AI in data and pricing

How do we balance opportunities and complexities?

Actuarial Congress 2026

Inge Lodder AAG, MSc – Founder of ELAVV

ELAVV

We stand for innovation, digitalisation & advising for complex financial systems

- Technology
- Data
- Artificial Intelligence
- Actuarial & Risk

Inge Lodder AAG, MSc
Inge.lodder@elavv.nl

Short introduction story

ELAVV

2014-2018 - Amazon’s Selection Policy

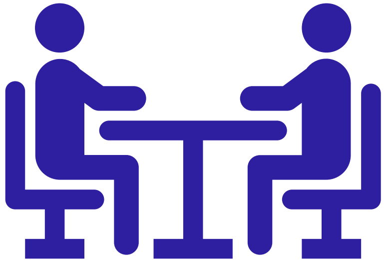

2014-2018 - Amazon’s Selection Policy

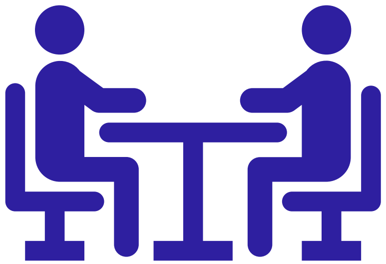

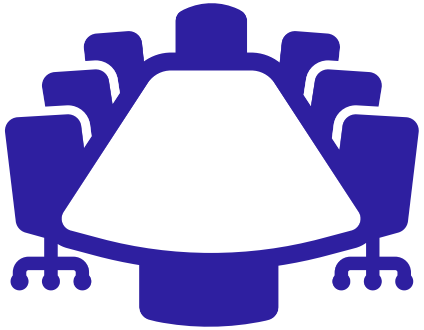

2014-2018 - Amazon’s Selection Policy

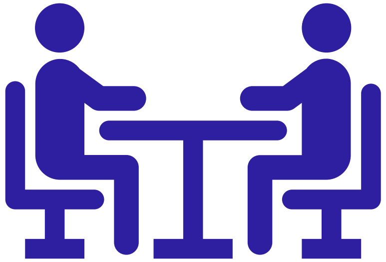

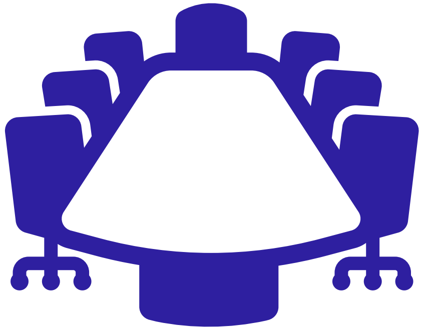

Today:

3 practical cases considering AI

The cases

1. Pricing & model bias
2. The GenAI trainingsloop
3. Privacy issues of external algorithms

1. Pricing & model bias

Opportunity in the Pricing of Non-Life Insurance Products

- The application of ML-techniques can prove to increase the preciseness of price estimates.
- Less data might be needed to make a more precise estimate.
- Less human interference might be needed, speeding up the process.

Discrimination in the Pricing of Non-Life Insurance Products

Legal	Ethnicity, Gender, Sexual Orientation, Political views, etc.
Social / Sales	Dutch consumers value their privacy more than those in other countries (Source: AFM). Data sharing is not taken lightly (!) .
Ethics / Image	Would an insurer actually want this from an ethical perspective? Does it show innovation or will it scare customers?

Pricing & AI in Practice (CRISP-DM model)

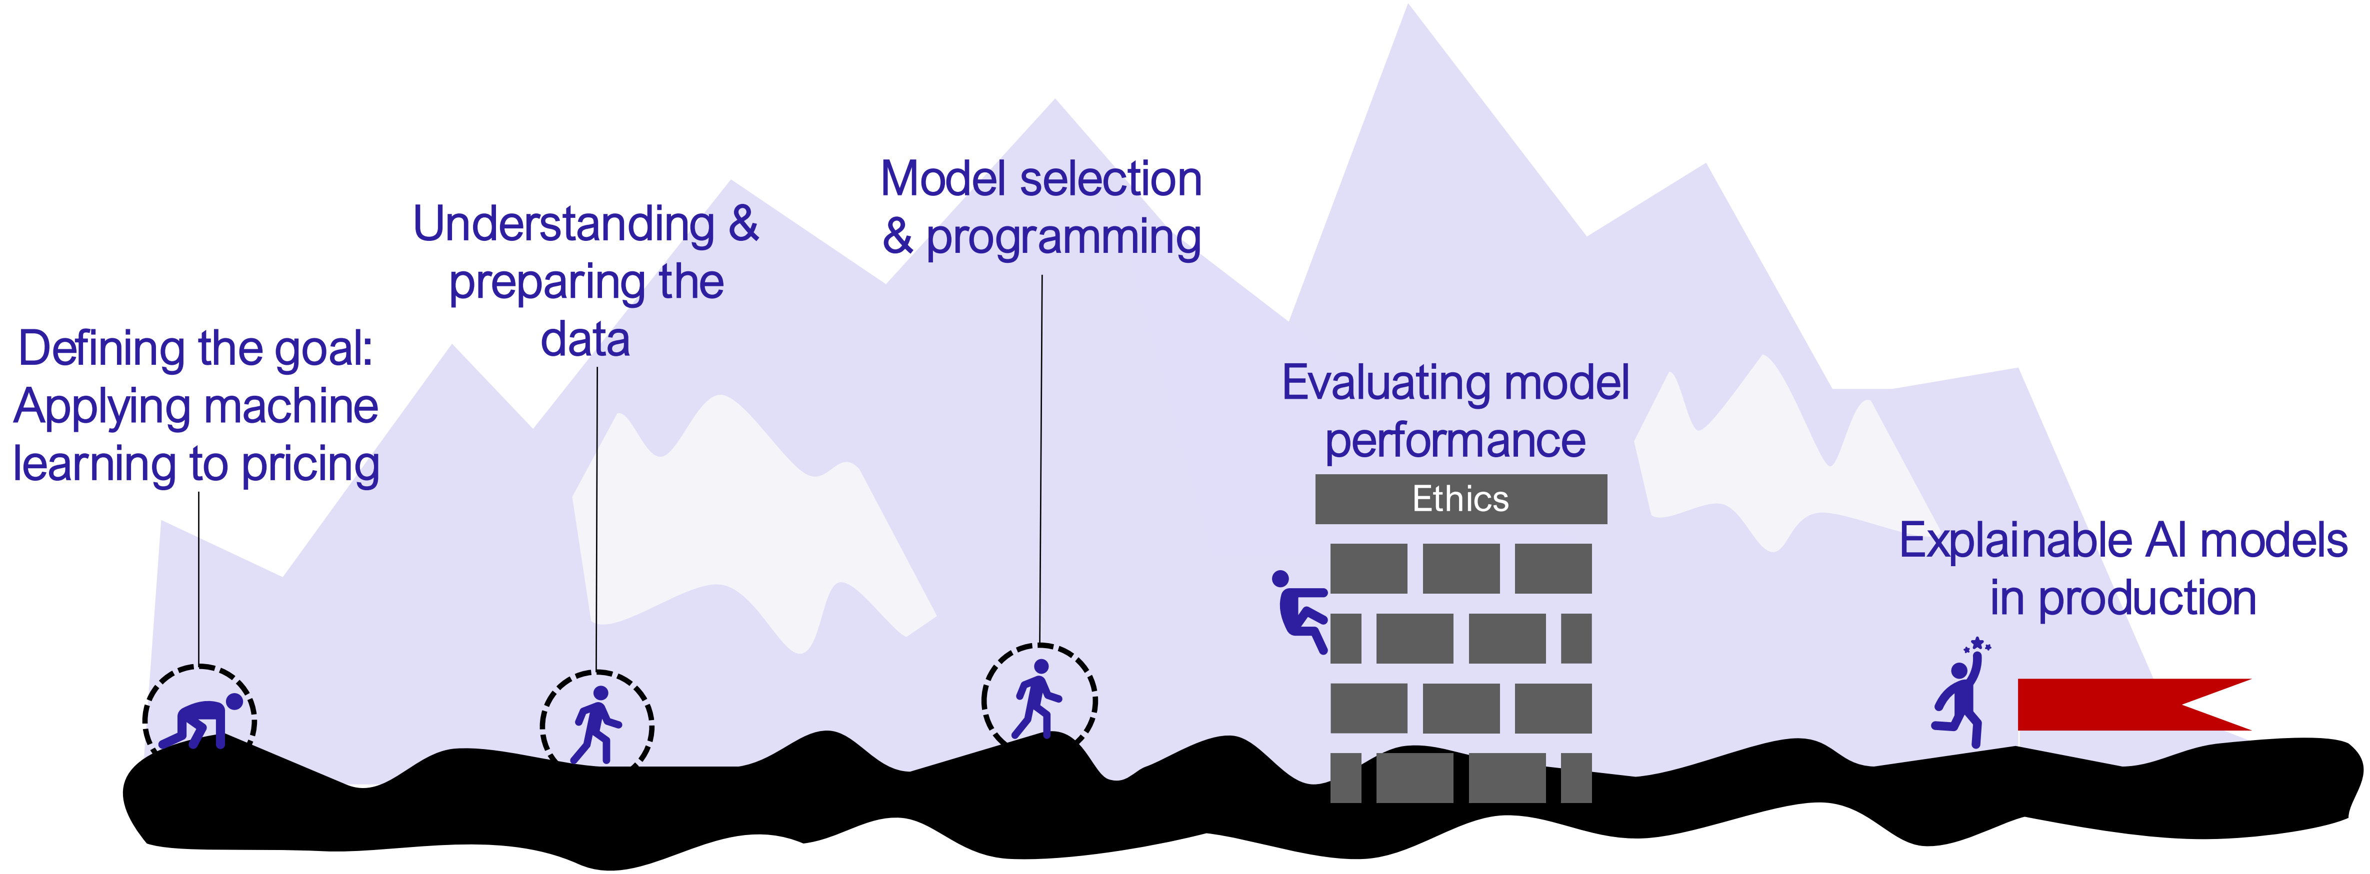

Pricing & AI in Practice (CRISP-DM model)

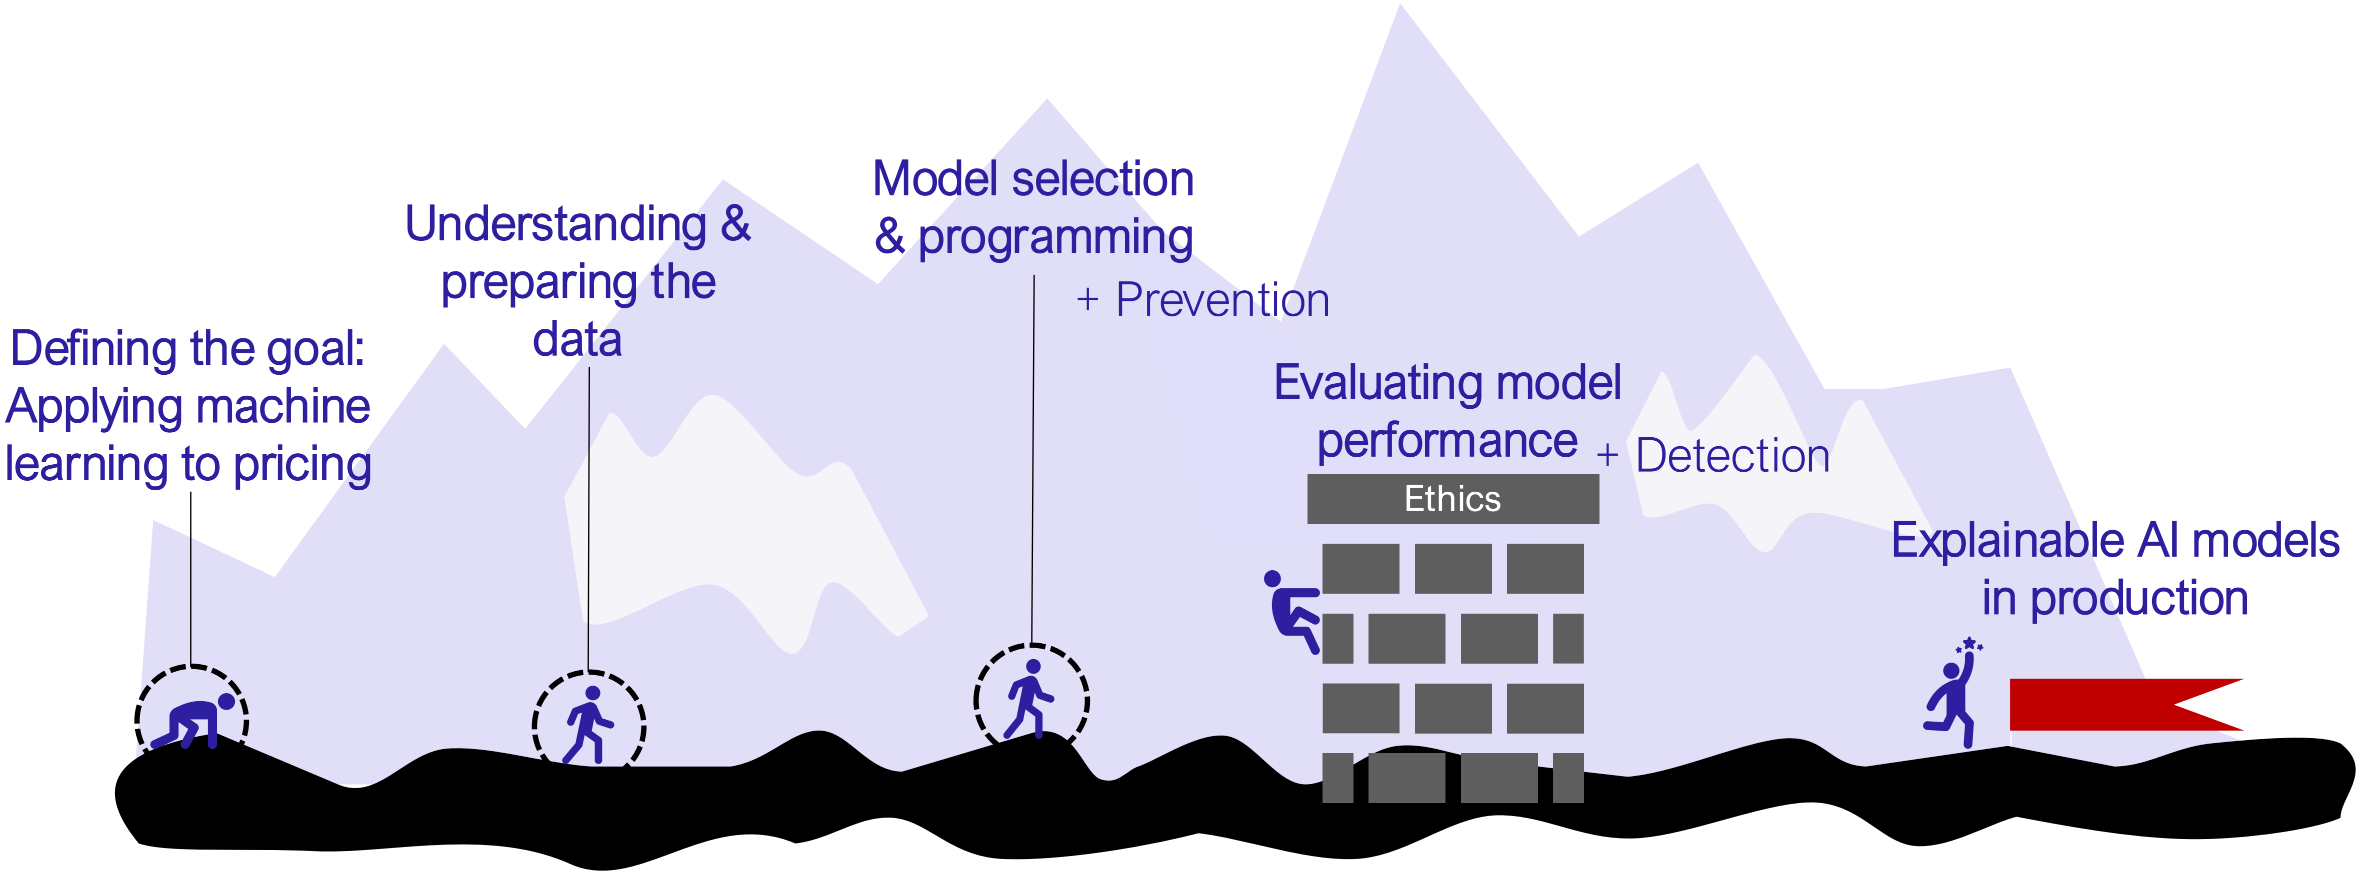

How do you prevent discrimination in ML models

What are preventative measures one can take?

- Model adjustments – See academic research
  - *Reading tip: Machine Learning with Multitype Protected Attributes: Intersectional Fairness through Regularisation (Oct 2025)*
- Reverse ML – Predict discriminative components with ML itself
- Causal AI - Beyond making predictions, causal AI helps explain "why" outcomes occur.

Contribution	Categorical treatment	Model architecture	# Data sets	Case study	Interpretation tools
Dugas et al. (2003)	–	LR, GLM, DT, NN, SVM	1	Tech. tariff	–
Yang et al. (2018)	–	TDBoost	1	Tweedie compound	PDP, VIP
Henckaerts et al. (2018)	–	GLM	1	Freq, sev	–
Wüthrich (2019)	Dummy encoding, embedding layers	GLM, NN, CANN	1	Freq	Avg. neuron activation
Schelldorfer and Wüthrich (2019)	Embedding layers	CANN	1	Freq	–
Ferrario et al. (2020)	One-hot encoding	Boosted trees, NN	1	Freq	–
Noll et al. (2020)	Dummy encoding, empirical means, one-hot encoding	GLM, DT, Boosted trees, NN	1	Freq	Loss-per-label
Henckaerts et al. (2021)	–	DT, RF, GBM	1	Freq, sev, tech. tariff	PDP, VIP, ICE
Kuo and Richman (2021)	One-hot encoding, embedding layers, attention layers	GLM, NN, Transformer, TabNET	1	Sev	–
Delong and Kozak (2023)	Autoencoder	NN	1	Freq	–
Meng et al. (2022)	Convolutional autoencoder	GLM	1	Freq	–
Henckaerts et al. (2022)	–	GBM	6	Freq	PDP, SHAP, Surrogates
This paper	Autoencoder	GLM, GBM, NN, CANN	4	Freq, sev, tech. tariff	PDP, VIP, surrogates, Shapley

How do you detect discrimination in ML models

Machine Learning models 'learn' rational effects from the data itself, making it difficult to determine exactly what the pricing is based on.

But it is possible to gain (partial) insight into the explanatory variables with these two statistics:

- Partial Dependence Plots: Plots of average impact
- Shapley Additive explanations: Impact per characteristic

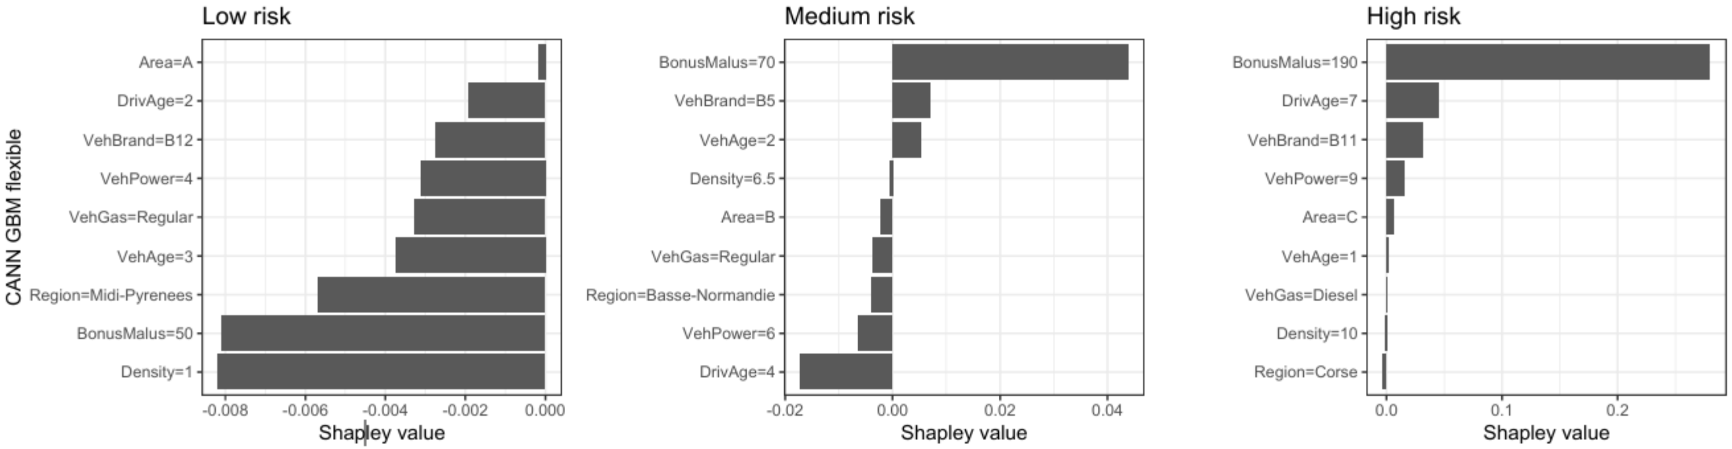

Note:

- *Using sensitive information to detect discrimination is permitted (AI act art. 10-5)*

Summary

Pricing, AI & Discrimination

	1. Prevention		2. Detection
Direct	Remove the variables		Regular checks
Indirect	Remove the variables	Correcting estimations	Partial Dep. Plots
	Constraints	Hybrid Models	Shapley Additive explanations
	Reweighting dataset	Causal AI	

3. Explainability / Auditability

- Proof of checks in step 1 -2
- Model validations
- Test & controls

Discussion questions

- Do current GLM completely eliminate the possibility of bias/discrimination?
- Is anyone using ML to enhance pricing as a main model? Why?
- Other questions?

2. The GenAI trainingsloop

The definition

GenAI – Generative AI is a type of artificial intelligence that can create new content such as text, images, audio, code, or video by learning patterns from existing data.

What is happening?

GenAI models are trained on data scraped from the internet

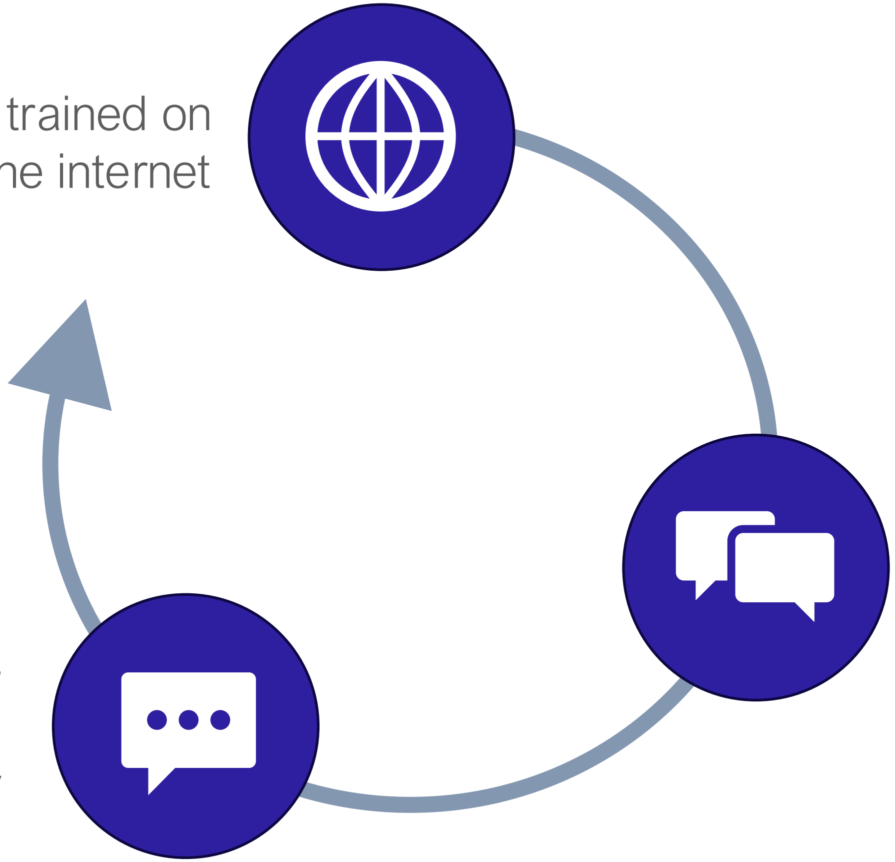

Users interact with GenAI models and correct them when necessary

Feedback is integrated, or AI-generated content is released publicly

What could happen?

You can enter a loop driven by ”fake” data

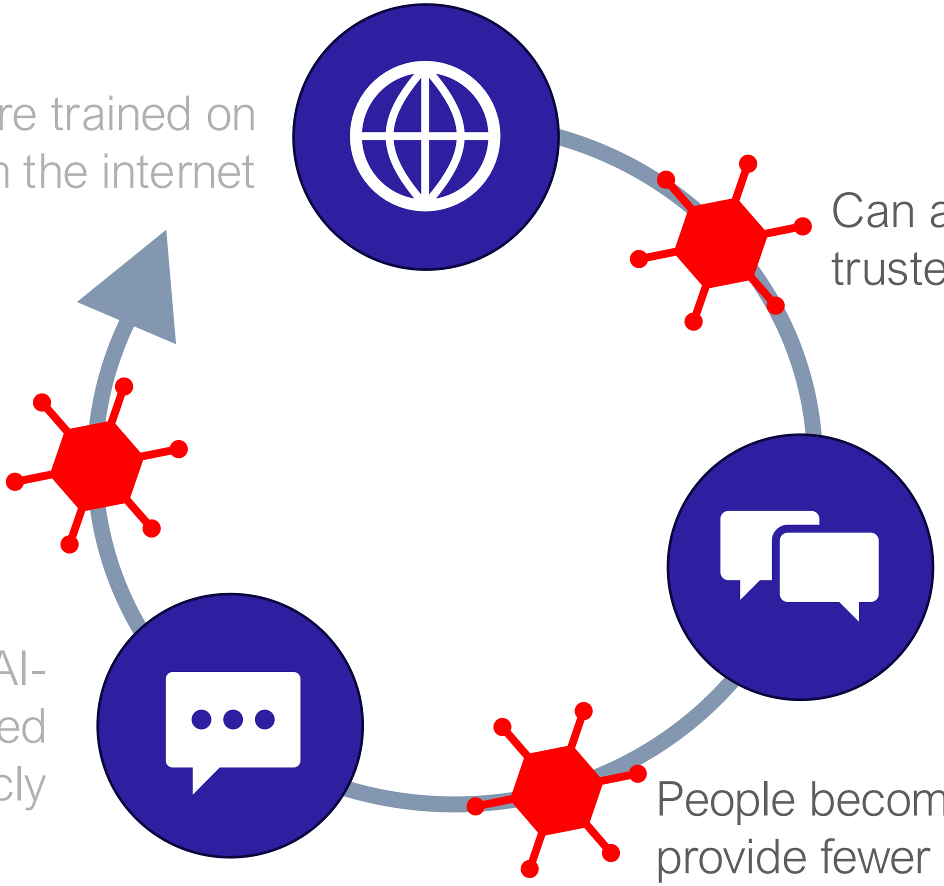

GenAI models are trained on data scraped from the internet

Can all internet sources be trusted?

Users interact with GenAI models and correct them when necessary

People become less skeptical and provide fewer corrections

Feedback is integrated, or AI-generated content is released publicly

Published information has lower quality - quantity over quality?

Discussion question

- Should AI-generated data be used to train other LLM models?
- How could we separate AI-generated data from ‘real’?
- At what point does synthetic data start to undermine model reliability?

3. Privacy issues of external algorithms

Opportunities - External LLMs

More insights, faster insights and even independent decision making (AgenticAI)

Complexities - Privacy Sensitivity

Key complexity comes down to Knowledge vs. Privacy

Data Leakage:

- Using Public LLM's can lead to accidental leaks to competitors and malicious users.
- Data leakage of e.g. salary data amongst employees

Shadow AI:

- Employees using unauthorized AI tools to handle company data create blind spots for security teams, bypassing corporate firewalls.

Strict Enforcements:

- Regulations like the EU AI Act and GDPR mandate requires "Privacy by Design" for all generative systems.
- If a company accidentally leaks customer data through AI, they face large fines.

The solution - Building “In-House” AI?

- Private AI models
  - Mitigates leakage or Privacy risk

- Reading vs. Learning (RAG)
  - AI learns and answer question based on private and reliable documentation you provide
  - No learning applications of the algorithm.

- Specialized tool:
  - Instead of one large general AI, companies are developing specialized AI tailored to in their specific business

Advantages

- Data Control
- Customization (start small)
- Compliance
- Future-proof due to flexibility

Disadvantages

- Higher setup cost
- Maintance (you are responsible for debugging and updating)
- Knowledge

Discussion question

- Is the productivity gain of public LLMs worth the potential IP and privacy risks?
- Are public LLMs like ChatGPT/CoPilot trustworthy for private data?
- What minimum controls should be in place before employees can use public LLMs?

Hopefully, this provides new insights and ideas for applying AI!

Contact us if you are looking for more information.

Appendix; Pricing & AI in Academic Research

Paper: Neural networks for insurance pricing with frequency and severity data (aug. 2024)

CANN (combined actuarial neural network)

For example: GLM + tree-based model & neural network adjustments

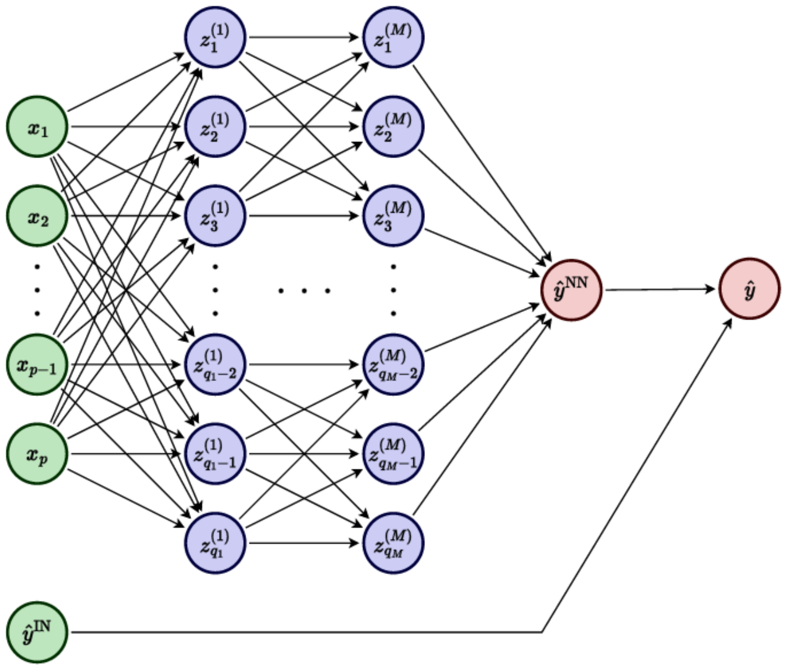

GLM as a global surrogate for a deep learning model

"Black box"-algorithm summarized in a table

GBM model as black box

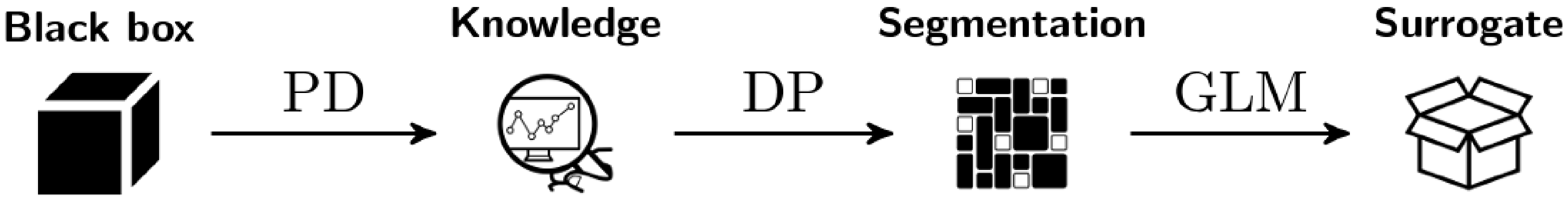

Appendix; Or: Could causal AI be the solution?

What is Causal AI?

- Designed to understand and quantify causal relationships between variables.
- Beyond making predictions, causal AI helps explain “why” outcomes occur.

Methods / Functions under the causal AI umbrella:

- ***Causal Bayesian Networks*** *(CBNs): Causal structure analysis.*
- ***What if scenario analysis:*** *ORSA or other risk analyses*
- ***Causal Forests:*** *A random forest aimed to minimize causal effects.*

Advantages	Disadvantages
Fairer Pricing Improved accuracy Greater pricing transparency	Increased model complexity (requiring sufficient knowledge) Higher variance than ”regular” ML Greater data required