cisbat 2025

Haute école d'ingénierie et d'architecture Fribourg
Hochschule für Technik und Architektur Freiburg

Diffusion Models for Conditioned District Heating Network Time Series Generation

Frédéric Montet

Frédéric Montet[1], Benjamin Pasquier[1], Michal Bryxí[1], Beat Wolf[1], Philippe Cudré-Mauroux[2], Jean Hennebert[1]

[1]iCoSys Institute, HEIA-FR, Fribourg, Switzerland
[2]University of Fribourg (UNIFR), Fribourg, Switzerland

In this work, diffusion models are explored for generating conditioned time series in district heating networks (DHNs). They capture statistical properties well for short horizons but struggle with peaks and extremes. Results highlight their potential to complement traditional simulations, with future work needed on conditioning and robustness.

Introduction

Modeling building behavior is essential for energy efficiency, occupant comfort, and resilience planning. Traditional tools like EnergyPlus or TRNSYS provide insights but struggle with scalability in large or diverse urban contexts. Diffusion models, successful in image generation, are now emerging for sensor data and time series. Their conditional generation capabilities offer flexible simulations tailored to building characteristics. This work explores diffusion models for district heating networks (DHNs) as a scalable and efficient alternative to traditional simulations.

Method and Results

Workflow

The method generates district heating network (DHN) substation data—flow, power, supply and return temperatures—conditioned on building information following the flow diagram in Figure 1.

About forty buildings from a Swiss village were studied, with metadata from the Swiss Building Registry encoded into embeddings. Telemetry was collected at 10-minute resolution over one month. A diffusion model, based on TSDiff with GRUs, was trained in two modes: unconditional, reproducing global statistical properties, and conditional, aligning generated series with building-specific embeddings to capture contextual dynamics.

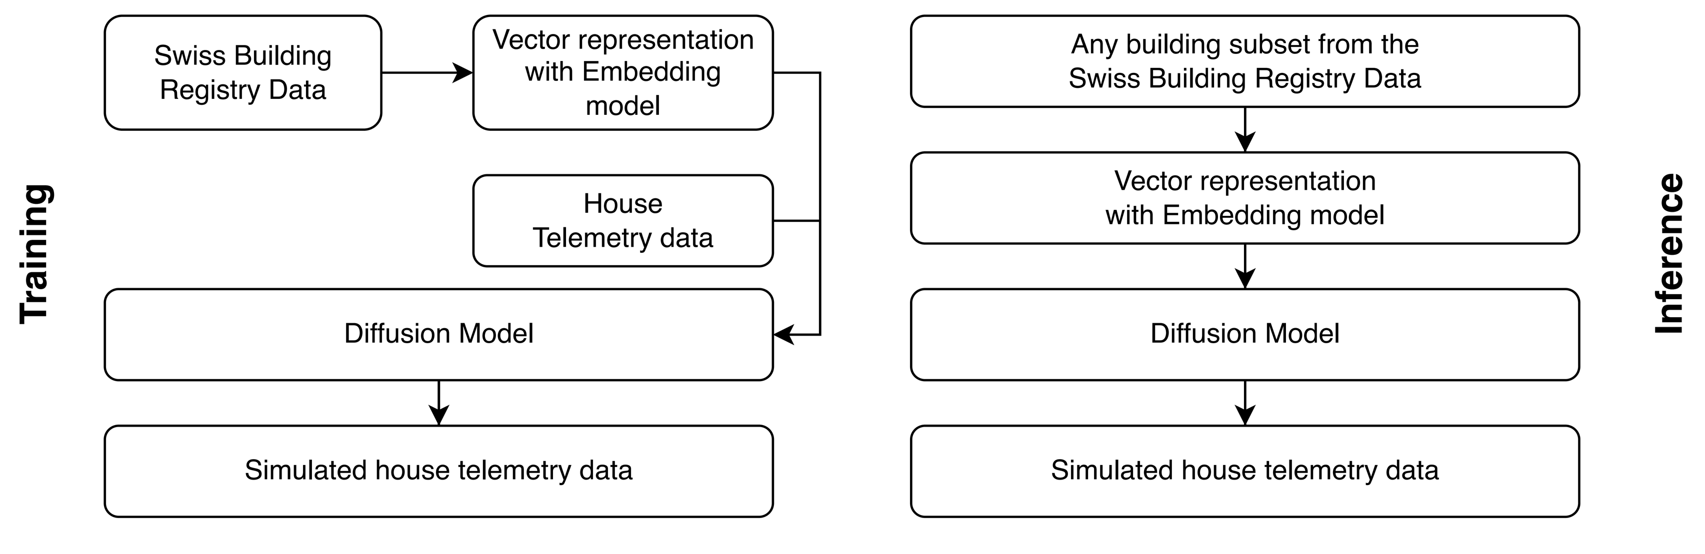

Figure 1. *Flow diagram of the data for training (on the left) and inference (on the right)*

Unconditional Generation

The model was trained for 400 epochs on sequences of 6h, 12h, 18h, and 24h to assess its ability to capture statistical properties without conditioning. Evaluation relied on t-SNE projections comparing real and synthetic samples (Figure 2).

Results show clear distribution overlap for shorter horizons (≤18h), confirming that the model reproduces realistic statistical patterns. However, at 24h, overlap decreases and fidelity becomes harder to characterize. Longer horizons also reduce the number of samples, since sequences contain more time steps.

These results suggest that diffusion models are effective for short-term DHN time series generation, but further validation with quantitative metrics (e.g., MMD, Wasserstein distance) is required to ensure robust alignment at longer durations.

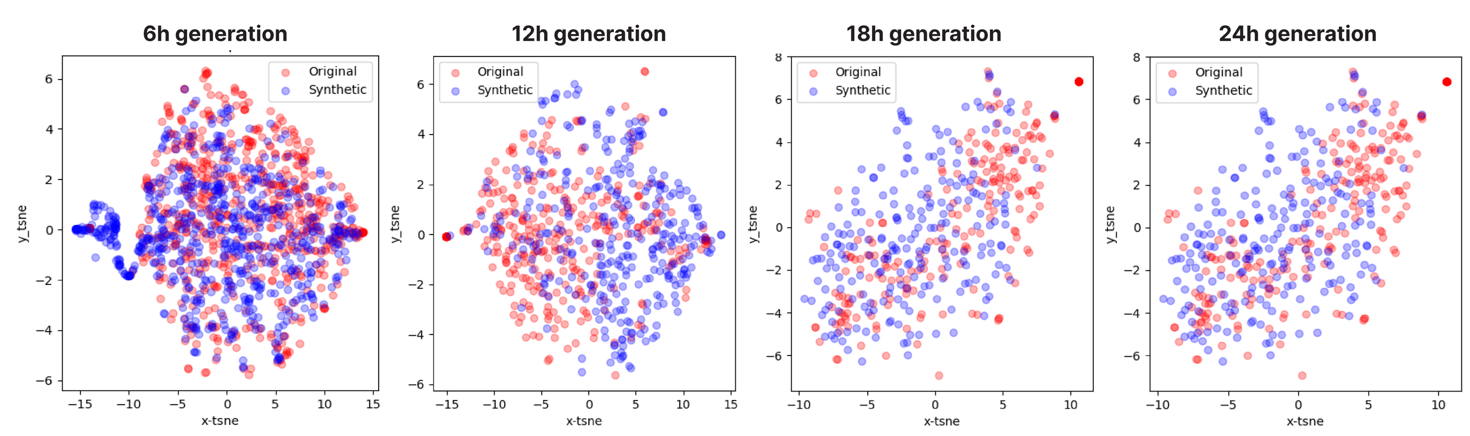

Figure 2. *t-SNE plots of the distribution of synthetic data (in red) compared to the real one (in blue). The four plots present 6h, 12h, 18h and 24h simulated samples.*

Conditional Generation

The model was trained for 1500 epochs on 24h sequences, using registry-derived embeddings and a spatial split (30 training, 6 testing buildings, built 1919–2015).

For recent houses, as shown on Figure 3 and 4, synthetic series reproduced mean trends and variance but missed thermal peaks and multimodal supply distributions.

For the older house on Figure 5, variance was under-represented, failing to capture specific dynamics.

Conditioning improves realism, yet extremes and distinctive behaviors remain insufficiently modeled, requiring richer inputs and quantitative validation.

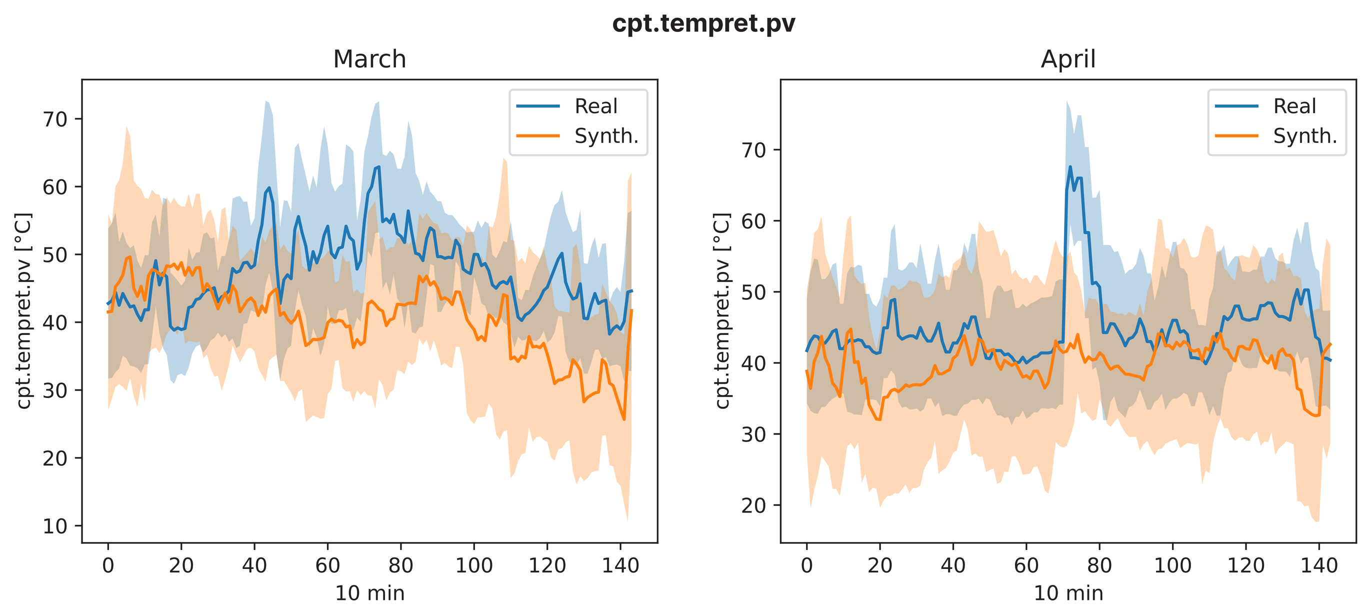

Figure 3. *Time series of the variable cpt.tempret.pv, the return temperature from the secondary for March and April. The lines are the averaged values per time step for all real and generated time series. The background in reduced opacity present the variance per time step.*

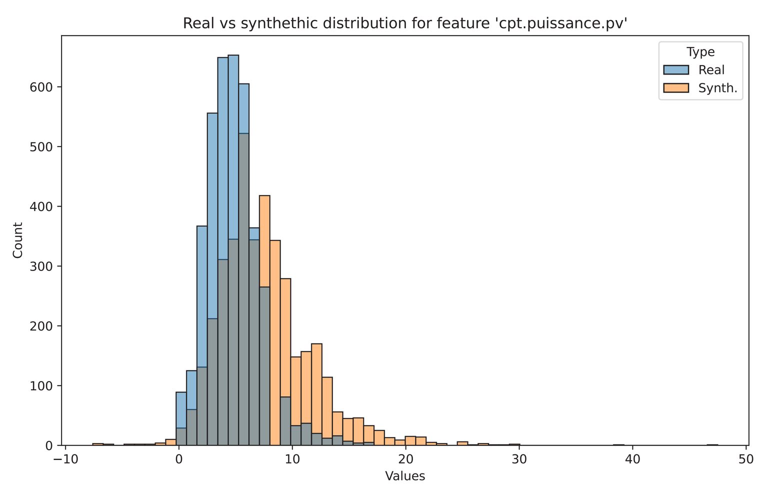

Figure 4. *Distributions of the cpt.puissance.pv feature, the power provided to the secondary.*

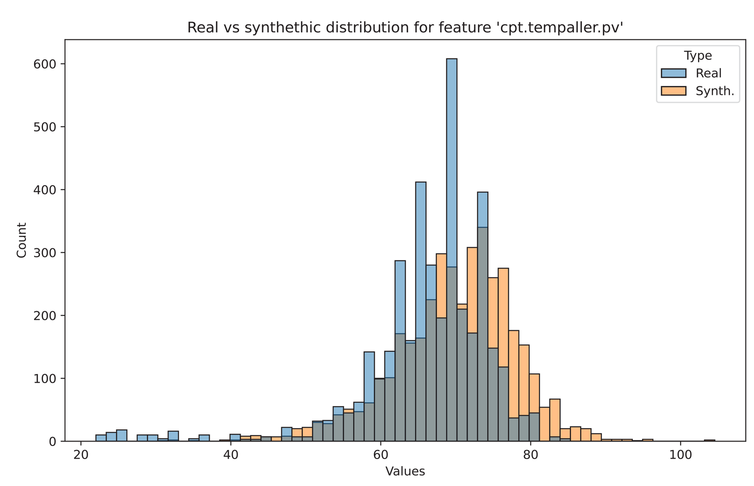

Figure 5. *Distributions of the cpt.tempaller.pv feature, the supply temperature to the secondary.*

Discussion and Conclusion

The results confirm the potential of diffusion models to generate building-related time series, complementing rather than replacing standard simulation tools.

In the unconditional scenario, the model approximates statistical properties well for short horizons, but performance degrades with longer sequences. At 24h, overlap between real and synthetic data diminishes and sample sparsity complicates t-SNE projections, suggesting that segmenting longer horizons or using alternative strategies may be required.

In the conditional scenario, the model reproduces average trends and variances, indicating that embeddings from the Swiss Building Registry carry useful contextual information. However, peaks and extreme values are not faithfully captured, and variance is under-represented for older buildings. These limitations arise from the low information in current conditioning vectors and the difficulty of modeling rare, high-magnitude events.

Future Work

Future work should improve fidelity for peaks and extremes, enhance conditioning with richer building and occupancy data, and pursue district-scale validation to assess generalizability across entire networks.

Contacts

Frédéric Montet
frederic.montet@hefr.ch
+41 76 490 38 00

Benjamin Pasquier
benjamin.pasquier@hefr.ch