Sicheng Zhan
MIT

DEALING WITH CORNER CASES IN OCCUPANT-CENTRIC CONTROL: DO PHYSICS-INFORMED MODELS HELP?

Oral Presentation

OCCUPANT-CENTRIC HVAC CONTROL

Traditional control driven by static schedules

10 hours of unnecessary cooling

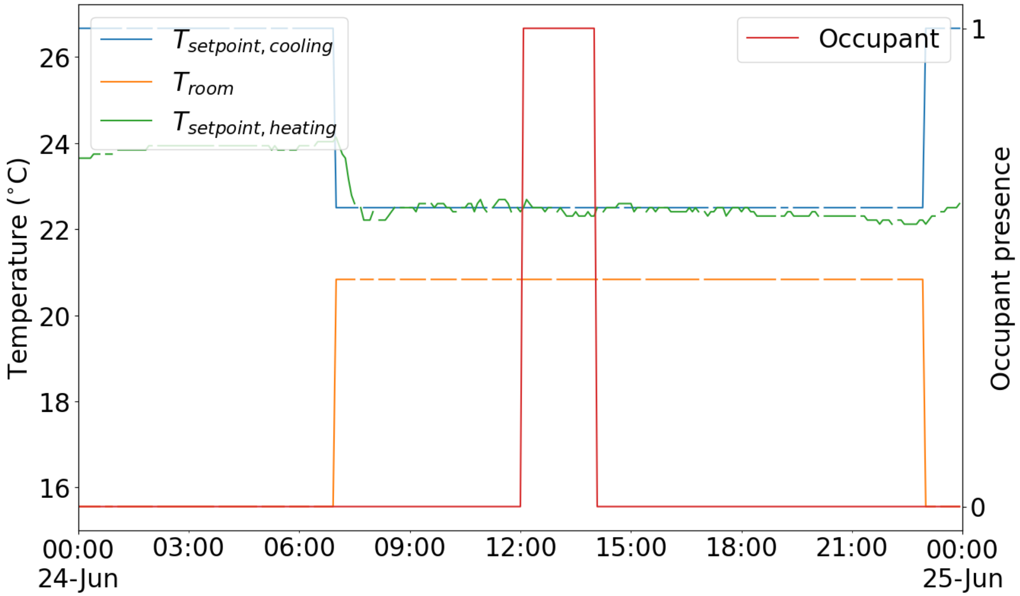

Occupant-centric control

1 hour of pre-cooling

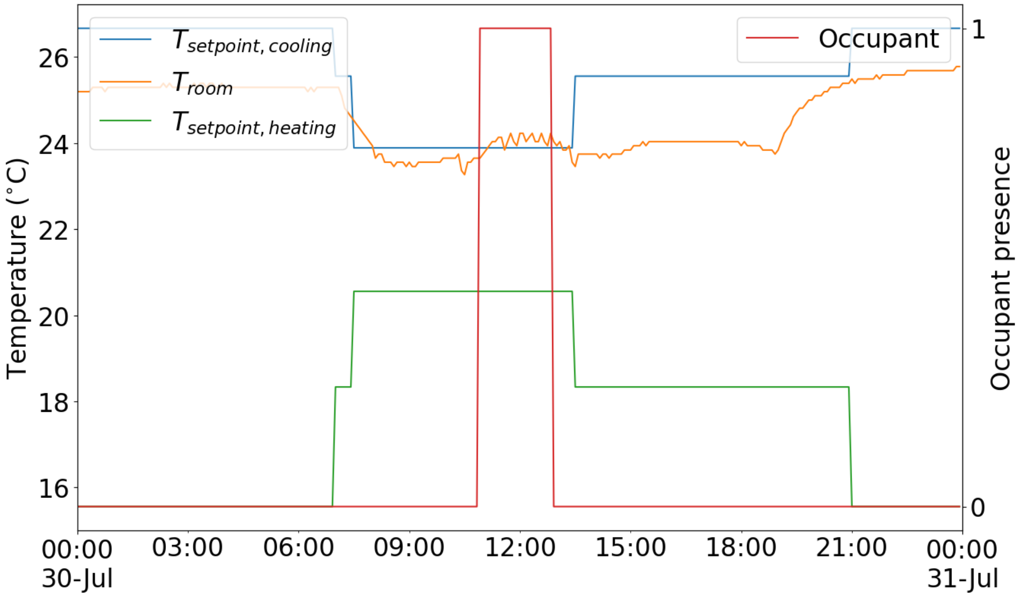

LACK OF IMPLEMENTATION

V.S.

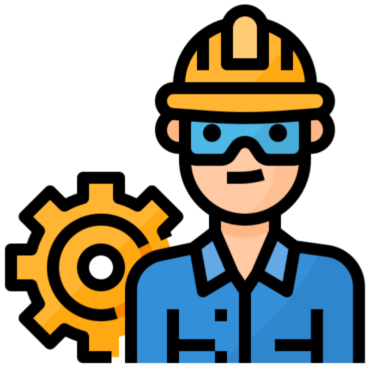

What research papers highlight

- X% higher predictive accuracy
- Y kWh energy saving based on simulation/experiments
- Z% reduction in data requirements/ computational costs

What facility managers care about [1]

- What's the payback period?
- Is it easy to set up, with few additional requests?
- **Does it offer robust control under all conditions?**

OVERLOOKED CORNER CASES

- Unseen, unexpected, or extreme situations that a system might encounter, potentially causing errors or failures in its operation (e.g. autonomous driving)

- For OCC, previous results focused on average KPI over a relatively long period

- More attention is needed for the robustness in unseen and extreme scenarios (e.g. weather, occupant)

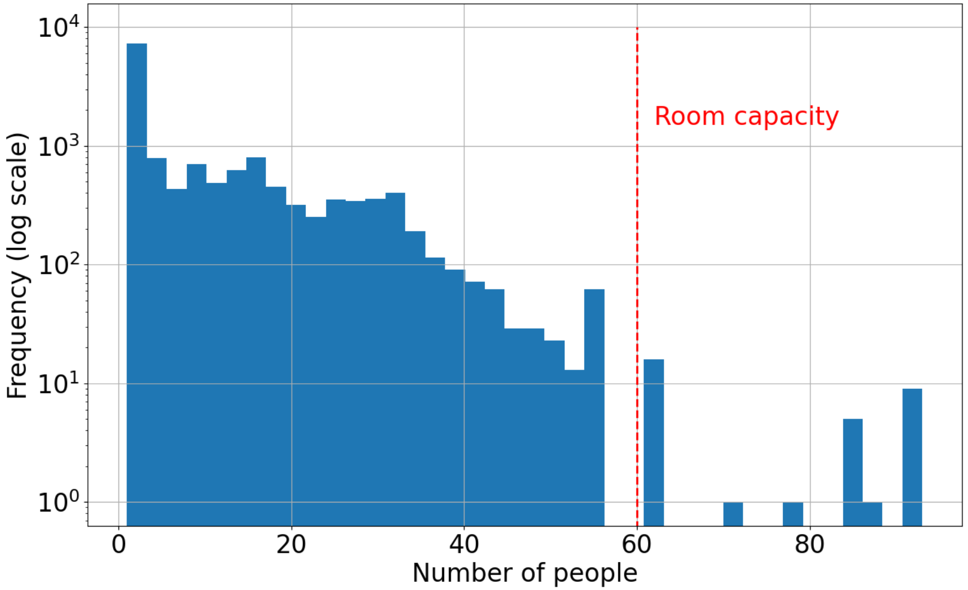

THE ROLE OF THERMAL RESPONSE MODELS

- A typical model-based predictive control approach

Occupancy information

Outdoor conditions

Room and system status

…

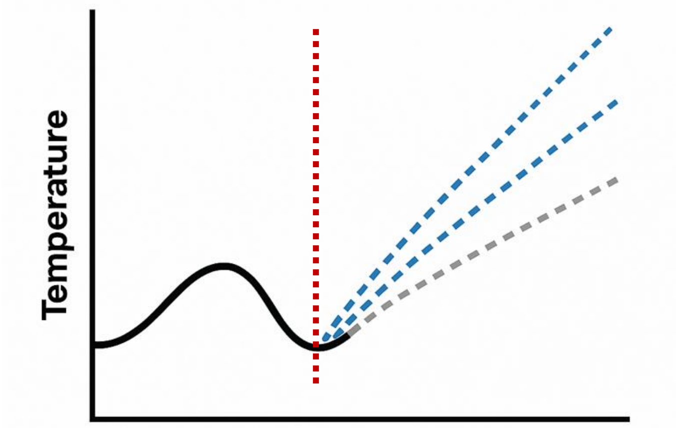

Predicted temperature response in the next control horizon given certain actions

→

Optimal dispatch of heating/cooling power

PHYSICS-INFORMED PREDICTIVE MODELS

- White-box models better describe the physics but usually less accurate
- Black-box models are more capable of fitting data but less reliable in extrapolation
- Hybrid models to combine the advantages of both sides
- Emerging trend of physics-informed
  - Still based on data-driven models, but integrating physics-informed constraints
  - More physics-consistent than black-box and easier to learn than white/gray-box

Black-box
- ANN
- Long Short-Term Memory (LSTM)

→

Physics-informed learning
- Physics-informed Neural Network (PiNN)
- Physics Consistent Neural Network (PCNN)

→

Gray-box
- Resistor-capacitor (RC)
- calibration

→

White-box
- EnergyPlus
- Modelica

RESEARCH OVERVIEW

- Objectives
  - How does the physics integration improve the robustness of predictive models?
  - How does the change in predictive behavior impact the decision-making process?

- Research design
  - Four models with **incrementally added physics**
  - Three predictive tests using **real-world data**

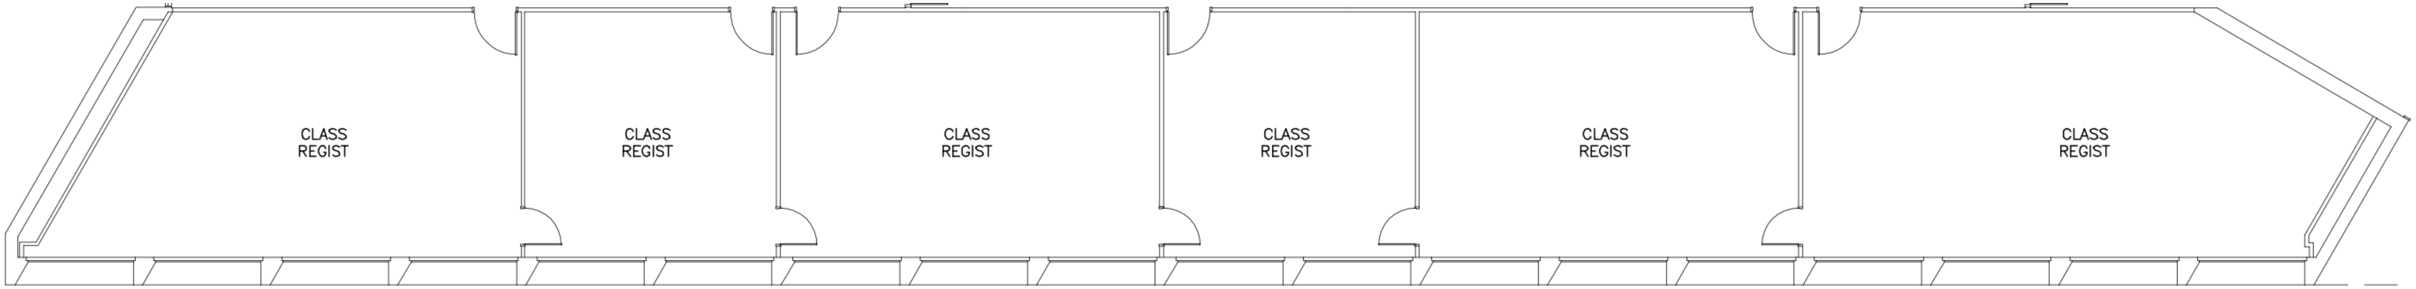

LSTM -> PINN -> PCNN -> RC

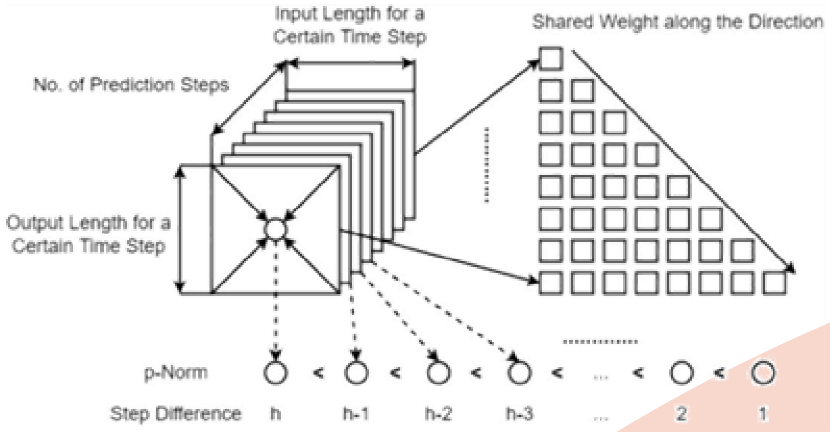

Increasing levels of physics

RC [4]

PCNN [3]

PINN [2]

LSTM

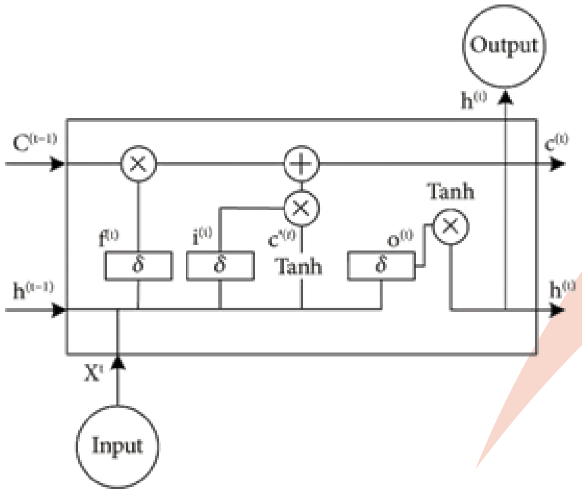

$$\mathcal{L}_{PINN} = \mathcal{L}_{MSE} + \lambda \mathcal{L}_{grad},$$

$$\mathcal{L}_{grad} = \frac{1}{l} \sum_{k=0}^{l-1} \left[\frac{1}{m} \sum_{z=1}^{m} g_k^z \right],$$

$$g_k^z = \sum_{y=1}^{m} ReLU\left(-\frac{\partial \hat{T}_l^z}{\partial u_k^y}\right) + ReLU\left(-\frac{\partial \hat{T}_l^z}{\partial T_k^{out}}\right)$$

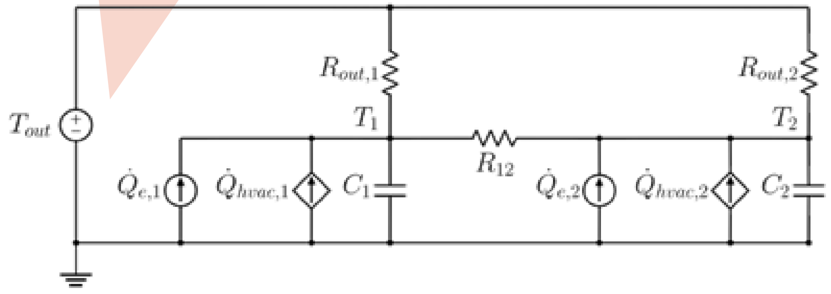

DATA AND TESTING CASES

- Inputs: current temperature, HVAC power, CO_2, outdoor temperature, solar irradiance
- Training and first testing dataset from 2023 winter: standard operation with static operating schedules
- Second testing dataset from 2024 winter: occupant-centric control driven by class schedules[1]

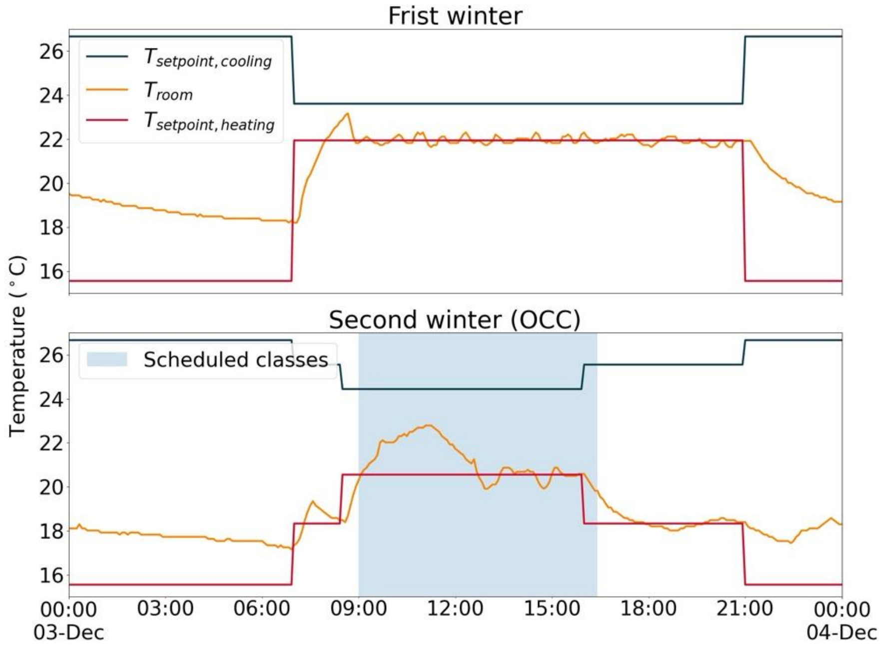

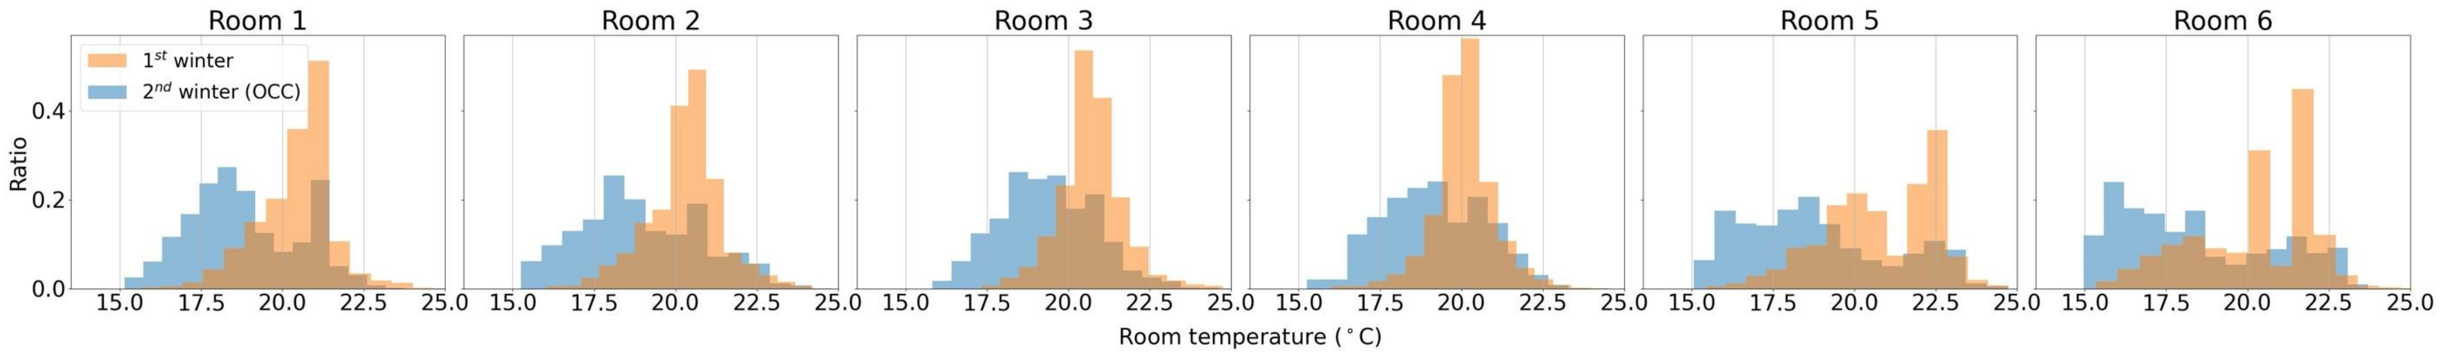

[1] Data from November-December 2024 were collected from experiments conducted by You Lin, MIT LIDS.

24 HOUR OPEN-LOOP PREDICTION

- Predictive accuracy at a glance

Model	Mean absolute errors (°C)								
	Training[a]			1st case (standard)			2nd case (extrapolation)		
	by room	*std*	**mean**	by room	*std*	**mean**	by room	*std*	**mean**
LSTM	0.6, 0.65, 0.77, 0.88, 1.29, 0.79	*0.23*	**0.83**	0.69, 0.7, 0.82, 1.11, 1.25, 0.81	*0.21*	**0.89**	1.04, 0.99, 1.17, 0.88, 2.18, 1	*0.44*	**1.21**
PINN	0.54, 0.82, 0.62, 0.51, 1.44, 1.02	*0.33*	**0.82**	0.66, 0.93, 0.7, 0.76, 1.36, 0.95	*0.23*	**0.89**	0.99, 0.88, 0.95, 0.58, 2.19, 1.13	*0.51*	**1.12**
PCNN	0.7, 0.57, 0.66, 0.56, 1.08, 1.59	*0.37*	**0.86**	0.76, 0.65, 0.74, 0.74, 1.16, 1.22	*0.22*	**0.88**	0.69, 0.92, 0.6, 0.64, 1.38, 1.01	*0.27*	**0.87**
RC	1.46, 1.41, 1.19, 1.16, 1.59, 1.73	*0.2*	**1.42**[b]	1.34, 1.53, 1.31, 1.14, 1.48, 1.59	*0.15*	**1.39**	1.78, 1.97, 1.9, 1.62, 1.6, 1.82	*0.13*	**1.78**

24 HOUR OPEN-LOOP PREDICTION

- LSTM and PINN yielded comparable results in both tests

Model	Mean absolute errors (°C)								
	Training[a]			1st case (standard)			2nd case (extrapolation)		
	by room	*std*	**mean**	by room	*std*	**mean**	by room	*std*	**mean**
LSTM	0.6, 0.65, 0.77, 0.88, 1.29, 0.79	*0.23*	**0.83**	0.69, 0.7, 0.82, 1.11, 1.25, 0.81	*0.21*	**0.89**	1.04, 0.99, 1.17, 0.88, 2.18, 1	*0.44*	**1.21**
PINN	0.54, 0.82, 0.62, 0.51, 1.44, 1.02	*0.33*	**0.82**	0.66, 0.93, 0.7, 0.76, 1.36, 0.95	*0.23*	**0.89**	0.99, 0.88, 0.95, 0.58, 2.19, 1.13	*0.51*	**1.12**
PCNN	0.7, 0.57, 0.66, 0.56, 1.08, 1.59	*0.37*	**0.86**	0.76, 0.65, 0.74, 0.74, 1.16, 1.22	*0.22*	**0.88**	0.69, 0.92, 0.6, 0.64, 1.38, 1.01	*0.27*	**0.87**
RC	1.46, 1.41, 1.19, 1.16, 1.59, 1.73	*0.2*	**1.42**[b]	1.34, 1.53, 1.31, 1.14, 1.48, 1.59	*0.15*	**1.39**	1.78, 1.97, 1.9, 1.62, 1.6, 1.82	*0.13*	**1.78**

24 HOUR OPEN-LOOP PREDICTION

- PCNN had slightly larger fitting error, performed similarly in standard tests, but significantly better in extrapolation

Model	Mean absolute errors (°C)								
	Training[a]			1st case (standard)			2nd case (extrapolation)		
	by room	*std*	**mean**	by room	*std*	**mean**	by room	*std*	**mean**
LSTM	0.6, 0.65, 0.77, 0.88, 1.29, 0.79	*0.23*	**0.83**	0.69, 0.7, 0.82, 1.11, 1.25, 0.81	*0.21*	**0.89**	1.04, 0.99, 1.17, 0.88, 2.18, 1	*0.44*	**1.21**
PINN	0.54, 0.82, 0.62, 0.51, 1.44, 1.02	*0.33*	**0.82**	0.66, 0.93, 0.7, 0.76, 1.36, 0.95	*0.23*	**0.89**	0.99, 0.88, 0.95, 0.58, 2.19, 1.13	*0.51*	**1.12**
PCNN	0.7, 0.57, 0.66, 0.56, 1.08, 1.59	*0.37*	**0.86**	0.76, 0.65, 0.74, 0.74, 1.16, 1.22	*0.22*	**0.88**	0.69, 0.92, 0.6, 0.64, 1.38, 1.01	*0.27*	**0.87**
RC	1.46, 1.41, 1.19, 1.16, 1.59, 1.73	*0.2*	**1.42**[b]	1.34, 1.53, 1.31, 1.14, 1.48, 1.59	*0.15*	**1.39**	1.78, 1.97, 1.9, 1.62, 1.6, 1.82	*0.13*	**1.78**

24 HOUR OPEN-LOOP PREDICTION

- RC model has too simplified model structure
- Minimized the predictive error with smoother profiles
- Robustly captured the (slow) free-floating temperature
- But couldn't (quickly) respond to HVAC inputs and disturbances

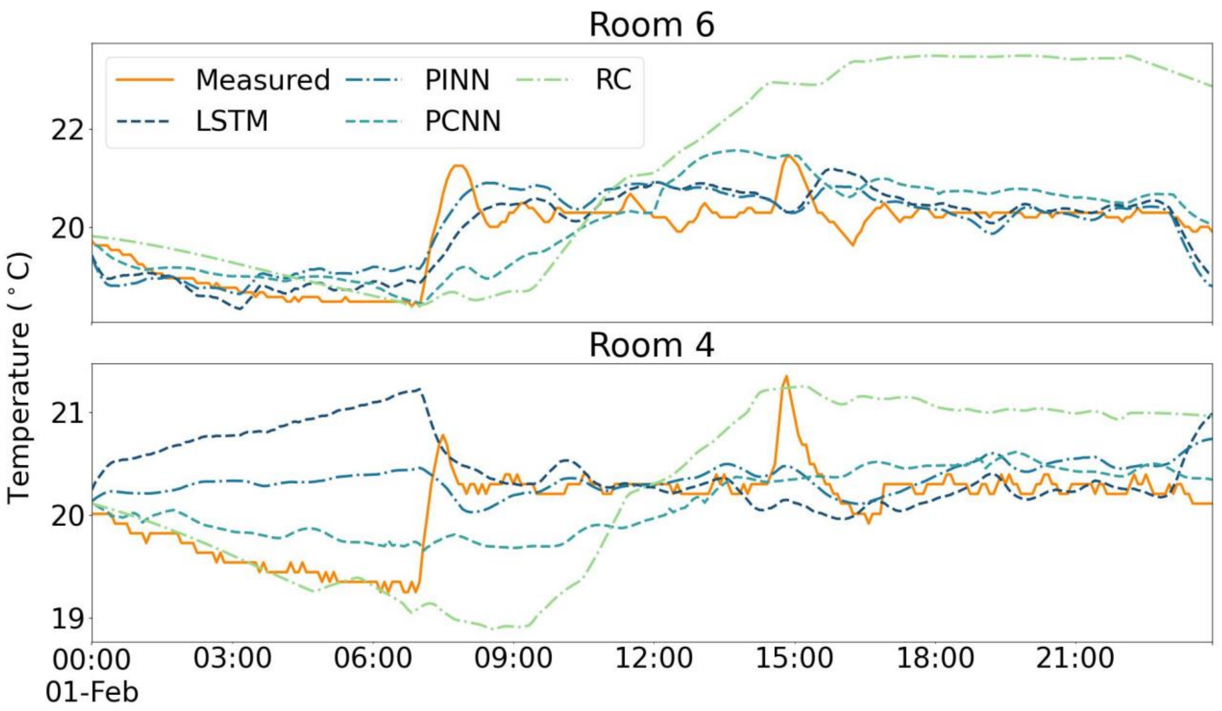

INPUT INTERPRETATION -> PHYSICS CONSISTENCY

- LSTM overestimated the impact of slightly increased CO_2 level
- Projected the room temperature to increase when it should decrease (0°C outside)
- PINN and PCNN with more physics gradually improved the situation
- Change in physics consistency not reflected in long-term predictive errors

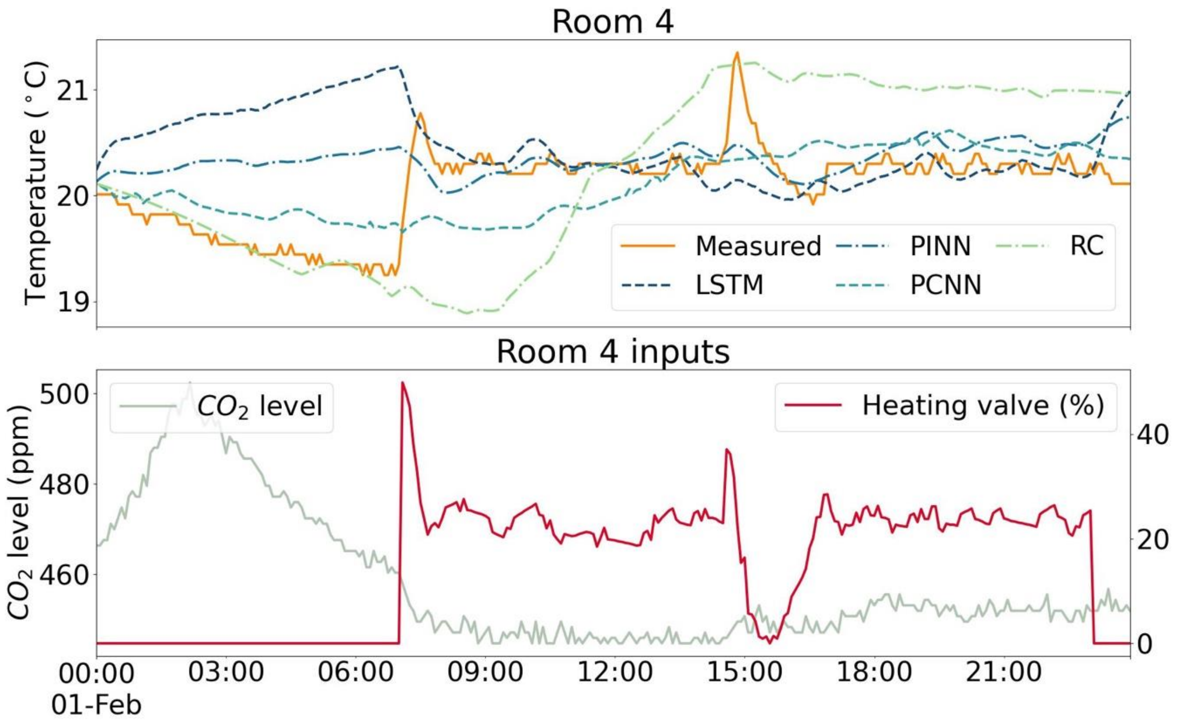

CHALLENGES IN EXTRAPOLATION

- LSTM and PINN wrongly predicted abrupt changes, affected by irrelevant inputs
- PCNN and RC were more consistent
- The challenges came from the change in operational patterns rather than the temperature range

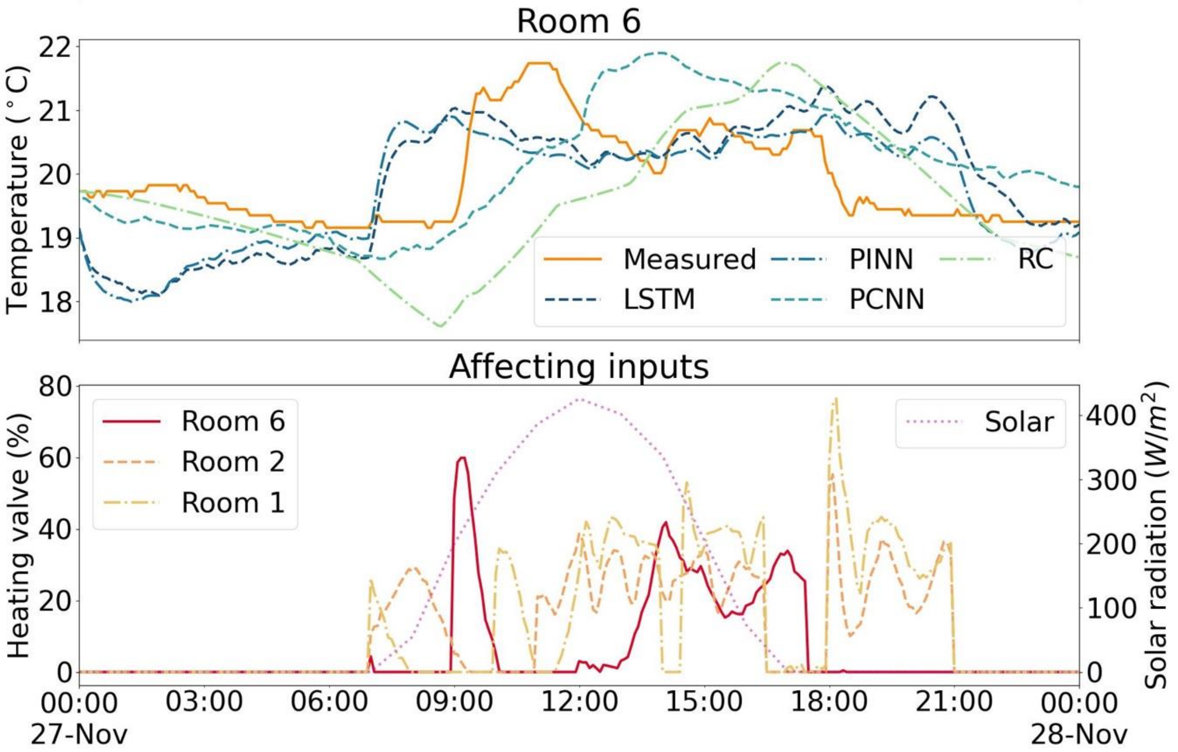

THE CORNER CASE

- Temperature spike caused by an unexpected large group around noon
- Thermal-response models are required to foresee the change in temperature given occupancy signals

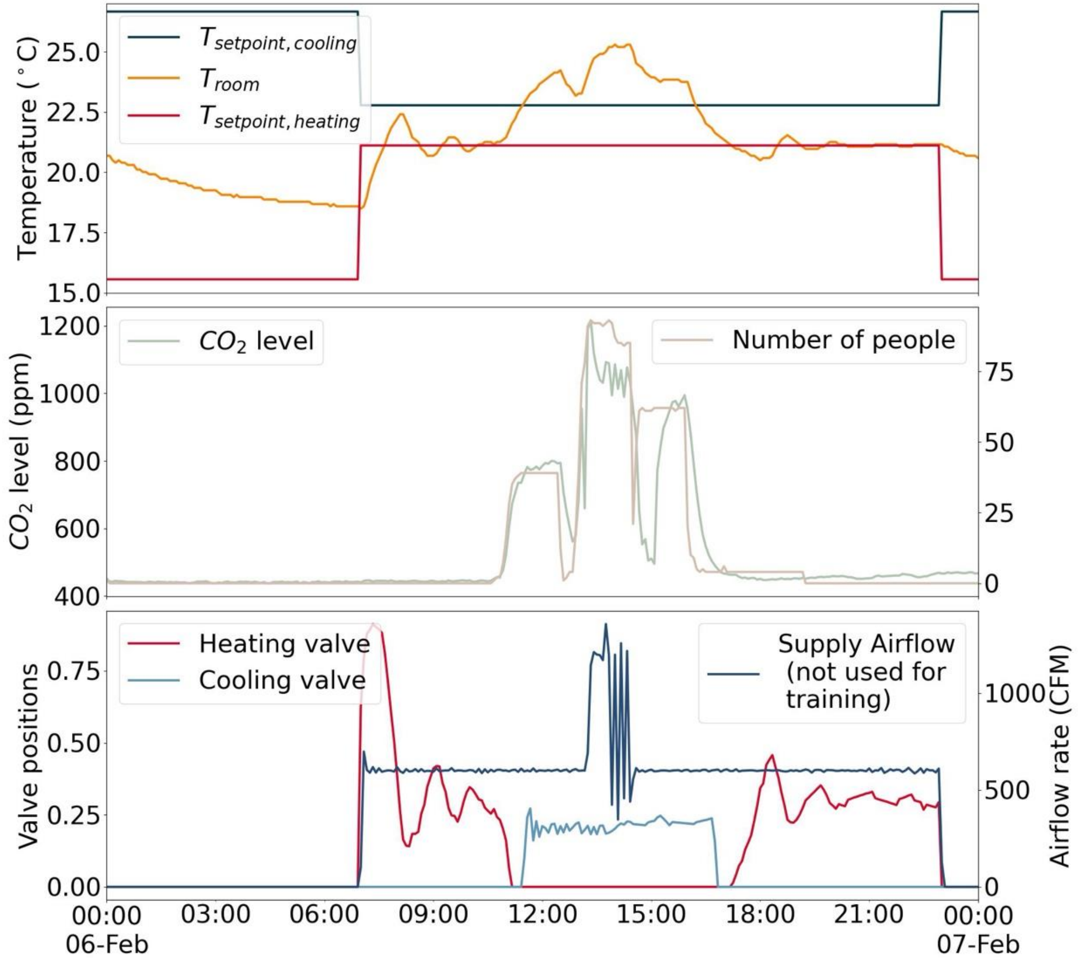

THE CORNER CASE

- None of the model could capture the spike in 24-hour ahead predictions
- Many applications could be based on shorter prediction horizon -> 3-hour ahead prediction
- PCNN captured the spike with a closer starting point, which could inform control precaution

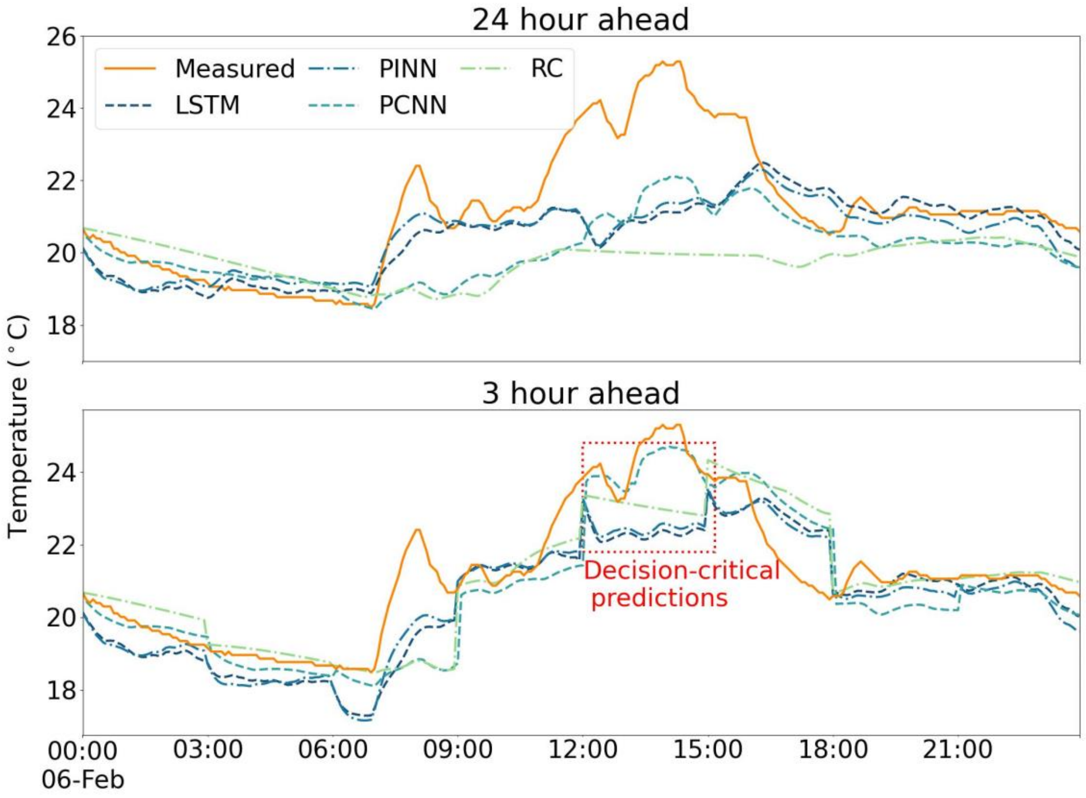

SUMMARY

- Real-world implementation requires more attention for control robustness and corner cases

- Physics integration could improve model robustness, but requires careful balance between physical constraints and model expressiveness

- A models’ deficiency in extrapolation and corner cases could be disguised by lower errors in standard tests, calling for more comprehensive evaluation framework

- Future work to quantitatively translate the predictive performance to control performance

THANK YOU!

References

[1] Royapoor, M., Antony, A., & Roskilly, T. (2018). A review of building climate and plant controls, and a survey of industry perspectives. Energy and Buildings, 158, 453-465.

[2] Di Natale, L., Svetozarevic, B., Heer, P., & Jones, C. N. (2023). Towards scalable physically consistent neural networks: An application to data-driven multi-zone thermal building models. Applied Energy, 340, 121071.

[3] Wang, X., & Dong, B. (2023). Physics-informed hierarchical data-driven predictive control for building HVAC systems to achieve energy and health nexus. Energy and Buildings, 291, 113088.

[4] Green, D. H., Lin, Y., Botterud, A., Gregory, J., Leeb, S. B., & Norford, L. K. (2024). Pareto-optimized thermal control of multi-zone buildings using limited sensor measurements. IEEE Transactions on Smart Grid, 15(5), 4674-4689.