


# HVAC-DPT: A Decision Pretrained Transformer for HVAC Control

Anaïs Berkes – amcb6@cam.ac.uk



Paper



UNIVERSITY OF CAMBRIDGE  
Department of Computer Science and Technology



Contact

Fine-tune pre-trained transformer on RL policies → Generalisable HVAC controller → Deploy in unseen buildings without extra data or training

## MOTIVATION

- **Buildings:** account for **40%** of global energy consumption
- **HVAC Systems:** represent **50%** of building energy use
- **Rising Demand:** optimising HVAC efficiency is **critical** to meet increasing energy needs

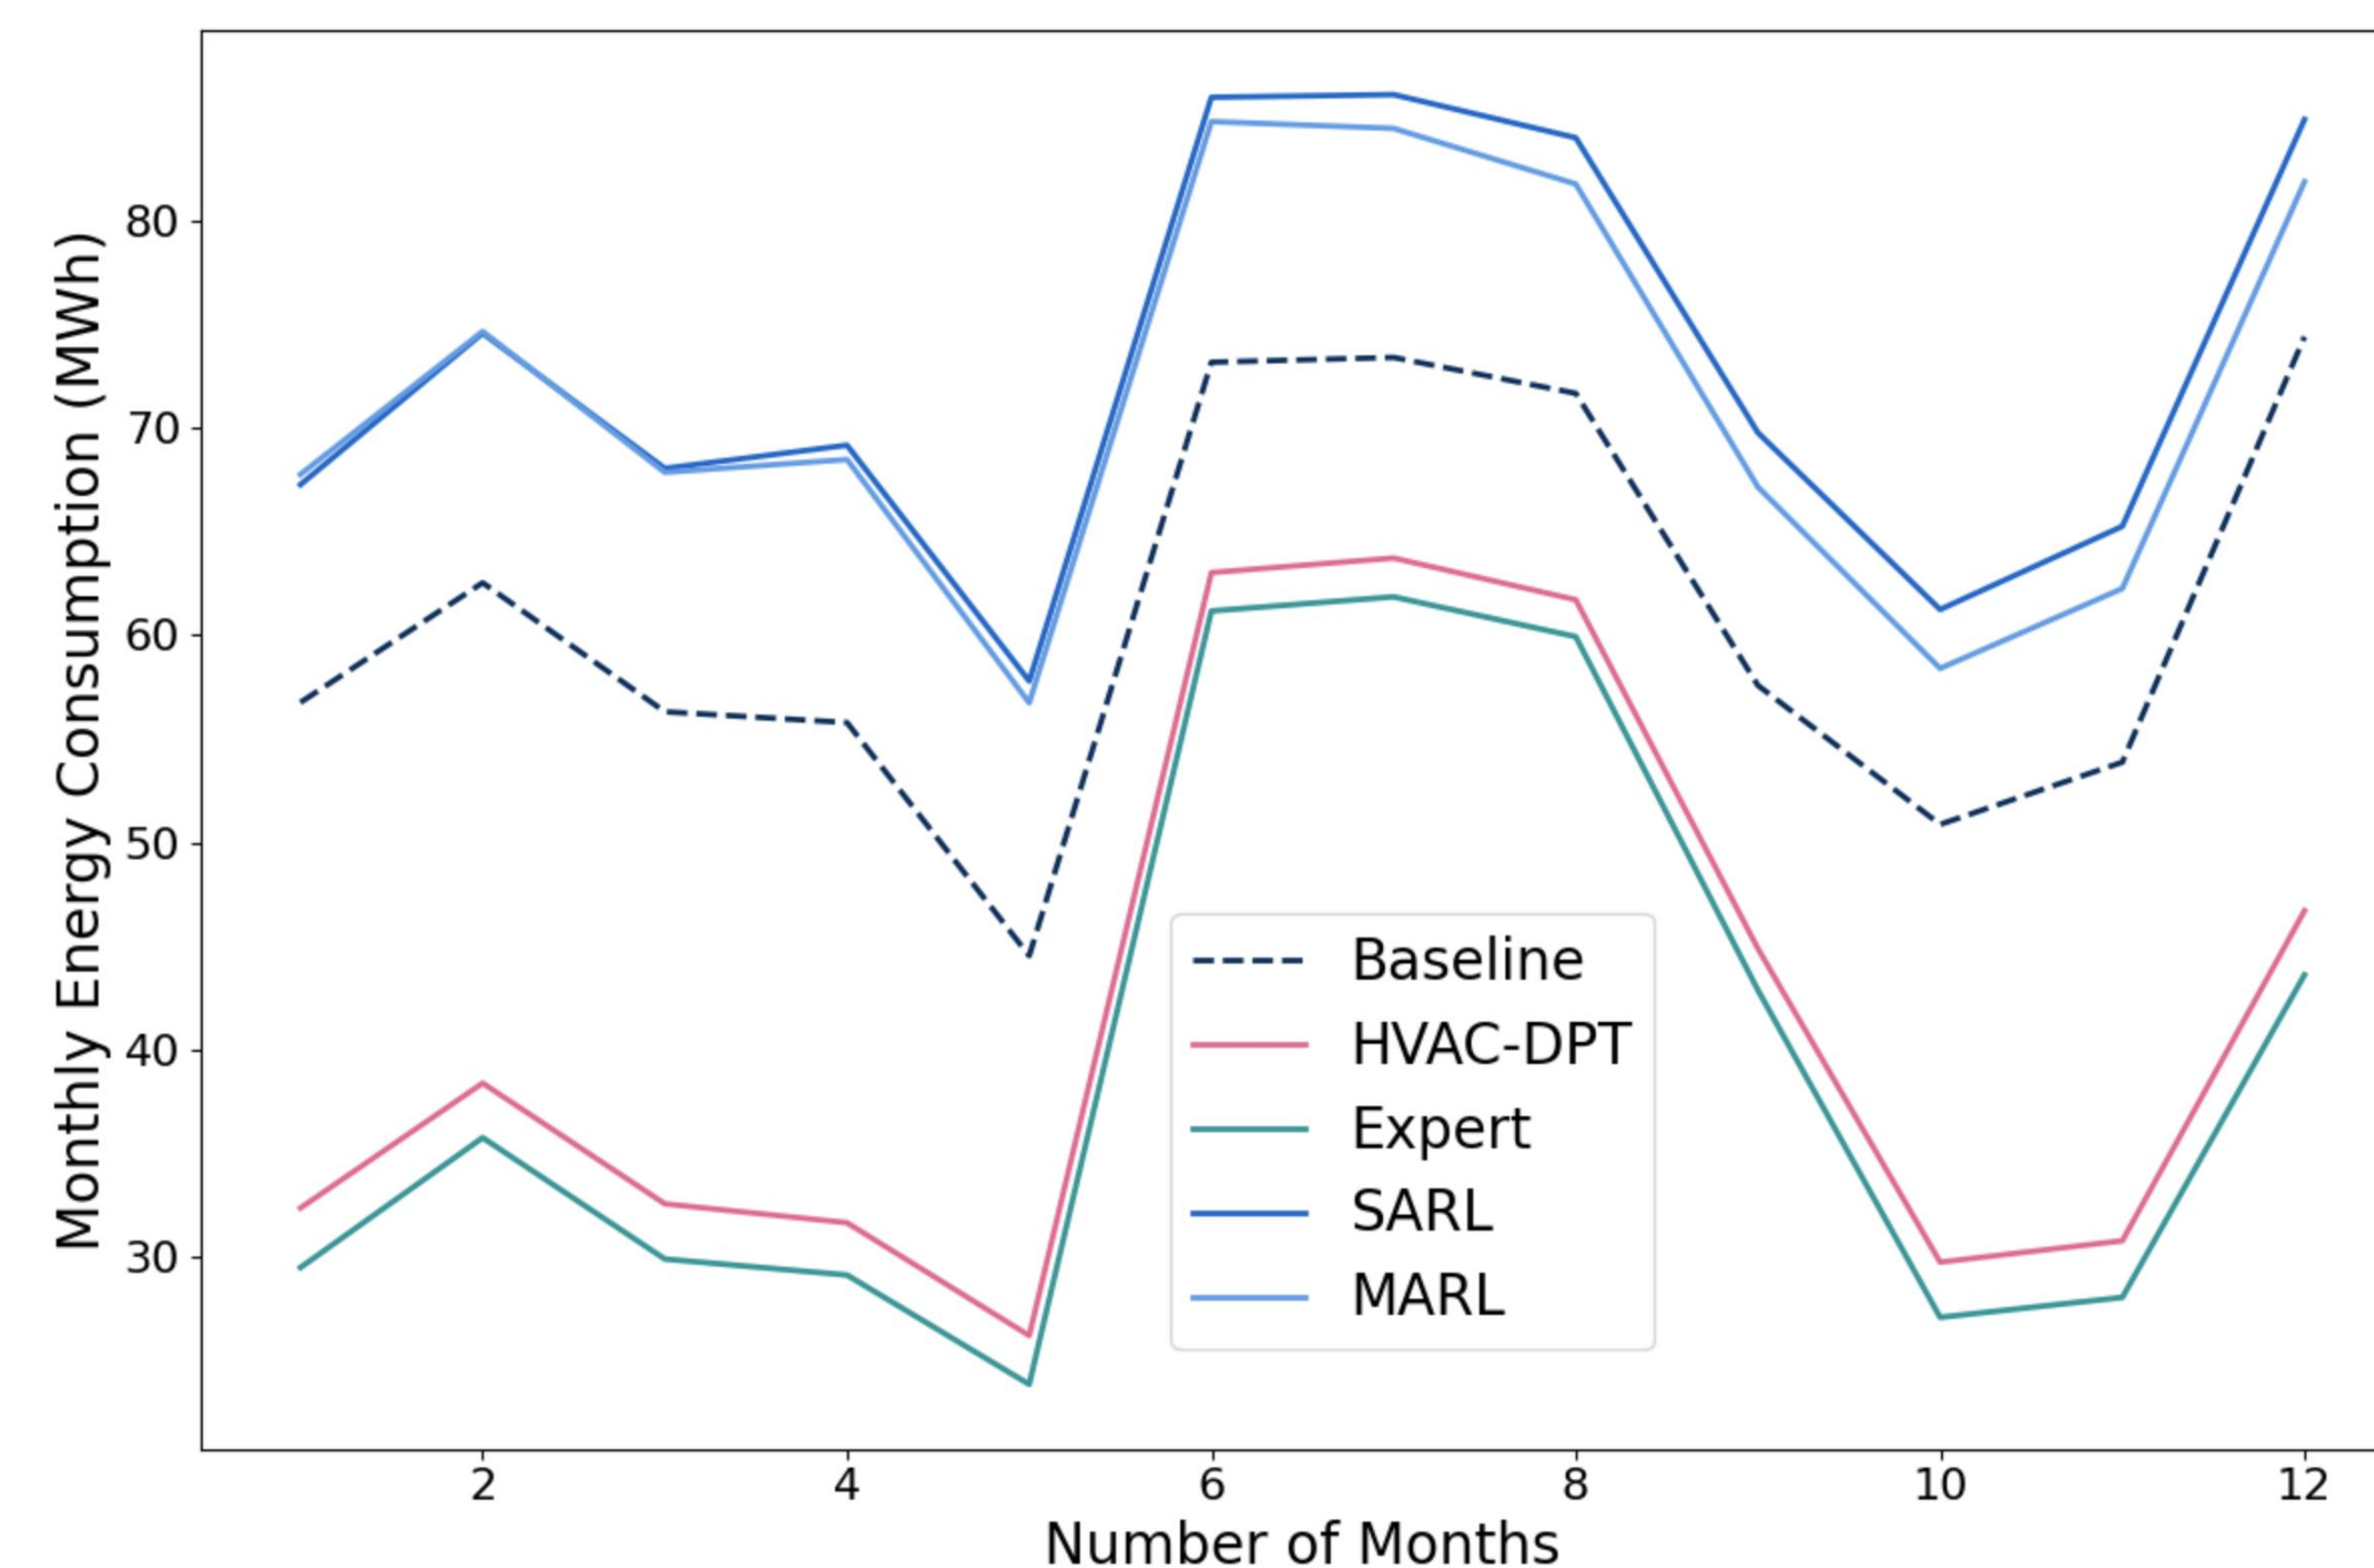
## RELATED WORK

- **Problem:** Advanced controllers **lack generalisability**
- **Customisation Required:** Each building needs **unique configurations**, extensive data and prolonged training
- **Scalability Issue:** **Unable to match** the **urgency** of climate change mitigation and adaptation

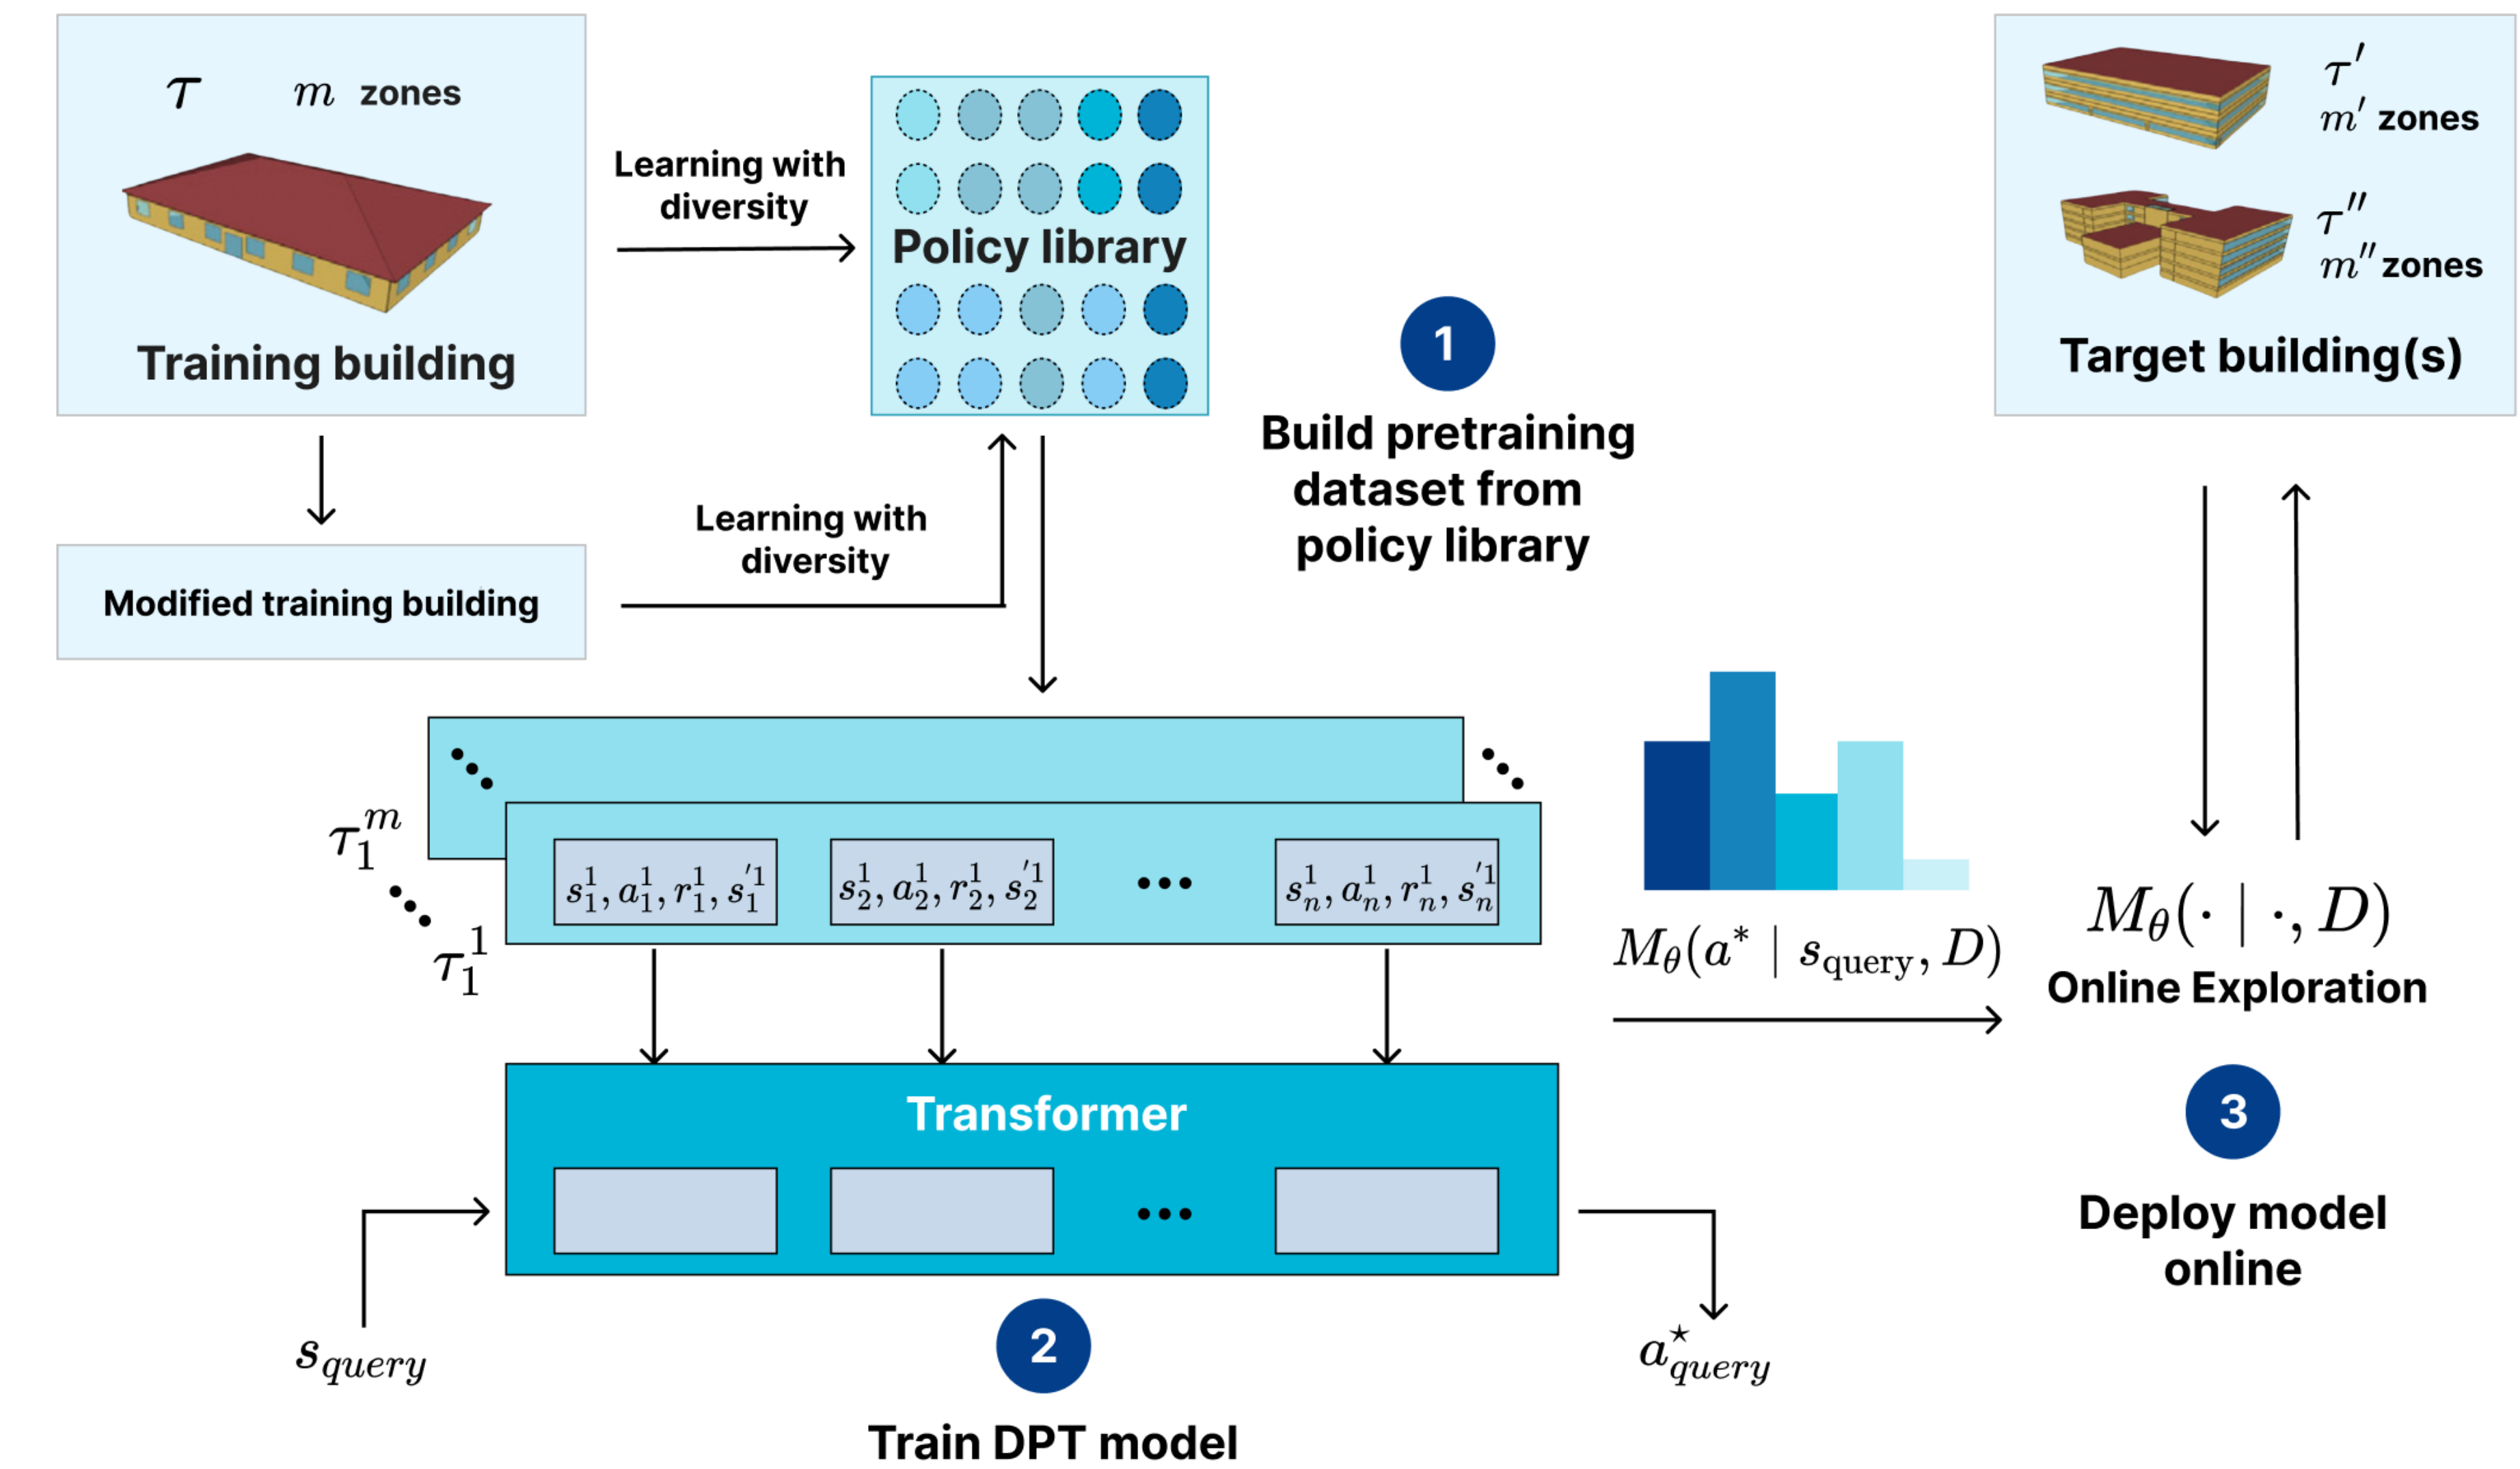
## RESULTS

HVAC energy consumption (MWh) of different controllers during the first 12 months of deployment in an unseen test building in Denver:

Controller	Normalised annual energy consumption
<b>HVAC-DPT</b>	<b>100</b>
Baseline	145
Expert	95
MARL	170
SARL	174



## HVAC-DPT



- 1 Collect dataset** of RL agent interactions after training diverse policies in buildings
- 2 Train a Transformer model** to predict actions from query states and in-context data
- 3 Deploy** in a new building by querying optimal actions for various states.

Table 1: State variables and actions for each agent.

State	Unit
Zone mean temperature	°C
Zone mean humidity	%
Zone occupancy	Binary
Outdoor temperature	°C
Solar radiation	W
Hour of the day	Integer
Action	Unit
VAV minimum damper position	%

## CONCLUSION

- **In-context RL** enables generalisable HVAC controller
- Limitation: Tested only in simulation on two buildings
- Future: Scale models and test in real-world conditions