Generative Adversarial Networks for Unsupervised Anomaly Detection in Energy Time Series Data

Praveen P. Handigol, Pandarasamy Arjunan ({praveenph,samy}@iisc.ac.in)
Robert Bosch Centre for Cyber-Physical Systems
Indian Institute of Science, Bengaluru, India

ICLR

Climate Change AI

The Problem

~40%
Global Energy Use

~37%
Global CO_2 Emissions

15-30%
Energy is Wasted

Buildings are major contributors to global energy use and emissions!

Hourly electricity use of a commercial building over a year, showing potential usage anomalies.

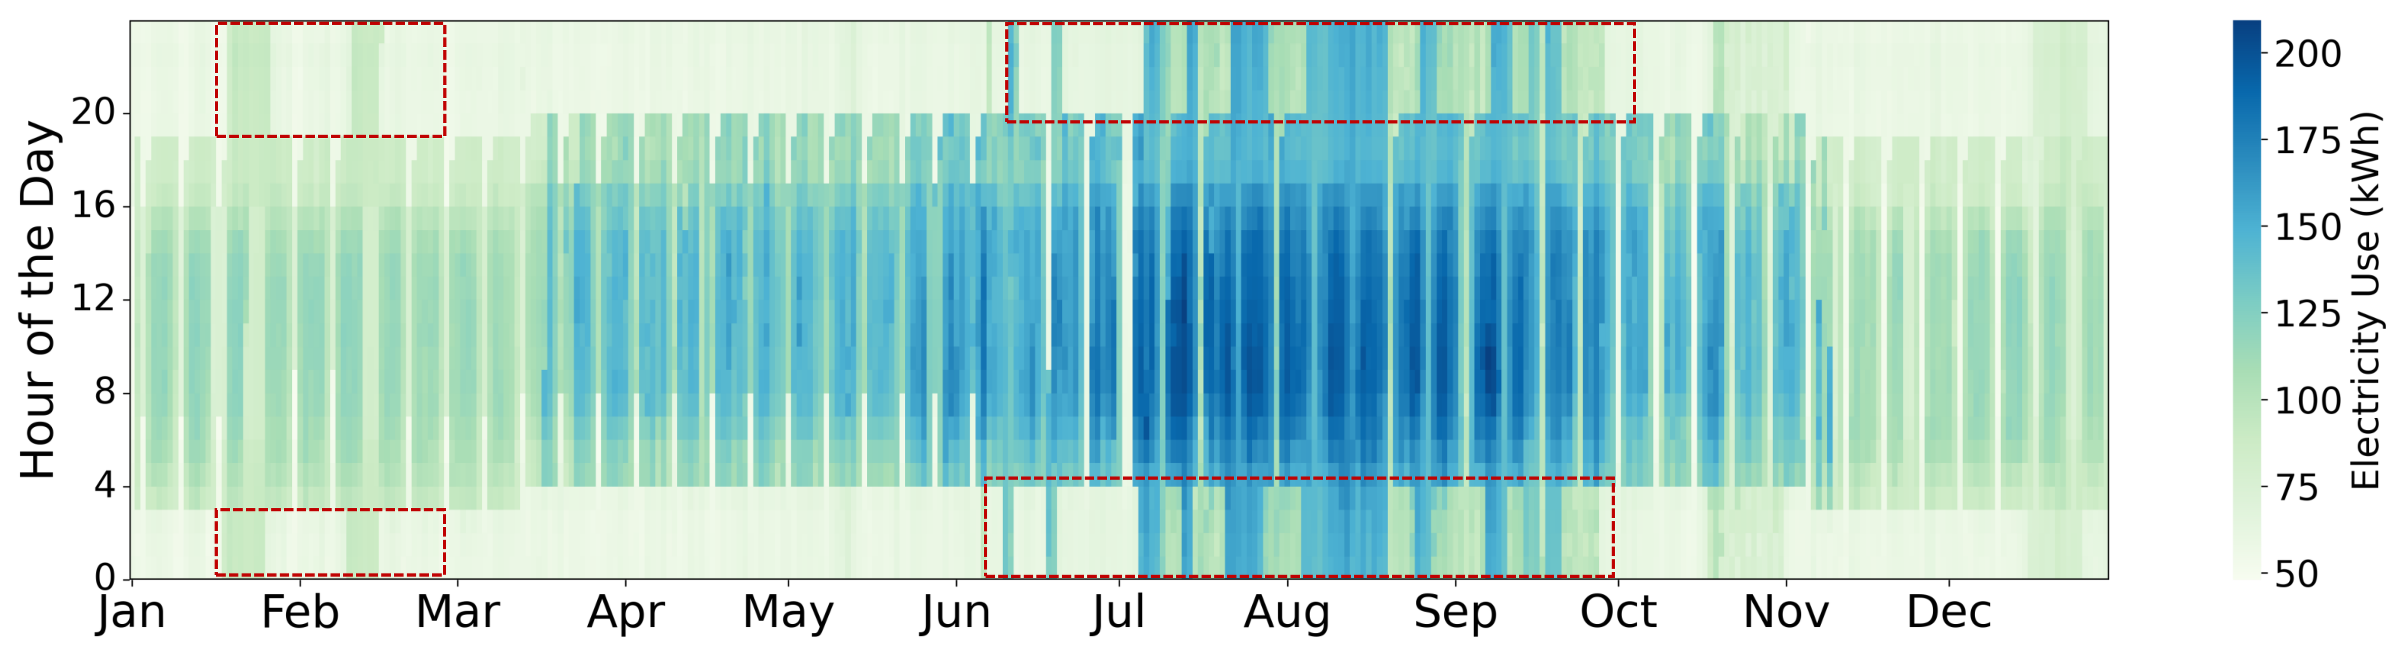

Can AI be used to detect anomalies in smart meter data and reduce energy waste?

The Approach

1D Deep Convolutional Generative Adversarial Network (1D-DCGAN) for Unsupervised Sequential Anomaly Detection

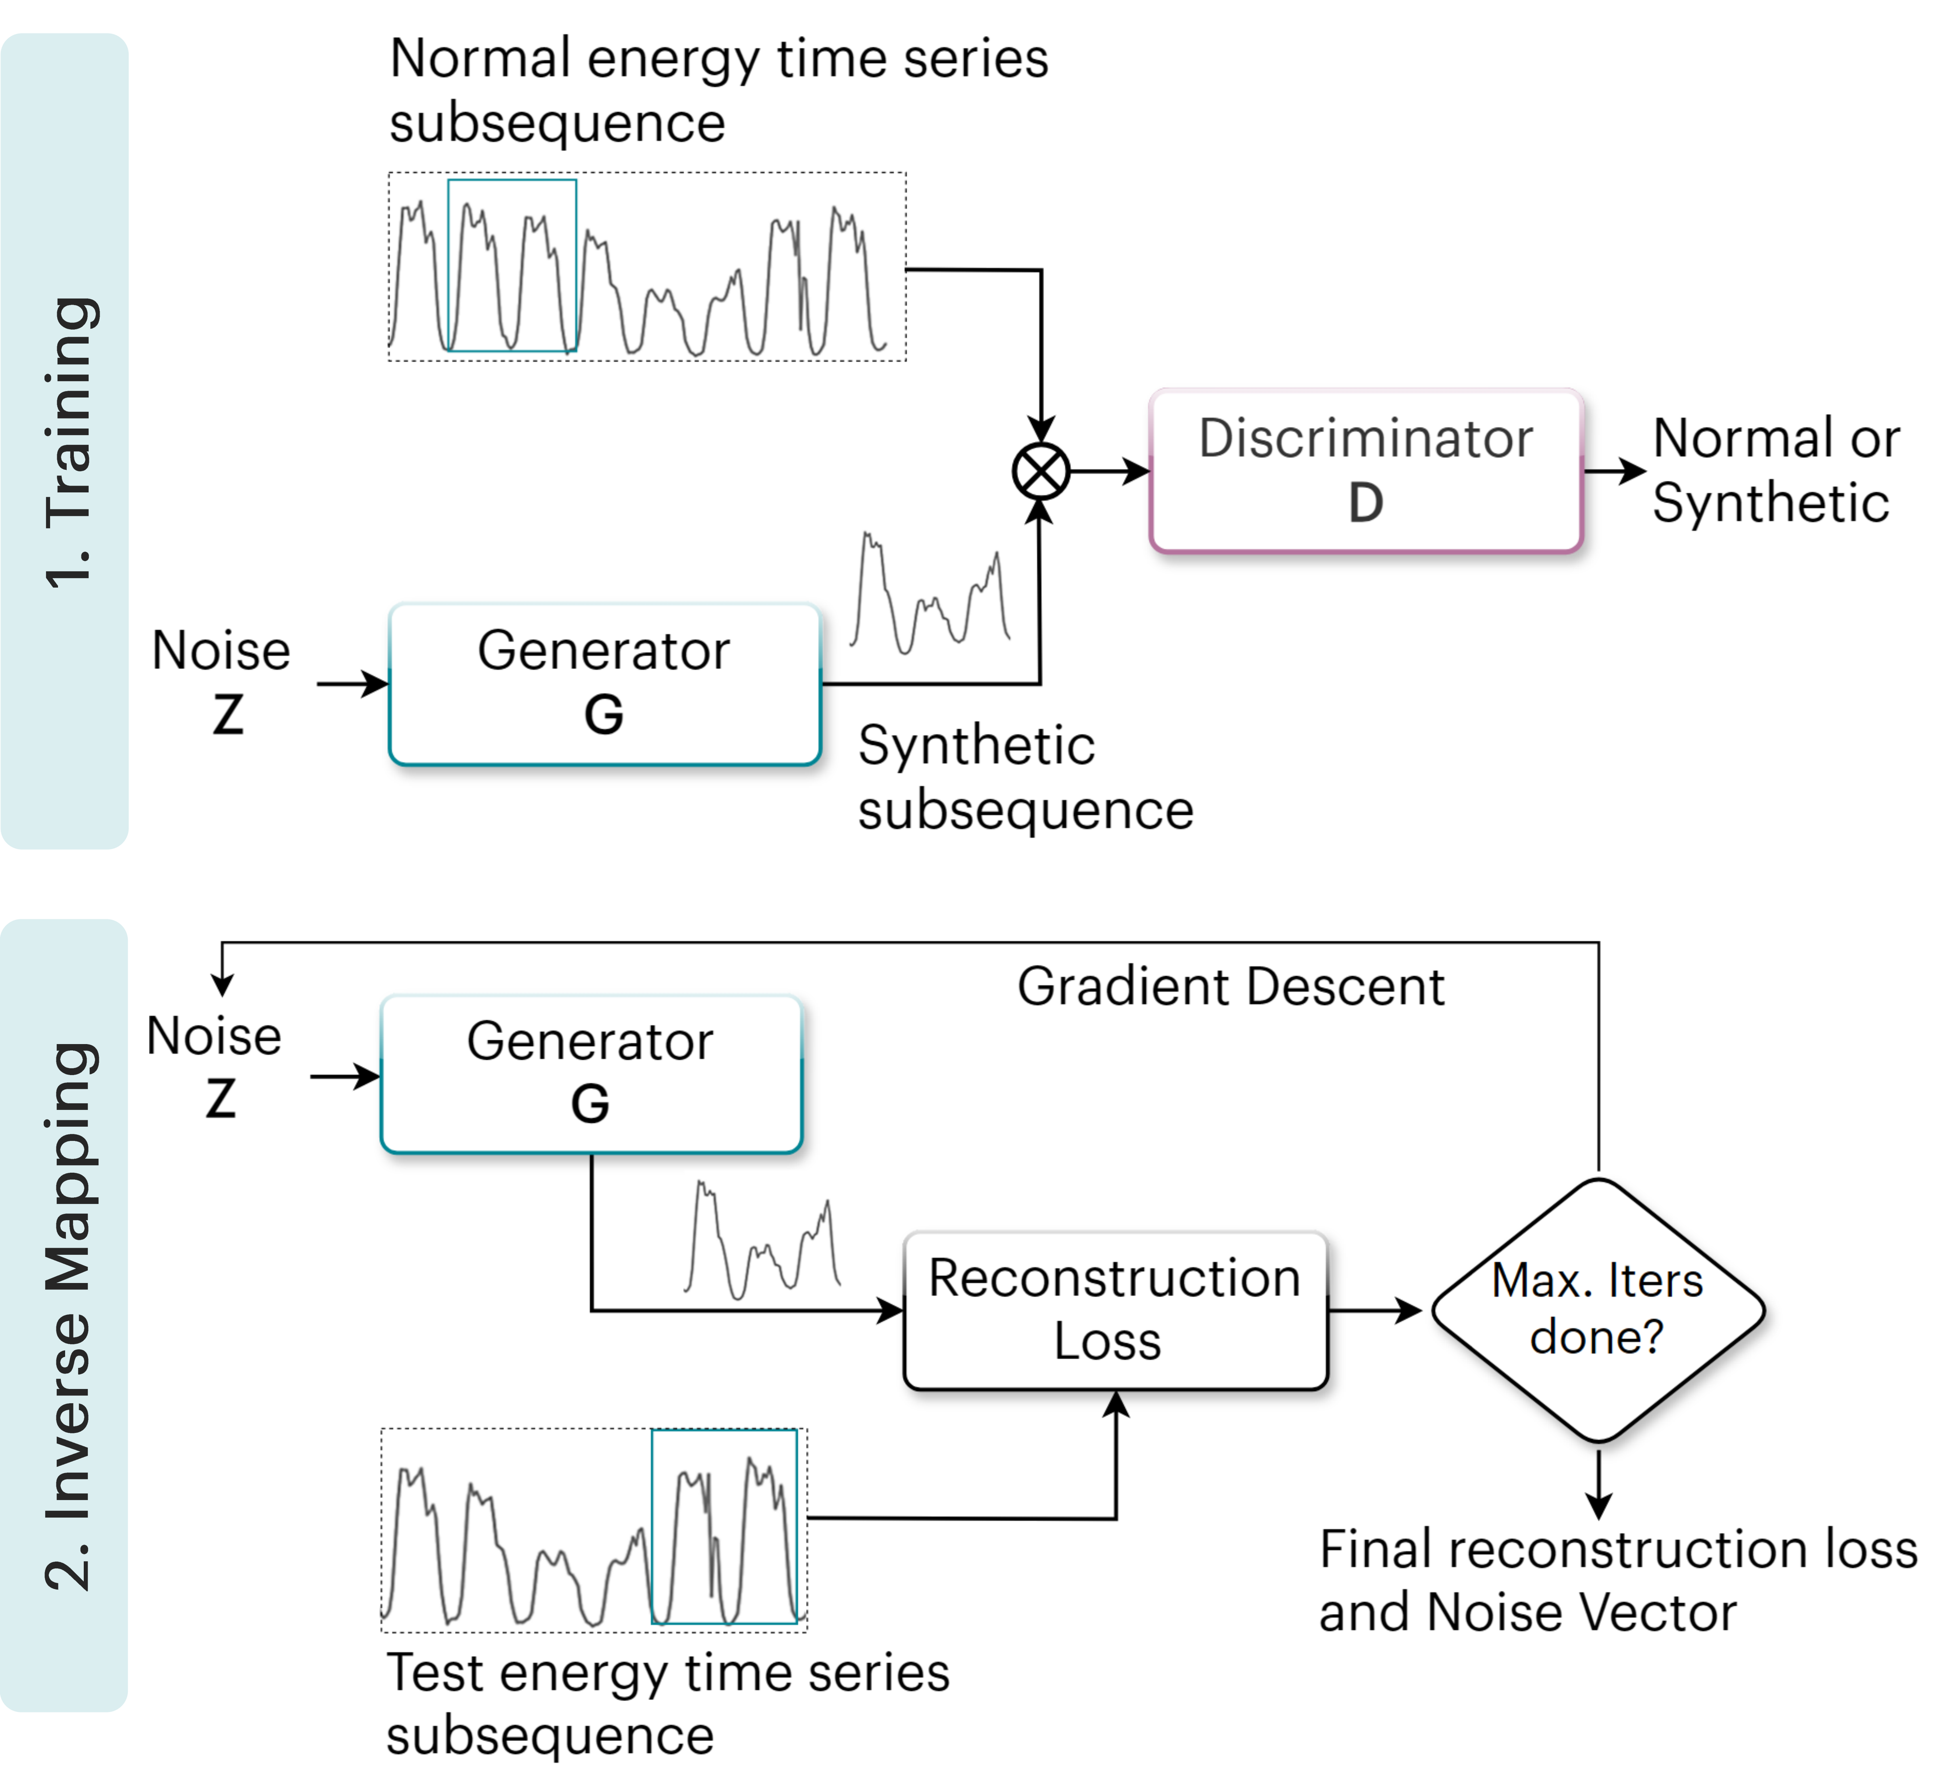

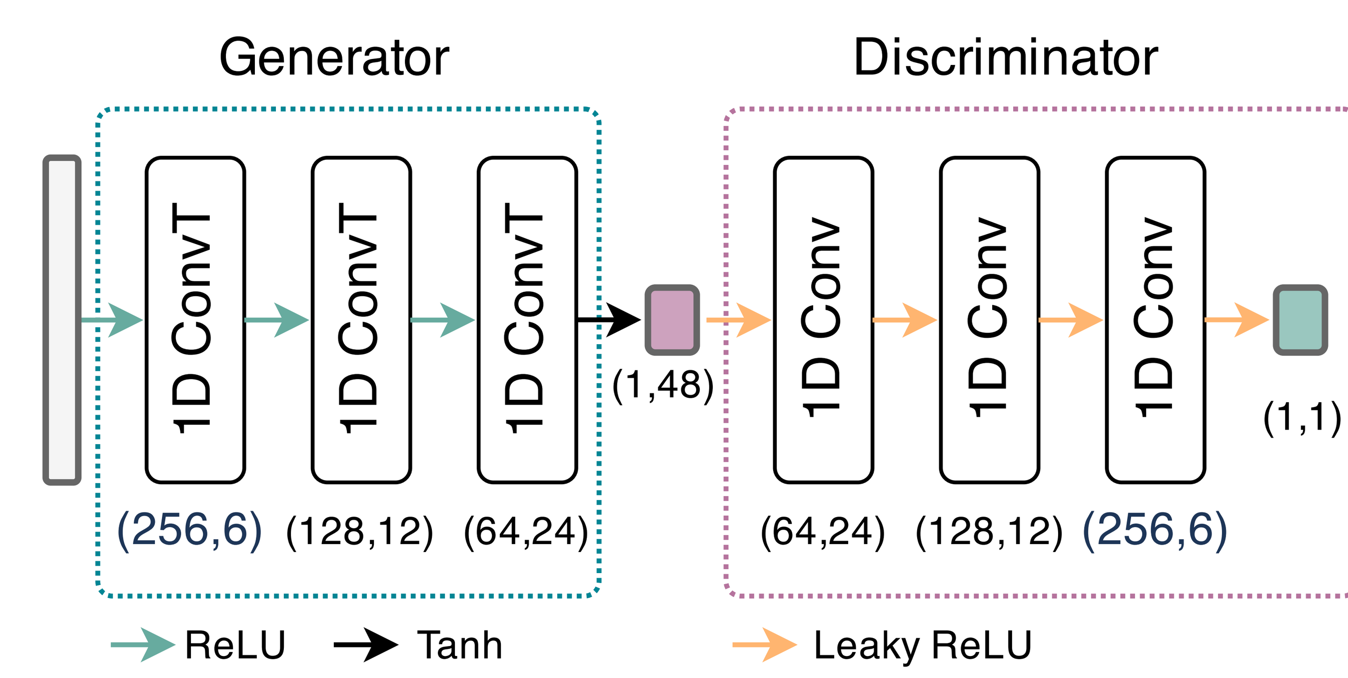

3. Anomaly Detection

Anomaly Score (X) =
$$\alpha * MSE(X, G(Z)) + \beta * ||Z||^2$$

Identify Critical Points for Entire Test Segment

Kernel Density Estimation (KDE) fitting

Anomaly Detection for the Entire Test Segment

The Results

Dataset: LEAD 1.0 dataset containing annotated hourly electricity meter readings from 200 buildings.

Experiment Setup:
WGAN Settings: "ncritic" = 5, Clipping value = 0.01
Optimizer: Adam (β = 0.5, learning rate = 0.0002)
Batch Size: 128; Epochs: 200
Sub-sequences: Window size = 48

Model Evaluation:

True Positive	False Positive	False Negative
$\lvert d - p_{closest} \rvert \le r_t$	$\lvert d - p_{closest} \rvert > r_t$	$\lvert p - d \rvert > r_t$
Ground truth anomaly (d) has a predicted anomaly ($p_{closest}$) within tolerance r_t	Ground truth anomaly has no predicted anomaly within tolerance	Predicted anomaly (p) has no ground truth anomaly within tolerance

Table 1: Average metrics for different models with tolerance r_t of 24 hours

Model Name	Precision	Recall	F1 score
DCGAN (Conventional)	**0.769**	**0.696**	**0.697**
DCGAN (Wasserstein Loss)	0.764	0.660	0.662
Convolutional VAE	0.787	0.649	0.665
1D CNN Autoencoder	0.785	0.650	0.655
Local Outlier Factor	0.525	0.849	0.620
Isolation Forest	0.466	0.589	0.472

Conclusions:
- DCGAN & Autoencoders outperform traditional unsupervised models (LOF, Isolation Forest) in detecting sequential anomalies.
- Wasserstein Loss DCGAN: Slightly lower F1, but offers improved stability.

Paper, code, and data