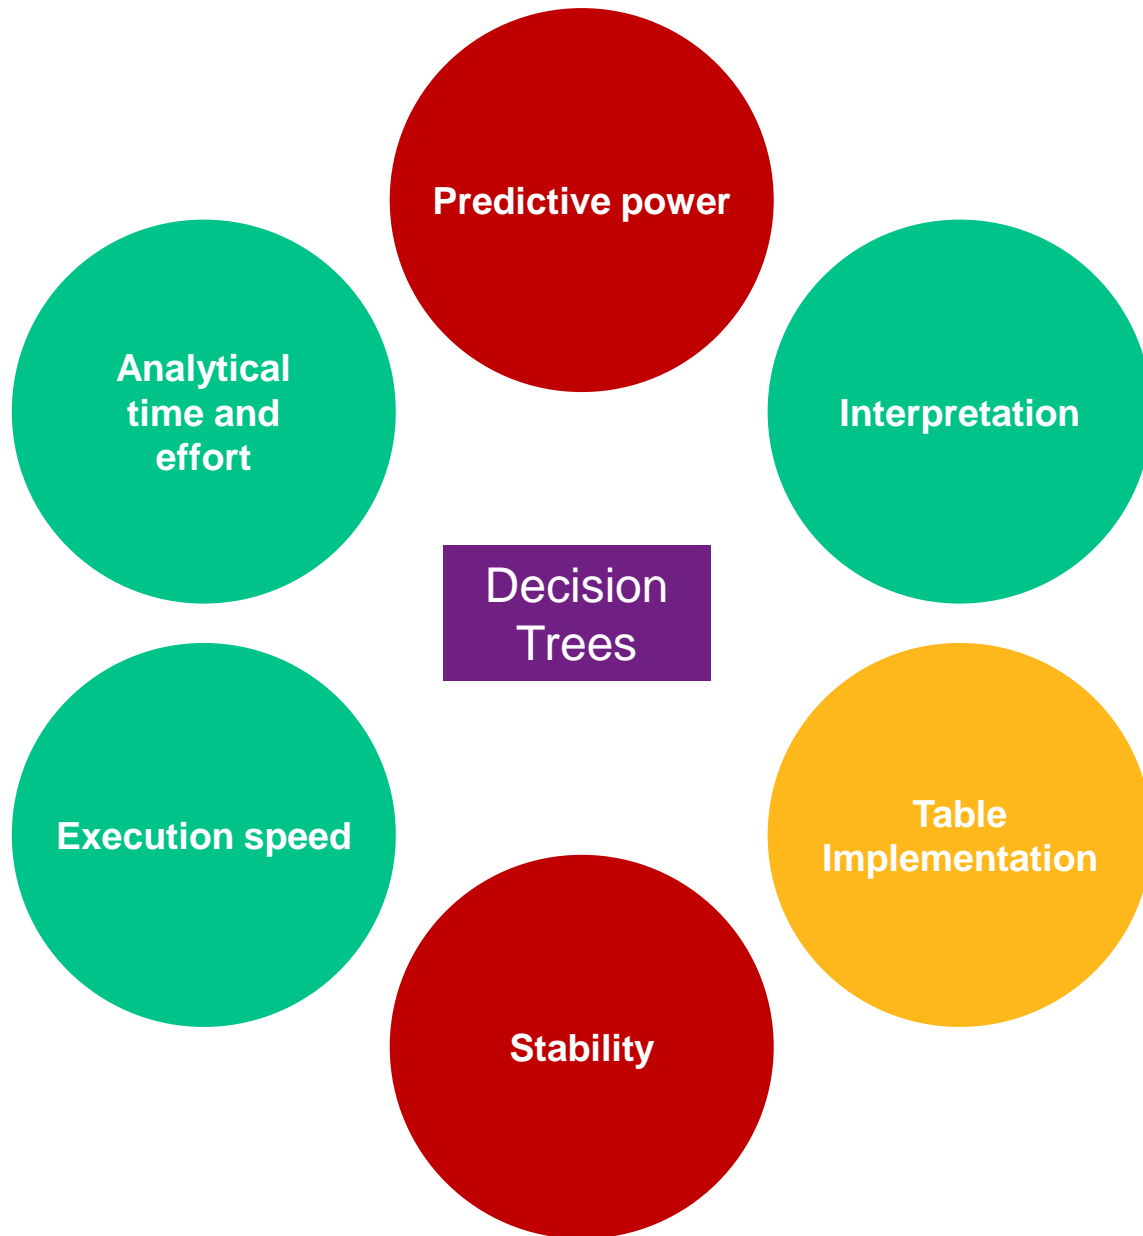


...and they can be bad at turning points

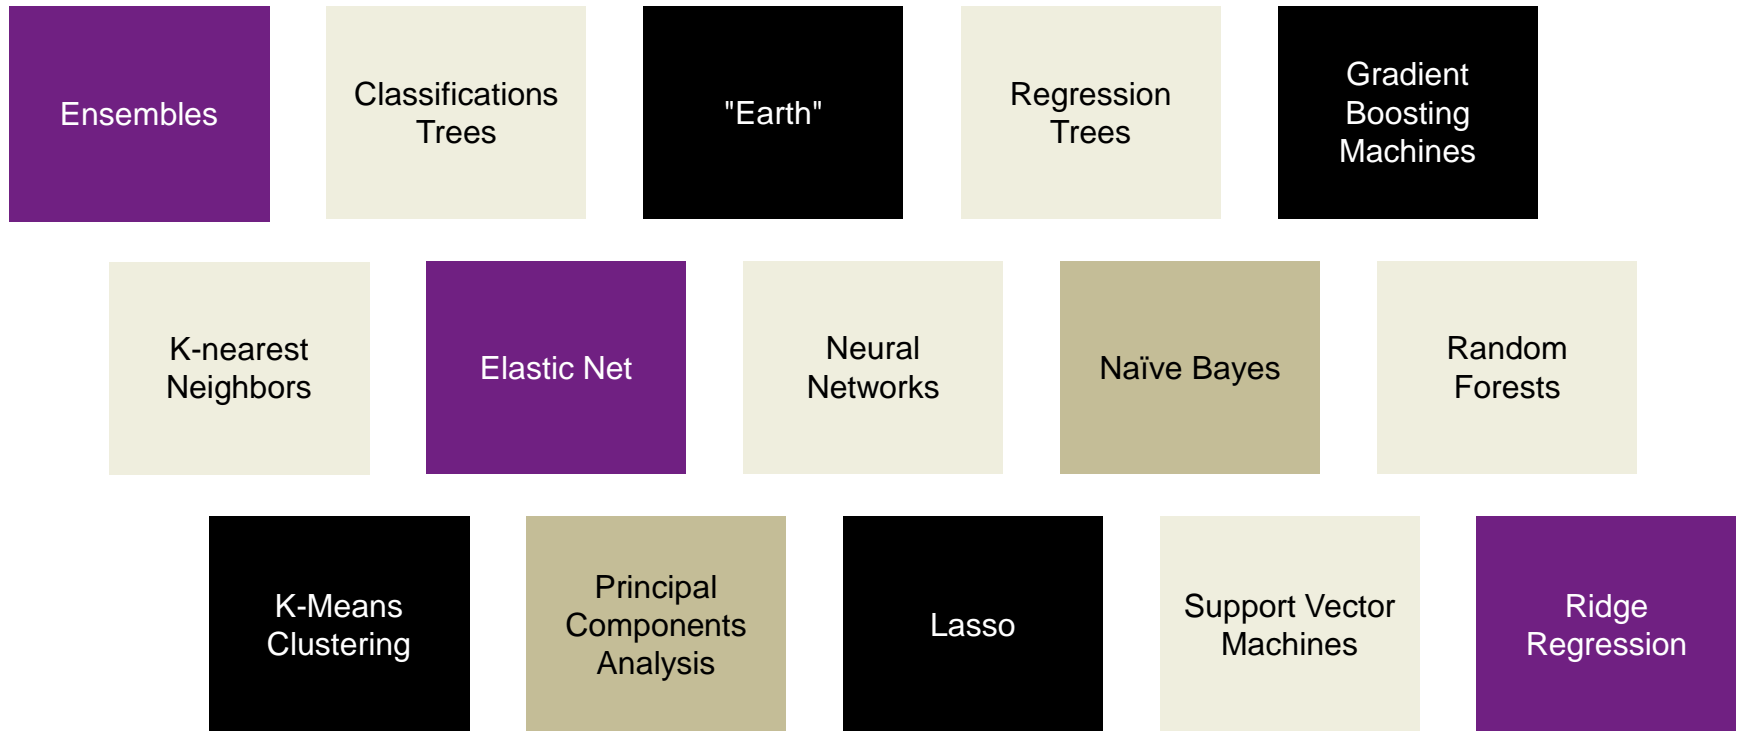


... they may struggle with categorical variables....






# Some machine learning methods



# 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:


Prediction = AVERAGE(Tree Predictions)

= AVERAGE(Tree Signal) + AVERAGE(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

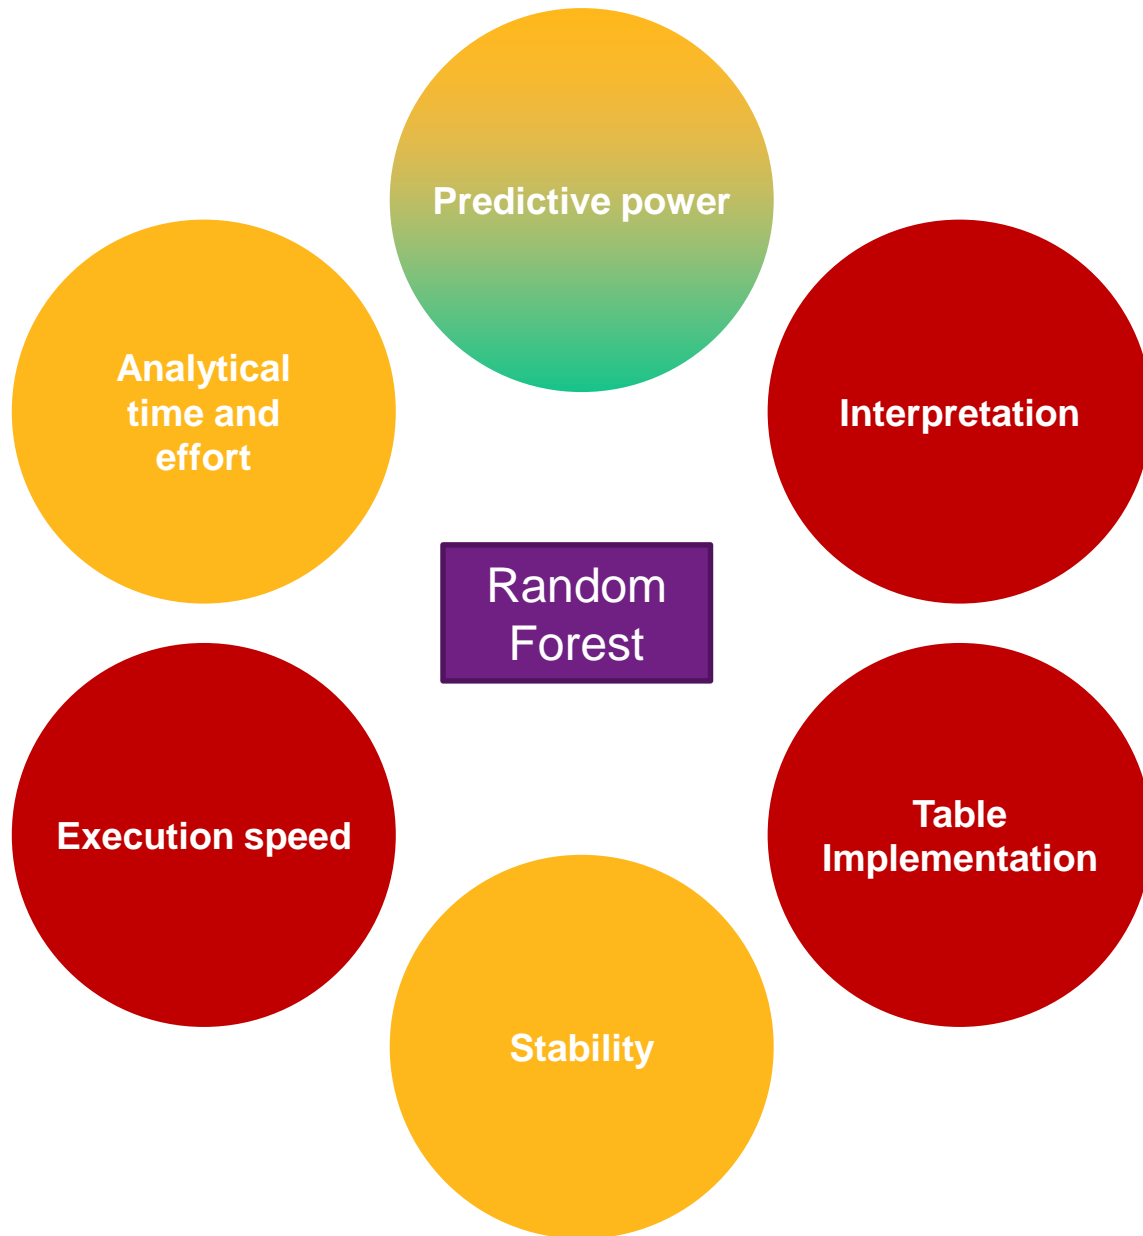


# A simple Random Forest example




# A simple Random Forest example






# Some machine learning methods



# Focus on Gradient Boosting Machines






# Gradient Boosted Machine or “GBM”


**A tree**

$f_i(x)$




**A GBM**

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




# Four main assumptions

- $\lambda$  **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




# A simple GBM example

GBM results at iteration 0




# A simple GBM example

GBM results at iteration 0




# A simple GBM example

GBM results at iteration 0




# A simple GBM example




# A simple GBM example

GBM results at iteration 1




# 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




# A simple GBM example

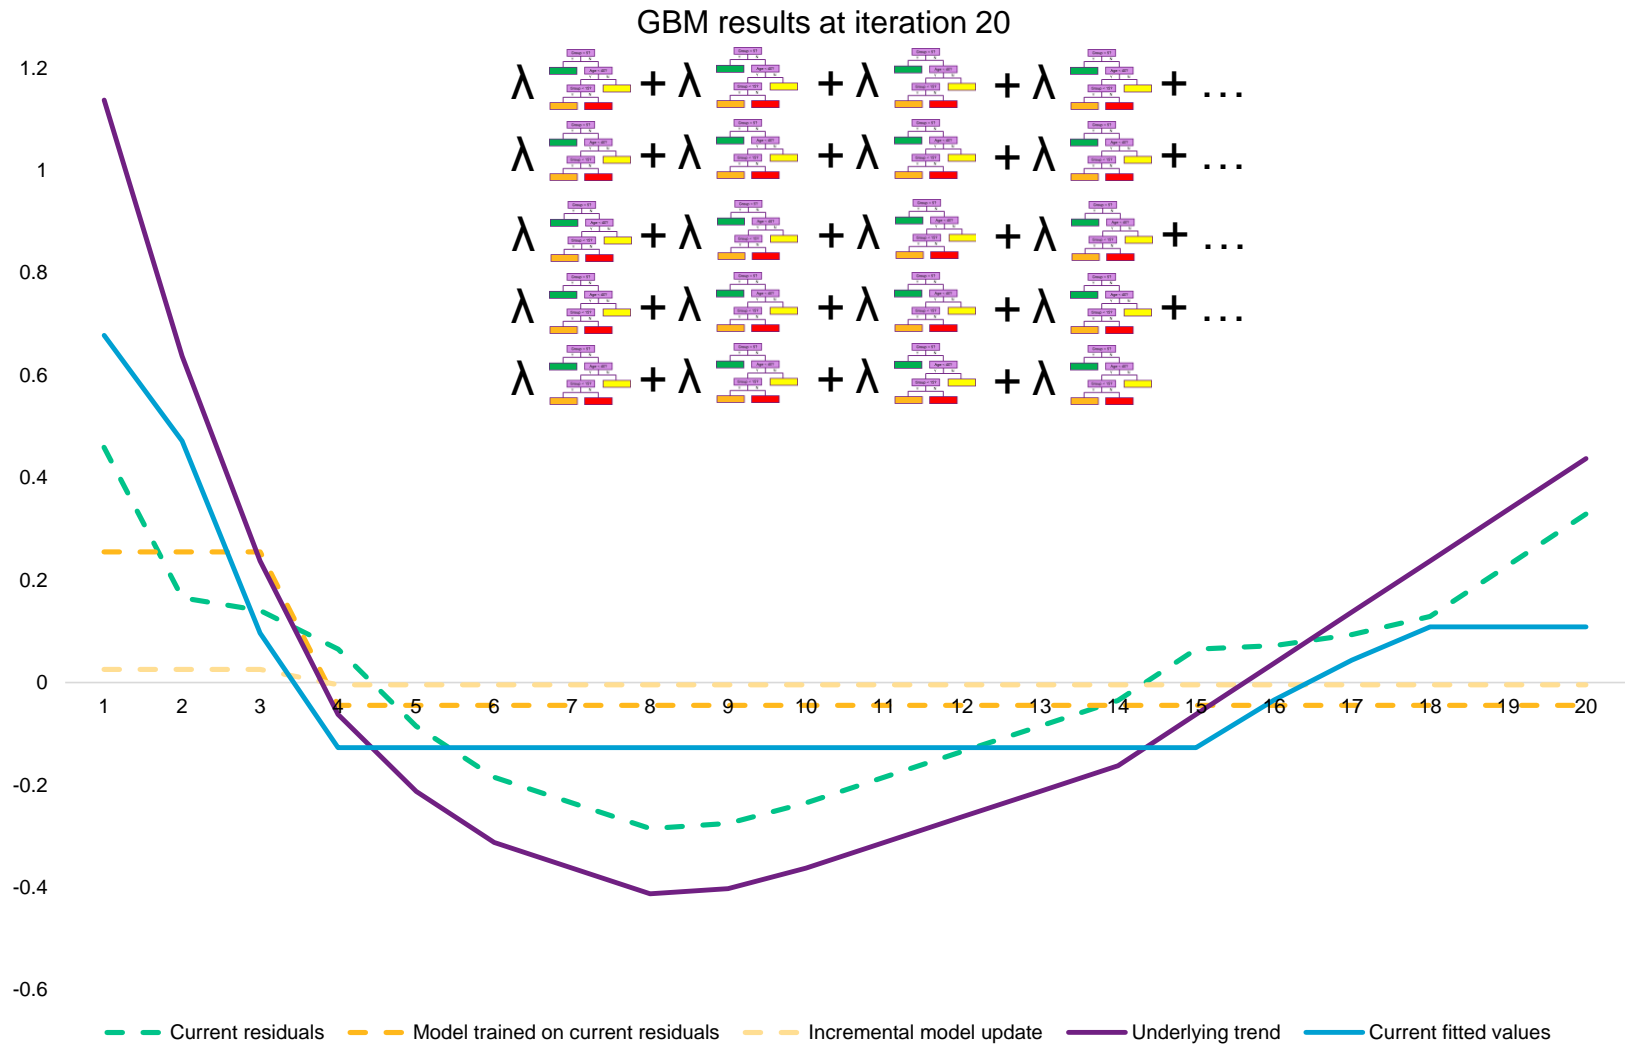





# A simple GBM example



# A simple GBM example




# A simple GBM example




# A simple GBM example

GBM results at iteration 40




# A simple GBM example

GBM results at iteration 50




# A simple GBM example

GBM results at iteration 100



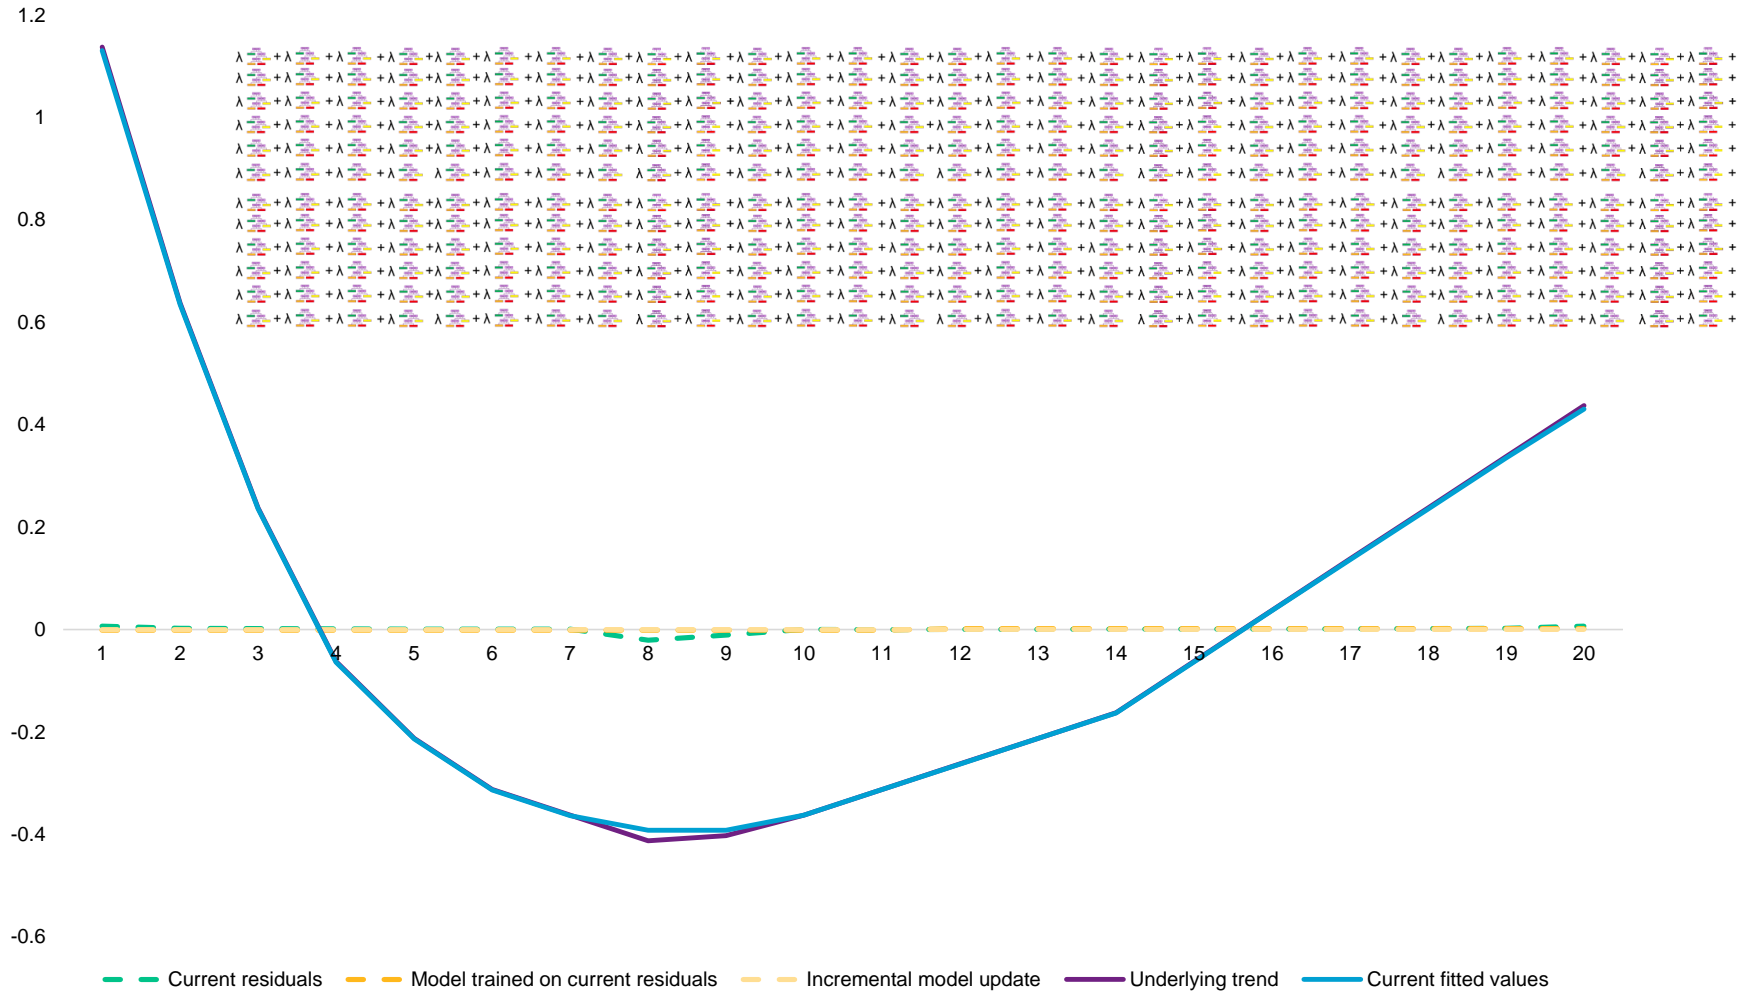
# A simple GBM example

GBM results at iteration 200



# A simple GBM example


GBM results at iteration 300





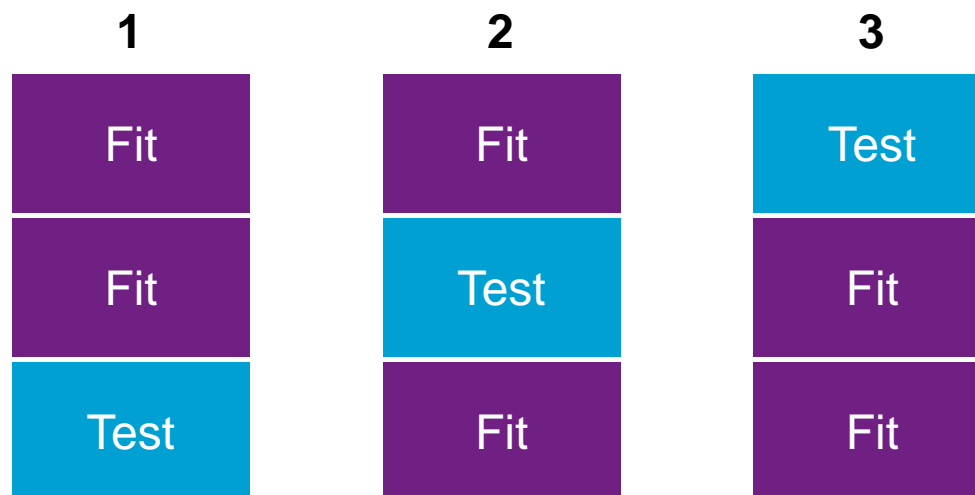
# A simple GBM example

GBM results at iteration 1,000



## 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

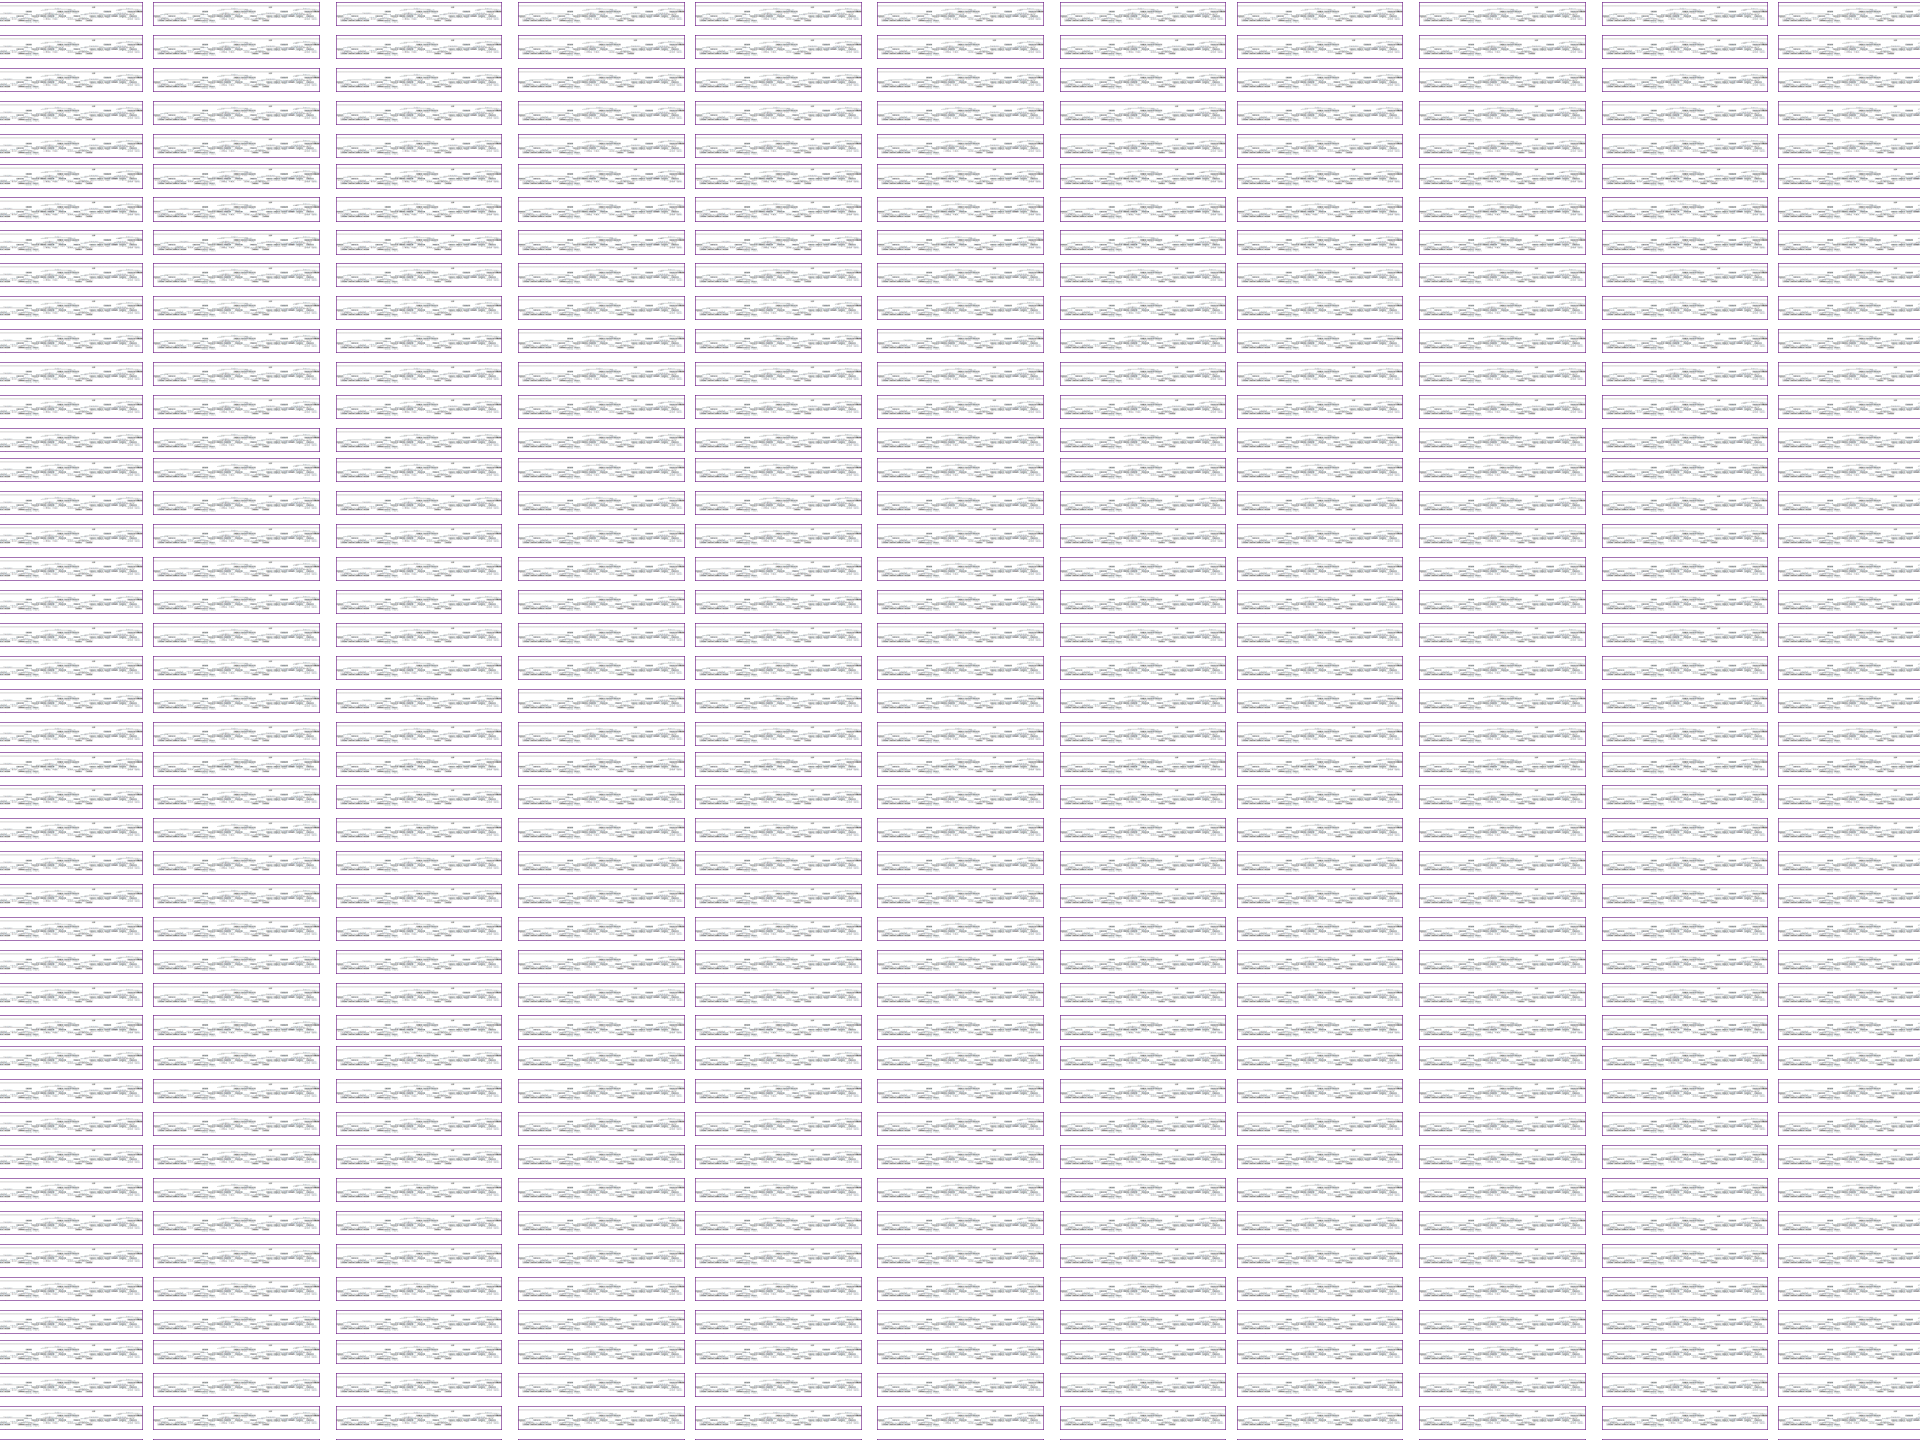


# What does a GBM look like?



# What does a GBM look like?






- Does it work?
- How does it work?

# Factor importance – relative influence

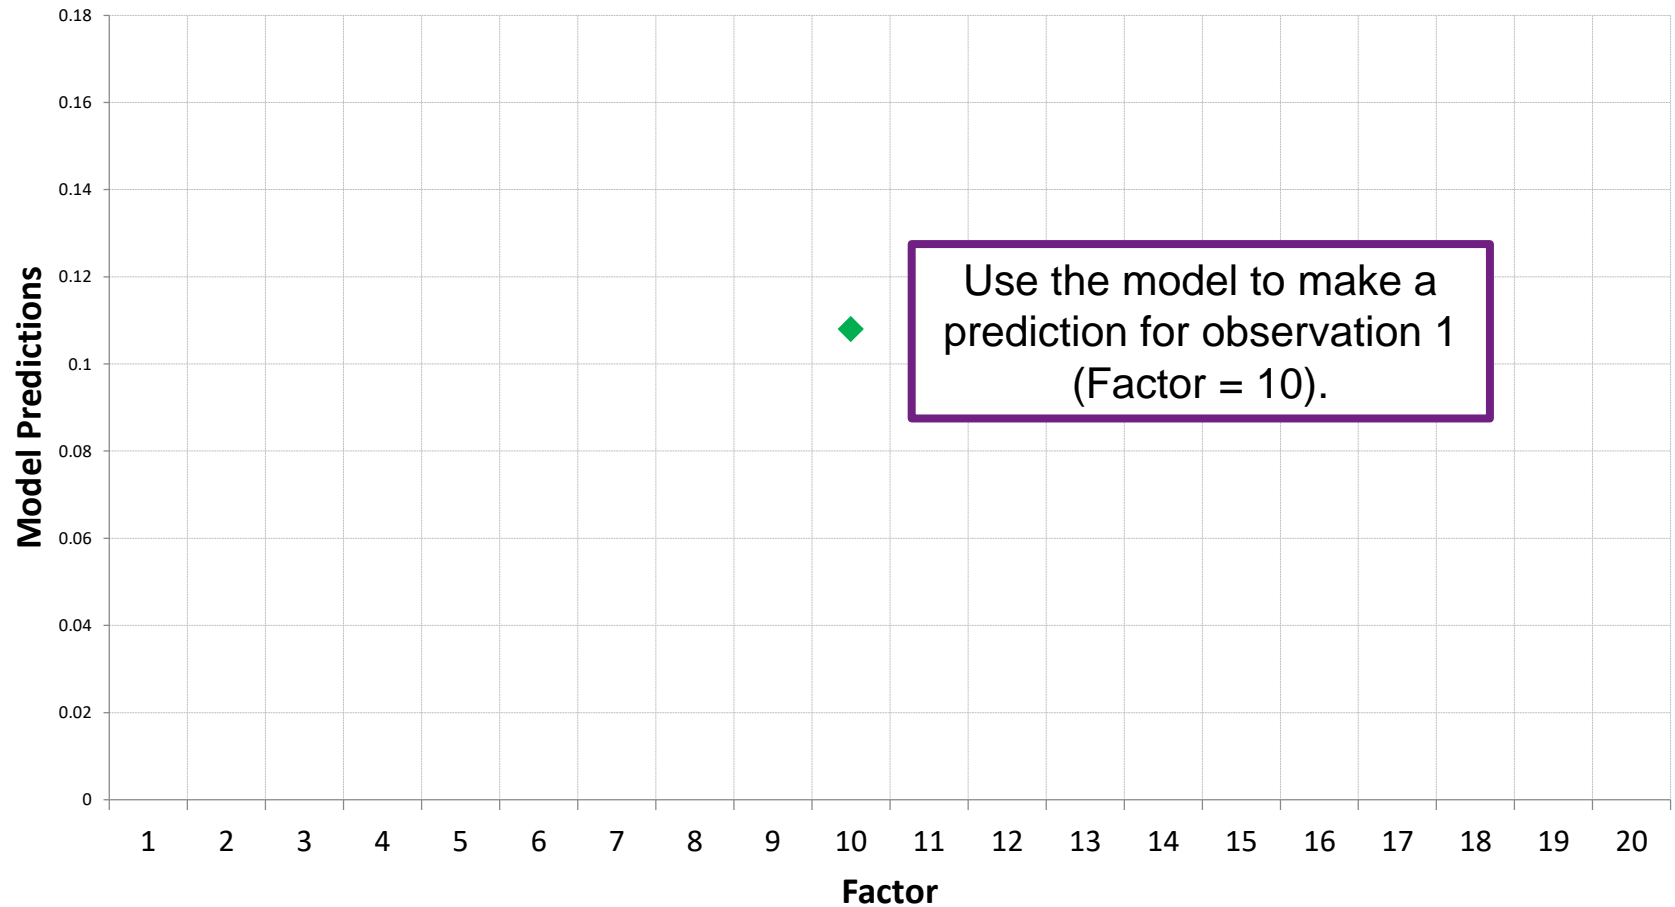
The relative influence of a factor can be measured as the total reduction in error attributable to splits by that factor, across all trees in the GBM






# Partial dependency plots

## Example




# Partial dependency plots

## Example



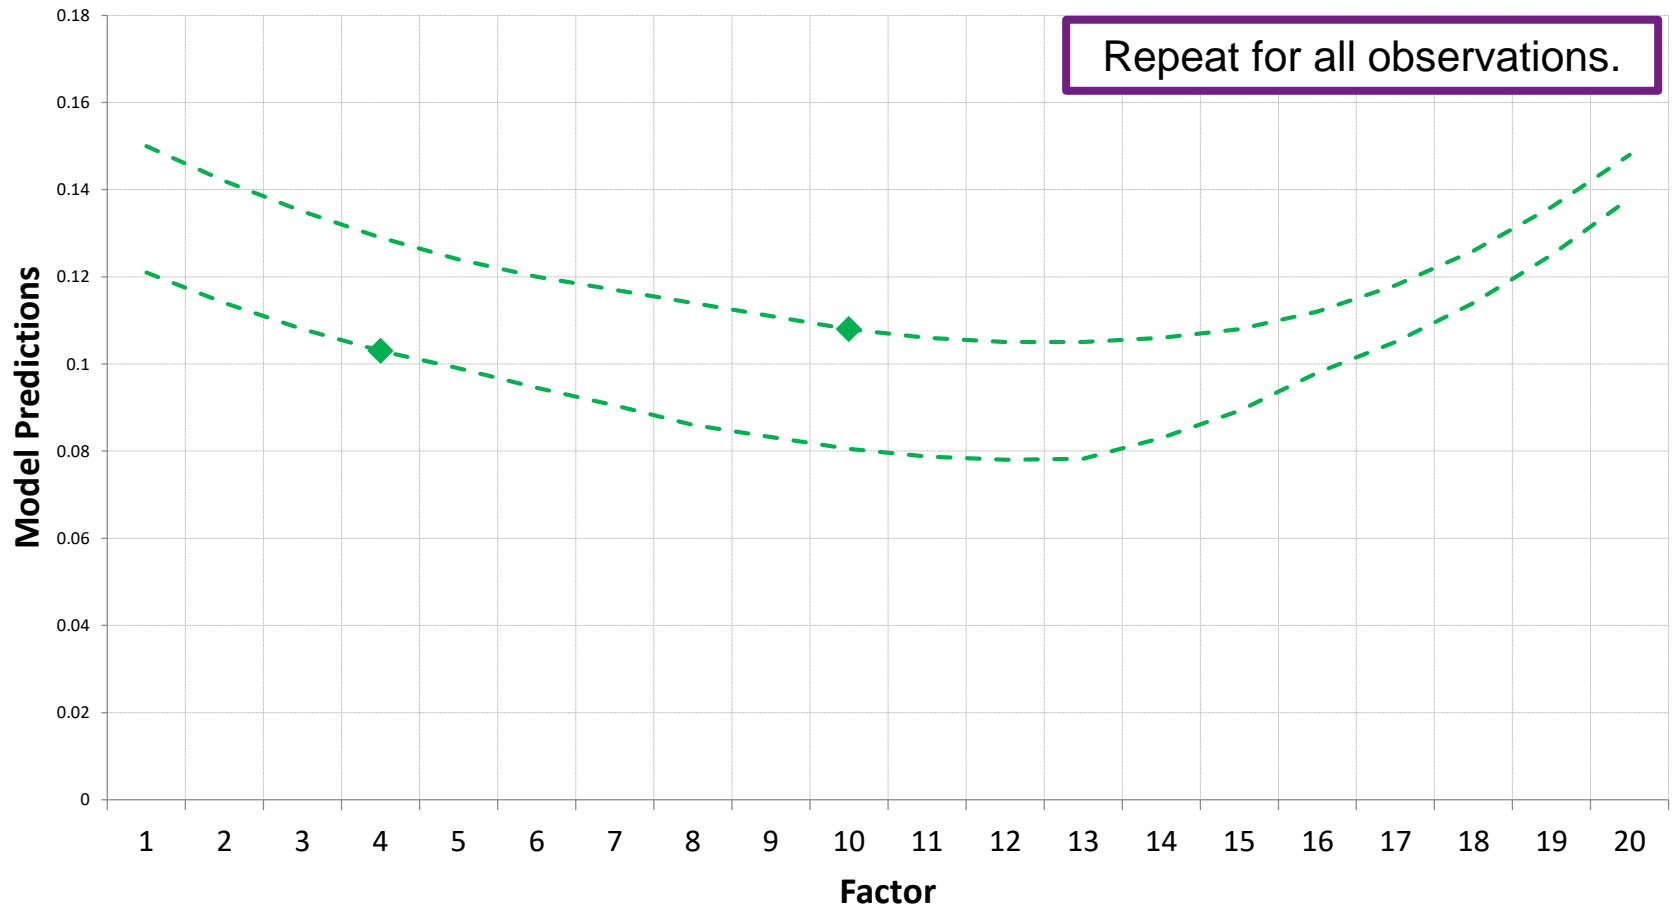
# Partial dependency plots

## Example




# Partial dependency plots

## Example




# Partial dependency plots

## Example




# Partial dependency plots

## Example



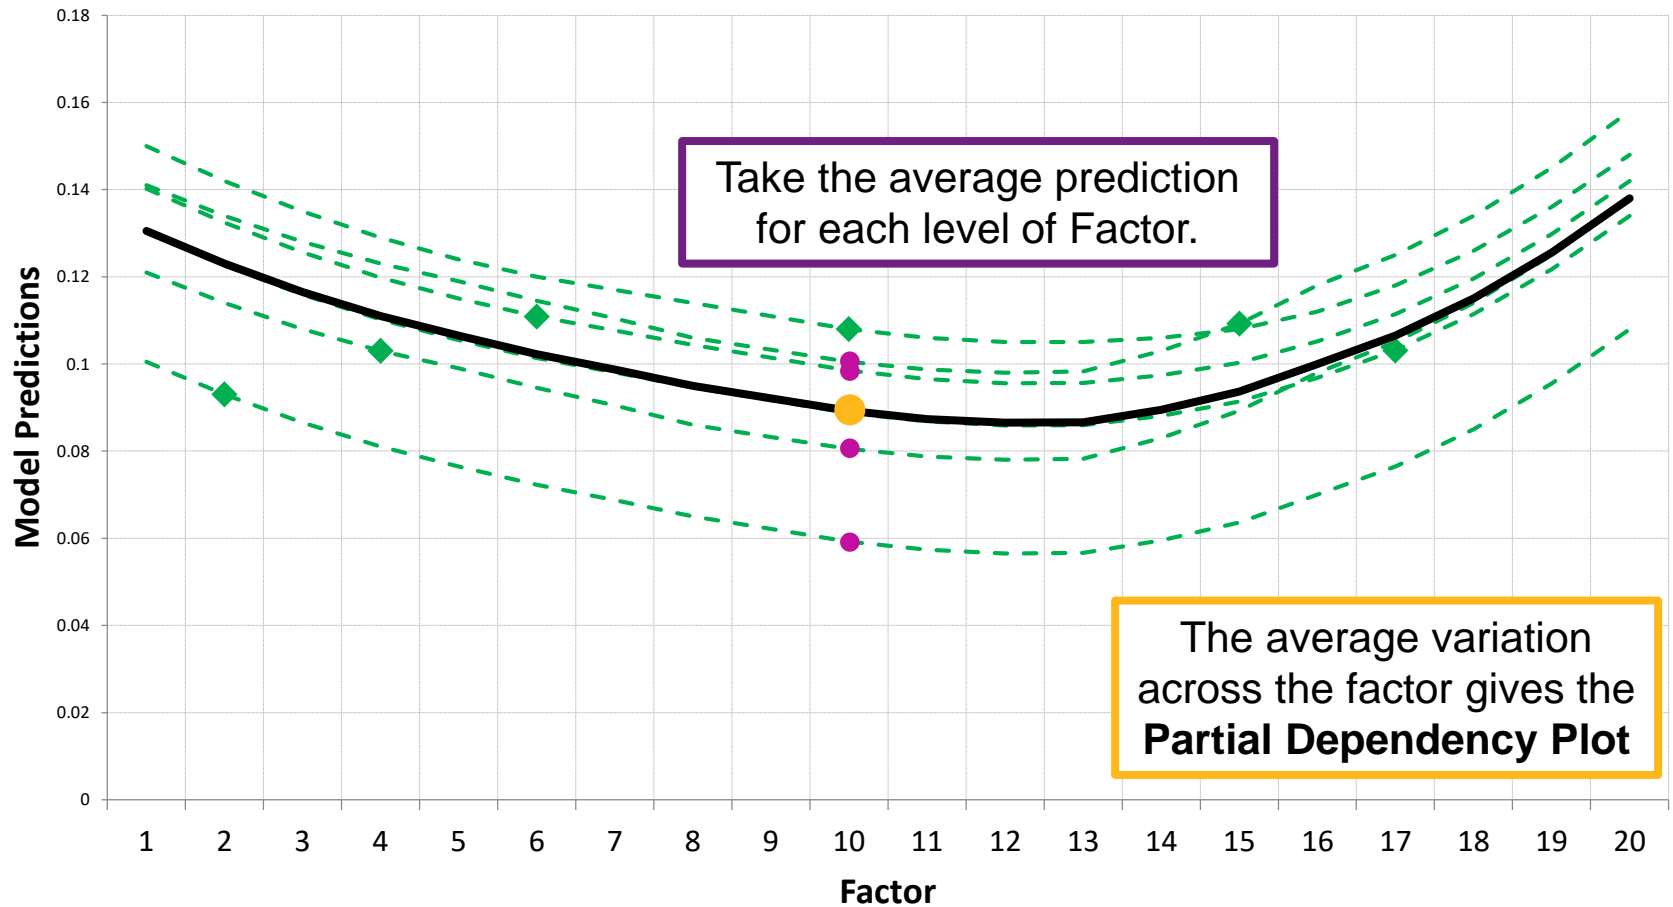
# Partial dependency plots

## Example



# Partial dependency plots


## Example






# Partial dependency plots

## Example




# Partial dependency plots

## Example




# Partial dependency plots

## Example










# Partial dependency plots


## Example



# Partial dependency plots etc



# Partial dependency plots



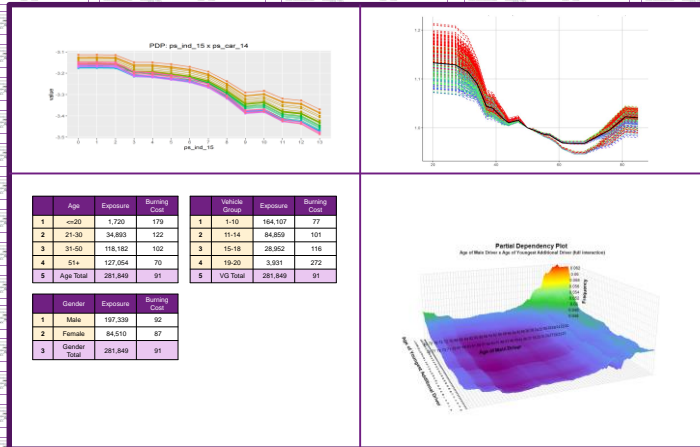
## Advantages

- Qualitative description of properties of relationships
- Most revealing of additive and multiplicative relationships

## Disadvantages

- “GLM view of a non-GLM thing”
- Interaction effects outside of the chosen subset may be obfuscated
- eg if  $X_1X_2$  is important and  $X_2$  is averaged out in the partial dependence plot,  $X_1$  may show as being heterogeneous, thus obfuscating the complexity of the modelled relationships

# So what?



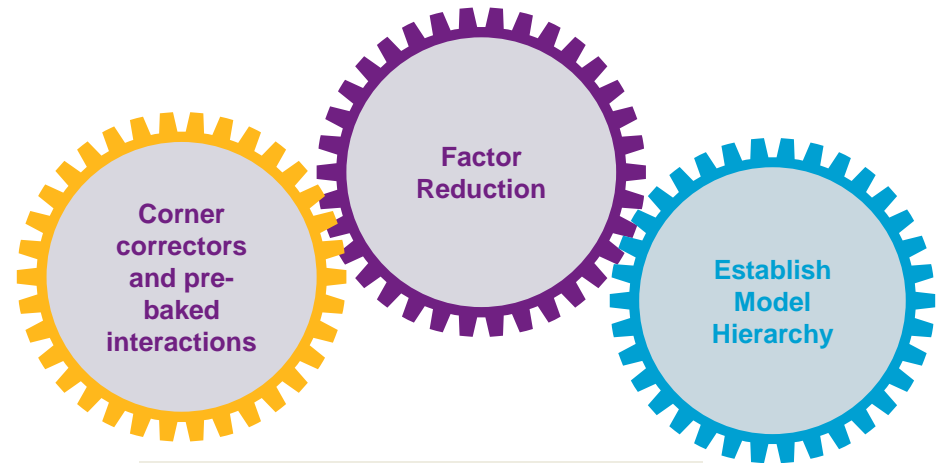
# Deploying GBMs

Model down into multiplicative tables via GLMs

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




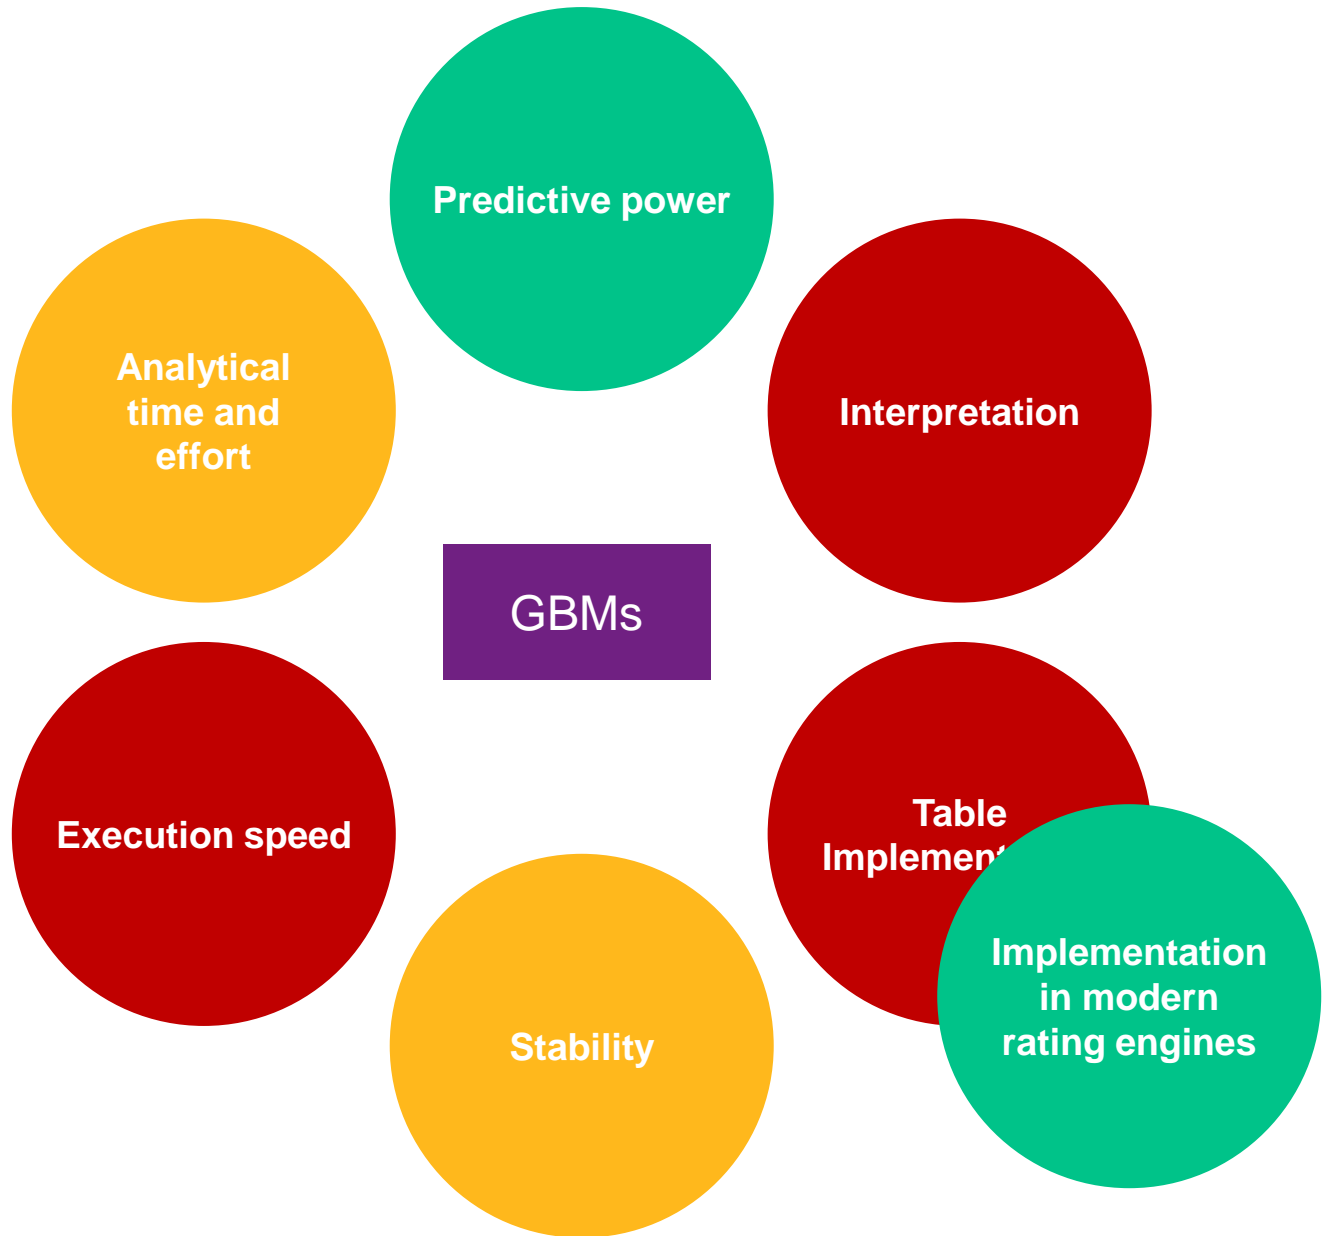
Use insights to guide GLM





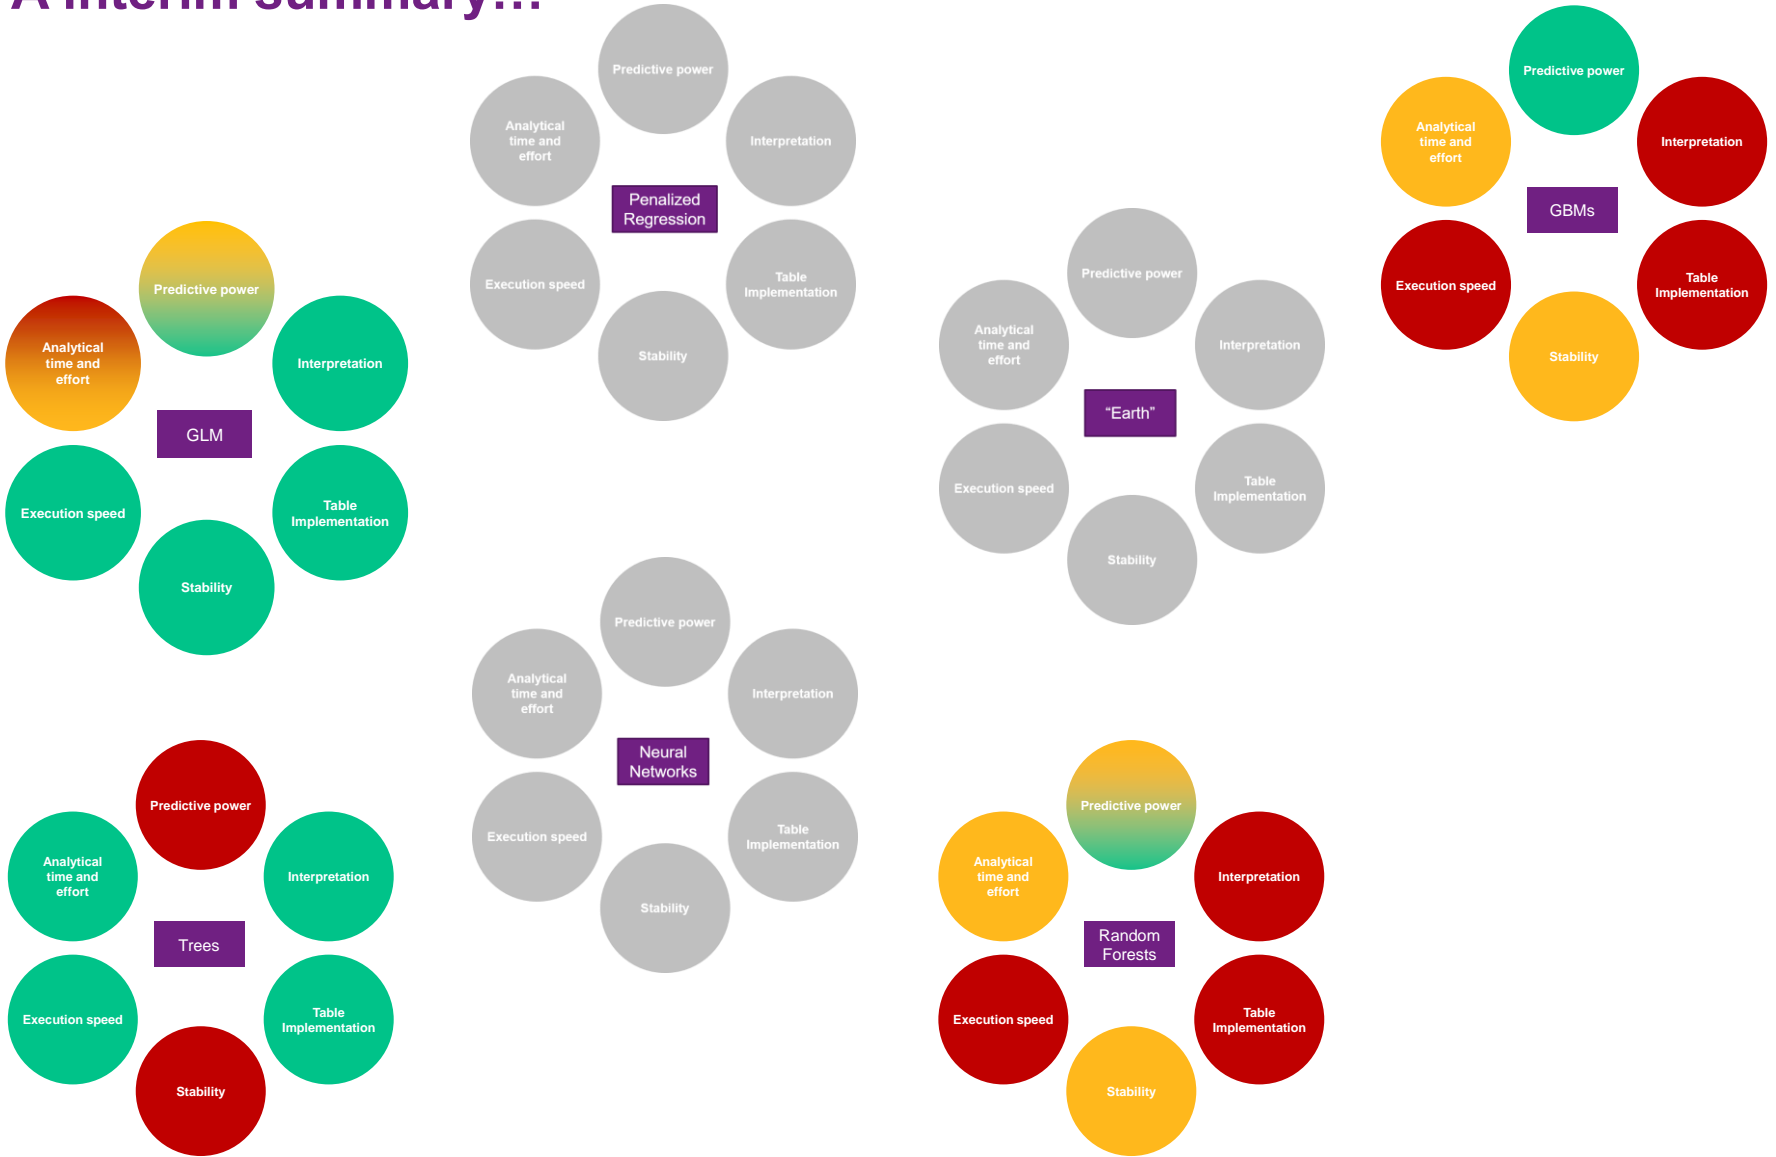
# Deploying GBMs







# A interim summary...



# 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, insights and conclusions to follow in Part 2...

# 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”  
Penalized regression  
Neural networks

---

### 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 & Product Management Seminar**  
Overview and Practical Application of Machine Learning  
Methods in Pricing – Part 2

Wednesday March 27, 2018

Ben Williams, Graham Wright



# Agenda

## Agenda

### Context of machine learning in pricing

---

#### Session 1:

---

Decision trees  
Random forests  
Gradient boosting machines

---

#### Session 2:

---

“Earth”  
Penalized regression  
Neural networks

---

### Conclusions

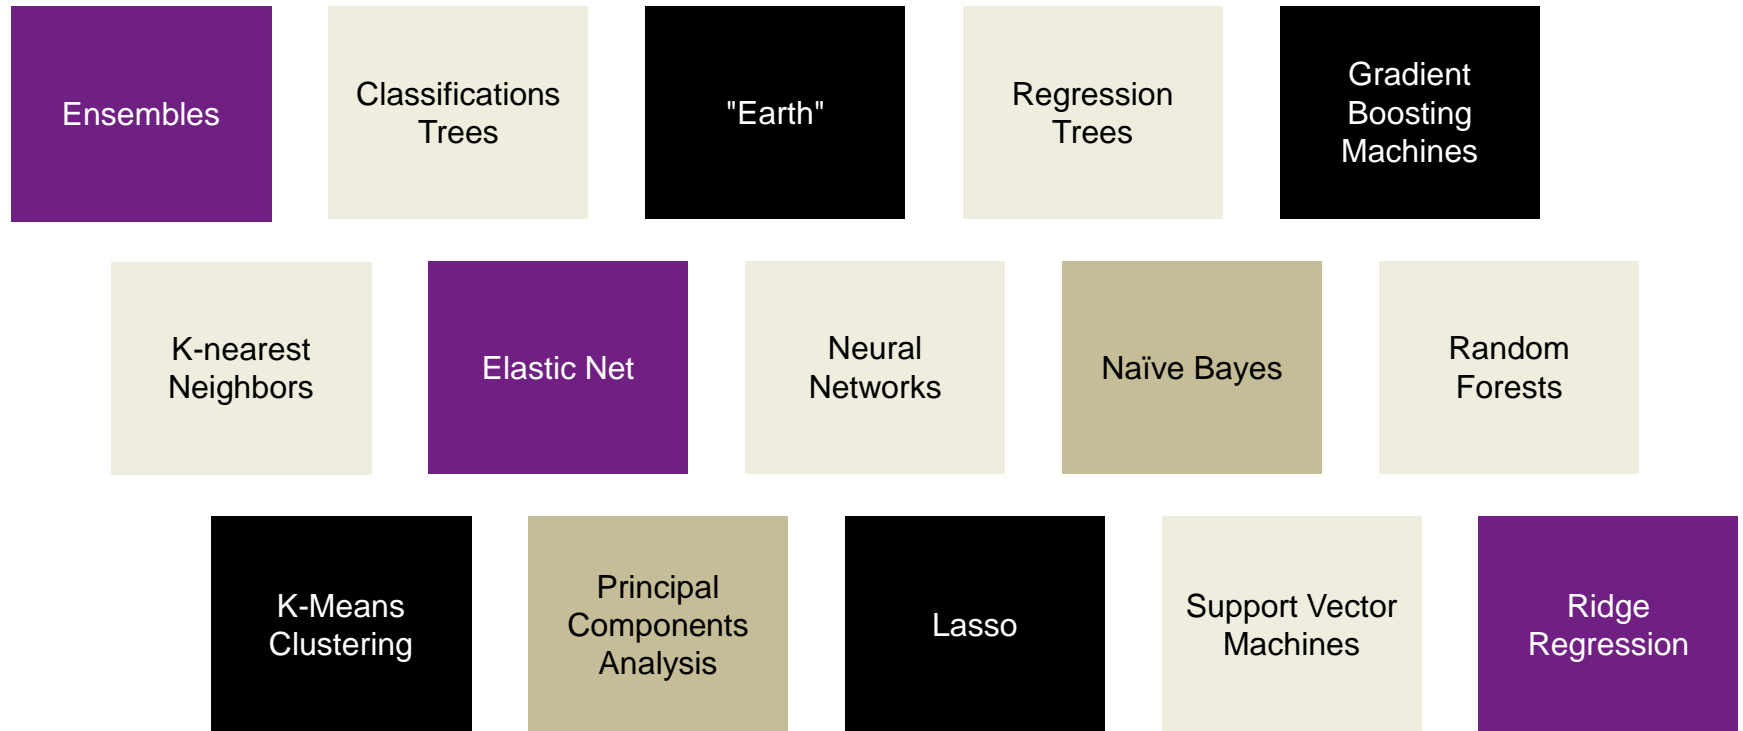
---

### Q&A

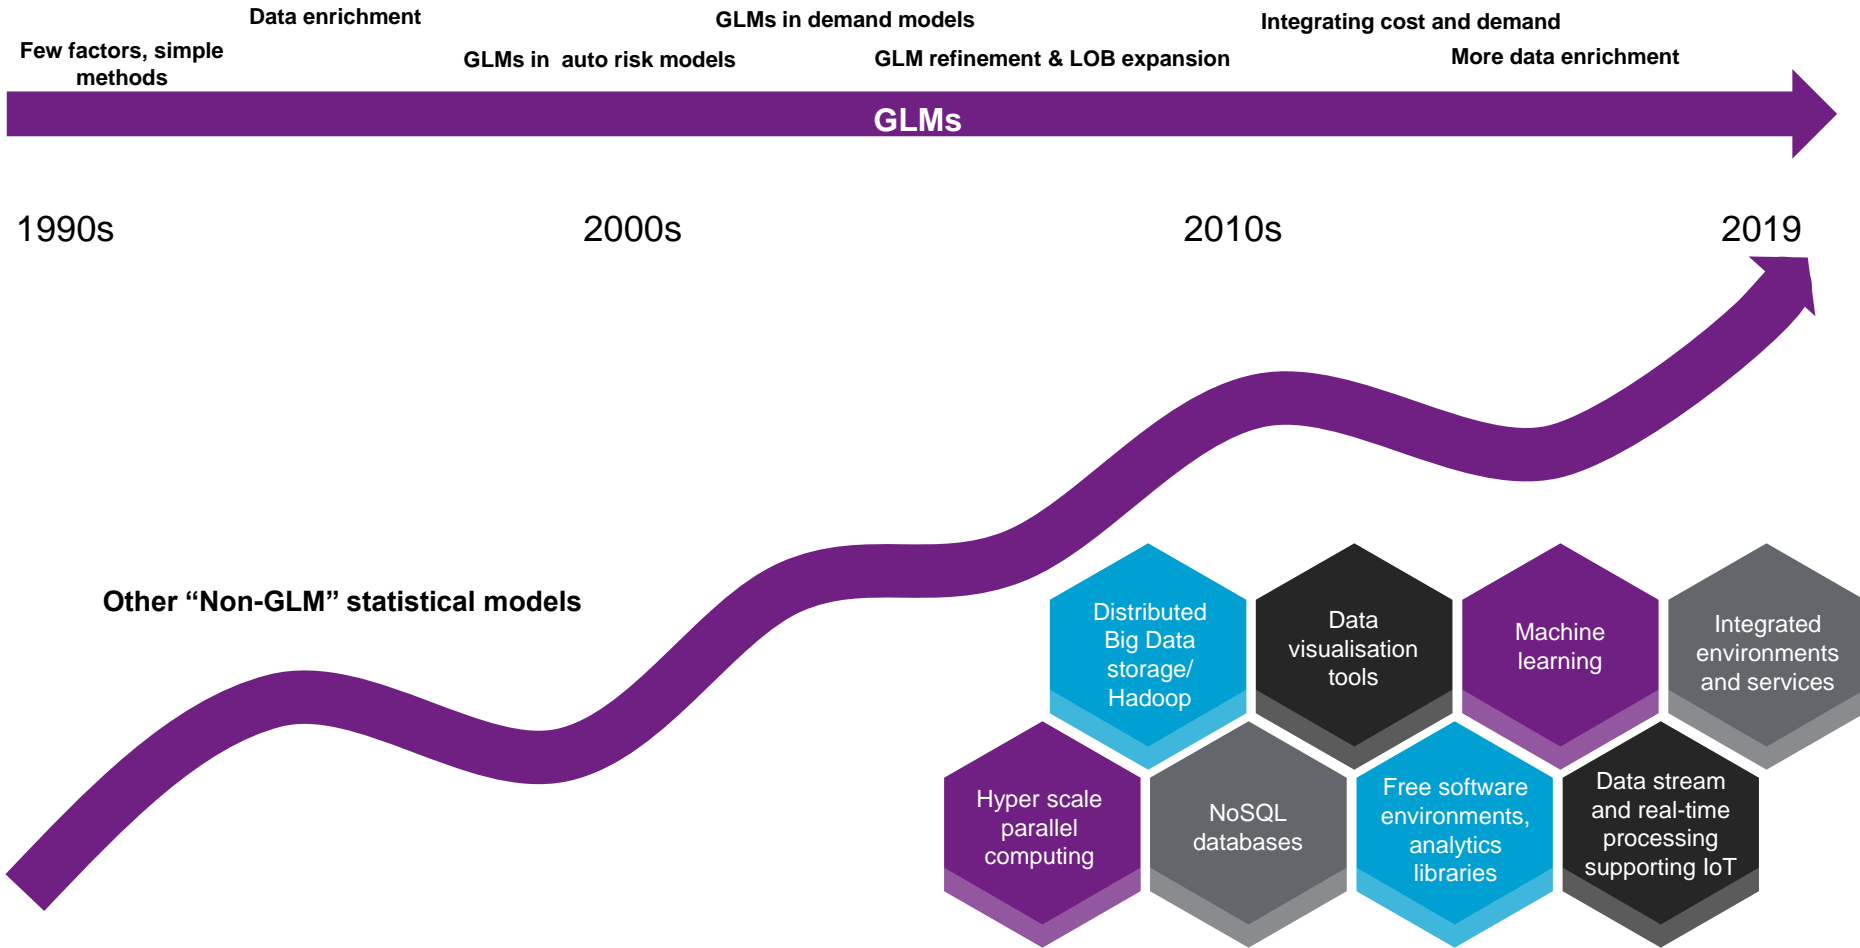
---

**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


# What are these machine learning methods?



# This is not new....

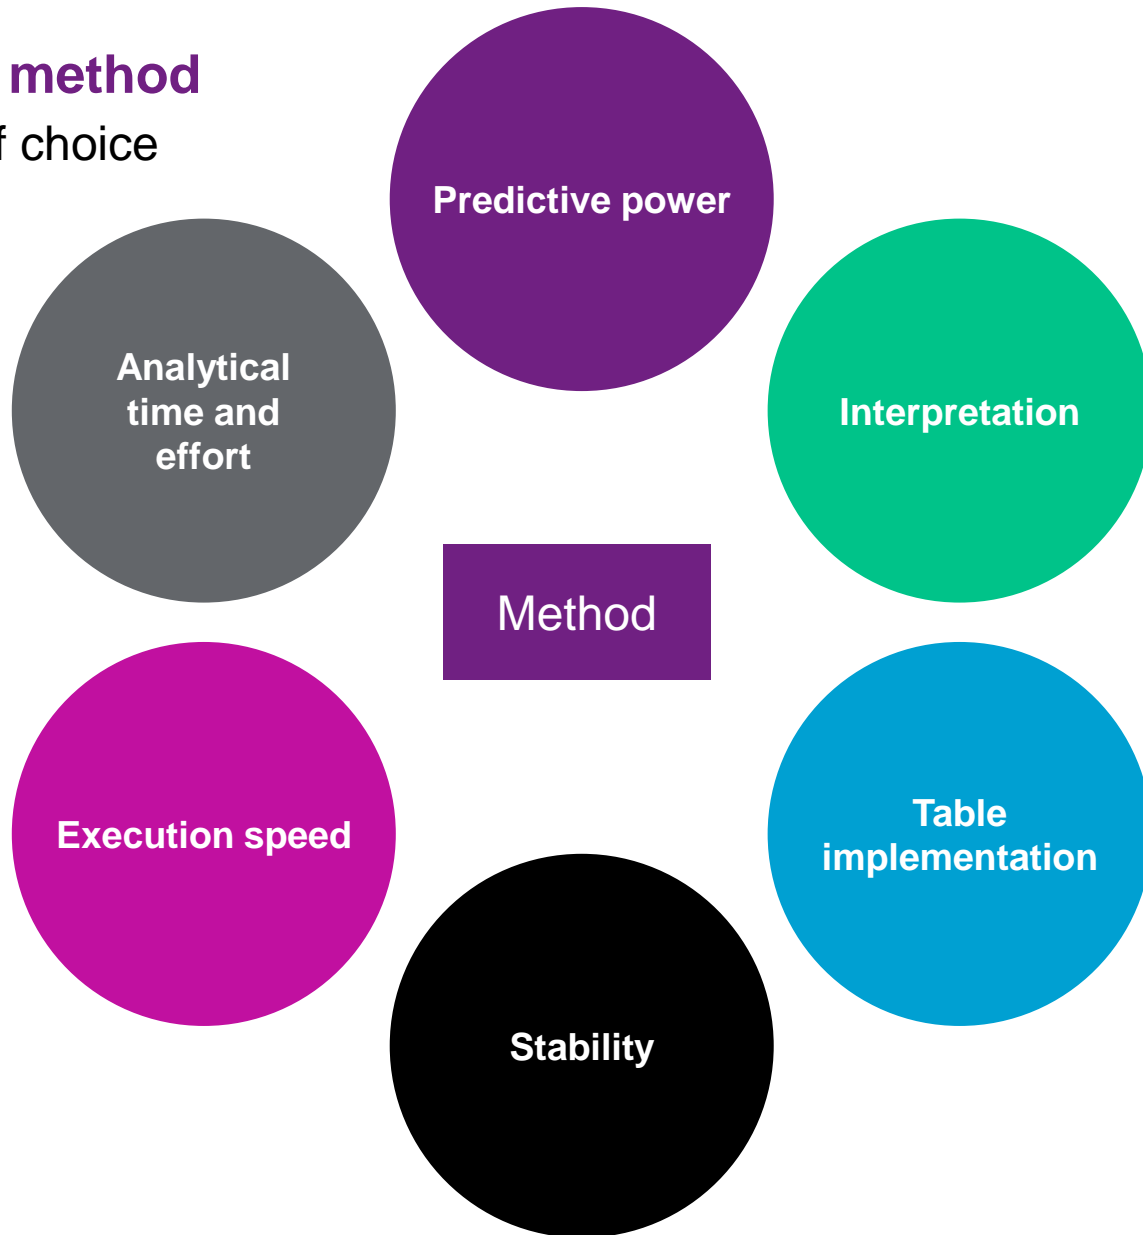


# Is it really all about the method?




# Choosing a method

Dimensions of choice



# It's domain expertise that helps decide



## Financial value estimate

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



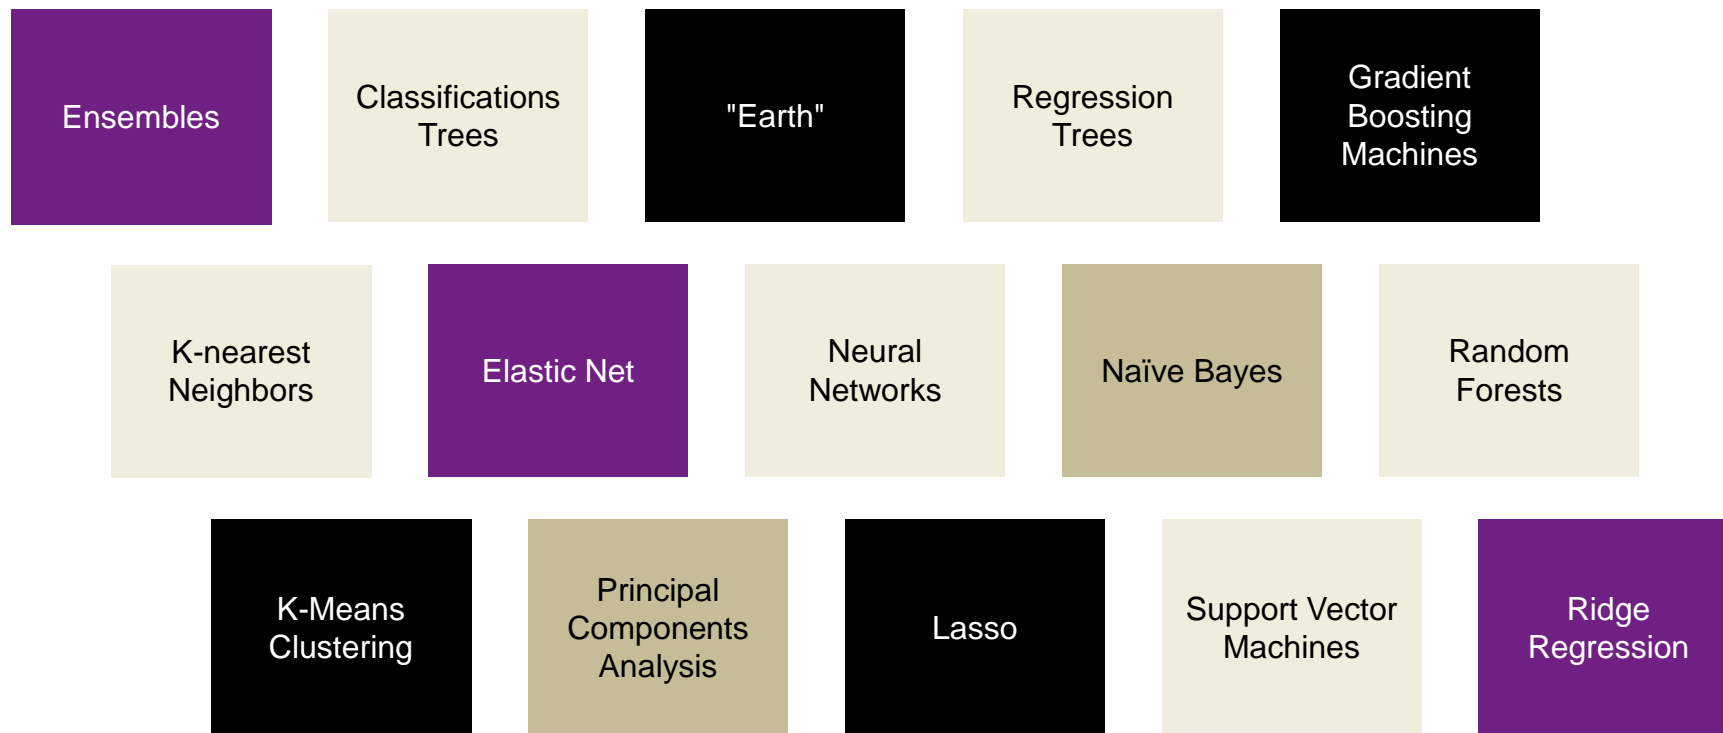
Example results redacted from printed version

# Illustrative results


Model	Gini	Gini improvement	Gini rank	Loss ratio @ elasticity 6	Loss ratio rank	Loss ratio @ elasticity 2	Loss ratio rank
GLM (main factor removed)	0.318	-2.6%	4	-0.9%	4	-0.4%	4
GLM (minor factor removed)	0.322	-1.3%	3	-0.4%	3	-0.2%	3
GLM	0.327	0.0%	2	0.0%	2	0.0%	2
New Model	0.330	1.0%	1	2.2%	1	0.5%	1




# Some machine learning methods



# Focus on “Earth”




# Multivariate adaptive regression splines (“Earth”)




# Multivariate adaptive regression splines (“Earth”)

## Categorical factors




# Multivariate adaptive regression splines (“Earth”)

## Categorical factors




# Multivariate adaptive regression splines (“Earth”)

## Categorical factors




# Multivariate adaptive regression splines (“Earth”)

## Categorical factors



# Multivariate adaptive regression splines (“Earth”)


## Categorical factors






# Multivariate adaptive regression splines (“Earth”)

## Categorical factors




# Multivariate adaptive regression splines (“Earth”)

## Numerical factors




# Multivariate adaptive regression splines (“Earth”)

## Numerical factors




# Multivariate adaptive regression splines (“Earth”)

## Numerical factors




# Multivariate adaptive regression splines (“Earth”)

## Numerical factors




# Multivariate adaptive regression splines (“Earth”)

## Interactions




# Multivariate adaptive regression splines (“Earth”)

## Interactions



# Multivariate adaptive regression splines (“Earth”)

## Interactions





# Multivariate adaptive regression splines (“Earth”)

## Advantages

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

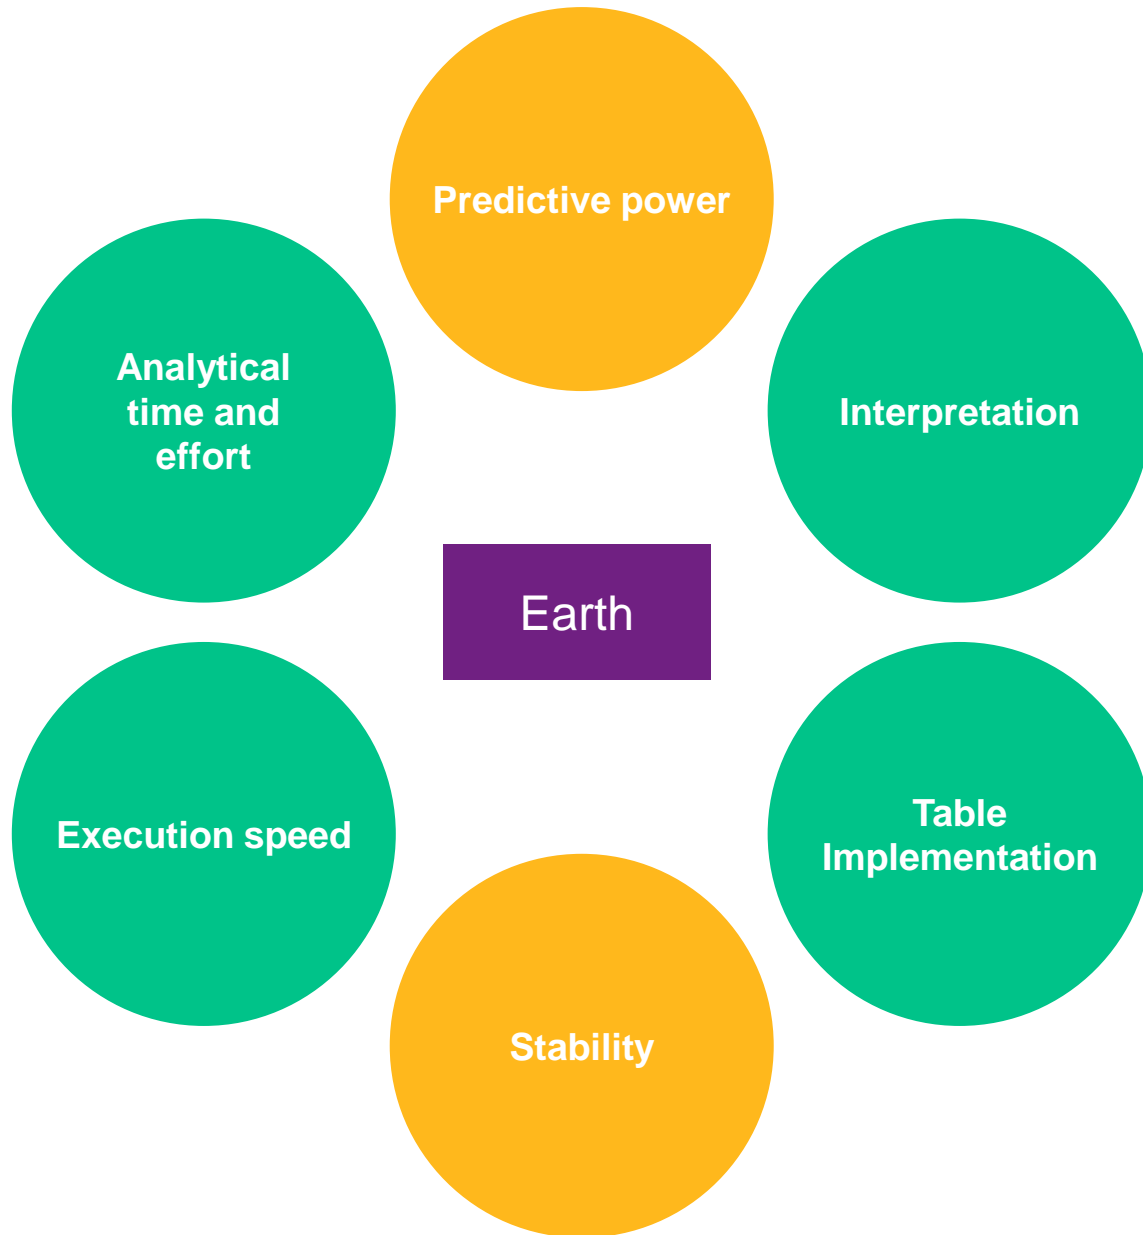
## Disadvantages

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

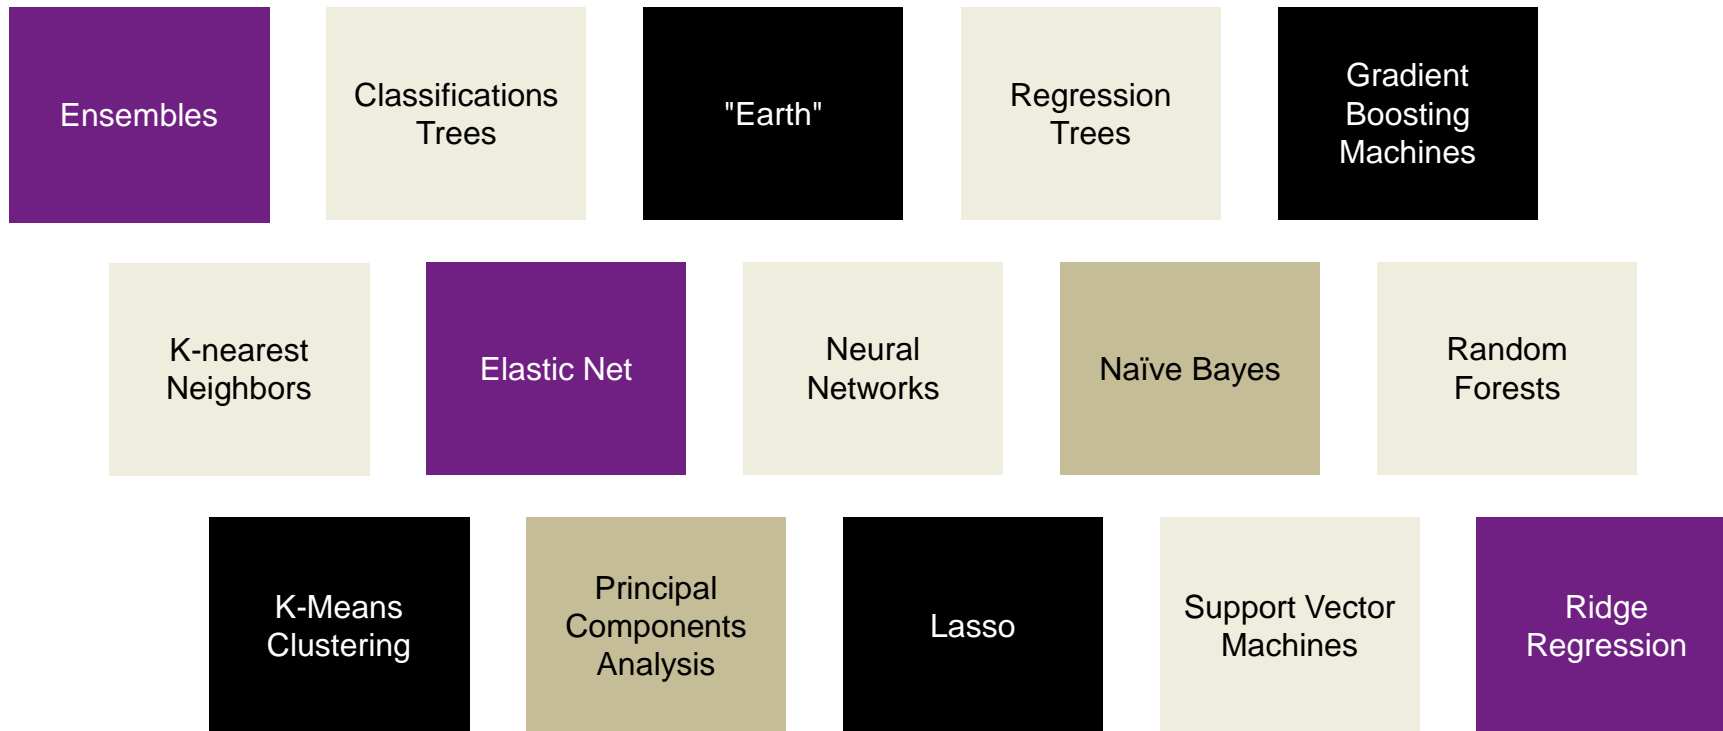
Intercept	0.412
UsuallyPayANNUAL	0.543
$h(\text{Log\_Premium} - 6.314)$	0.432
$h(\text{Age}-35)$	-0.329
UsuallyPayANNUAL * $h(\text{Log\_Premium}-6.5673)$	0.00654
Homeowner	-0.0291
etc	....

## How might “Earth” be applied?

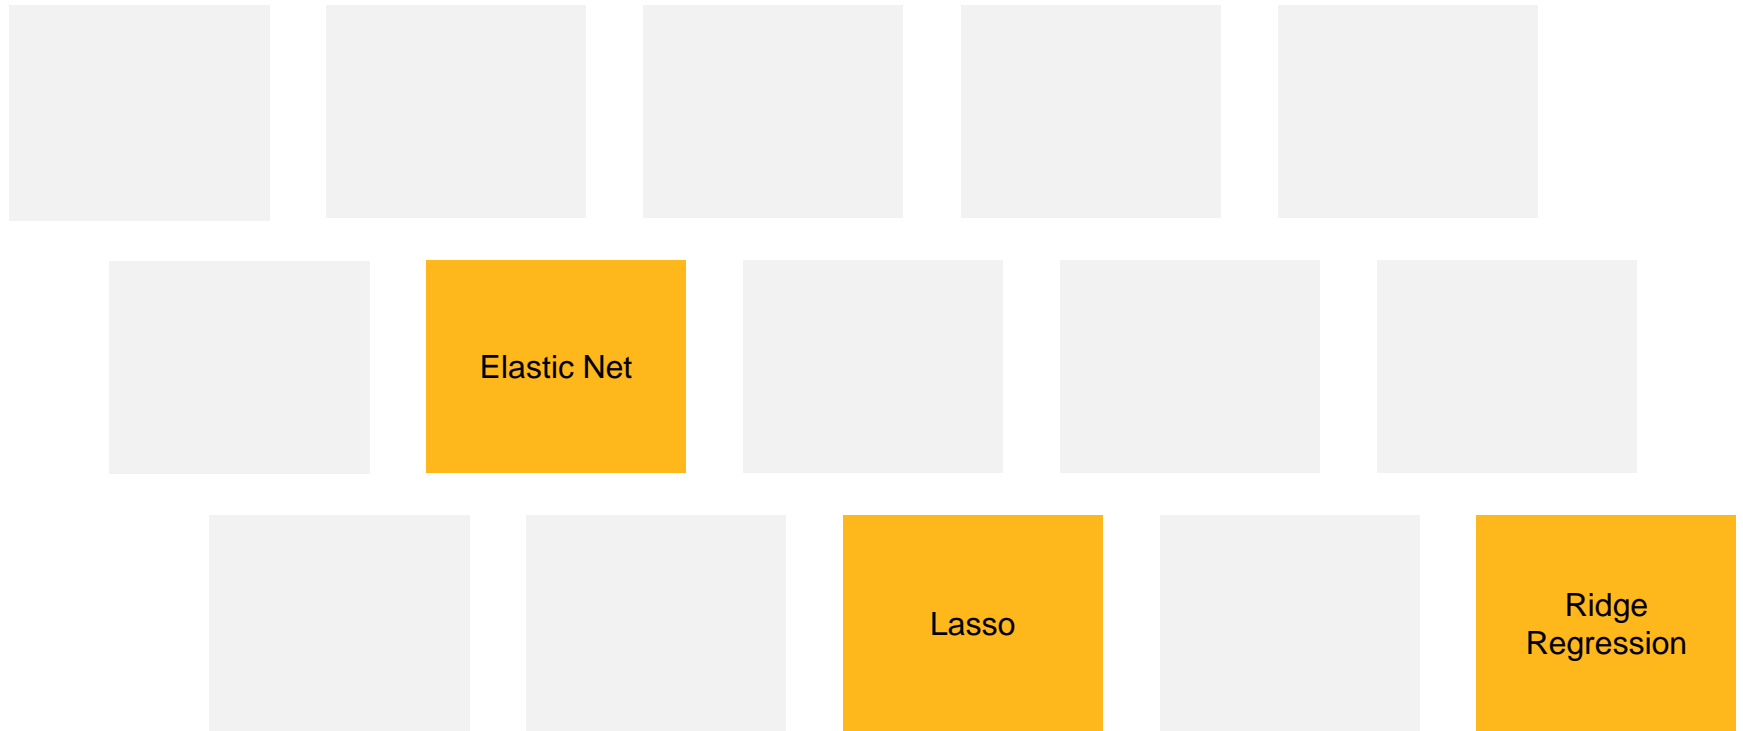
- Historically pricing models have been fit by coverage and/or peril – are these still the most suitable splits?
- When should models be split/combined? (e.g., homeowners and landlords policies or fire and lightning perils)
- How many models should we build and what should they predict?
- Increasing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



# Some machine learning methods



# Focus on Penalized Regression



# Penalized Regression

## Overview

### GLMs

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

### Penalized regression

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

### Elastic Net

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

**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

# Penalized Regression

## GLM

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



# Penalized Regression

$f(\underline{x}) = g^{-1}(\mathbf{X} \cdot \underline{\beta})$  where  $\underline{\beta}$  estimated by minimizing

GLM

Lasso


Ridge

$L(\beta|X, y)$


$$+ \lambda_1 \sum_i |\beta_i| + \lambda_2 \sum_i \beta_i^2$$

Elastic Net


Ridge  $\sum_i \beta_i^2$



Elastic Net



Lasso  $\sum_i |\beta_i|$



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



# Penalized Regression

$f(\underline{x}) = g^{-1}(\mathbf{X} \cdot \underline{\beta})$  where  $\underline{\beta}$  estimated by minimizing

GLM

Lasso


Ridge

$L(\beta|X, y)$


$$+ \lambda_1 \sum_i |\beta_i| + \lambda_2 \sum_i \beta_i^2$$

Elastic Net


Ridge  $\sum_i \beta_i^2$



Elastic Net



Lasso  $\sum_i |\beta_i|$



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

# Penalized Regression


## Parameter selection

- Minimize:  $L(\beta|X, y) + \lambda_1 \sum_i |\beta_i| + \lambda_2 \sum_i \beta_i^2$
- Penalty parameters can be re-written:  $\lambda_1 = \lambda\alpha, \quad \lambda_2 = \lambda \left(\frac{1-\alpha}{2}\right)$
- $\alpha$  controls the mixture between Lasso ( $\alpha = 1$ ) and Ridge ( $\alpha = 0$ )
- $\lambda$  controls the overall size of the penalty
- $\lambda, \alpha$  selected using cross-validation
- Factors automatically selected from initial set!



# Penalized Regression

## Parameter selection - example




Models range from Lasso ( $\alpha = 1$ ) to Ridge ( $\alpha = 0$ )

# Penalized Regression

## Parameter selection - example

- The fitting process can be investigated to help with feature selection



As size of penalty decreases, parameters begin emerge as non-zero

Parameters that are still zero at the optimal lambda could be discarded

# 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


## Vehicle classification – categorical 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
...	...	...	...	...	...	...

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	...	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	...	0
0	...	1
0	...	0
0	...	0
0	...	0
1	...	0
0	...	1
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





# 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
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



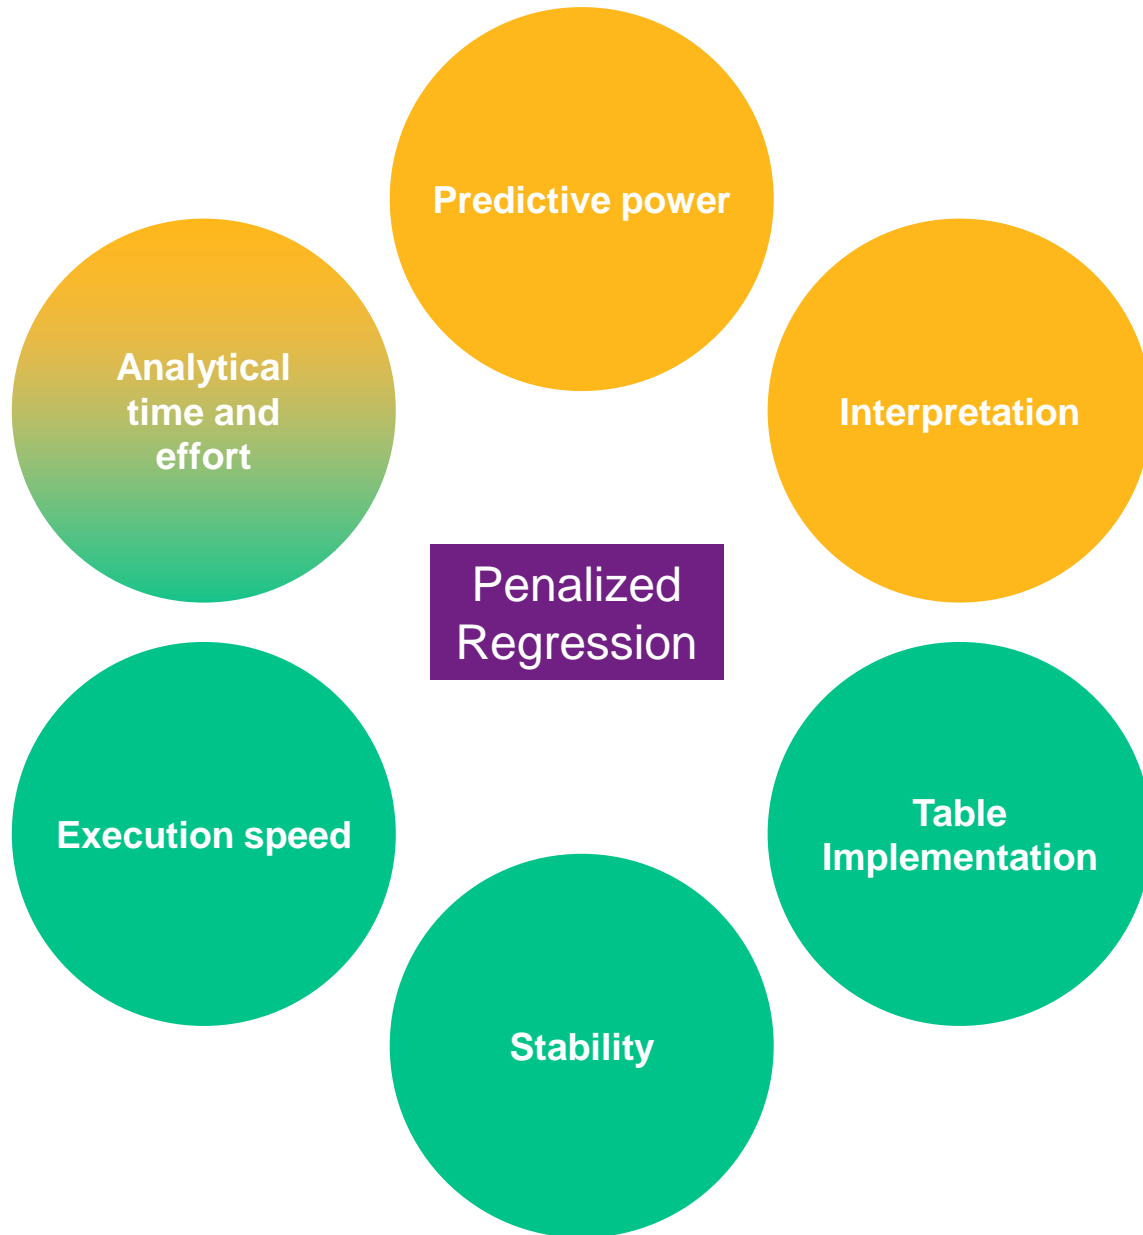
# Deploying Penalized Regression

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




# Practical applications of regression methods in pricing

Streamlining factor selection



Geodemographic information




Vehicle clustering

	Age	Exposures	Summly Cost
1	<=20	1,720	179
2	21-30	34,803	122
3	31-50	118,182	102
4	51+	127,054	70
5	Age Total	281,849	91

	Vehicle Group	Exposures	Summly Cost
1	V10	164,107	77
2	V1-14	84,859	101
3	V15-18	28,952	116
4	V19-20	3,931	272
5	VG Total	281,849	91


	Gender	Exposures	Summly Cost
1	Male	197,239	92
2	Female	84,610	87
3	Gender Total	281,849	91




Modelling hierarchies

Simplification of ML models to improve interpretability

# Some machine learning methods




# Focus on Neural Networks



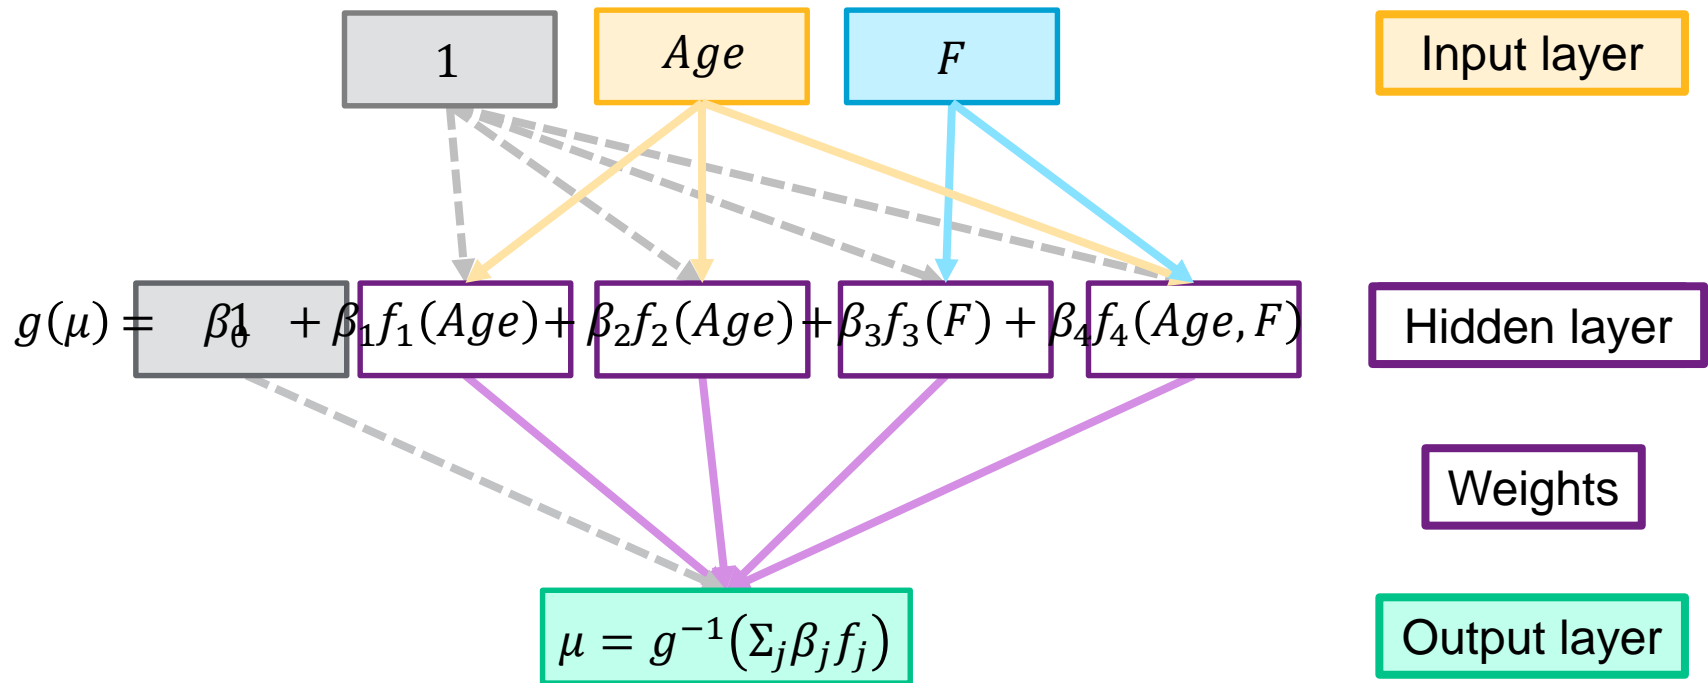
## Start with a simple GLM...

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

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

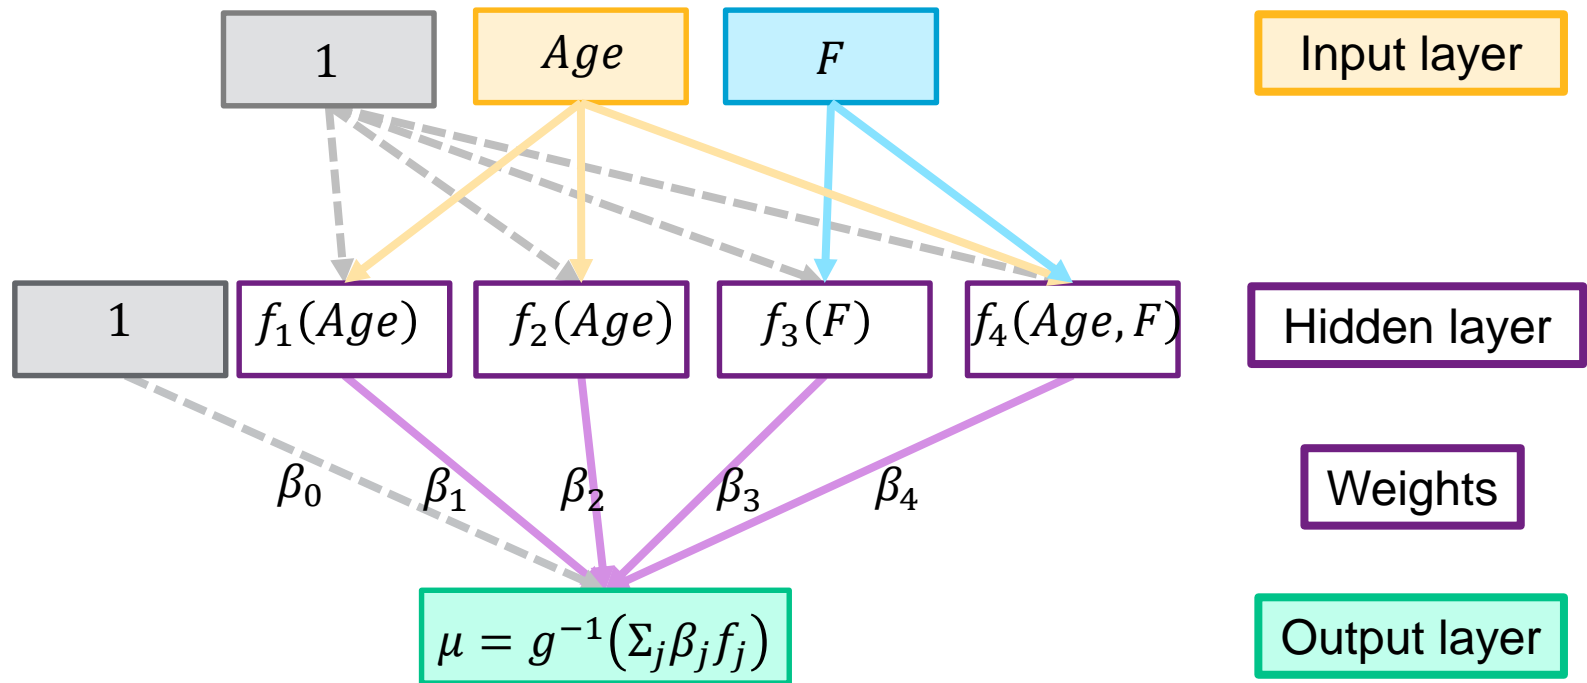


# We can represent GLMs as a network...

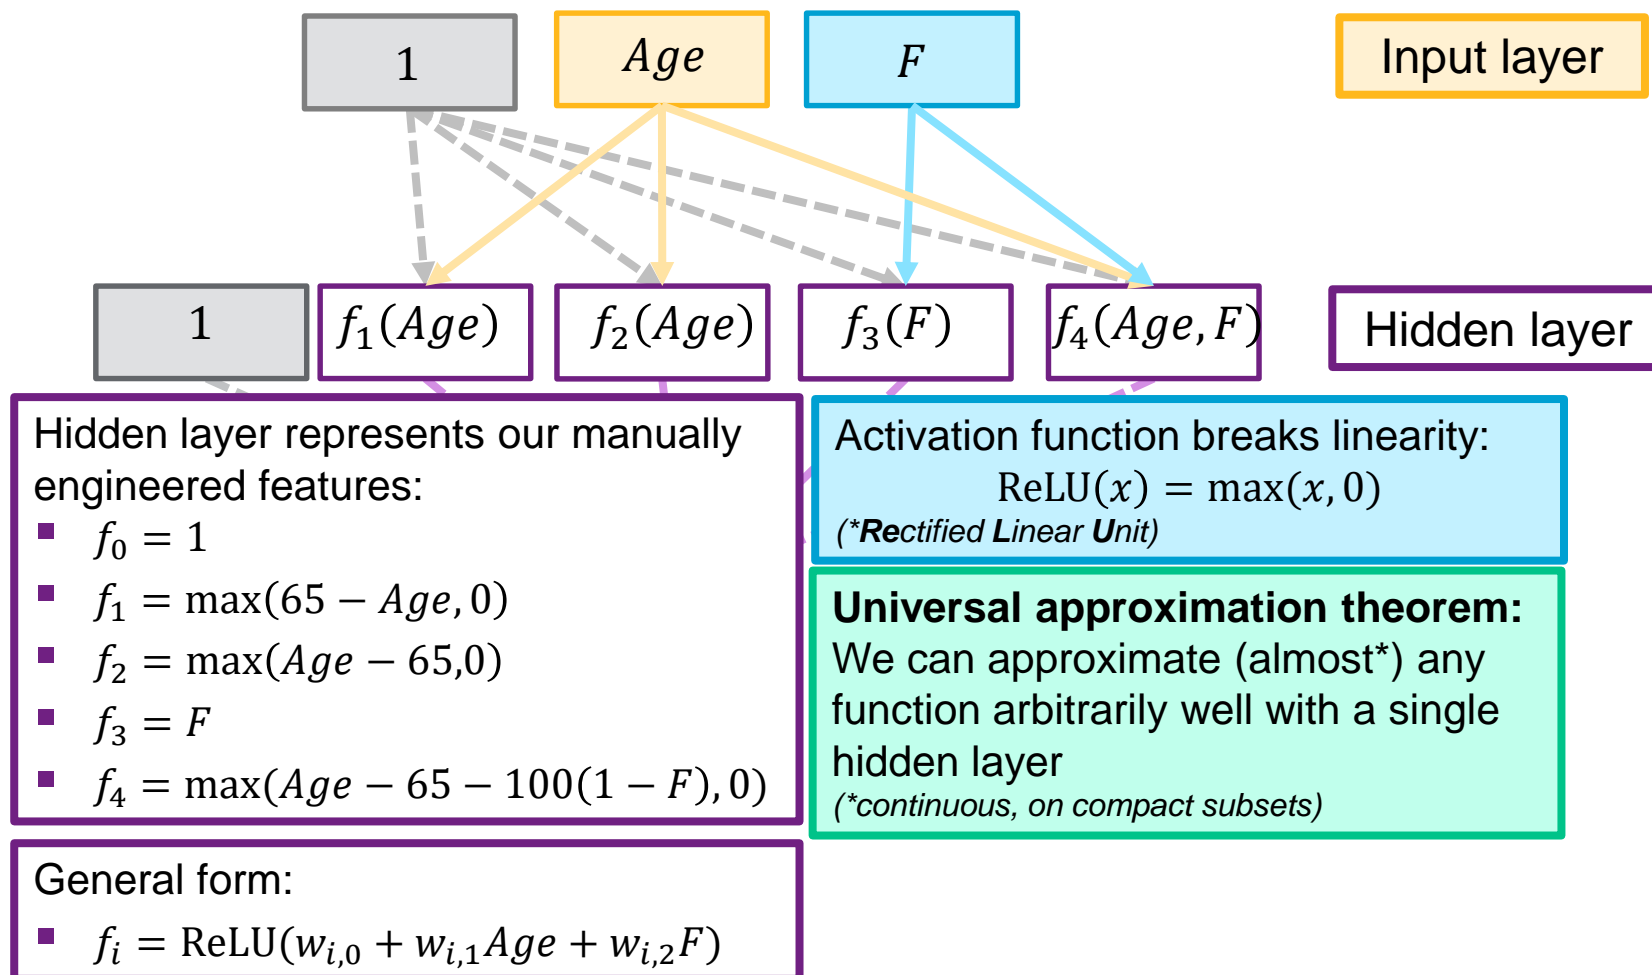




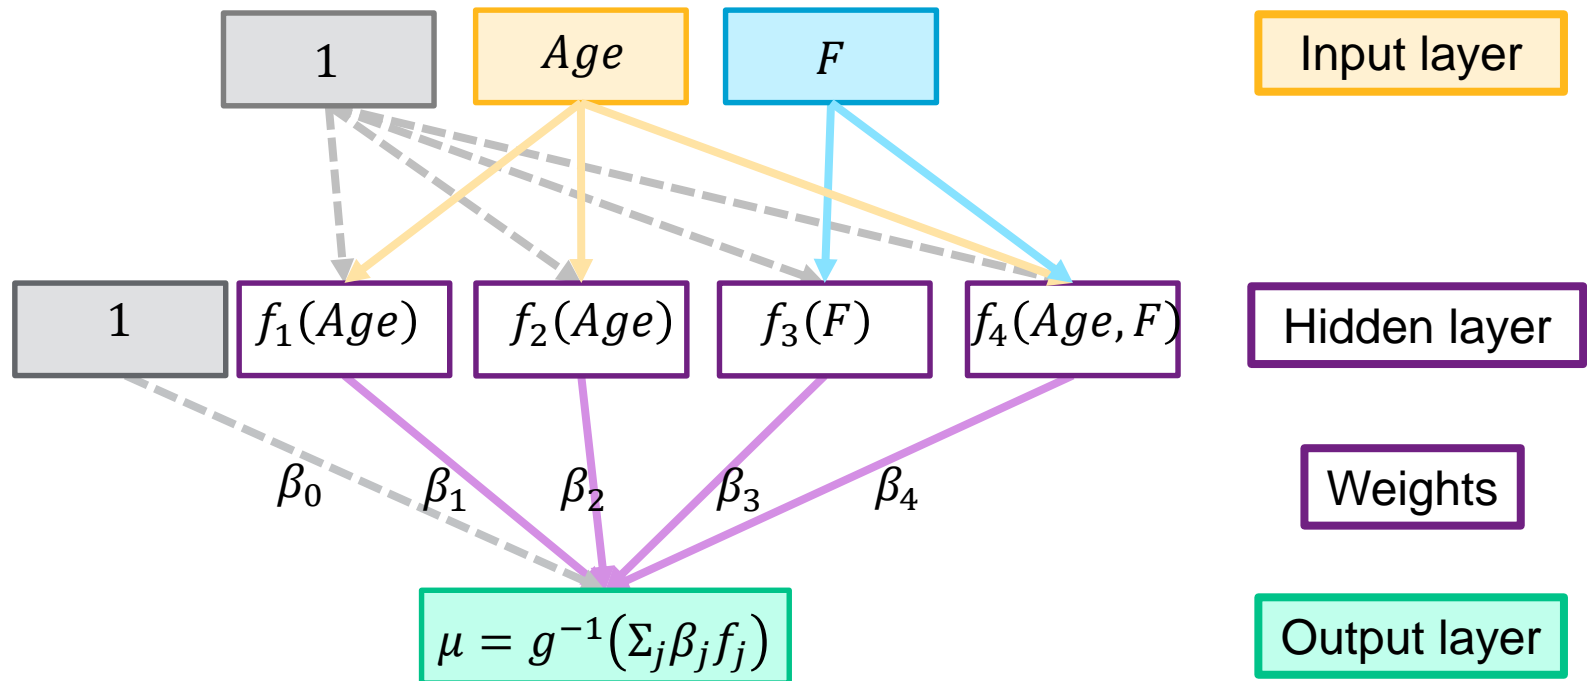
# We can represent GLMs as a network...



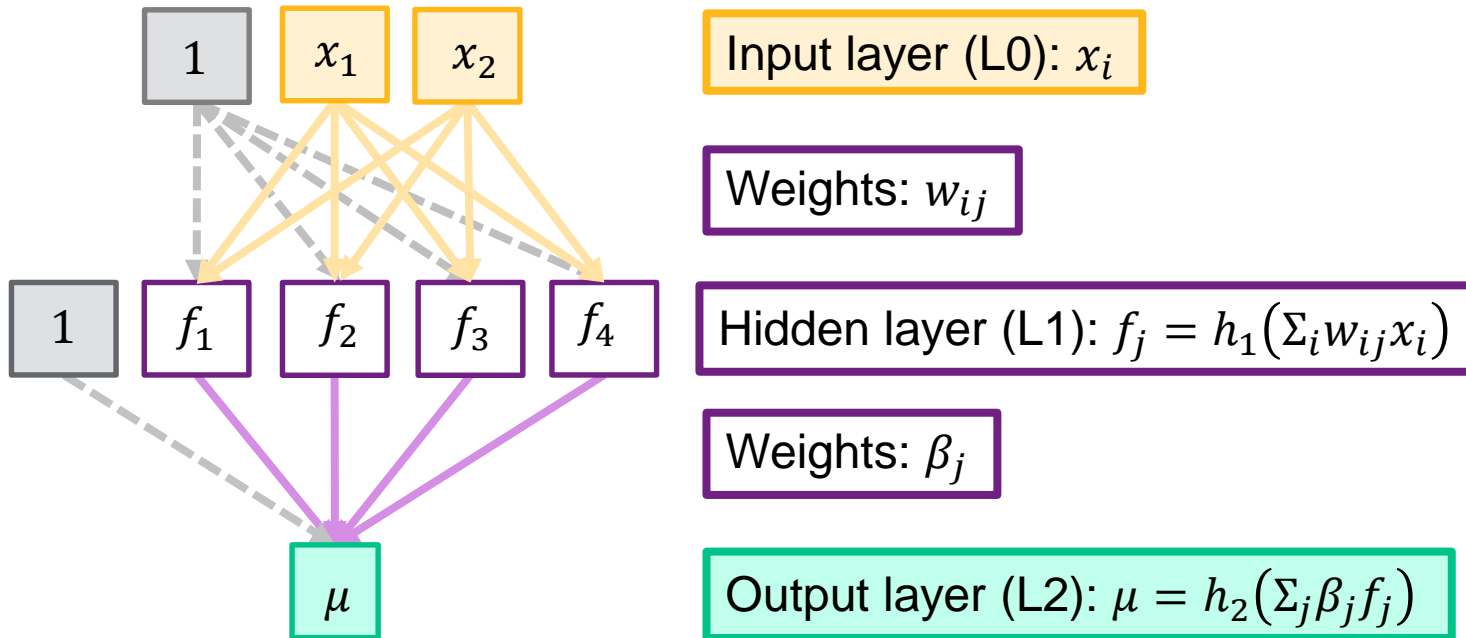
## We can represent GLMs as a network...



# 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

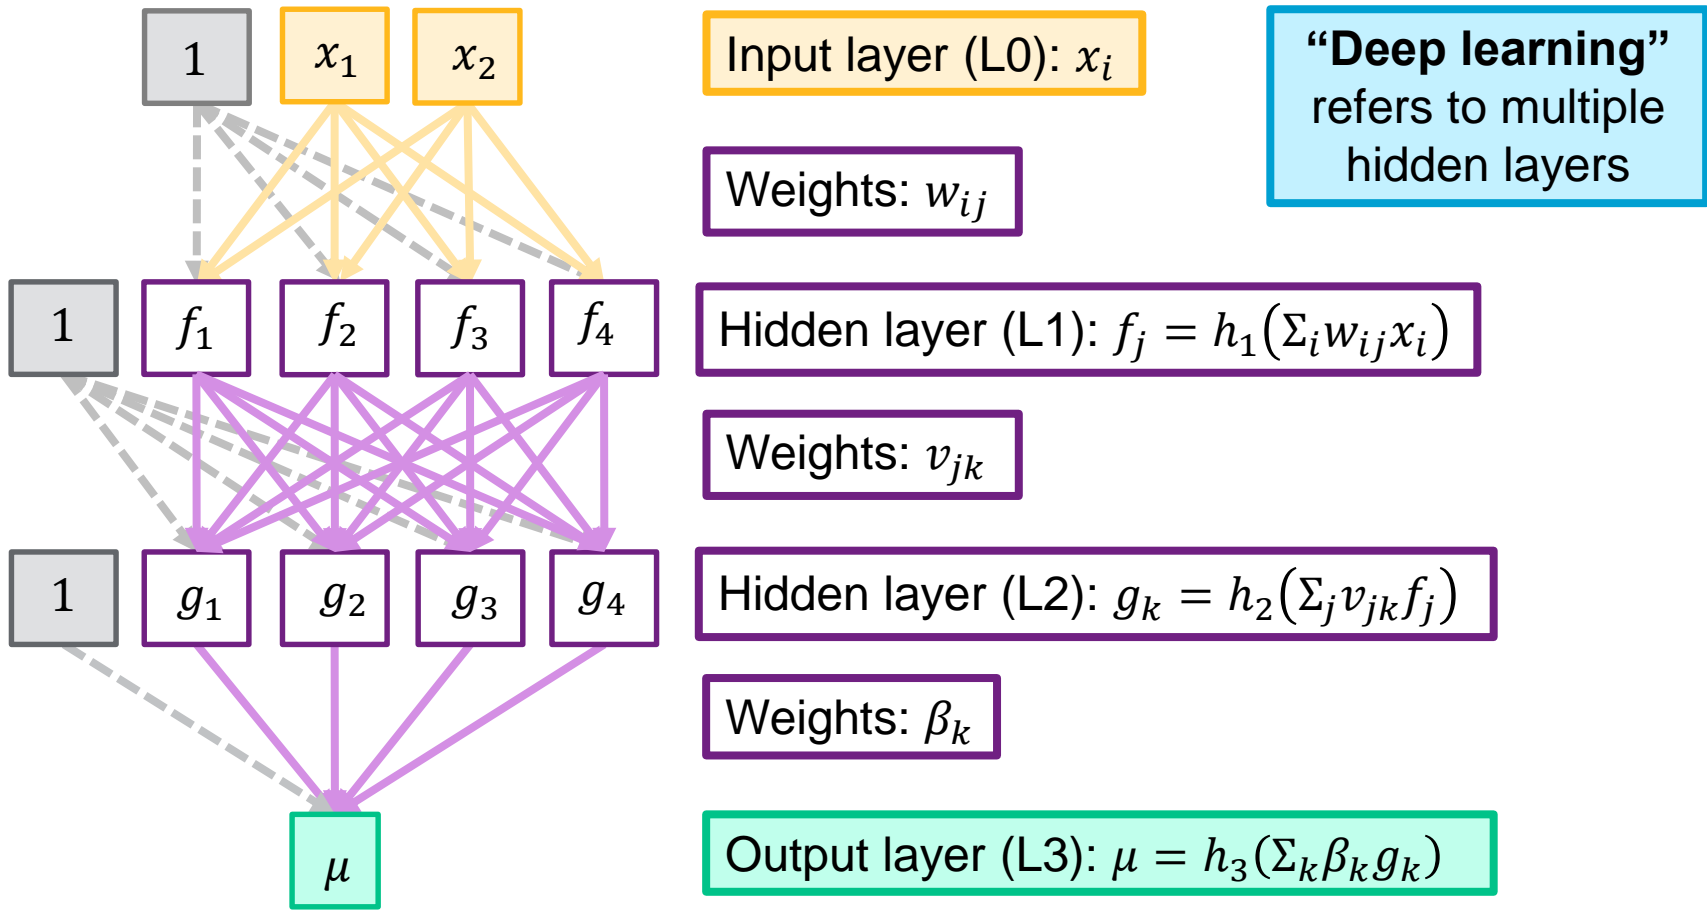
# Generalizing to neural networks

## 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



# 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?



## Where is the value?




Which picture is more likely to be of a cat?



# Neural networks

Evolution or revolution?



# 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

# Neural networks

## Case study – GLM benchmark

Model	Test error	Training error
GLM	34.7%	34.0%

# Neural networks

Require some work!

Input layer

Dropout

Optimization algorithm

Output layer

Hidden layers

Epochs

Learning rate



Regularization

Batch size

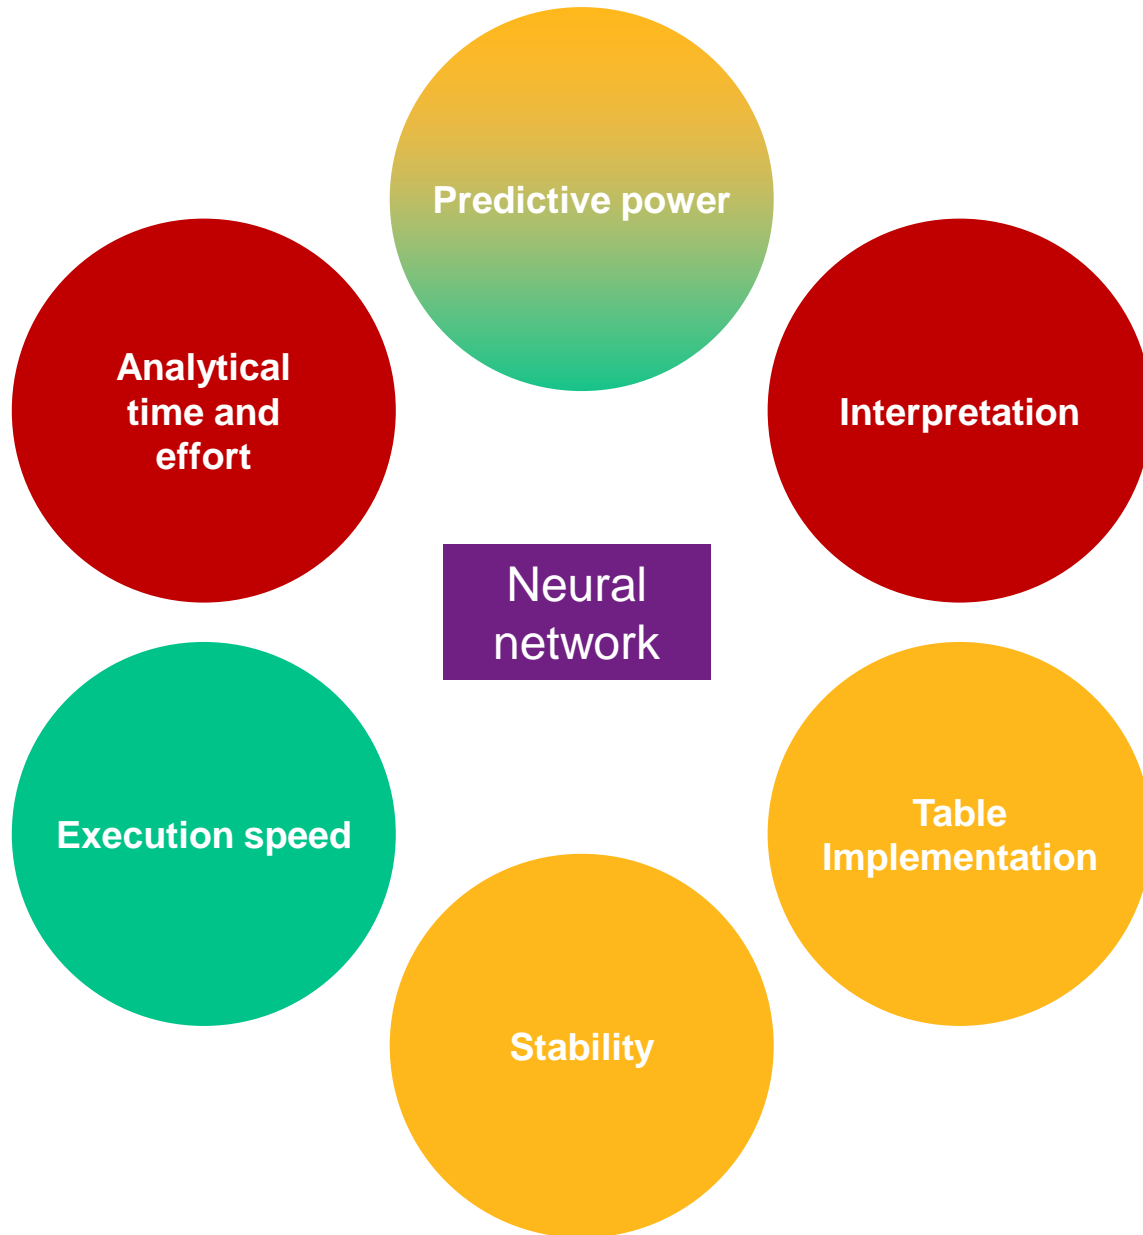
Initial weights

Activation functions

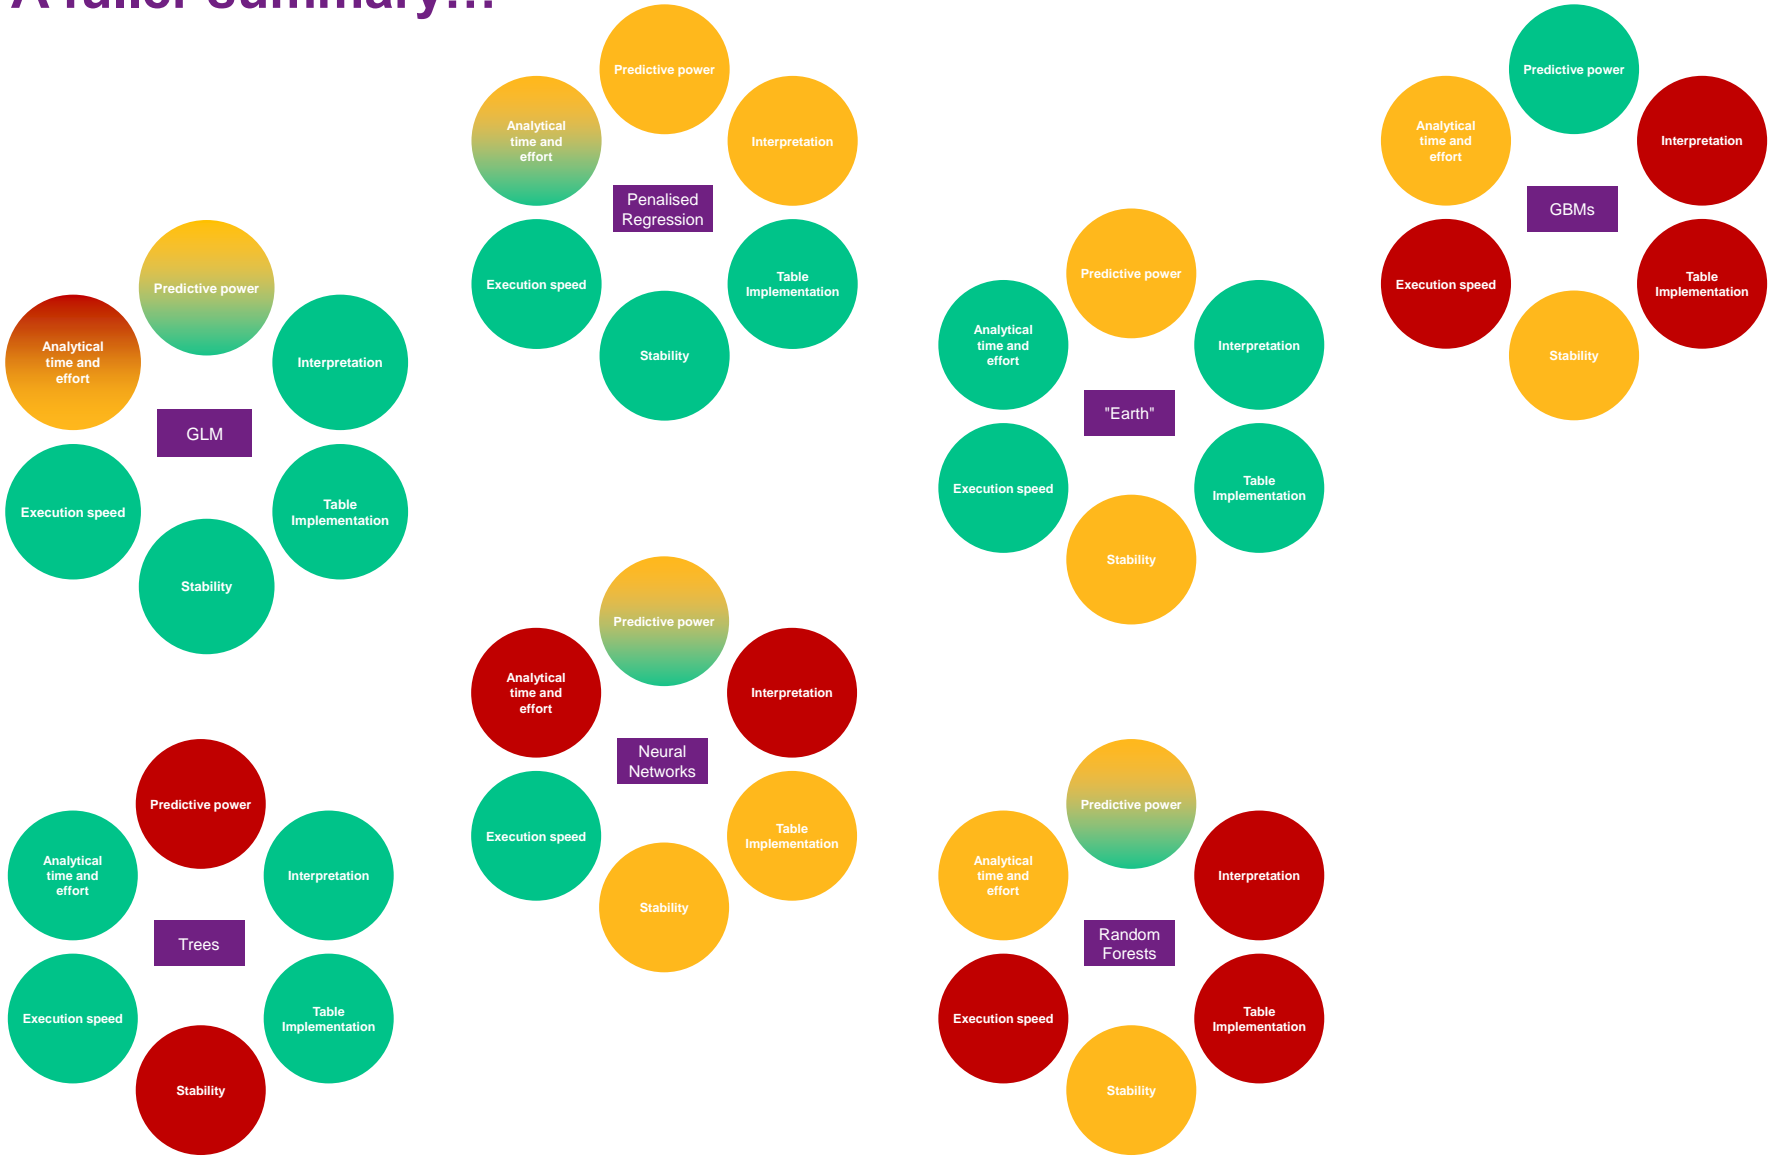
# Practical applications of neural networks in pricing





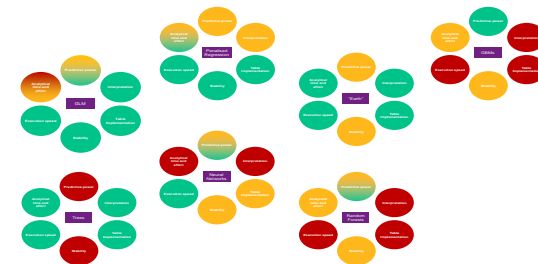


# A fuller summary...



# 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

# 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		
Modelling tools and platforms		
Internal skills sets		
Measuring value		
Application		

# 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 methods		Slowly upward
Modeling tools and platforms		
Internal skills sets		
Measuring value		
Application		

# 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)

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

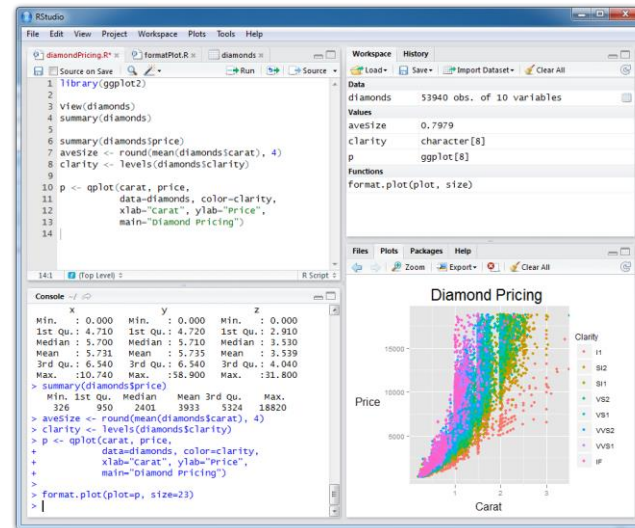
Some critical success factors

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

## Price assessment – scenario testing



VARIABLES, DATA SOURCES AND TECHNICAL ISSUES





## 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%

# 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		
Modeling tools and platforms		
Internal skill sets	?	Slowly upward
Measuring value		
Application		

*“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		
Modelling tools and platforms		
Internal skills sets		
Measuring value		Static
Application		

# 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)

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

Some critical success factors

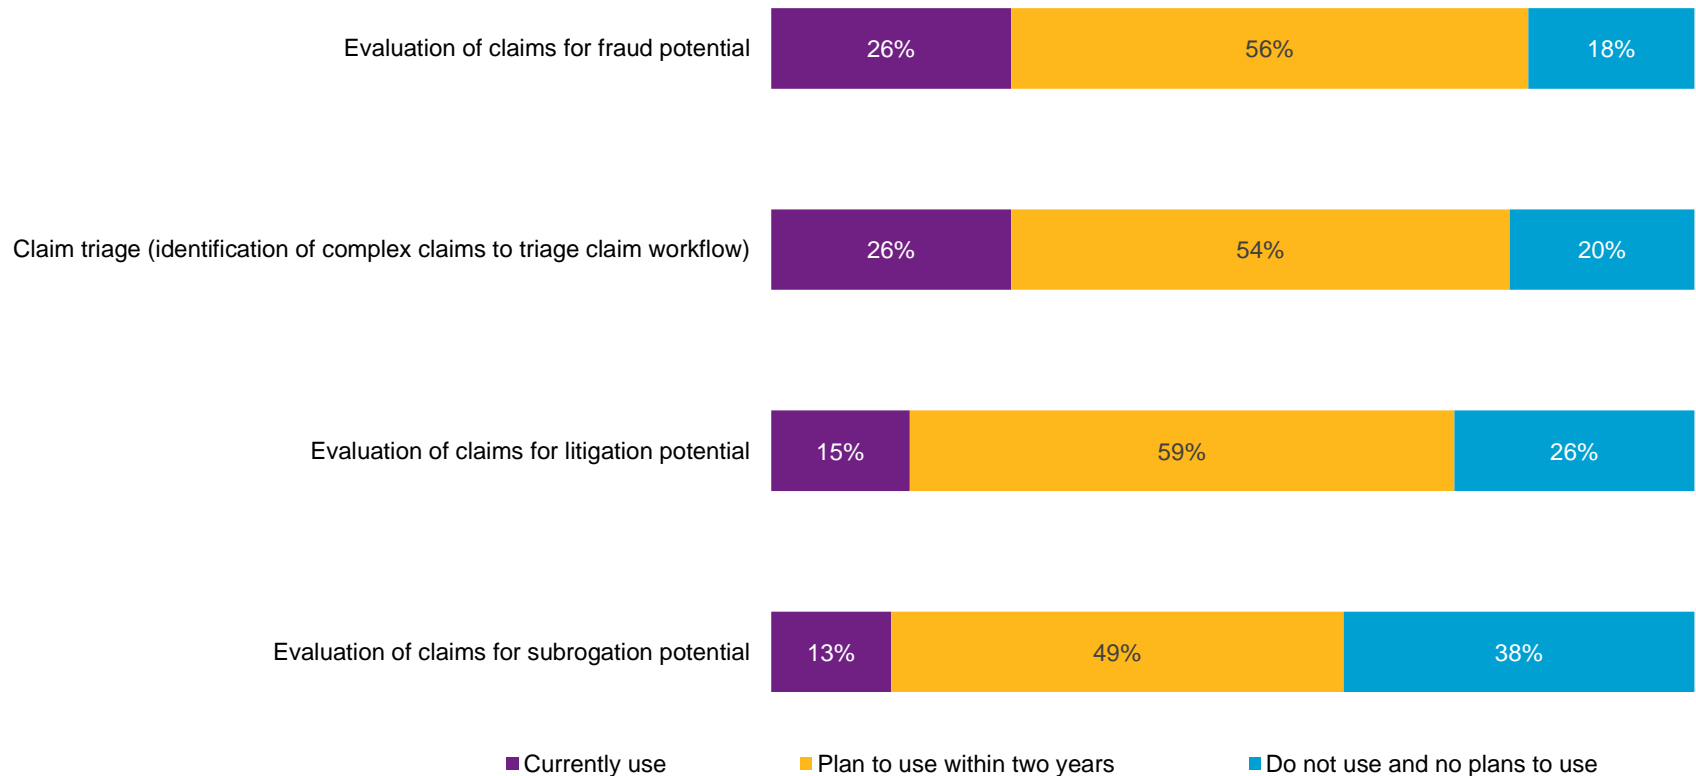
Component	Rating	Directional trend
Data availability		Static
Appetite to try new approaches		
Modelling tools and platforms		
Internal skills sets		
Measuring value		Slowly upward
Application	?	Slowly upward

# 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)

# 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)




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


## Reserving



## Expense management



## Marketing



## Agency/Broker management



■ Currently use
 ■ Plan to use within two years
 ■ Do not use and no plans to use

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

# 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

