EESTI
AKTUAARIDE
LIIT

Forward Together

25th Anniversary Conference

Tallinn
September 27, 2024

Artificial Intelligence in Insurance

EESTI
AKTUAARIDE
LIIT

Charles Cowling
International Actuarial Association

Prof. Meelis Kull
University of Tartu

Bogdan Tautan
Actuarial Association of Europe

Jan Küthe
Akur8

Matvei Mirosnikov
Allianz Commercial

UNIVERSITY OF TARTU

Estonian Centre of Excellence in AI

UNIVERSITY OF TARTU Institute of Computer Science

Risk assessment inside AI: current methods and future outlook

Meelis Kull

University of Tartu

Head of the Estonian Centre of Excellence in AI

Estonian Actuarial Society 25th Anniversary Conference

Estonian Centre of Excellence in AI

3 institutions
13 research groups

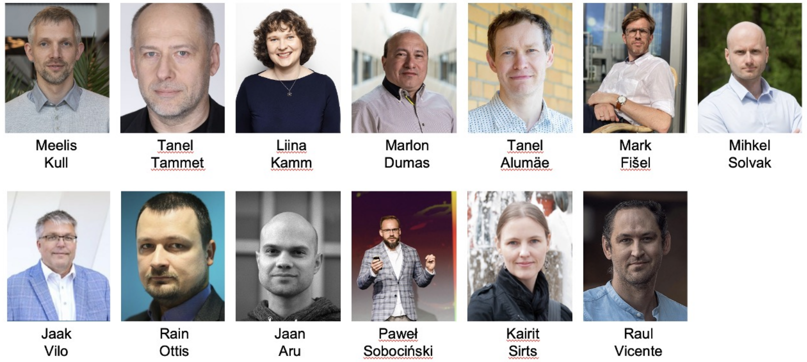

UNIVERSITY OF TARTU
TAL TECH
CYBERNETICA

Foundations of AI:

- Hybrid AI Pipelines
- Adaptation of Foundation Models
- Safeguards and Trust in AI
- Privacy and Security in AI

Applications of AI:

- AI for E-governance
- AI for Healthcare
- AI for Business Processes
- AI for Cybersecurity
- AI for Education

AI and probabilities

AI and probabilities

- Large Language Models output probabilities
- Image classifiers output probabilities
- Speech recognition systems output probabilities

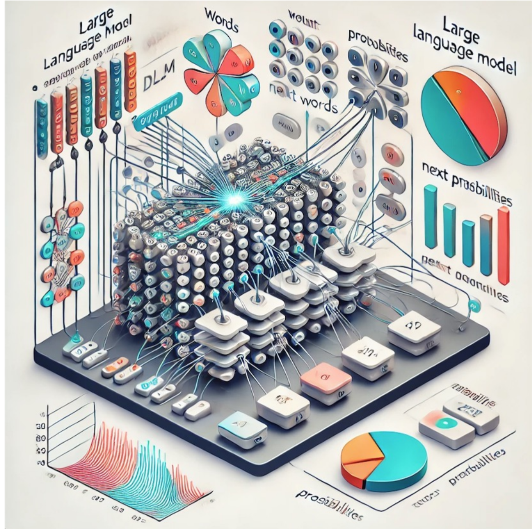

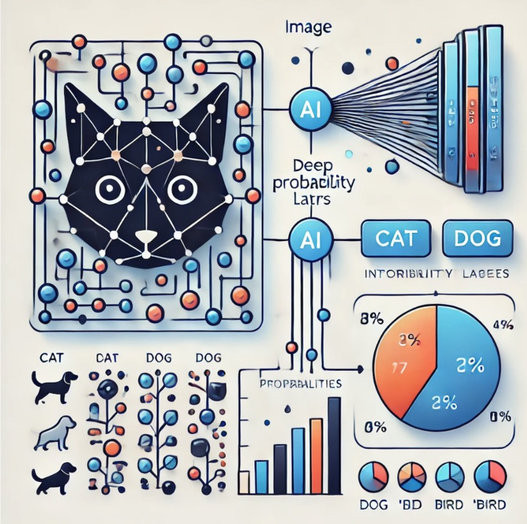

What do these probabilities really mean?

- The probabilities in machine learning are often training-time frequencies, not genuine uncertainty estimates

- Example:
  Coin tossed twice;
  Head and Tail does not imply it is a fair coin

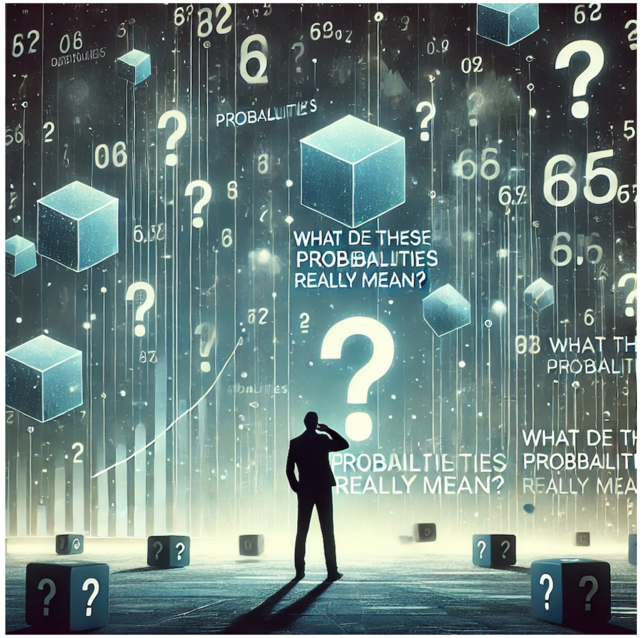

$\not\Rightarrow$ fair coin

Over-confidence in machine learning

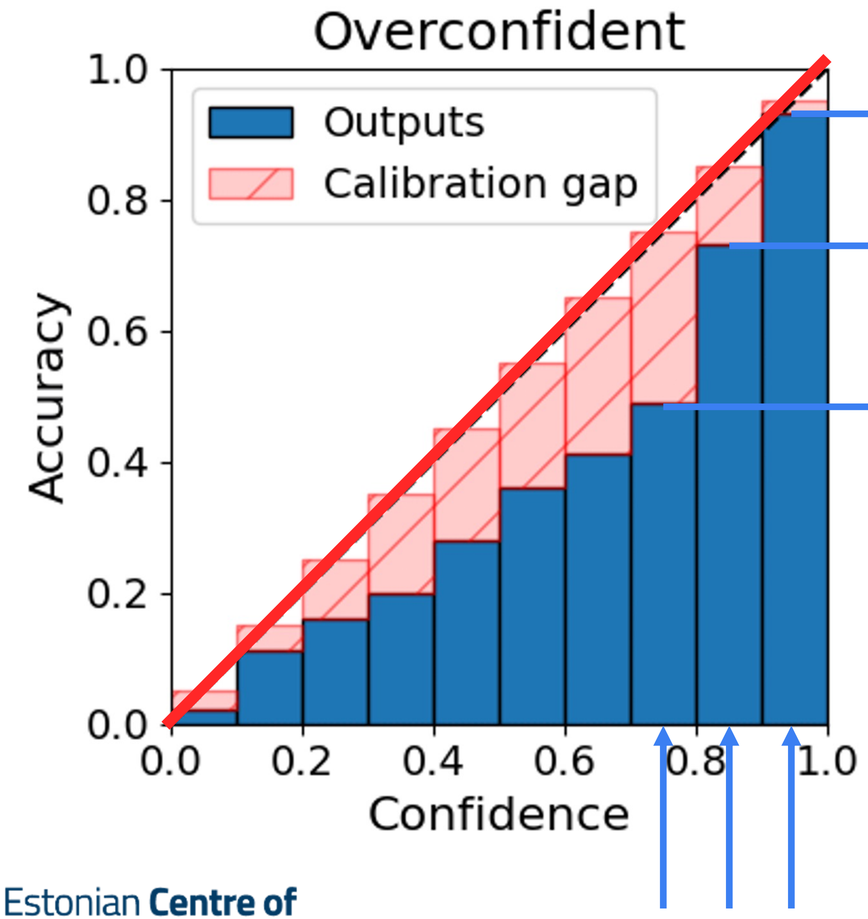

$$P(\text{Accurate} \mid \text{Confidence} = 0.95) = 0.90$$

$$P(\text{Accurate} \mid \text{Confidence} = 0.85) = 0.70$$

$$P(\text{Accurate} \mid \text{Confidence} = 0.75) = 0.50$$

$$\forall p \in [0,1] : \quad P(\text{Accurate} \mid \text{Confidence} = p) = p$$

How do we get calibrated probabilities

- Treat the existing model as a black box
- Use a validation dataset to learn a ‘calibration map’
- On new data:
  use the obtained transformation on top of the black box

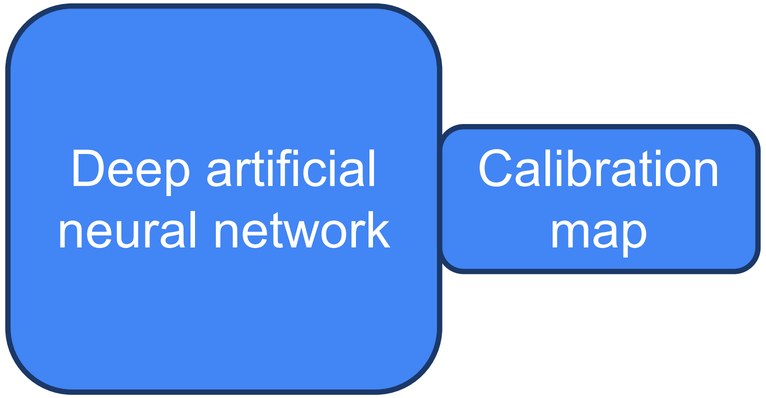

Probabilities in AI –
just existing methods from statistics?

Probabilities in AI – just methods from statistics?

- AI uses, adapts, and advances many methods from statistics:
  - Linear regression, logistic regression
  - Generalised linear models (GLMs)
  - Maximum likelihood methods
  - Variational inference
  - Bayesian statistics
  - ...
- New statistics emerging in the AI field (or between stats and AI)?
  - Deep learning
  - Uncertainty quantification:
    - E.g., scientific conferences: AISTATS, UAI

Uncertainty quantification in AI

- Aleatoric uncertainty - inherently stochastic input-output relationships
  - Will this car have an accident in the next year (given full information about the driver and car)?
- Epistemic uncertainty – lack of knowledge about input-output relationships
  - Will the new material melt at 100 degrees (given full structural information about the material)?
- Epistemic uncertainty is risky for insurance:
  - Someone might have more accurate knowledge than the insurance company
  - Can lead to mispricing
- Most AI methods:
  - Output predictive uncertainty – probability distribution of output given input
  - Do not distinguish aleatoric and epistemic
- Some recent methods aim to quantify epistemic uncertainty separately
  - My research group has an ongoing project on categorising types of uncertainty

Uncertainty in AI - Relevance for Actuarial Science?

- AI is unreliable, how can it estimate risks reliably?

- Use uncertainty calibration methods
  - Are these methods reliable enough?
  - Potential solution – calibration with probabilistic guarantees (e.g. conformal prediction)

- Take AI-extracted features, then apply classical actuarial techniques
  - Adversarial attack – carefully crafted small changes to inputs lead to very different outputs
  - Potential solution – more constrained AI methods

- Separate estimation of epistemic uncertainty
  - Perhaps helps to identify cases where insurance should not be given

Future outlook for AI + Actuarial Science?

- AI as a tool to find more complex patterns
- Dynamic and real-time enhanced risk assessment
- Personalised insurance products
- Automation and efficiency

Estonian Centre of Excellence in AI

Thank you!

Summary

Relevance of AI for Actuarial Science?

- AI is unreliable, how can it estimate risks reliably?
- Use uncertainty calibration methods
- Take AI-extracted features, then apply classical actuarial techniques
- Separate estimation of epistemic uncertainty

Future outlook for AI + Actuarial Science?

- AI as a tool to find more complex patterns
- Dynamic and real-time enhanced risk assessment
- Personalised insurance products
- Automation and efficiency

AKUR8

Company overview

Founded
2018

Global offices
Paris, NYC, London, Milan, Cologne, Tokyo, Atlanta, Montreal

Employees
150+

Nationalities
25+

Activity
Non-Life Insurance Pricing (e.g. P&C, Health, Travel, Pet)

Customers
130+ in 40+ countries

The challenge

Deliver a Pricing Process that is

fast, predictive and interactive

Common attempts to deliver pricing sophistication

Traditional manual pricing process is long (months), iterative and inefficient.

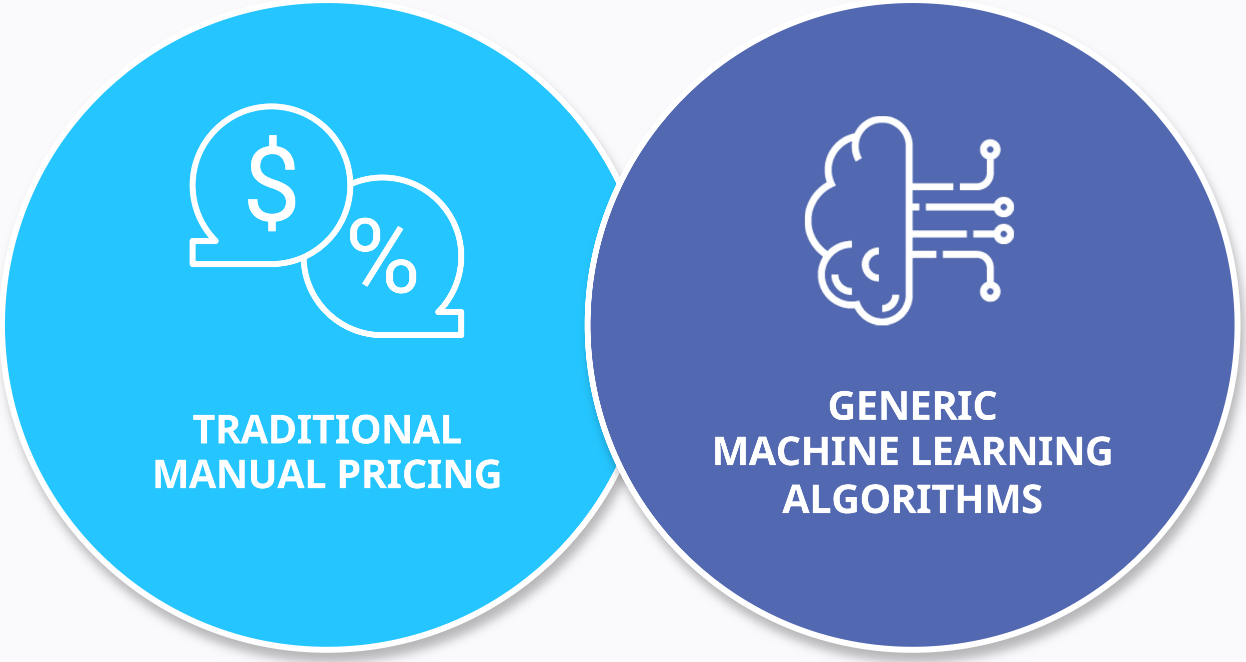

ML models can address those pains but are not explainable (black box), creating unacceptable adverse selection and regulatory issues.

The big picture

<table>
<tr><th></th><th>Levels Selection</th><th>Credibility</th><th>Ridge Regression</th><th>Lasso Regression</th><th>GBM</th><th>Derivative Lasso</th></tr>
<tr><td>Control low-exposure segments to prevent overfitting</td><td colspan="6">All the techniques presented today aim at controlling overfitting</td></tr>
<tr><td>Set coefficients of low-exposure segments at zero</td><td>Selection of effects</td><td colspan="2">No selection of effects</td><td colspan="3">Selection of effects, allowing binary decisions (if the effects are visualized - not always true for GBMs)</td></tr>
<tr><td>Shrink low-exposure segments</td><td>No</td><td colspan="5">This allows to tolerate segments with limited (yet usable) data</td></tr>
<tr><td>Work for multivariate models</td><td>Yes</td><td>No</td><td colspan="4">Yes; apply the same priors / rules for all levels</td></tr>
<tr><td>Creates transparent models (GLM or additive models)</td><td colspan="4">Designed for the GLM framework</td><td>Usually, output not transparent</td><td>Additive models</td></tr>
<tr><td>Natively manage non-linear effects</td><td colspan="4">These techniques work on “pure GLM” (linear or categorical effects)</td><td colspan="2">Yes</td></tr>
<tr><td>Coefficient depending on the robustness parameter</td><td>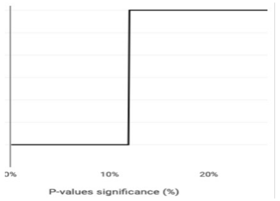</td><td>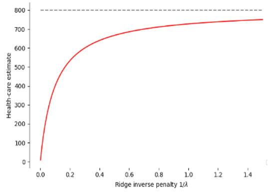</td><td>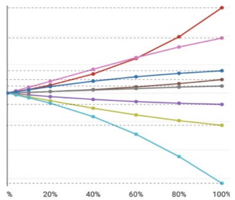</td><td>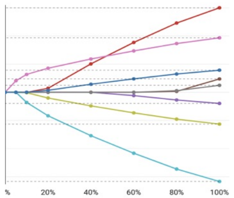</td><td>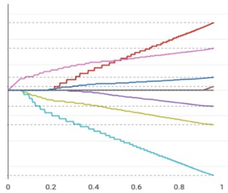</td><td>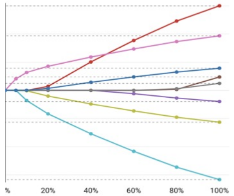</td></tr>
</table>

Constraining the number of variables

$$\begin{cases} \beta^* = Argmax_\beta \; LL(x,y,\beta) - \lambda \, |\beta_i - \beta_{i+1}| \\ \#Variables < N \end{cases}$$

This problem is equivalent to finding the optimal coefficients **placing a "cost" on the number of variables** in the model (using the optimization "lagrangian trick").

$$\beta^* = Argmax_\beta \; \underbrace{LL(x,y,\beta)}_{\text{Minimize the Training Error}} - \lambda \underbrace{|\beta_i - \beta_{i+1}|}_{\text{Maximize the smoothness (control for overfitting)}} - \lambda_v \cdot \underbrace{\#Variables}_{\text{Minimize the number of variables (control for quality)}}$$

The Lagrangian trick is a common optimization technique that allows to transform a problem with multiple constraints into a simpler optimization problem*.

* this optimization problem is actually quite complex to solve; several optimization tricks and approximations need to be leveraged.

Extending the framework

Akur8 vs. Black-box models: control of the understanding

Akur8 allows the creation of complex GLMs which can be compared to black-box models.
However, the main benefit of the GLMs approach is to provide a control over the complexity / performance trade-off.

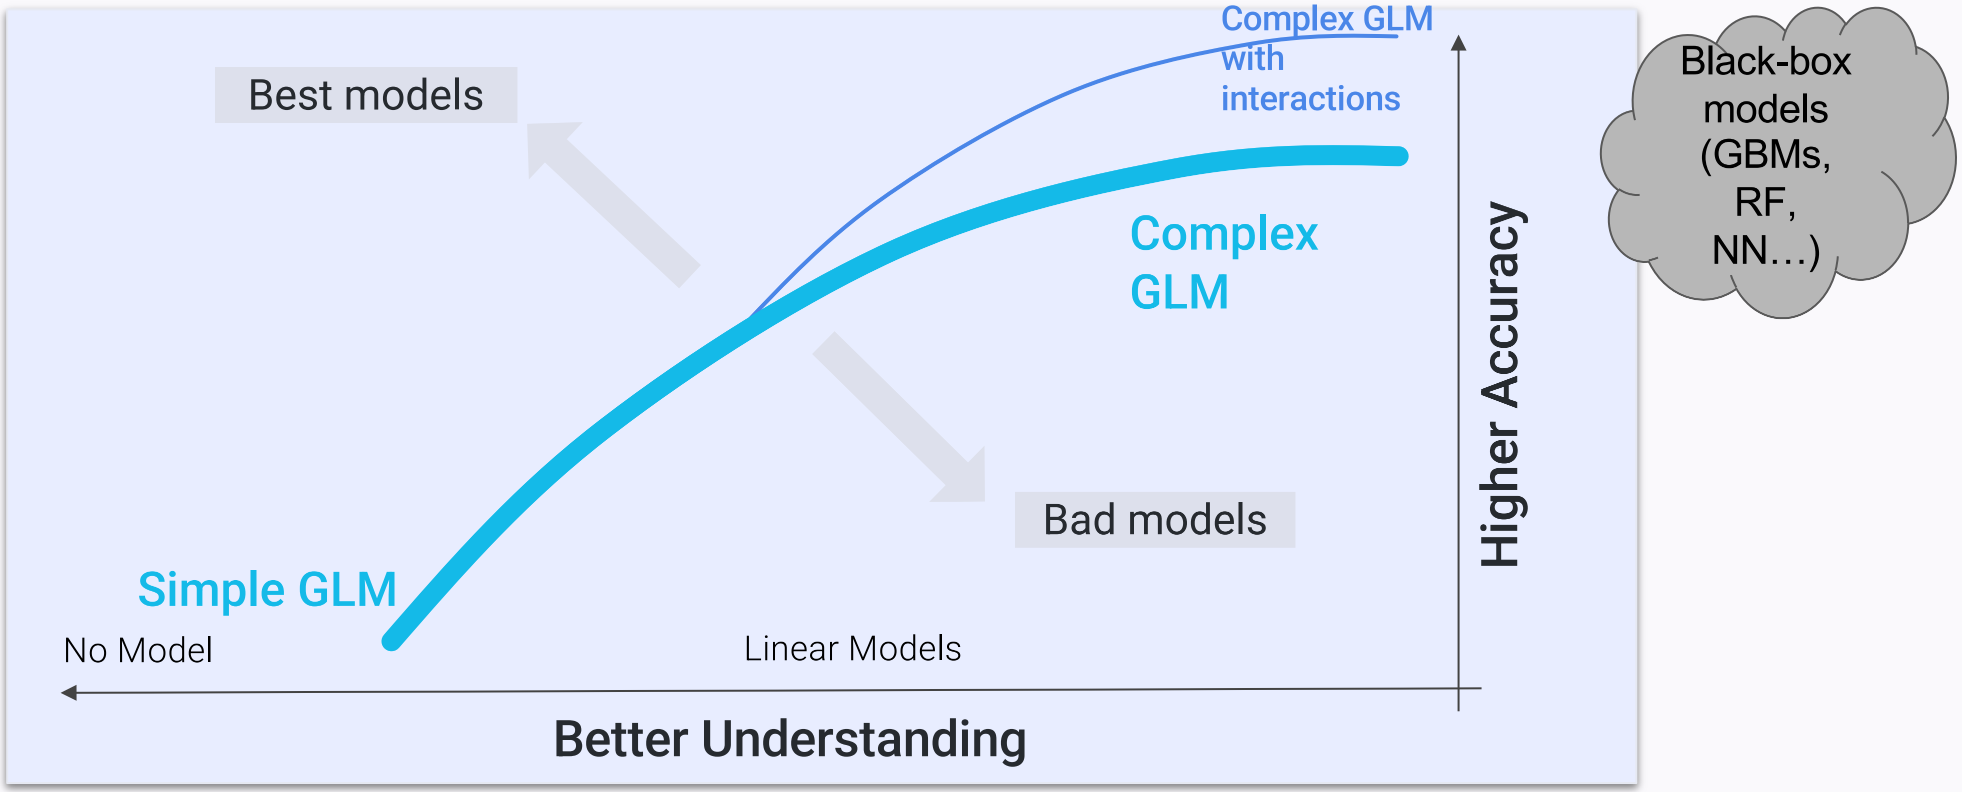

Applying to Interactions

The same principle can be applied in **two dimensions, to fit interactions**. The prior there is slightly different to take into account the 2-D nature of the problem.

For instance, on an interaction between two ordered variables, we could suppose as prior that the differences between all the "connected" levels are supposed to follow a Laplace distribution.

The prior term would become:

$$\begin{aligned} Penalty(\beta) = \cdots &+ \lambda\left|\beta_{18,2} - \beta_{19,2}\right| \\ &+\lambda\left|\beta_{18,2} - \beta_{17,2}\right| \\ &+\lambda\left|\beta_{18,2} - \beta_{18,1}\right| \\ &+\lambda\left|\beta_{18,2} - \beta_{18,3}\right| + \cdots \end{aligned}$$

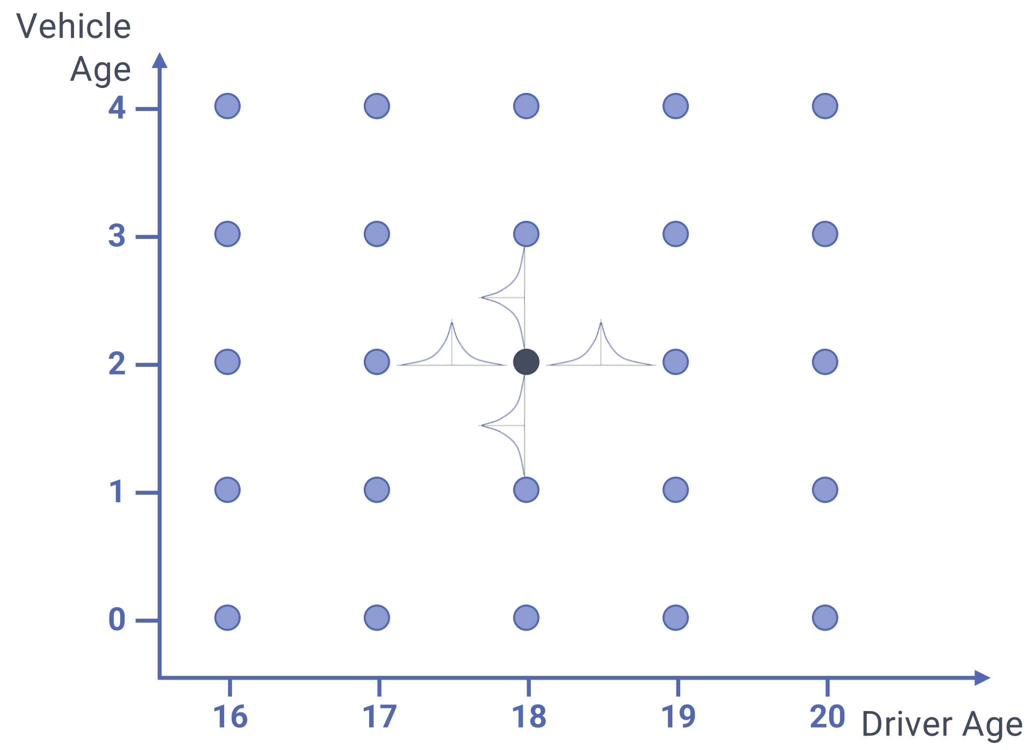

Applying to Interactions

The interactions generated by applying this kind of priors would naturally extend the properties of models to interactions, allowing to identify the relevant ones and fit them automatically.

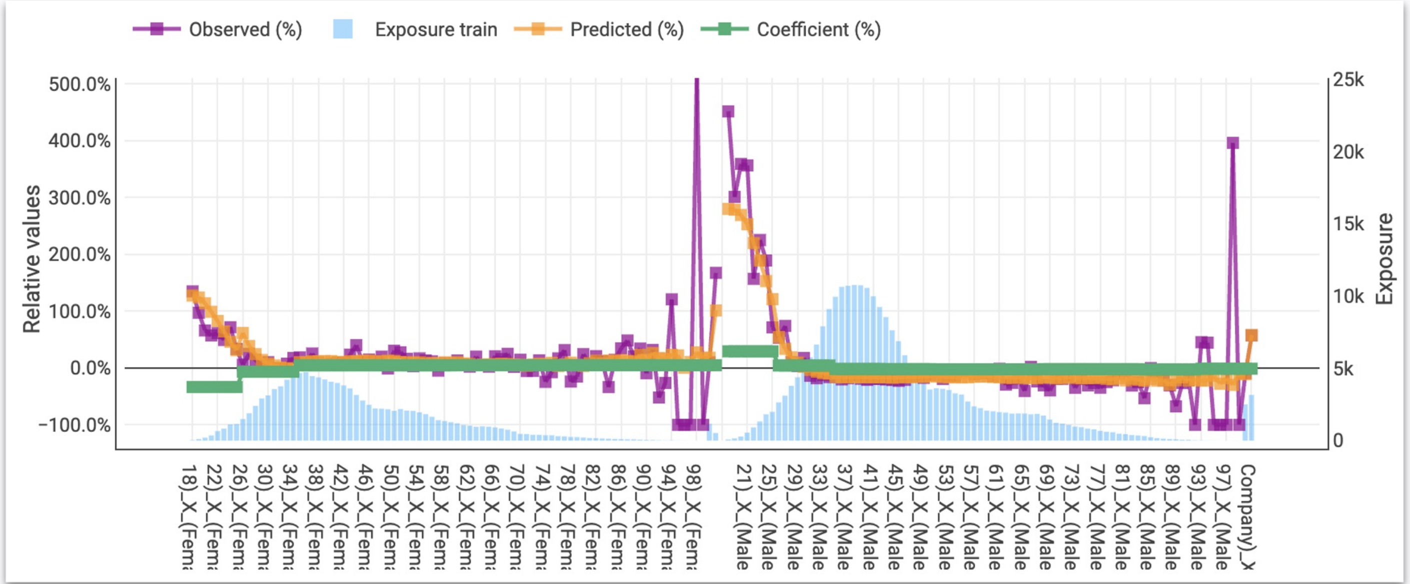

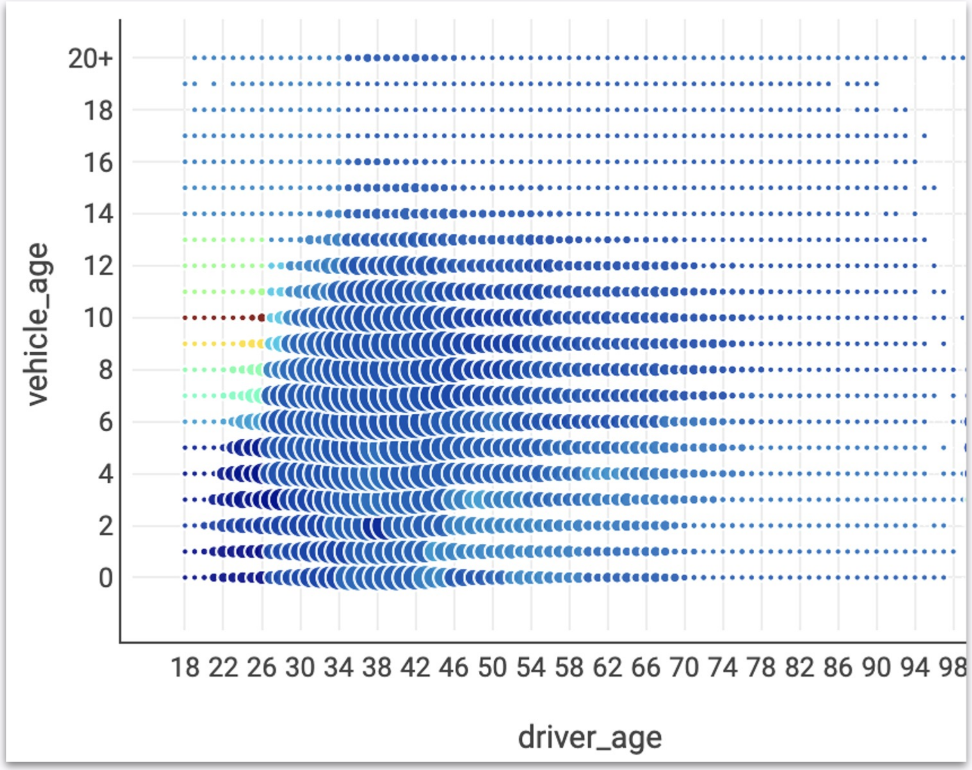

Applying to Geography

Geographic modeling can also be achieved with a similar method: the prior is that **nearby locations are expected to have similar risk levels**.

This has strong similarities to a **Gaussian Process** modeling.

Weak Prior

Intermediate Prior

Strong Prior

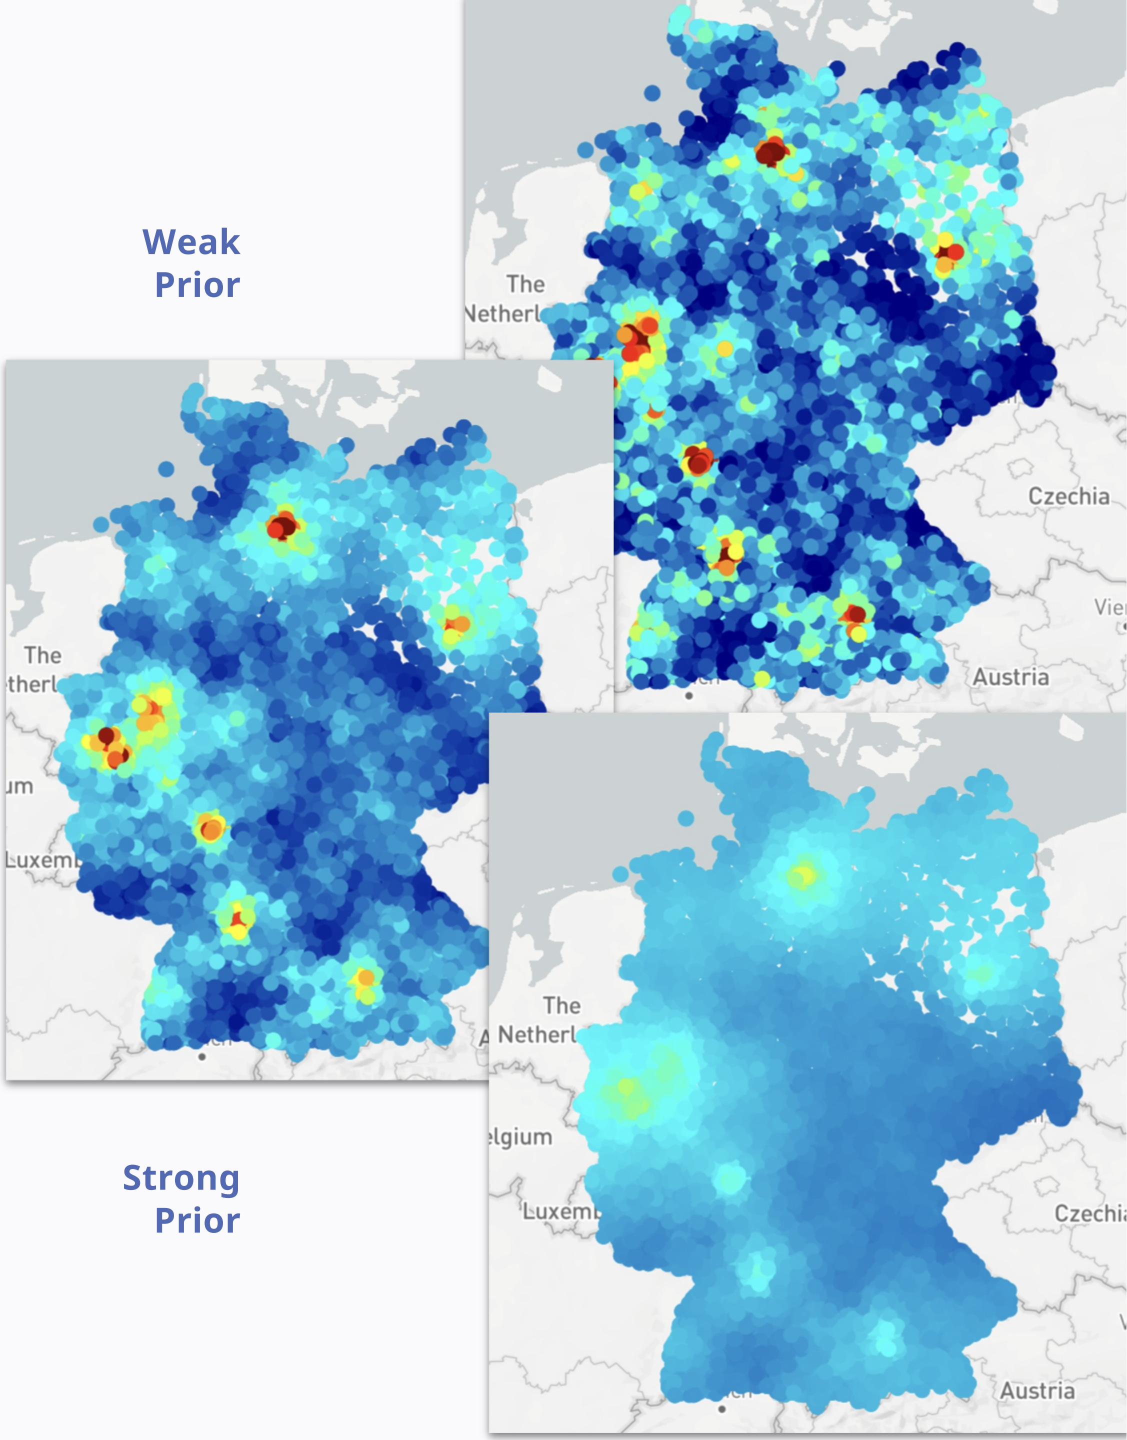

AKUR8

Summary of the Potential

Of efficient, data-driven GAMs

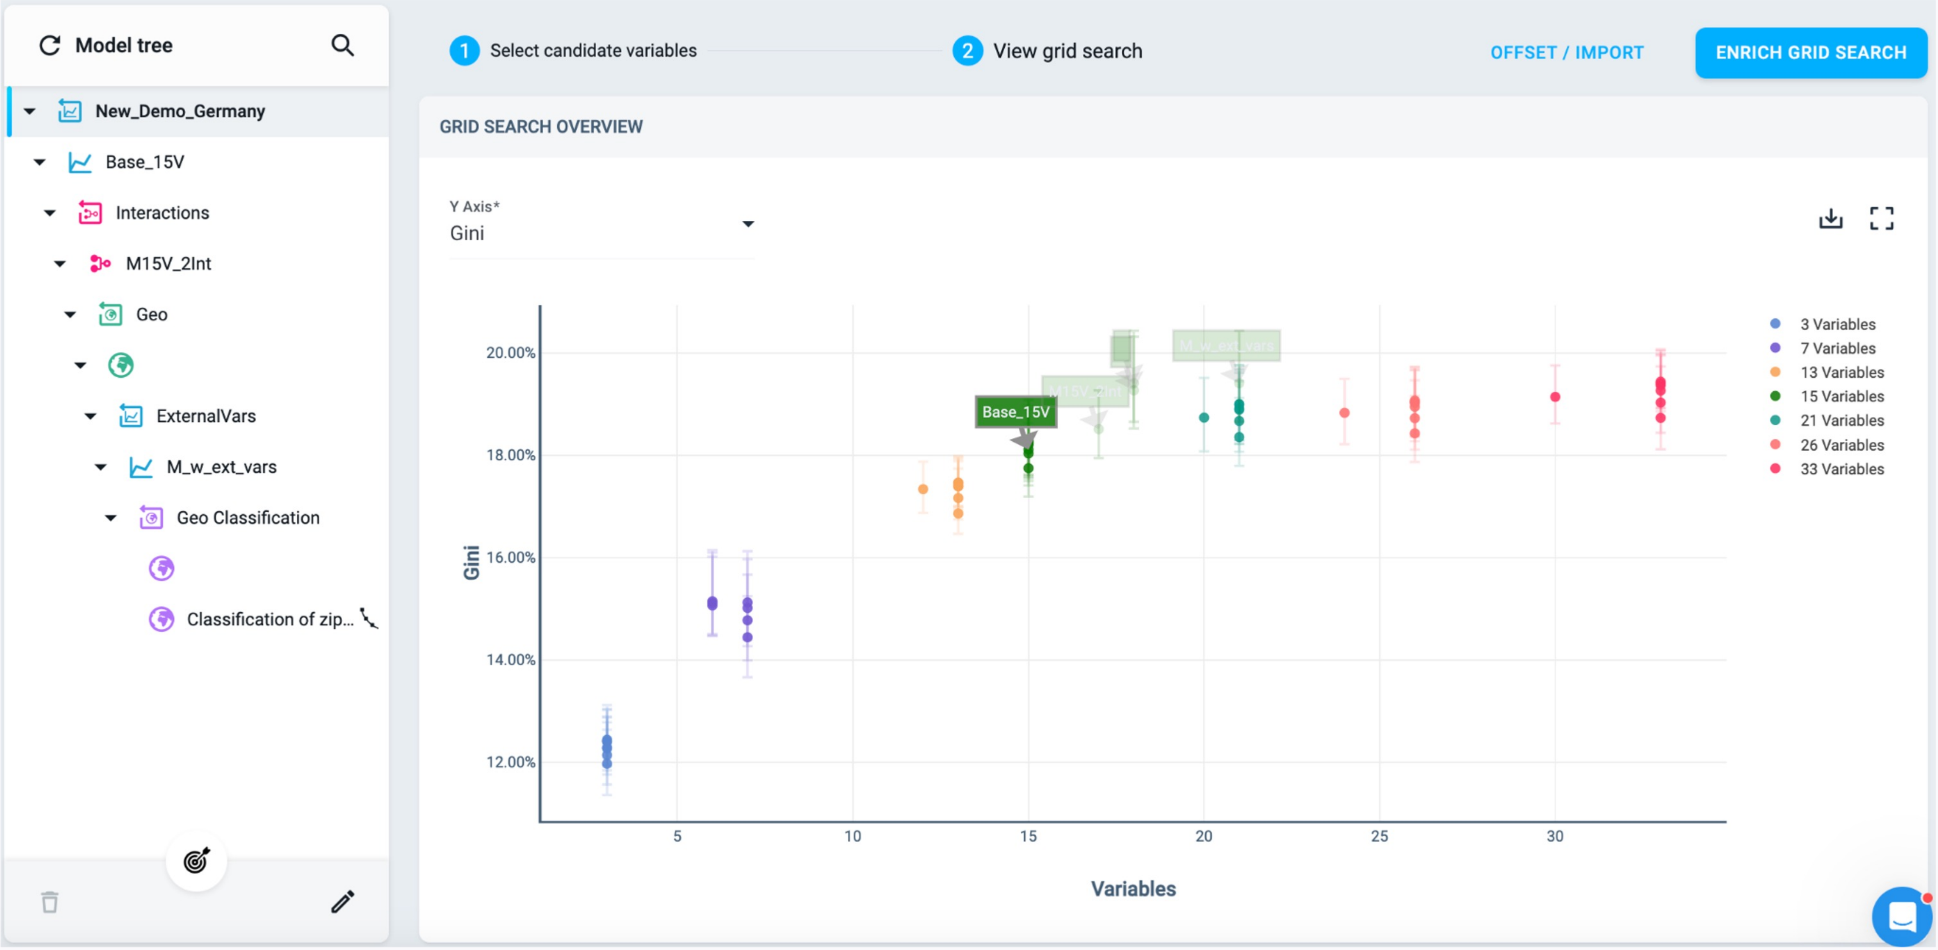

Thank you!

jan.kuethe@akur8.com

Addendum

The coefficient path graph

How to 'rescale' the impact of the penalty

It is possible to generalize this graph, tracking the **impact of penalty on several levels** simultaneously.

The **'coefficient path graph'** allows to globally analyse how the estimates /coefficient evolve when the smoothness increases:

- Y axis represents the value of the estimates.
- X axis represents the 'Empirical Credibility' - which is a 'Proportion of the GLM solution)

$$\text{Empirical Credibility} = \sum_{i \in \text{Classes}} \frac{|\text{Predicted}_i - \text{Grand Average}|}{|\text{GLM}_i - \text{Grand Average}|} \%$$

- Empirical Credibility = 100 % - Estimates match the observed
- Empirical Credibility = 0 % - Estimates match the Grand Average (or complement of credibility)

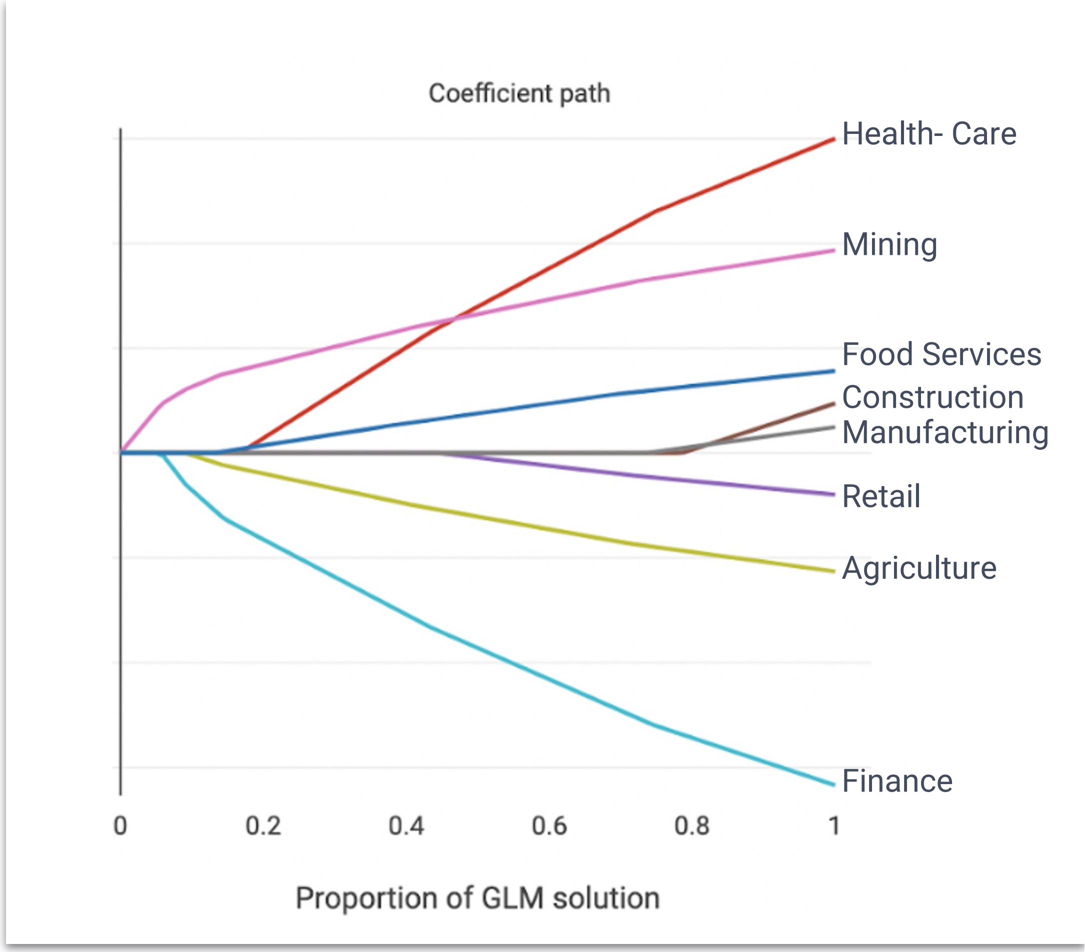

Coefficient path graph of the Lasso

Workers Compensation example

Lasso and Ordinal variables

Under these **"Lasso"** assumption on the **derivative**, penalized regression can **natively incorporate non-linear effects**.

Furthermore, the convergence result between GBMs and Lasso is still valid.

To control the training error and ability to generalise:

- Penalized Regression require the definition of a **single parameter**: the **smoothness**
- GBMs require to determine the combination of **several parameters:**
    - **number of trees**
    - **learning rate**
    - and other tree-related parameters

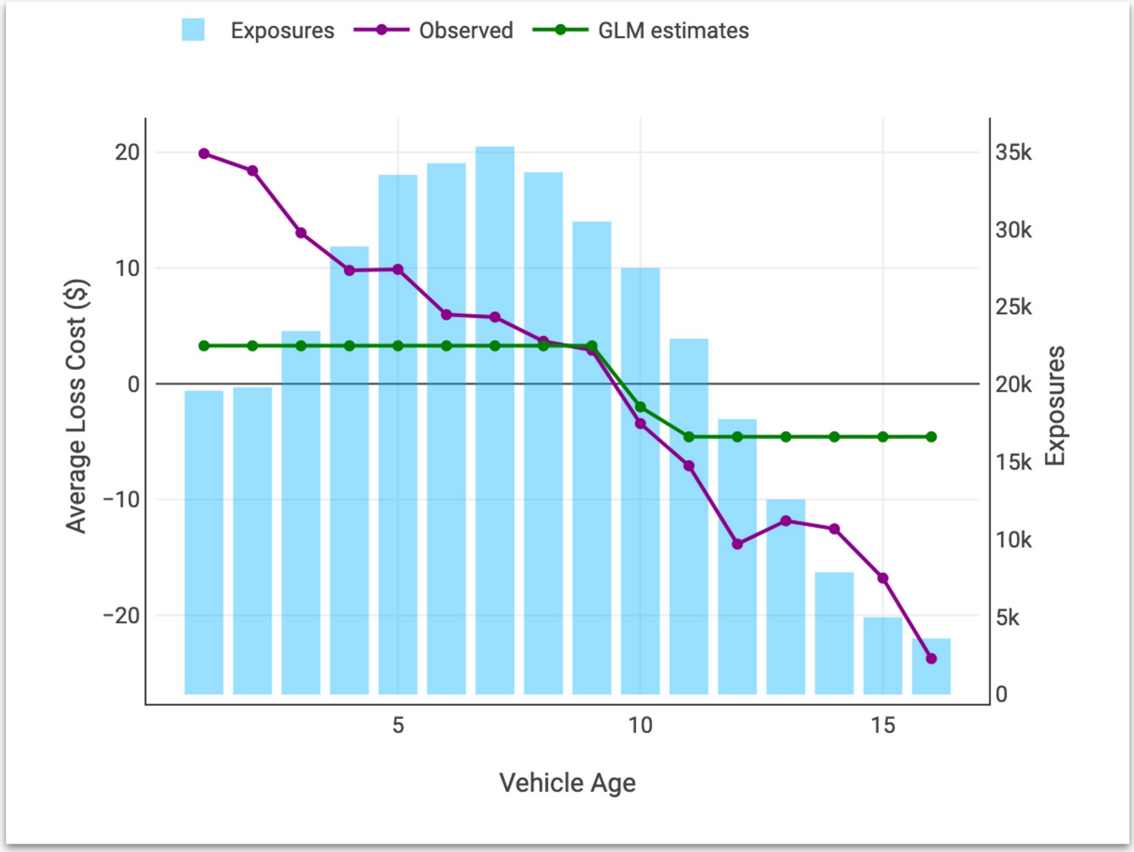

International Actuarial Association
Association Actuarielle Internationale

Safeguarding responsible AI

Implications of the AI Act on the actuarial work

Bogdan Tautan

Forward Together Conference

Agenda

- The global perspective
- Key information on the EU AI Act
- Risk management implications
- Key takeaways

The global perspective

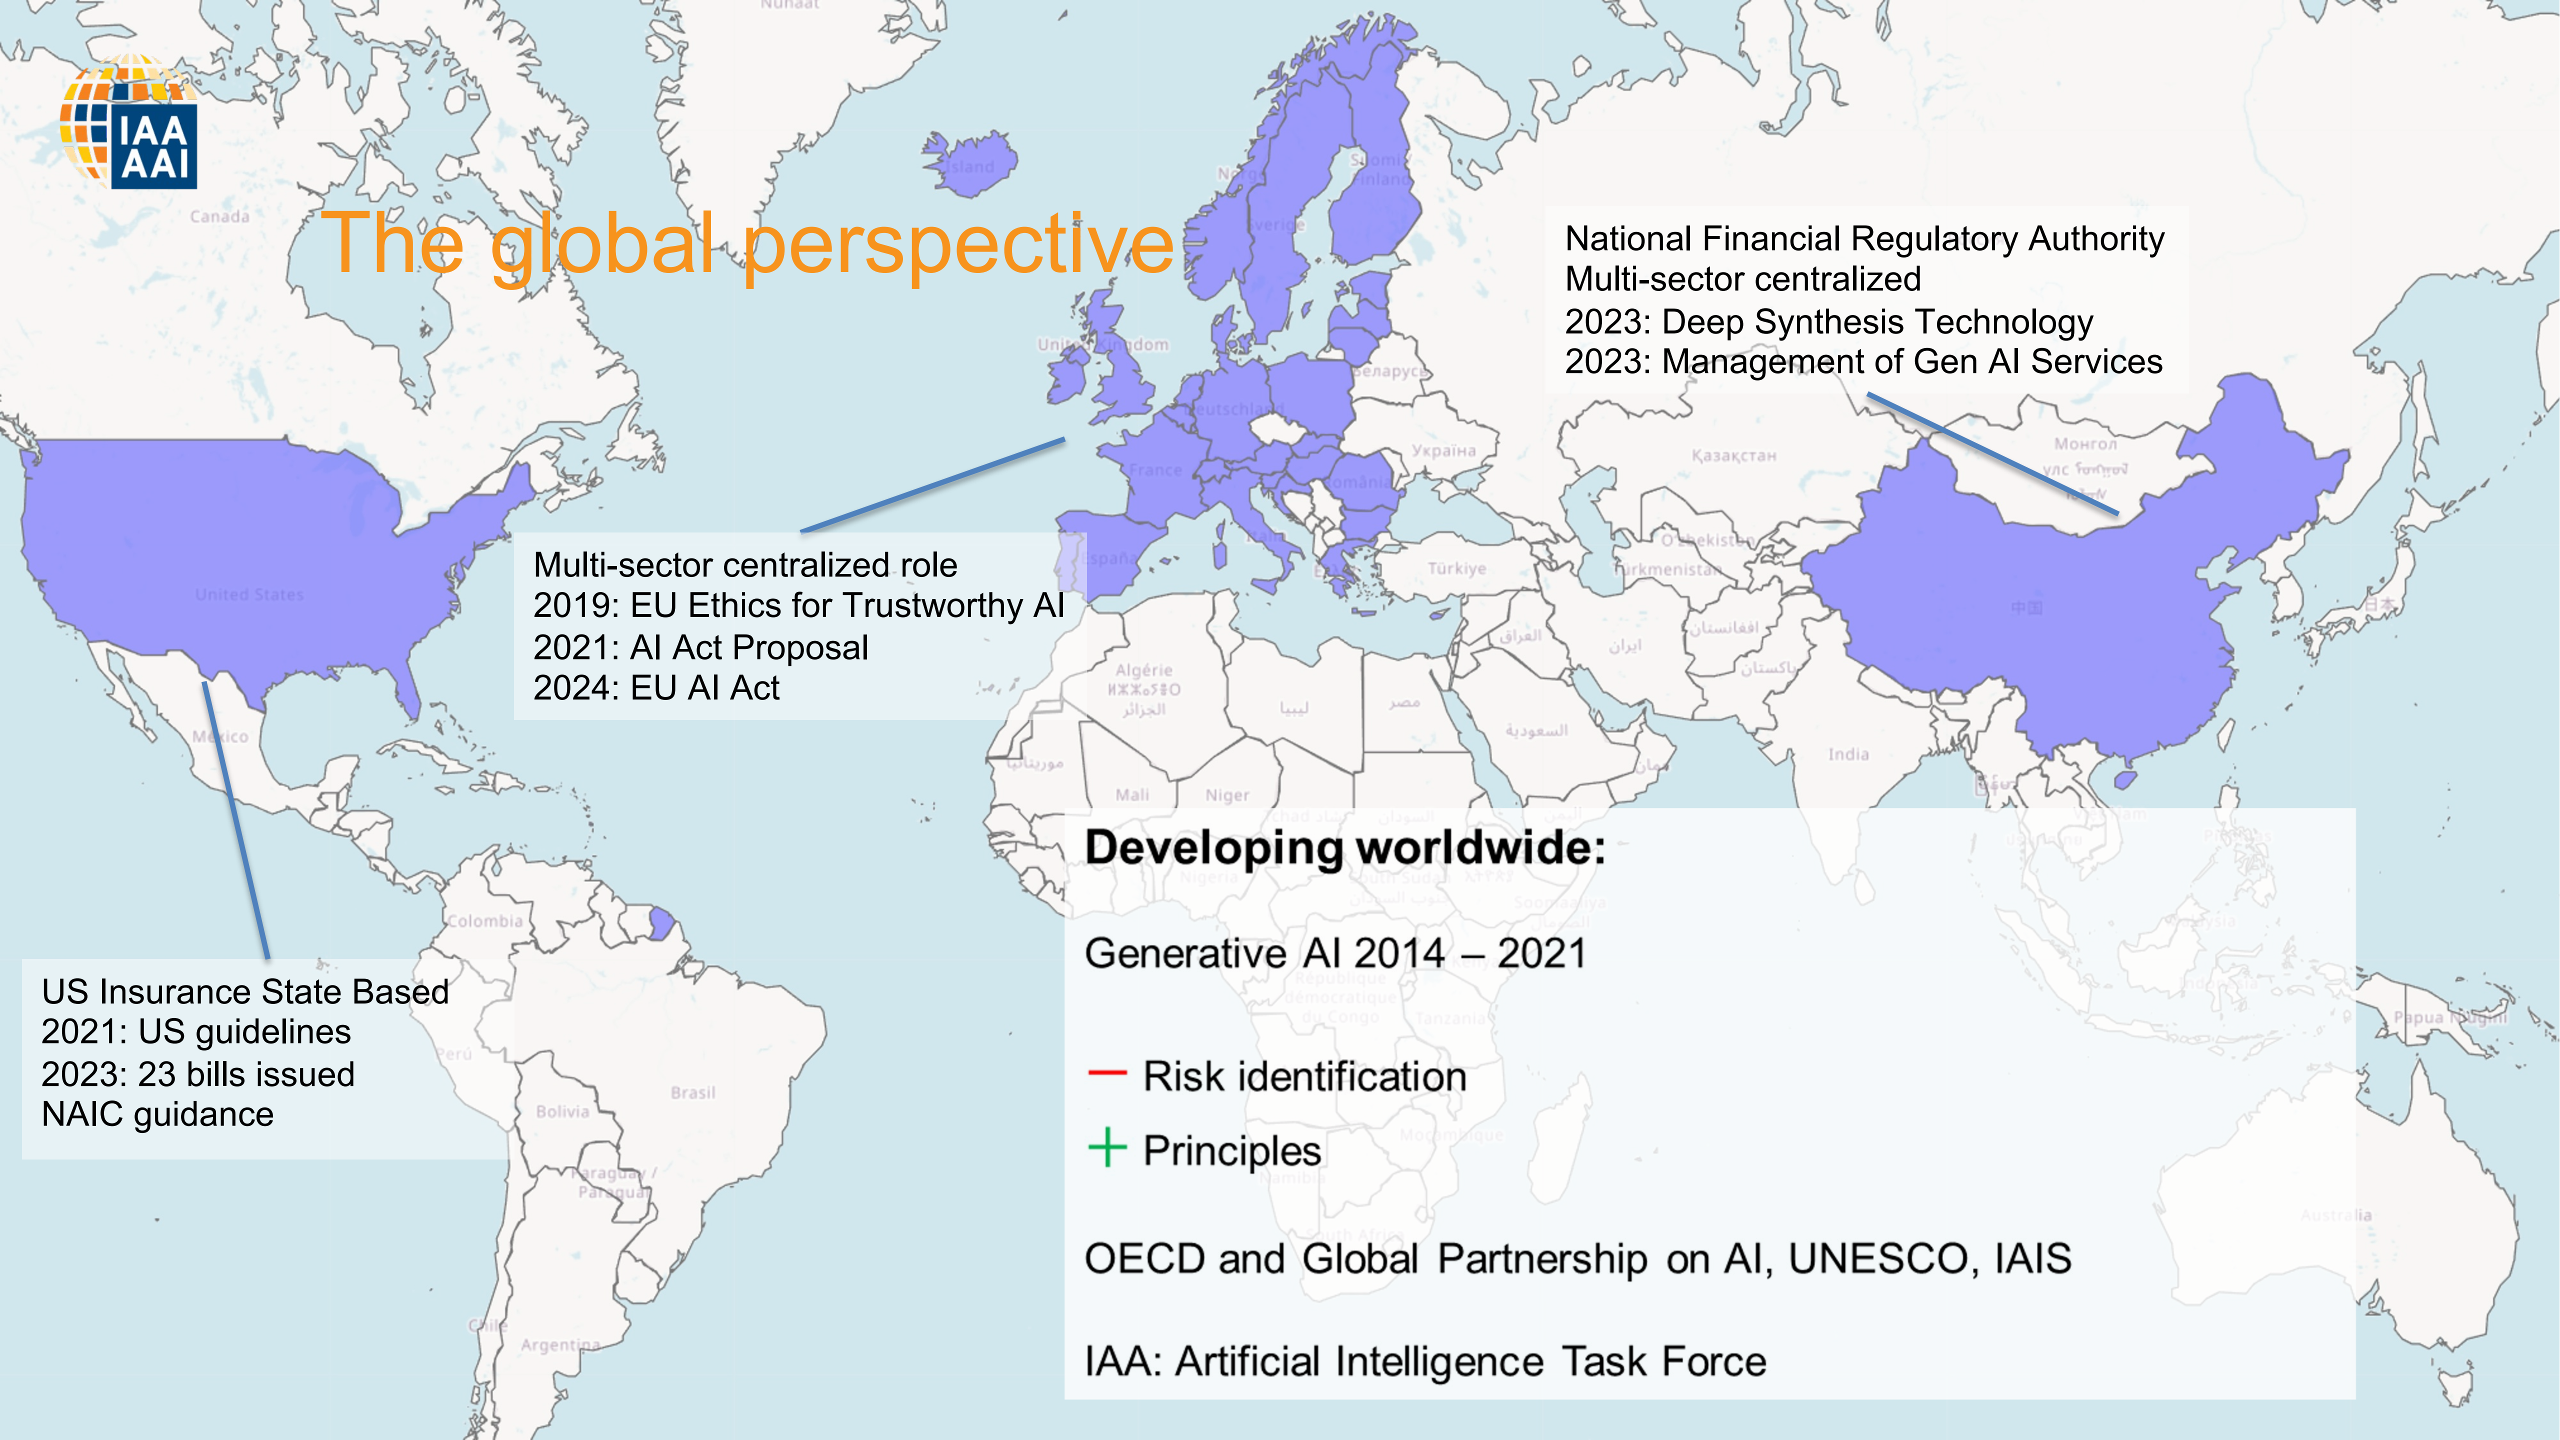

National Financial Regulatory Authority
Multi-sector centralized
2023: Deep Synthesis Technology
2023: Management of Gen AI Services

Multi-sector centralized role
2019: EU Ethics for Trustworthy AI
2021: AI Act Proposal
2024: EU AI Act

US Insurance State Based
2021: US guidelines
2023: 23 bills issued
NAIC guidance

Developing worldwide:

Generative AI 2014 – 2021

– Risk identification

\+ Principles

OECD and Global Partnership on AI, UNESCO, IAIS

IAA: Artificial Intelligence Task Force

EU AI Act

Final agreement on 13th of March 2024 agreement

- **Safeguards human oversight**
- **An iterative risk-management process**
- **Risk classification of AI systems**

European AI Office – the center of AI expertise across the Union:

- **AI Board**
- **Scientific Panel**
- **Advisory Forum**

Penalties up to:

- **Non-compliance:** 7% of annual turnover or 35 mln. EUR
- **Violations:** 3% of annual turnover or 15 mln. EUR
- **Misleading information:** 1% of annual turnover or 7,5 mln. EUR

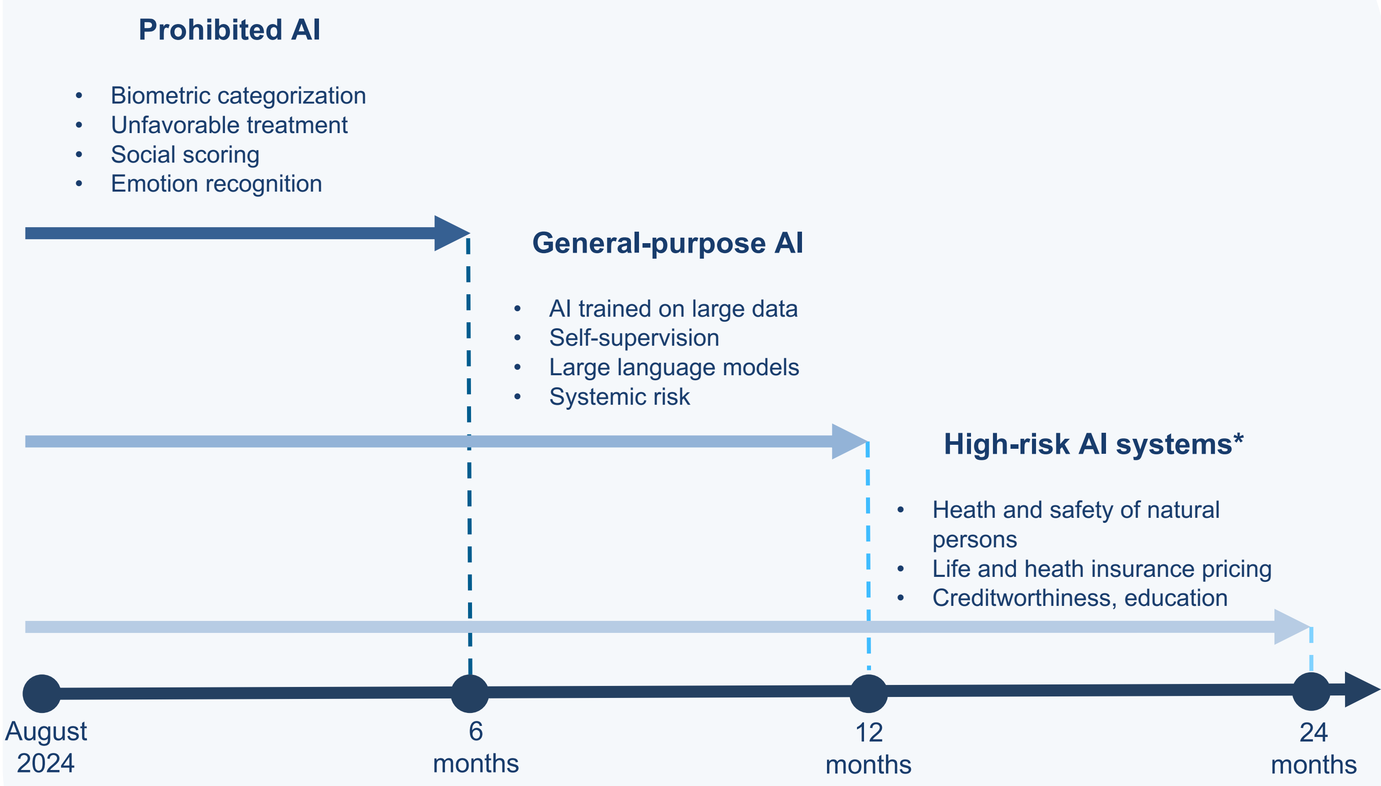

*Obligations for high-risks systems under Annex III; 36 months will be for those under Annex II

EU AI Act

Deployer

Data governance

Human oversight

Record-keeping

User of the AI system, ensuring right use of data, safeguarding human oversight and monitoring of the system

Distributor

Conformity

Compliance

Record-keeping

Makes the AI system available to the market, ensuring the standards, compliance and storage.

Provider

Data governance
- Security
- Logging
- Quality

Documentation
- Technical
- Key elements
- Compliance

Risk Management
- Identify
- Assess
- Monitor

Robustness
- Accuracy
- Cybersecurity
- Instructions

Human oversight
- Prevent risks
- Measures
- Implementation

Transparency
- XAI
- Information
- Accuracy

Record-keeping
- Monitoring
- Usage time
- Data

Compliance
- Registration
- Logging
- Quality

Develops the system and ensures the appropriate operation of an AI system, including the conformity and qualitative assessment, technical documentation and registration within the EU Database (high-risk).

Risk management implications

- Management
- Quality
- Completeness
- Accuracy
- Security
- Synthetic assumptions

- Scalability
- Reporting
- Robustness and Security
- XAI and Transparency
- Business alignment

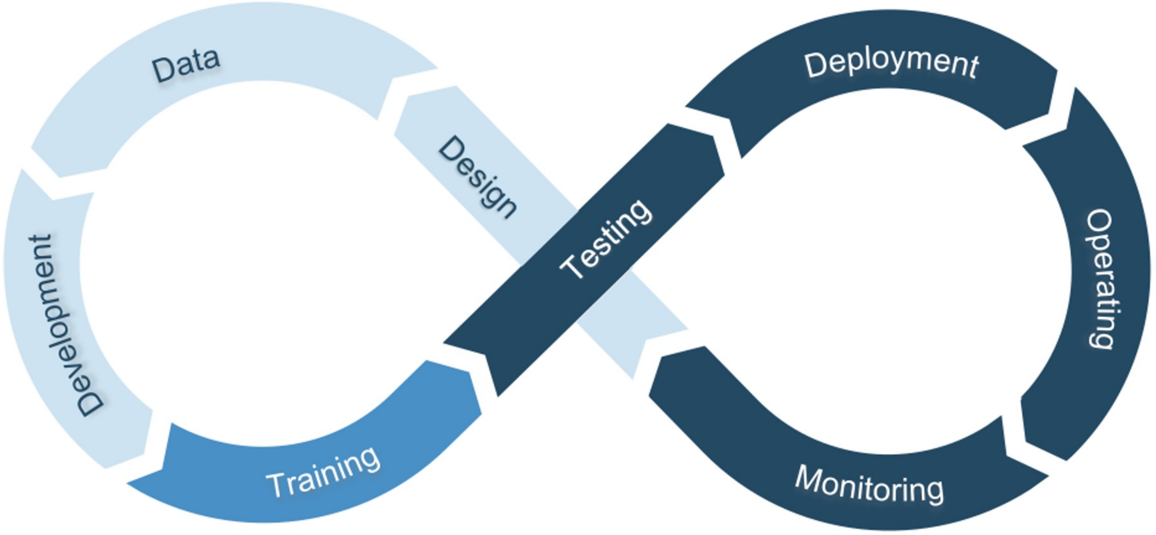

- Assumptions
- Parameters
- Bias and drift
- Model scorecards
- Fitting
- Test vs Training

Relying on existent frameworks

- GDPR
- Data Act, DORA
- Insurance Distribution Directive, Consumer Protection Code
- NIST (US), ISO RM practices
- Solvency II

New responsibilities for actuaries

- Understand current regulations and assess adaptiveness
- Address the skill gap within the organization:
  - Ontology and prompt engineering
  - Understanding programming languages
  - Infrastructure and environment
- Connect XAI to the business needs and the application

Key takeaways

AI Governance workstream's view

- Worldwide developments
- Governn responsibly
- Follow best practices
- Engage with stakeholders

1st line
Data
Development

2nd line
Compliance
Standards

3rd line
Audit
Reporting

- Definitions are (somehow) important
- Link-up disciplines – new skills and roles
- End user responsible for the deployment
- Enhance processes using Gen AI

Thank you!

25