# The New Actuarial Toolbox

From GAMs to Gaussian processes


3 March 2026

# Introduction

**Martin Tan**

- 32 years old, lives in Leiden
- Started as an offshore engineer, pivoted to actuarial consulting and currently working as a pricing actuary at Nationale-Nederlanden.
- Powerlifting, video games, robotics, card games (e.g. Yu-Gi-Oh!)

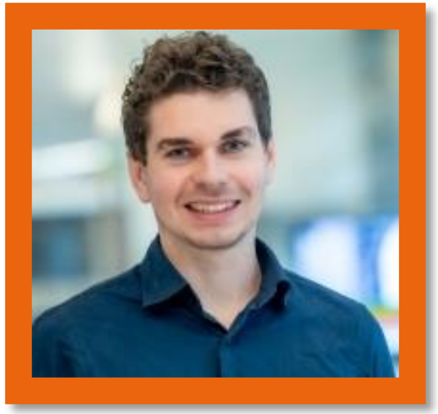

**Rob Schuitemaker**

- 31 years old, based in Amersfoort
- Nationale-Nederlanden
    - 2017-2021: Pricing Actuary at ABN AMRO Insurance
    - 2022-2026: Senior Pricing Actuary at Retail Non-life
- Running, volleyball

# Who we are

## Company profile

Founded in 1845, NN Group is a financial services company, active in Europe and Japan. With all its employees, the Group provides retirement services, pensions, insurance, banking and investments to approximately 19 million customers.
NN Group is listed on Euronext Amsterdam (NN).

## Our main brands

nationale nederlanden
azl
OHRA
movir
BE FRANK
woonnu

Our presence

## We operate in 10 countries

We provide retirement services, pensions, insurance, banking and investments to approximately 19 million customers.

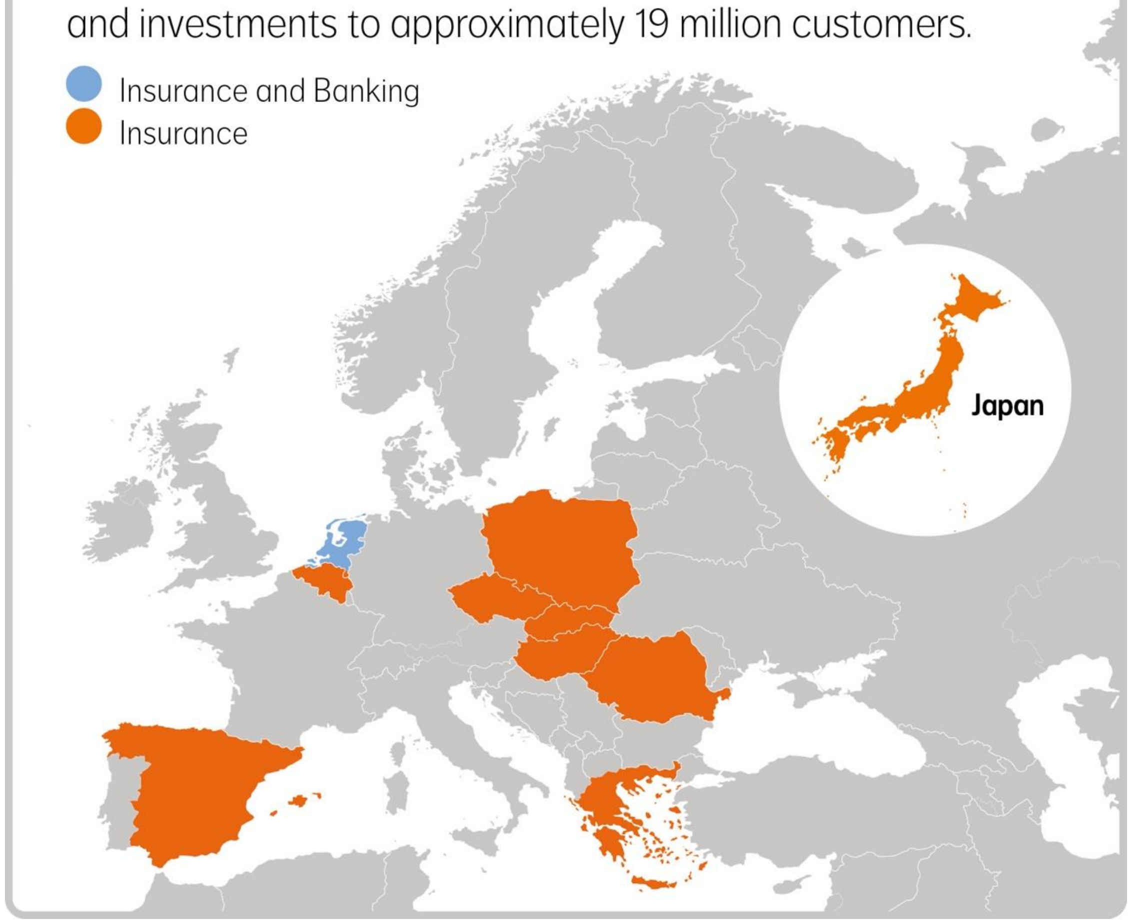

# Agenda

- NN & AI
- Trends in non-life pricing
- Pitfalls of spatial data
- Introduction to Gaussian processes
- Example: Enhancing weather data using Gaussian processes
- Wrap-up

# Our strategic framework

Our purpose

## We help people care for what matters most to them

Our ambition

**We want to be an industry leader, known for our customer engagement, talented people, and contribution to society**

Our values

- Care
- Clear
- Commit

Our brand promise

## You matter

Our strategic commitments

**Engaged customers**
We deliver an outstanding customer and distributor experience, and develop and provide attractive products and services.

**Talented people**
We foster a values-based culture and empower our colleagues to be their best.

**Contribution to society**
We contribute to the well-being of people and the planet.

**Financial strength**
We are financially strong and seek solid long-term returns for shareholders.

**Digital & data-driven organisation**
We use technology and data responsibly to transform our business and drive operational excellence.

# NN & AI

## Rapid AI developments

AI is developing at record speed, it’s not slowing down.

## Customer

Customers expect service that is easy, fast, and cost-efficient which can be realized by leveraging AI.

## Leader in AI

NN wants to lead the industry in applying AI to drive speed, precision, and customer impact necessitates NN to act now.

# Trends in non-life pricing

Pricing on peril level

Use of more flexibility in models

# Trends in non-life pricing

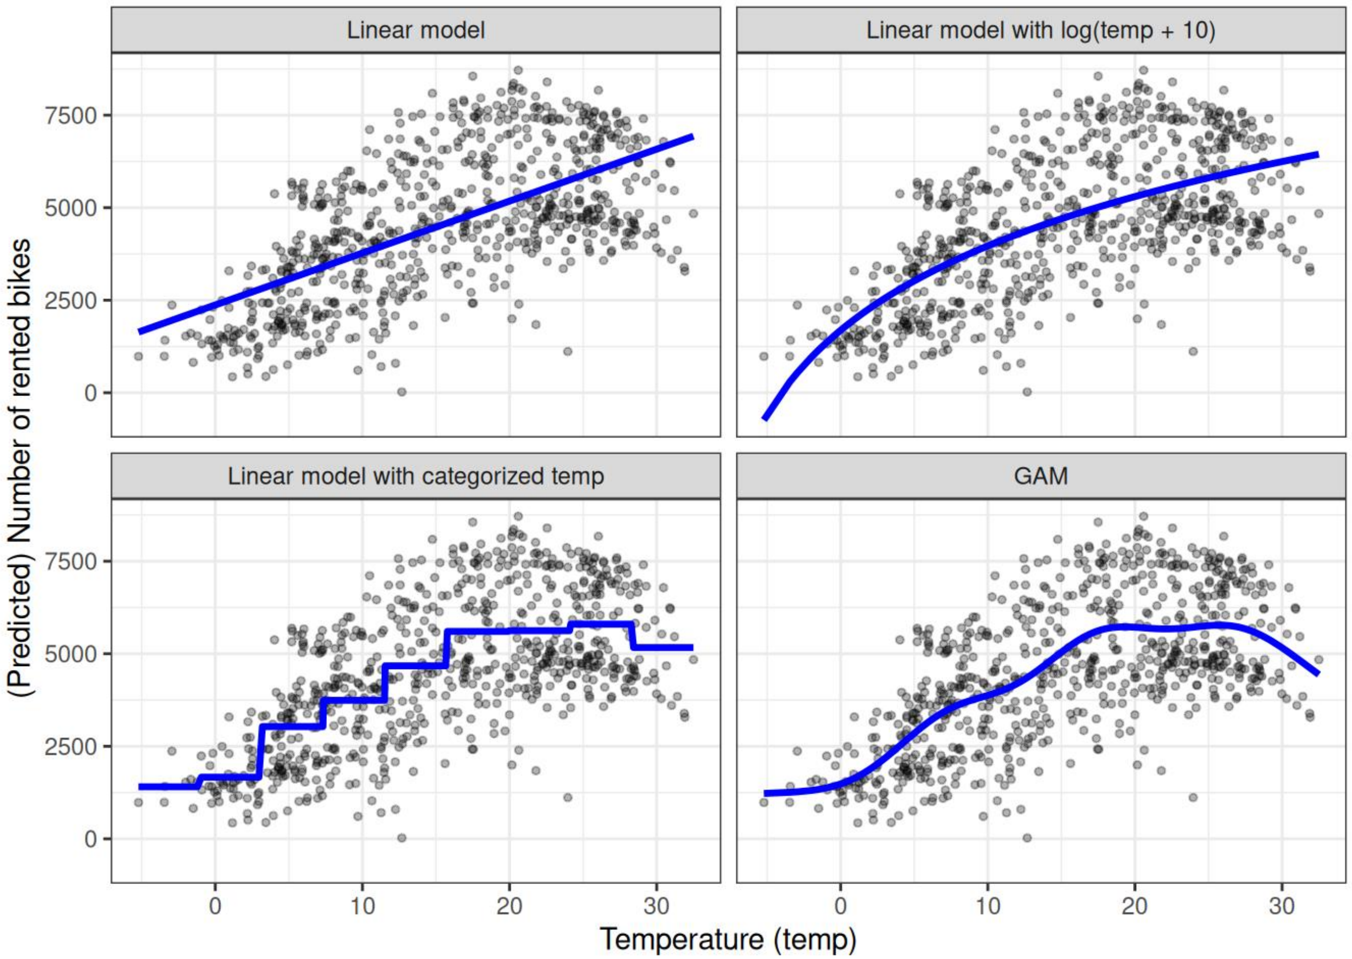


Source: 8 GLM, GAM and more – Interpretable Machine Learning

NN

# Trends in non-life pricing

- Pricing on peril level
- Use of more flexibility in models
- Fewer questions in sales funnel
- Use of more (external) data
- More granularity

# Pitfalls of spatial data

## Stylized example

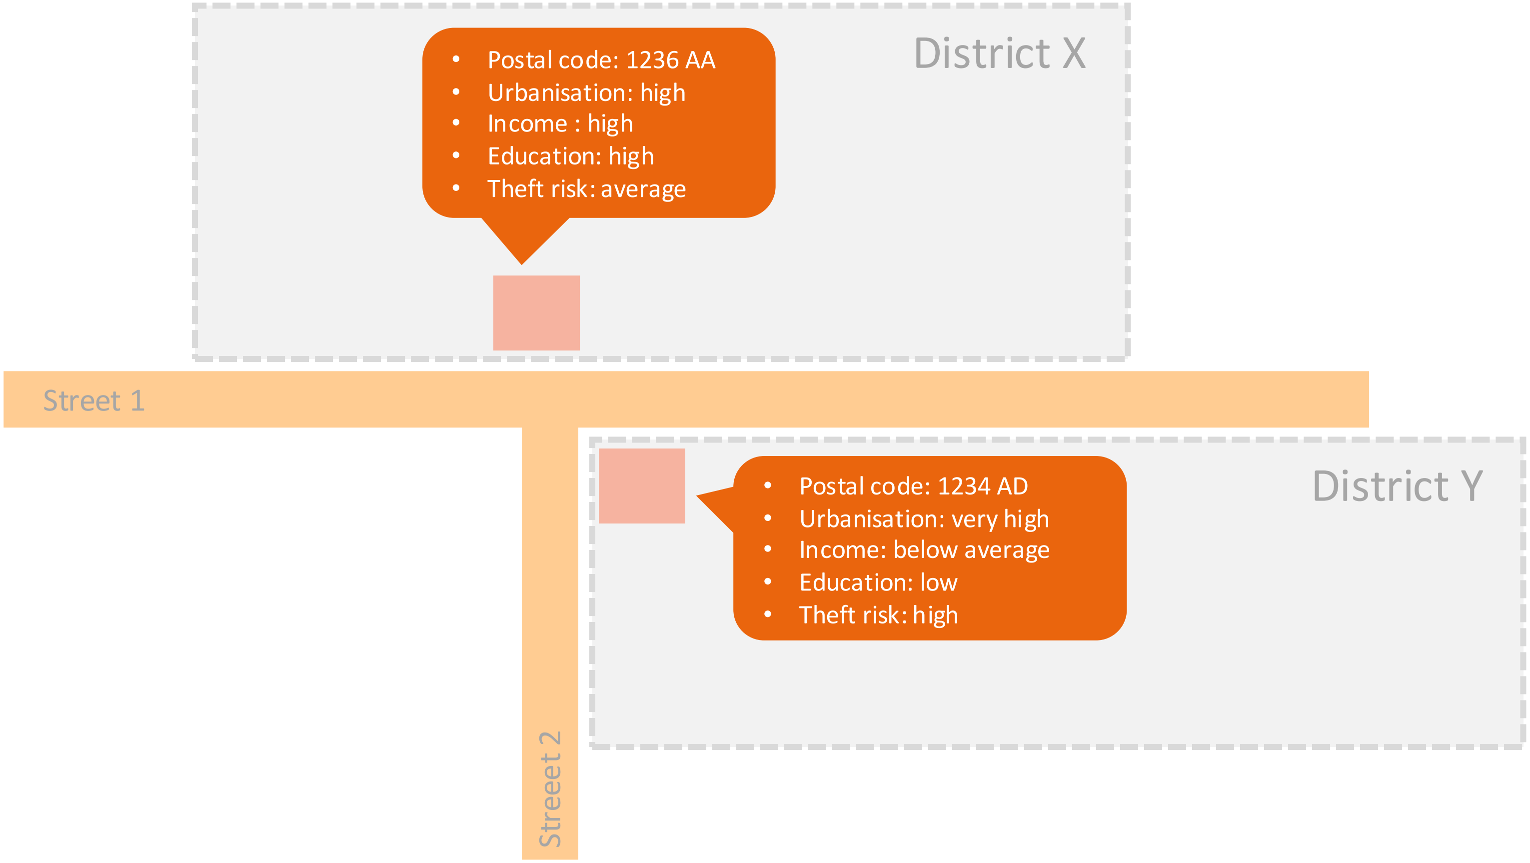


**Strengths**

- Many regional variables available via external sources without the necessity to ask the customer
- Spatial data has very high explanatory power on aggregated level

**Weaknesses**

- Region variables are often correlated
- Many sources do not have factual data for every address. Interview results are extrapolated using models
- Unlogical outcomes on individual level

# Pitfalls of spatial data

## Not fictional

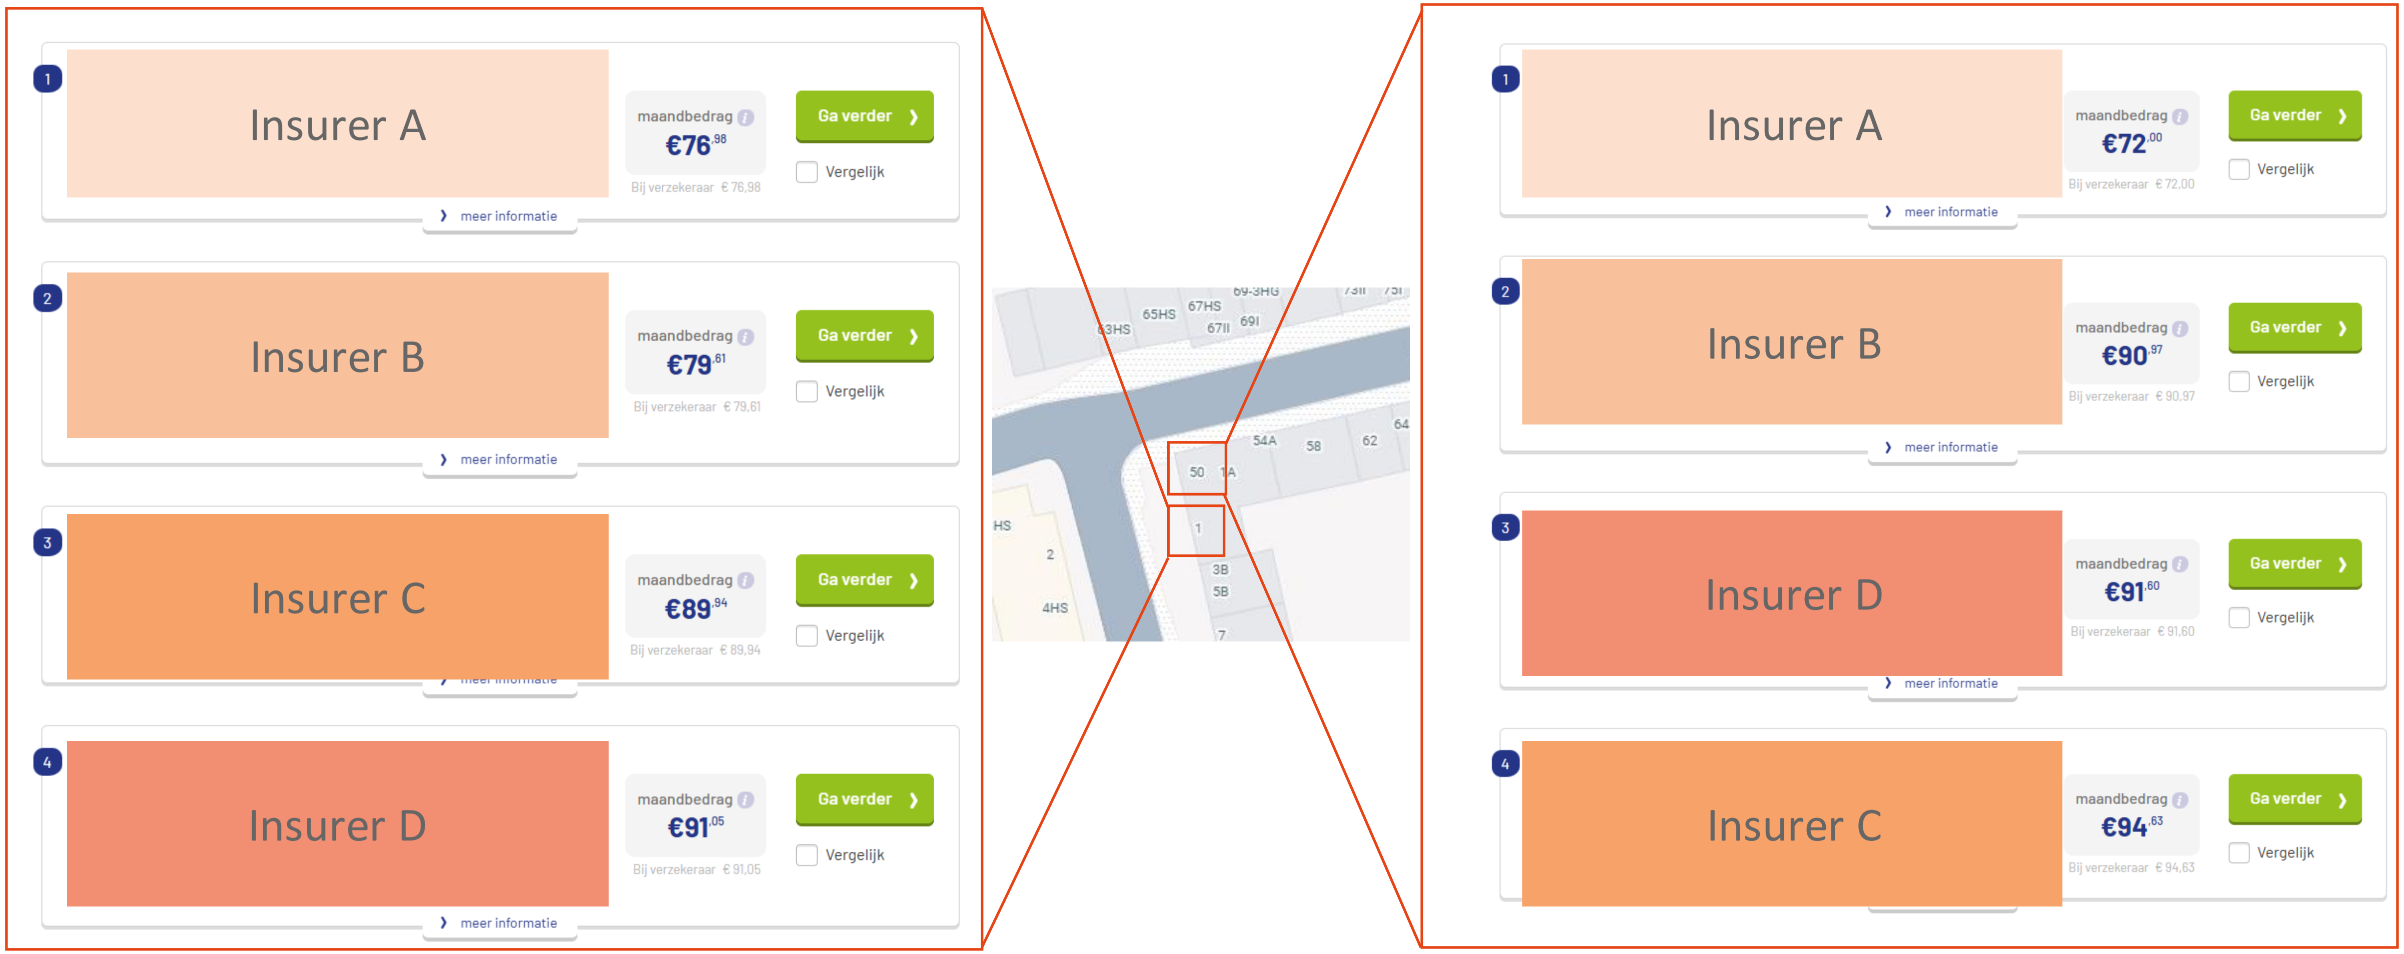


*Source: Pricewise, 27-01-2026*

nationale nederlanden

# Industry implications

## Impact on consumer confidence?

### Premie autoverzekering binnen dezelfde wijk tot honderden euro's duurder

Auteur Amanda Bulthuis | Bijgewerkt op 21 augustus 2023

Als je binnen je eigen woonwijk of straat verhuist, kan de premie van je autoverzekering flink stijgen. Dat blijkt uit onderzoek van de financiële vergelijkingssite Geld.nl. "Gemiddeld steeg de autopremie in de berekeningen die wij maakten bijna 22 euro per jaar, maar we zagen ook uitschieters waarbij de premie na de verhuizing ruim 380 euro per jaar hoger was. En dat in dezelfde woonwijk! Zelfs als je buurman precies even oud is als jij en precies dezelfde auto rijdt, dan nog betaal je dus misschien veel meer voor je autoverzekering dan je buurman", vertelt Amanda Bulthuis, expert geld & verzekeringen bij Geld.nl.

Autopremie in zelfde woonwijk soms fors hoger

| | Maandpremie adres A | Maandpremie adres B |
|---|---|---|
| Vogelwijk, Heerenveen | € 64,75 | € 81,83 |
| Oude Gracht-West, Eindhoven | € 71,20 | € 88,46 |
| Zevenkamp, Rotterdam | € 148,27 | € 180,08 |
| Gemiddeld verschil per jaar | € 262,40 | |

Geld nl

Amand
OUD

Amanda was van 2
productspecialist.
oorspronkelijk van
passen we het arti
laatste ontwikkeli
informatie in dit ar

Autoverzeker
Start ve

Grafiek 5: Ontvangen en behandelde klachten bij Kifid 2015-2024 (Kifid, 2024)

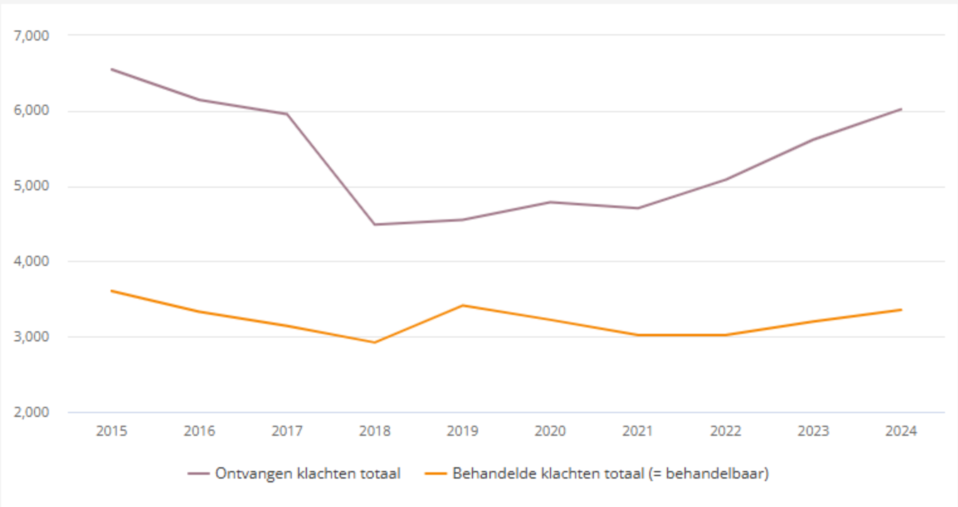


### Verschil premie autoverzekering op basis van huisnummer in dezelfde straat

Gepubliceerd op 01-11-2024

Het is een bekend gegeven dat de woonplaats invloed heeft op de premie van de autoverzekering. Dat geldt eveneens voor verschillen per postcodegebied in een plaats, maar er zijn zelfs verschillen in premies in dezelfde straat. Dat kan betekenen dat de ene autobezitter, die aan het begin van de straat woont een hogere premie betaalt dan de verzekerde aan het einde van de straat. Het huisnummer van het adres bepaalt dus mede de hoogte van de premie van de autoverzekering.


#### Premieverschil huisnummers vooral in grote steden

Het verschil in premie van de autoverzekering op basis van huisnummers is vooral merkbaar in grote steden met lange straten. Het verschil op basis van dezelfde gegevens kan tussen verre buren uit dezelfde straat wel oplopen tot meer dan €100 per jaar. Bijvoorbeeld als het om een WA verzekering gaat.

*Sources*:

- Premie autoverzekering binnen dezelfde wijk tot honderden euro's duurder | Geld.nl
- Verschil premie autoverzekering op basis van huisnummer in dezelfde straat
- Kifid - Jaarverslag 2024

nationale nederlanden

# Wooclap Question

Large premium differences between neighbours should be avoided, even if they are (technically) explainable based on external data sources.

1. Yes, because there are many concerns regarding the (external) data causing these differences.
2. Yes, because I do not believe the risk differs that much between neighbours.
3. No, if the data demonstrates this and is reliable, then it provides a better risk assessment.

# Patchwork

## Can we have a smoother pattern?

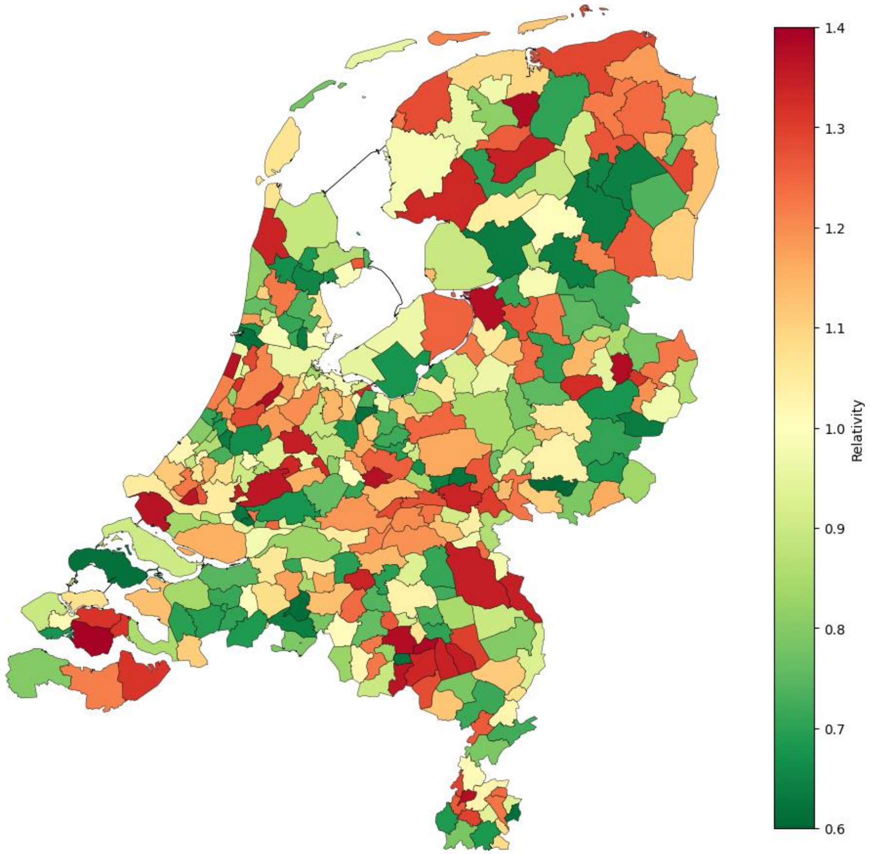


nationale nederlanden

# Gaussian Process

## A distribution over functions

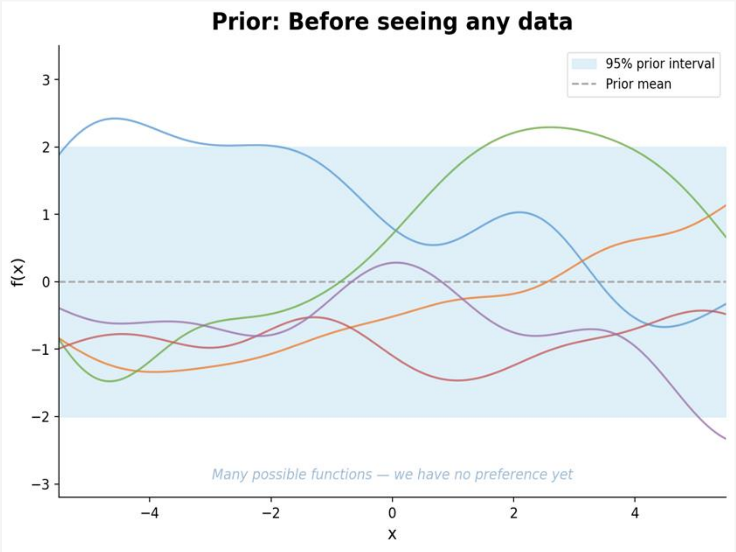


**1** **Before data:** many possible smooth functions

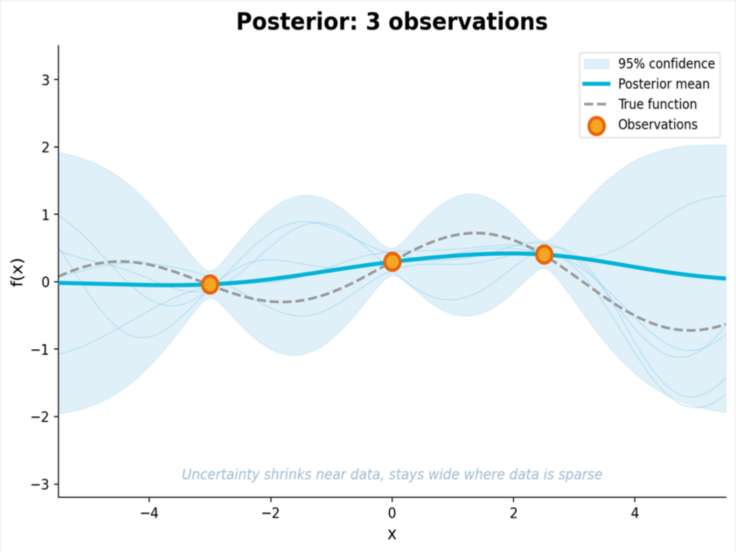


**2** **3 observations:** mean updates, bands narrow locally

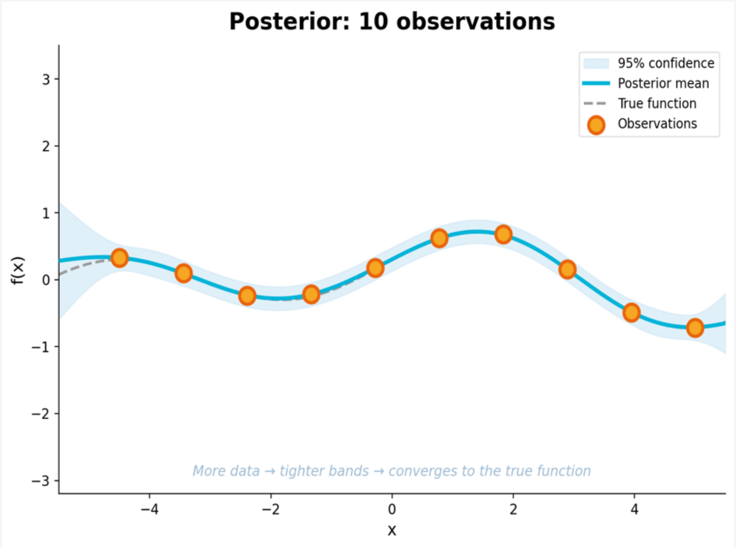


**3** **10 observations:** tight bands, recovers true function

**Key insight:** A GP starts with a prior belief (smooth functions), then updates as data arrives. Uncertainty shrinks near observations and stays wide where data is sparse, **exactly like credibility theory.**

# Gaussian Process

## Three hyperparameters, each interpretable

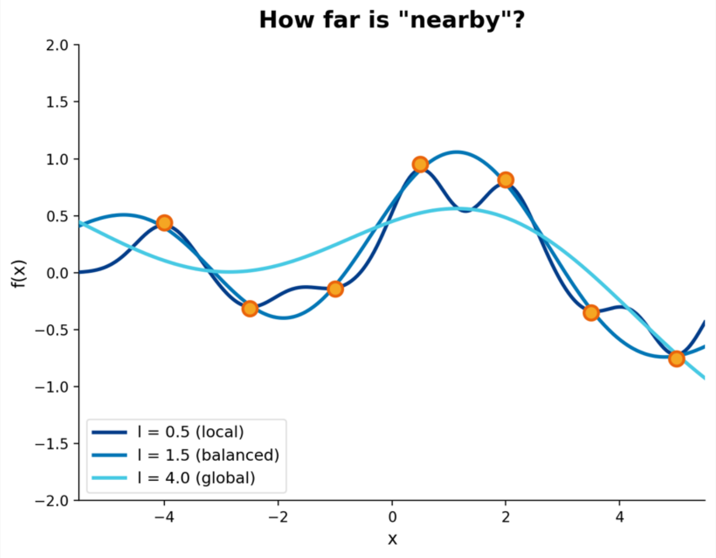


**$l$ Lengthscale**

How far = “nearby” (learned from data)

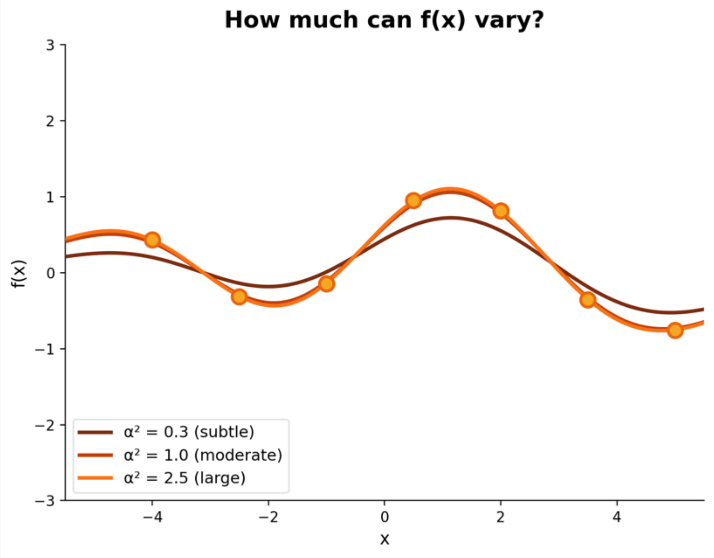


**$\alpha^2$ Output scale**

Amplitude of spatial differences

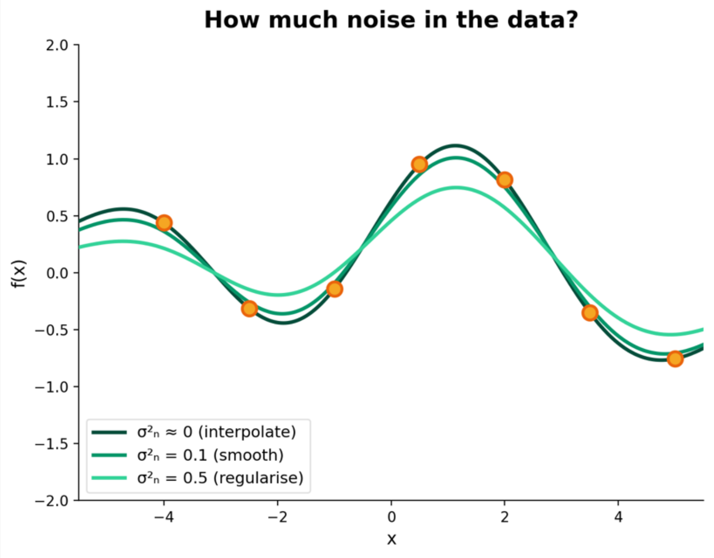


**$\sigma_n^2$ Noise (nugget)**

Observation error / micro-variation

**Key insight:** All three are learned from data via marginal likelihood no manual tuning needed.

# Gaussian Process

## The credibility interpretation

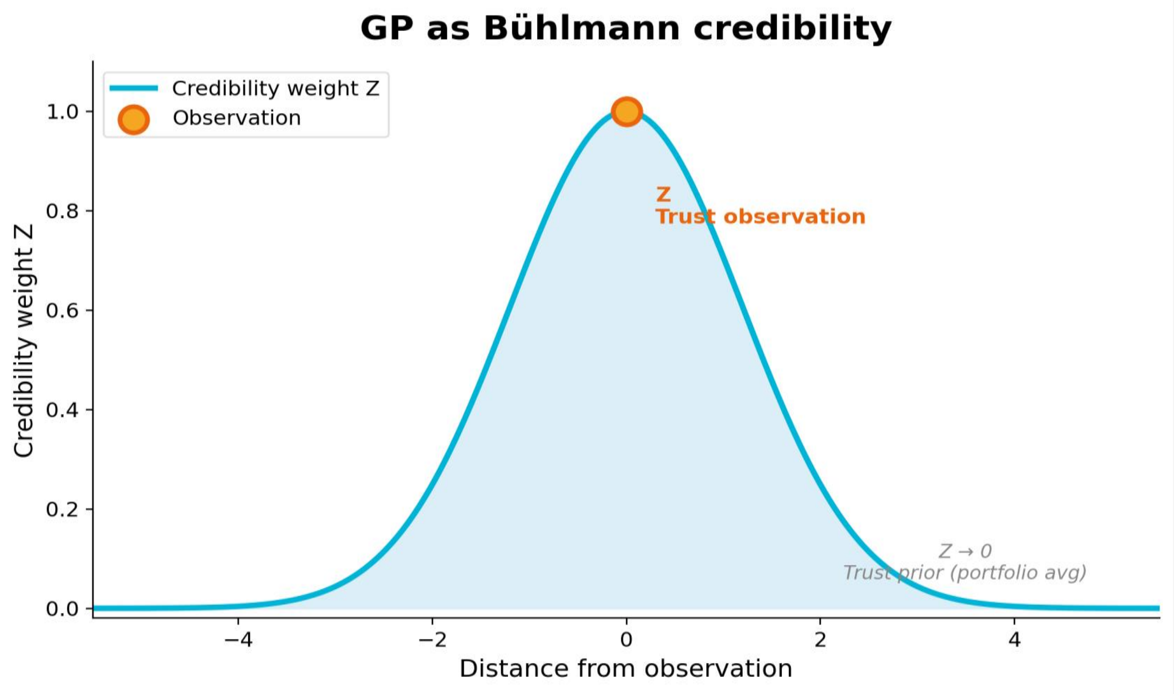


### The Parallel

$$\hat{\mu} = Z \cdot \bar{X}_i + (1 - Z) \cdot \mu$$

**$Z$ (credibility)** ↔ Kernel correlation to neighbors

**$\bar{X}_i$ (individual)** ↔ Nearby station observations

**$\mu$ (collective)** ↔ Prior mean (portfolio average)

### ✓ Why This Matters for Pricing

→ Sparse data zones get shrunk toward the spatial average

→ Data-dense zones keep their local estimate

→ Posterior σ gives confidence bands on every territorial factor

→ The lengthscale tells you the range of spatial correlation

*GPs generalize credibility to any number of dimensions and any spatial structure — it's what Bühlmann would have built if he had GPUs.*

# Gaussian Process

## The (historic) computational cost

**THE PITFALL**

$$O(n^3)$$

Inverting the $n \times n$ covariance matrix is the computational bottleneck.

For $n = 10{,}000$ observations: ~1 trillion floating-point operations

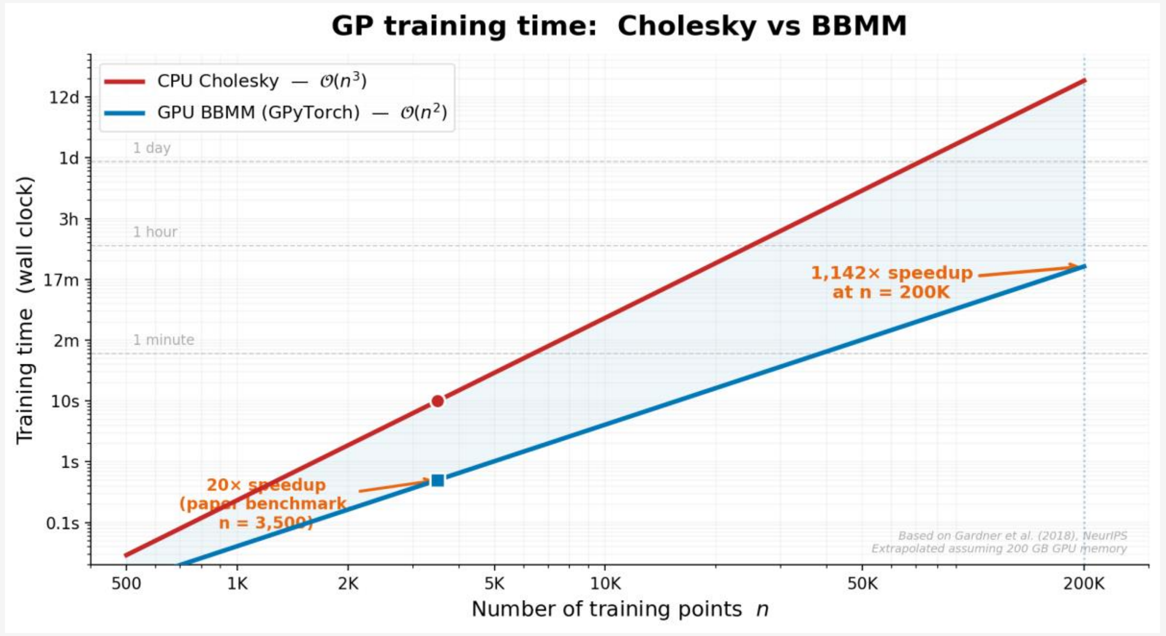


Scalability research has brought the cost down: **inducing-point methods** (FITC, variational) reduce to $O(nm^2)$, **structured kernel interpolation** (KISS-GP) exploits grid structure, and **GPU-native libraries** (GPyTorch) parallelize linear algebra on CUDA cores.

**Key insight:** With modern GPUs and scalable approximations, GPs could now handle the dataset sizes typical P&C pricing.

# Case study – rainfall interpolation

## THE DATA

37 weather stations (hourly) + ~300 volunteer rain stations (daily totals). Rain stations only have place-level coordinates.

## THE CLASSIFICATION

**No rain** <2mm/day **Light rain** 2–10mm/day **Heavy rain** >10mm/day

## THE CHALLENGE

**Predict the rainfall class label for any historic date and location in the Netherlands**

*Use both hourly and daily datasets; a classic multi-fidelity spatial interpolation problem.*

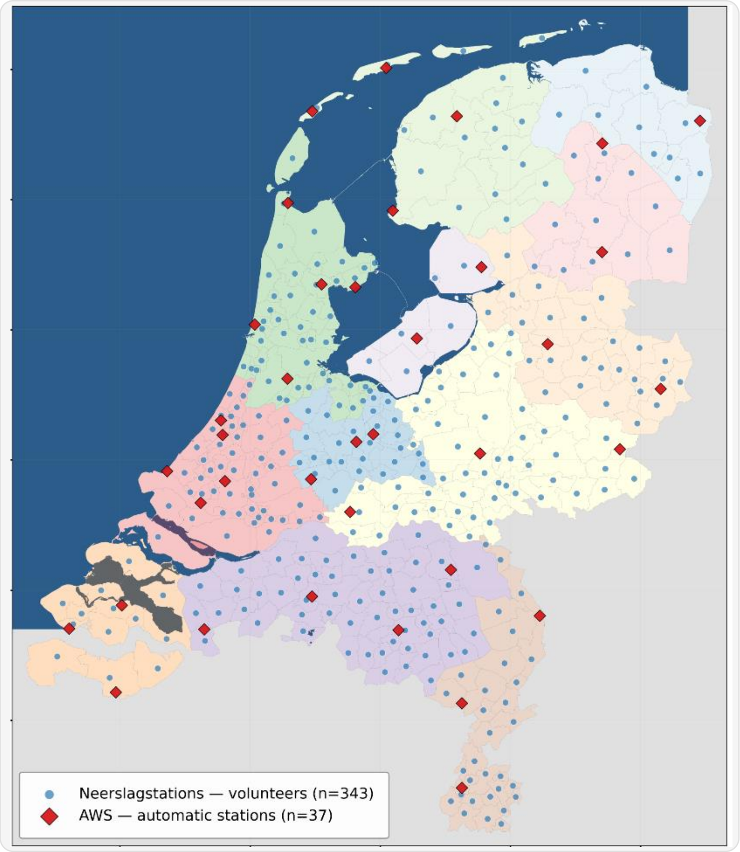


**Key insight:** Can we predict rainfall class for any location and date using only station data? **This is the GP's job.**

# Rainfall model

Leave-one-station-out cross-validation

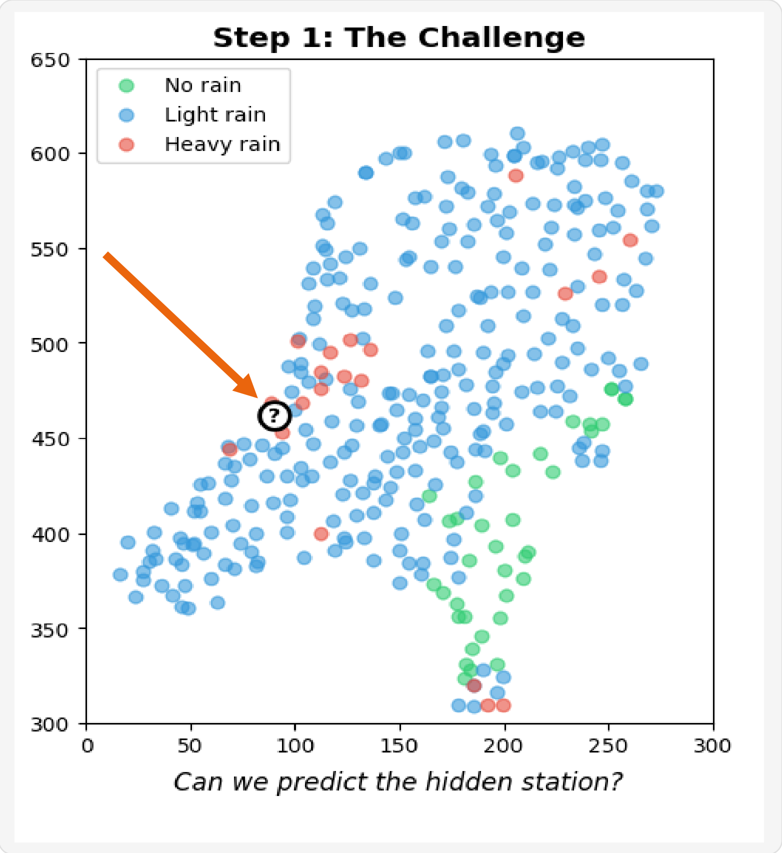


**Key insight:** Given only the surrounding stations, can the GP predict the rainfall class at the hidden location?

# Wooclap Question

What classification do you predict for this weather station?

1. No rain (green)
2. Light rain (blue)
3. Heavy rain (red)

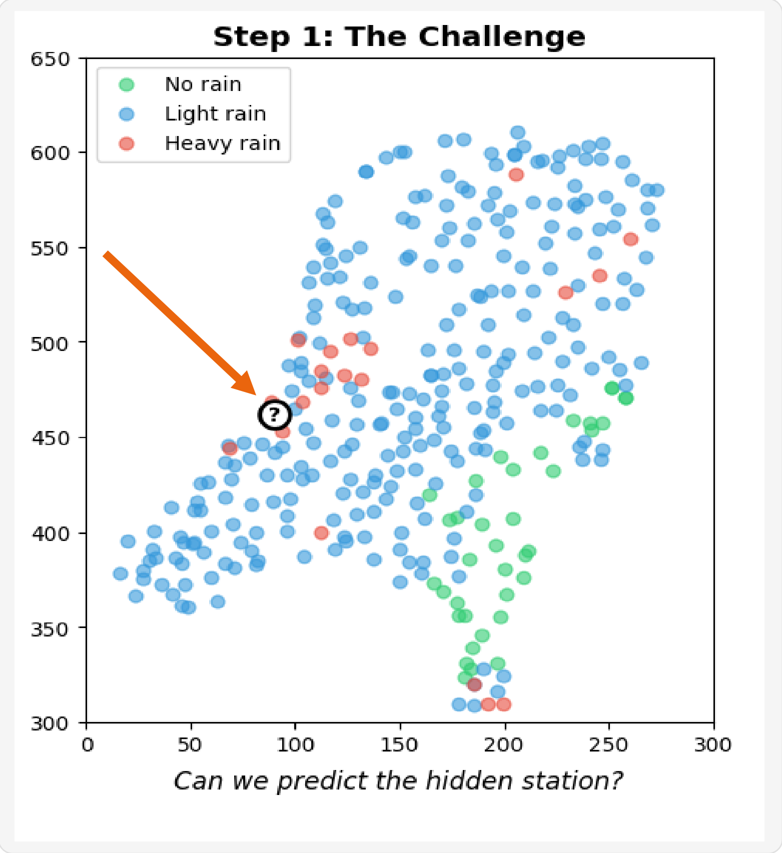


nationale nederlanden

# Rainfall model

Leave-one-station-out cross-validation

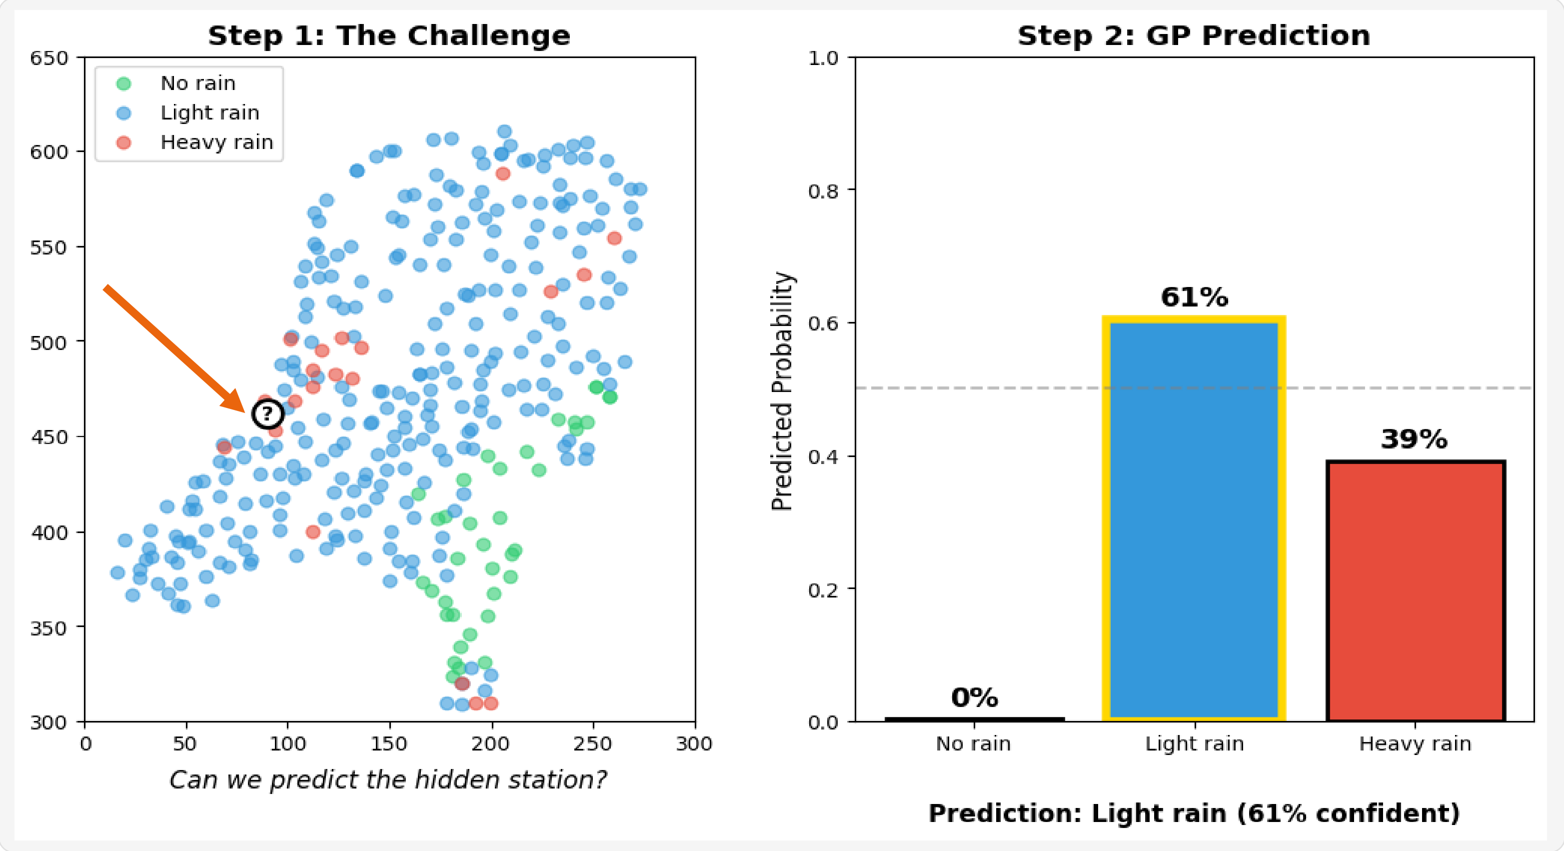


**Key insight:** The GP assigns class probabilities. It doesn't just guess, **it quantifies how confident it is.**

nationale nederlanden

# Rainfall model

Leave-one-station-out cross-validation

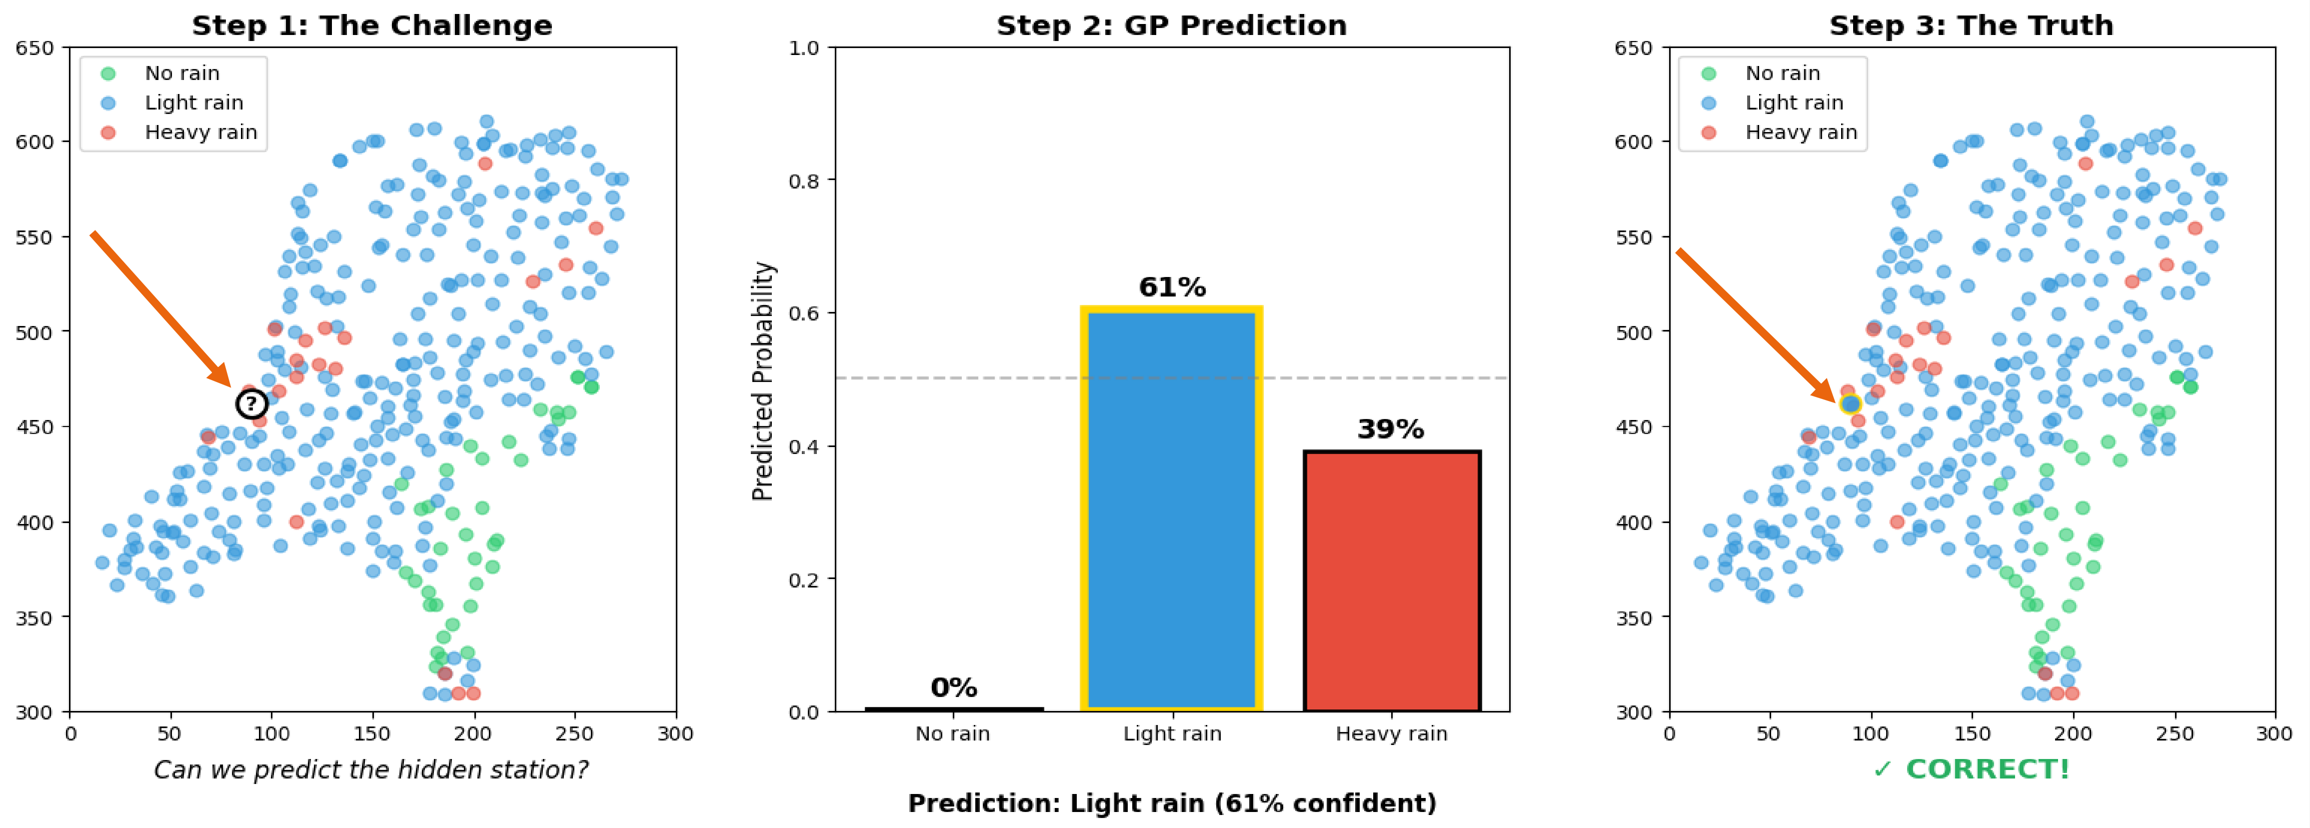


**Key insight:** The truth revealed: the GP got it right with calibrated confidence.

nationale nederlanden

# Rainfall model classification

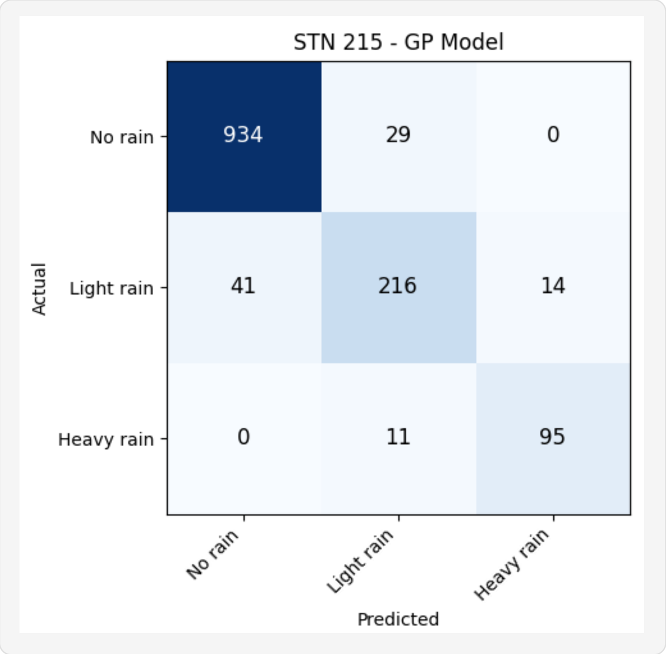


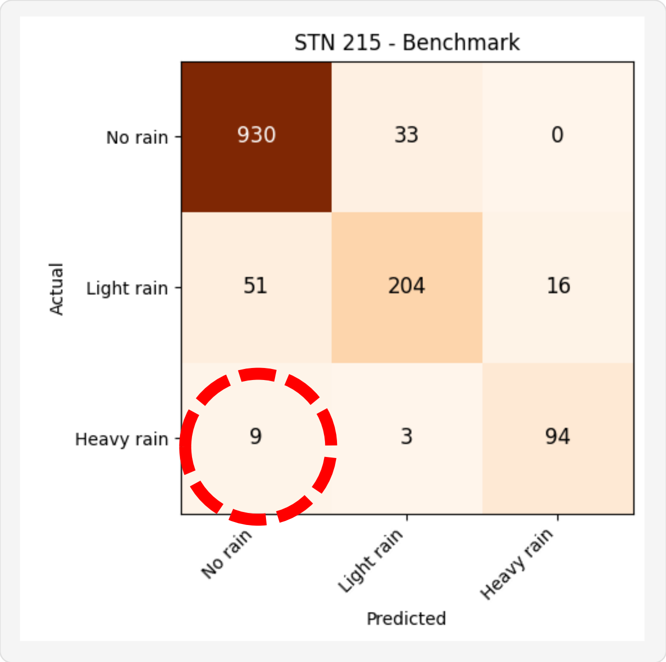


**92.9% overall accuracy** vs 91.6% for the 3-nearest-neighbour benchmark. Higher accuracy *and* safer predictions.

**0 dangerous misclassifications** — the GP does not predict"no rain" when heavy rain occurred. The benchmark makes this error 9 times.

**Key insight:** Accuracy alone does not tell the full story; **the GP eliminates the most dangerous errors.**

nationale nederlanden

# Rainfall model – from weather to pricing

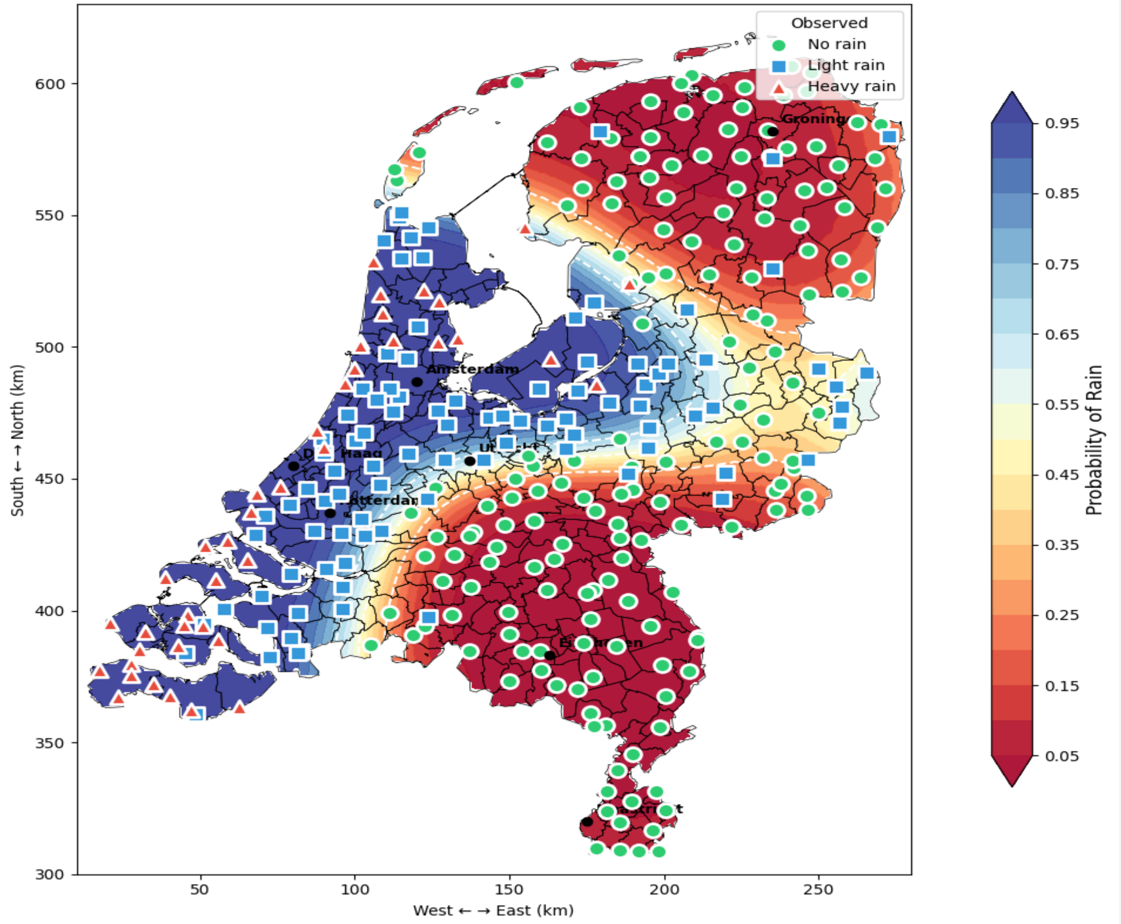


## FROM RAINFALL TO PRICING

| | |
|---|---|
| Rain stations → | Zipcodes |
| Rainfall class → | Premiums |
| Sparse observations → | Sparse claims data |

***Same sparse-data problem.***
***Same spatial solution.***

The smooth probability surface on the left is exactly how territorial pricing factors should behave: neighbouring areas get similar rates unless the data says otherwise.

**Key insight:** The GP probability surface smoothly interpolates between stations; **exactly how territorial pricing factors should behave.**

# Rainfall model – my model knows what it doesn't know

CALIBRATION

**92.9% accurate; and it knows when it is uncertain.**

A perfectly calibrated model has every bar on the diagonal.

**Green bars** on the diagonal → model confidence matches reality. When it says 80% confident, it's right ~80% of the time.

**Red bars** below the diagonal → low-confidence predictions. The model correctly flags these as uncertain rather than guessing.

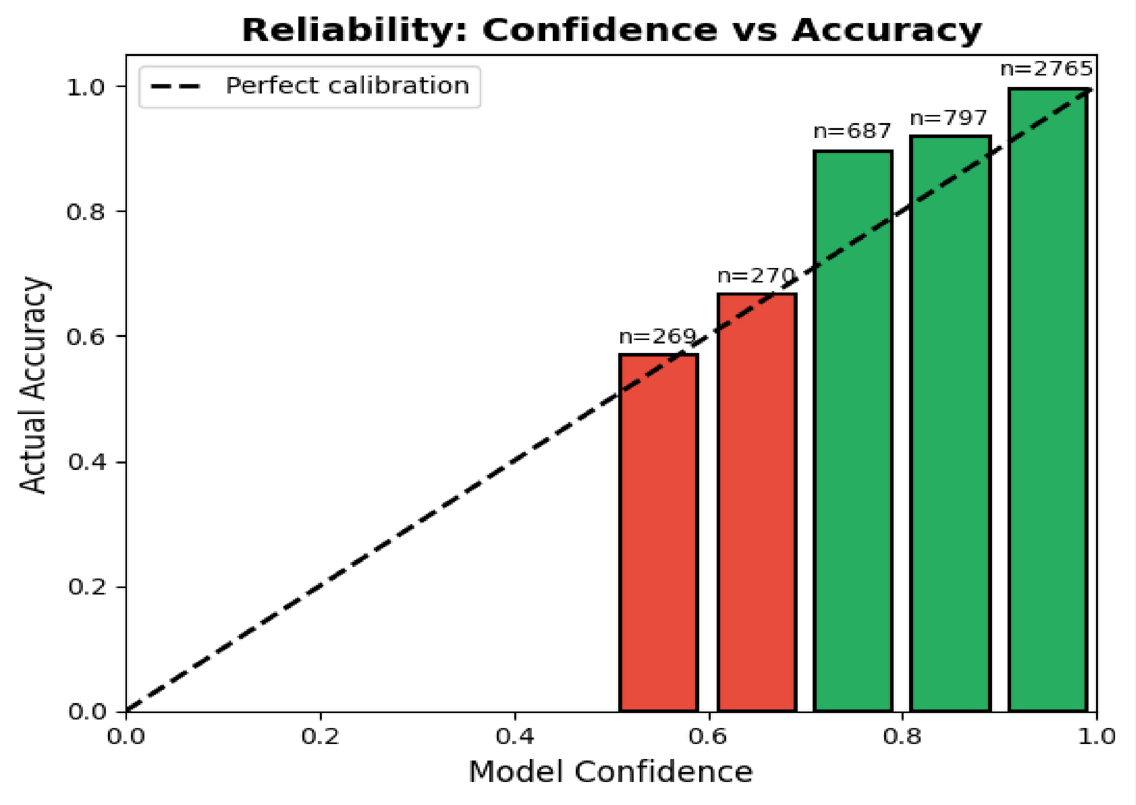


**Key insight:** A well-calibrated model gives actuaries the confidence to use its predictions; **the model knows what it doesn't know.**

# Conclusion

### THE CHALLENGE

Pricing demands granularity, but granularity creates sparse data and non-credible estimates.

### THE GP SOLUTION

GPs interpolate spatially, shrink sparse zones toward the mean, and quantify uncertainty on every prediction.

### NOT A BLACK BOX

The credibility interpretation means every actuary can understand and trust what the GP is doing.

### CASE STUDY

GP spatial interpolation on rainfall classification across the Netherlands; 92.9% accuracy.

### WELL-CALIBRATED

The model knows what it doesn’t know; it eliminates the most dangerous errors and gives actuaries confidence to use its predictions.

*Rainfall was our case study, but any spatial data applies.*

**What spatial problem in your portfolio would benefit from this?**

nationale
nederlanden