

# Probabilistic Geospatial Machine Learning as a Blueprint for Trustworthy AI in High-Stakes Domains

Andres Masegosa

Dept. of Computer Science, Aalborg University – Copenhagen

TrAI Workshop 1 — Pioneer Centre for AI — 3 March 2026

# Trustworthy AI: Beyond Fairness and Explainability

## The standard trustworthiness dimensions

Explainability, fairness, robustness, privacy, accountability — are all essential.

But in **high-stakes, decision-critical** domains, another dimension is often missing:

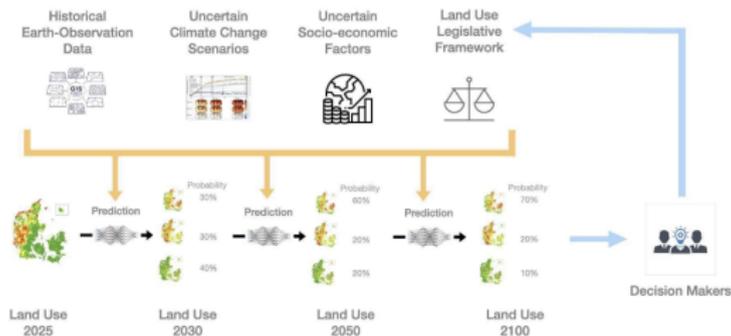
## What high-stakes domains additionally require

- **Calibrated uncertainty** — the system must communicate the full range of plausible outcomes and their probabilities.
- **Decision-relevant robustness** — reliable under the compound, non-stationary conditions where decisions are actually made.

A model can be explainable and fair, yet **dangerously overconfident** in a changing world.

# DK-Future: Land-Use Forecasting Under Compound Climate Impacts (Villum Foundation)

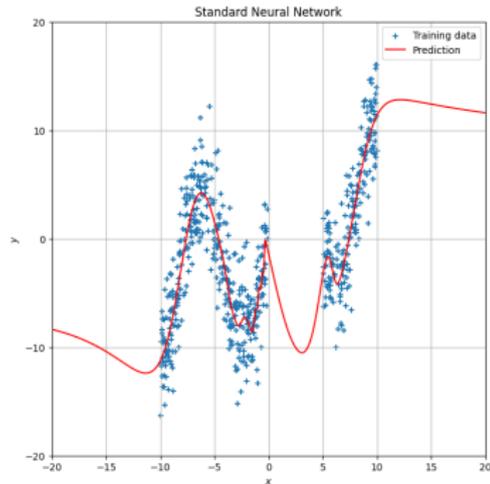
- Denmark's low-lying terrain and long coastline make land use highly sensitive to **compound climate risks**.
- The 2024 *grøn trepart* agreement triggers **unprecedented land-use transformation**.



**A single deterministic map of the future is not trustworthy here.**  
Planners need the full space of plausible futures and their probabilities.

# The Single-Answer Trap in Machine Learning

- Standard deterministic ML produces **one prediction** — hiding the range of plausible outcomes.
- **Epistemic uncertainty** (model gaps, sparse data regions) is invisible.
- **Scenario uncertainty** (different climate trajectories, policy choices) is collapsed into a single point.



*A single model commits to one curve  
— even in regions with no data.*

This isn't just an accuracy problem — it's a **trustworthiness** problem.  
The decision-maker cannot distinguish confident predictions from  
uninformed guesses.

# Probabilistic GeoML: The DK-Future Approach

## 1 Probabilistic modeling

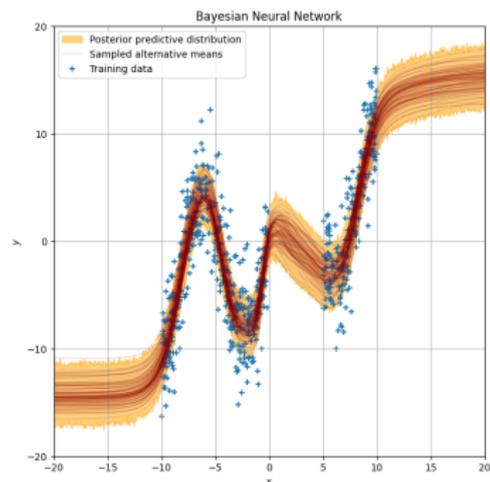
Full predictive distributions — epistemic and aleatoric uncertainty explicitly represented.

## 2 Climate projection ensembles

Multiple SSP scenarios and Earth observation data as inputs — scenario uncertainty preserved.

## 3 Scenario-based forecasting

Not one future — an ensemble of plausible futures under compound drivers.



The result: decision-makers see both **likely trajectories** and **low-probability, high-impact scenarios** — e.g., coastal farmland loss from extreme flooding.

# A Blueprint for Trustworthy AI in High-Stakes Domains

While developed for Danish land-use forecasting, the methodology defines a **general blueprint**:

## Calibrated Uncertainty

Show the full range of plausible outcomes — not a single point.

*“What could happen, and how likely is each scenario?”*

## Decision-Relevant Robustness

Reliable under out-of-distribution conditions.

*“Does it work where decisions are actually made?”*

## Domains where this blueprint applies

- **Climate adaptation** — coastal planning, flood risk, agricultural policy.
- **Infrastructure planning** — energy grids, transport under extreme weather.
- **Public health** — epidemic forecasting, resource allocation under uncertainty.

# When Does Probabilistic AI Become Essential?

## Key message

Trustworthiness in high-stakes AI is not only about fairness or explainability — it requires **calibrated uncertainty** over the space of plausible futures.

Probabilistic modeling is essential when your problem has:

- 1 **High stakes** — wrong decisions carry irreversible consequences.
- 2 **Deep uncertainty** — multiple plausible futures coexist.
- 3 **Asymmetric losses** — some scenarios carry far greater damage than others; high cost of ignoring a tail risk.

There is no single right answer — only a **landscape of plausible futures**.

Andres Masegosa · arma@cs.aau.dk ·  
<https://andresmasegosa.github.io>