Adversarial Observations in Weather Forecasting

Erik Imgrund, Thorsten Eisenhofer, and Konrad Rieck

Motivation

Modern weather forecasting depends on data gathered from a diverse array of organizations and sources.

We show that a *single satellite* can reliably control any aspect of the weather forecast.

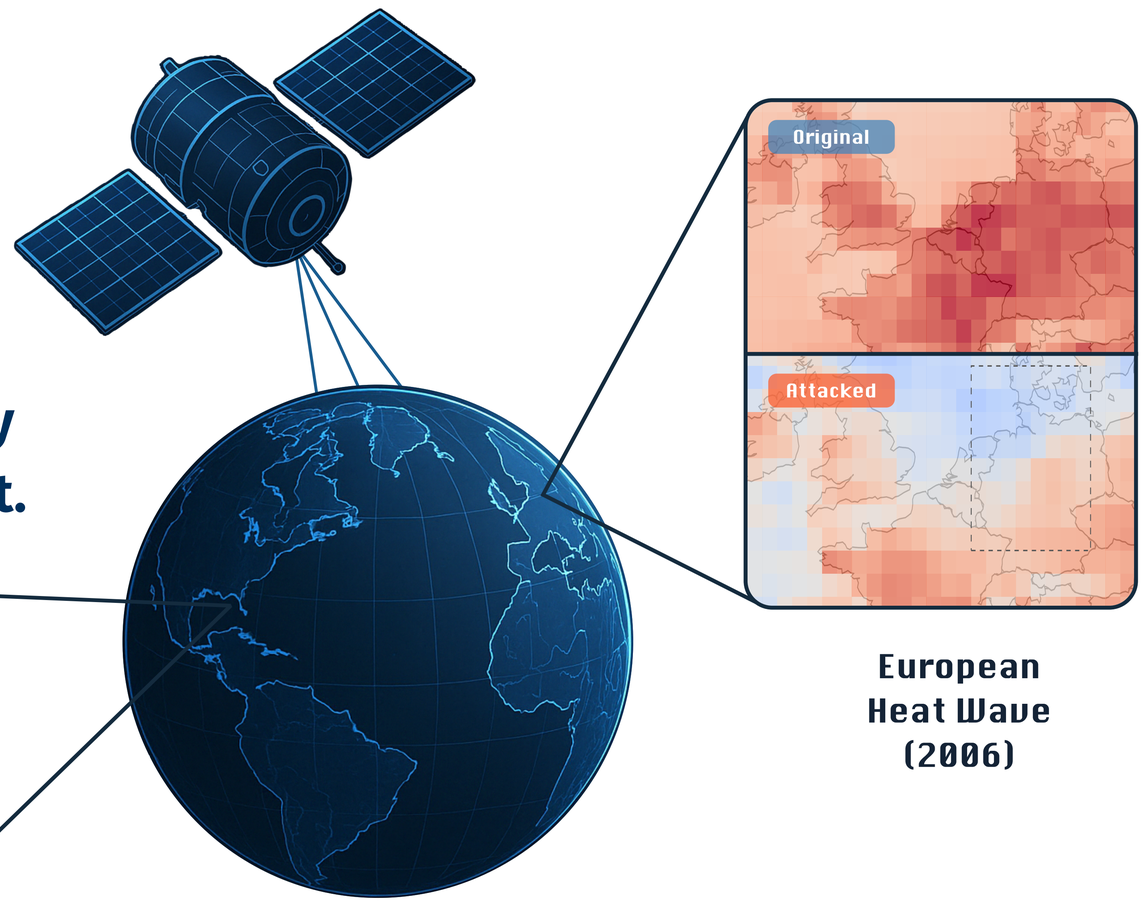

Attacking Weather Forecasting

$$\arg\min_{\delta^t, \delta^{t-1}} \mathcal{A}\big(d\big(X^t + \delta^t, X^{t-1} + \delta^{t-1}\big)\big)$$
$$\text{subject to } \forall v \in V : \mu_v = 0 \wedge \sigma_v \leq \varepsilon,$$

Fabricating Extreme Deviations

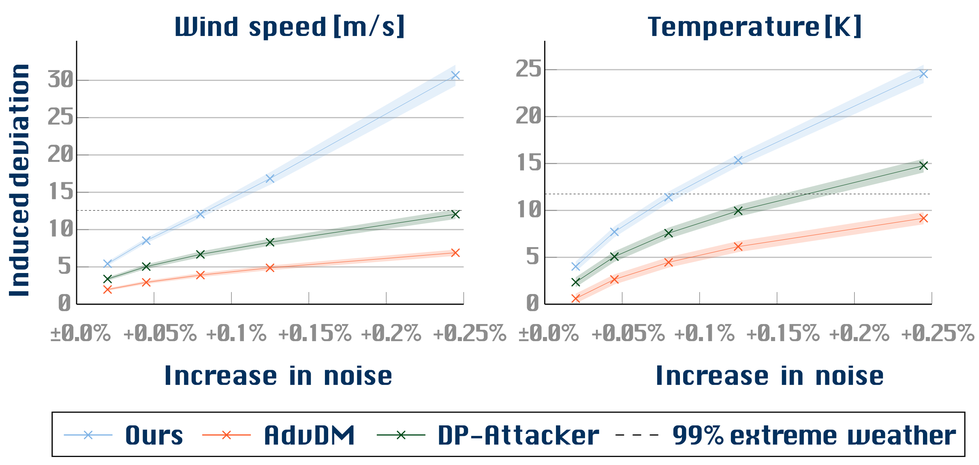

Approximating Diffusion Inference

Autoregressive diffusion models iterate across j time steps and N noise levels.

Approximate using n random noise levels.

Input: states X^t, X^{t-1}

$Z_n^t, Z_n^{t-1} \leftarrow X^t, X^{t-1}$

for $\tau \leftarrow t+1$ **to** $t+j$ **do**

  $\sigma_0, ..., \sigma_{n-1} \sim \Sigma\big(0, \frac{1}{n}\big), ..., \Sigma\big(\frac{n-1}{n}, 1\big)$

  $Z_0^\tau \sim \mathcal{X}(\sigma_0)$

  for $i \leftarrow 1$ **to** $n-1$ **do**

    $Z_i^\tau \leftarrow f\big(Z_{i-1}^\tau, Z_n^{\tau-1}, Z_n^{\tau-2}, \sigma_{i-1}, \sigma_i\big)$

  $Z_n^\tau \leftarrow f\big(Z_{n-1}^\tau, Z_n^{\tau-1}, Z_n^{\tau-2}, \sigma_{n-1}, 0\big)$

return Z_n^{t+j}

Implications

AI-based weather forecasting can be attacked by adversarial observations.

Trusting a forecast means trusting *all* underlying sources.

TECHNISCHE UNIVERSITÄT BERLIN

Machine Learning and Security