

A meta-learning approach to enable autonomous buildings

PI: Panagiota Karava; Co-PI's: Ilias Bilonis, Jianghai Hu

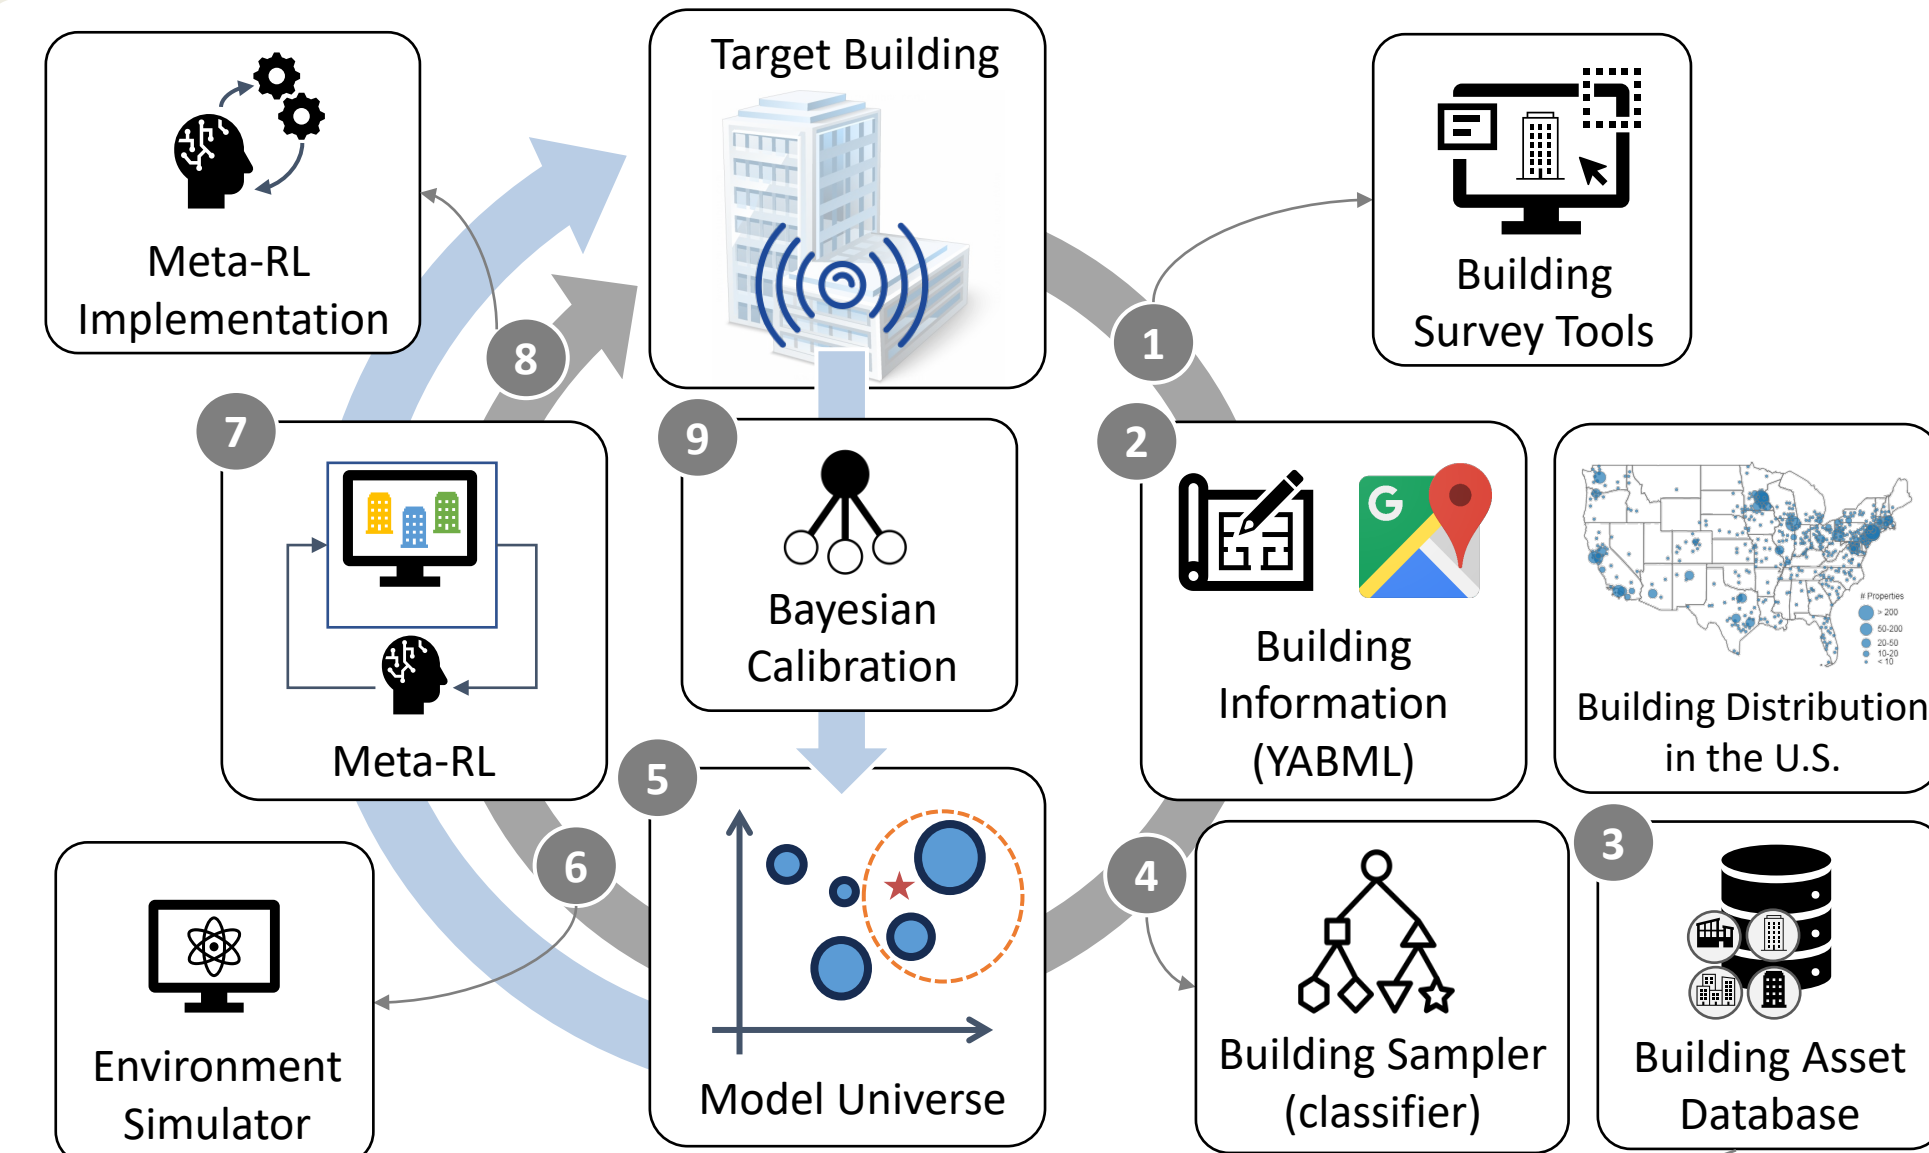
Students: Ting-Chun Kuo, Sreehari Manikkan, Hemanth Devarapalli

Challenge

Intelligent building controls require engineering customized, site-specific and costly solutions

Solution

Develop the AI-Enabled Building Energy Expert (AI-BEE) integrates data-driven and model-based learning and fuses different sources of information to automate the discovery of optimal policies for building control



- 1 Survey Tool
- 2 Information Schema
- 3 Asset Database
- 4 Prior Knowledge Sampler
- 5 Model Universe
- 6 Environment Simulator
- 7 Meta-Reinforcement Learning (Meta-RL)
- 8 Meta-RL Implementation
- 9 Bayesian Update of Model Universe

Scientific Impact

- Novel TD-based meta-RL algorithm for smaller-scale control problems
- Novel gradient-based meta-RL algorithm for complex control problems
- Automatic generation of differentiable building models sampled from the model universe
- Automated Bayesian update of model structure and parameters using building data

Broader Impacts

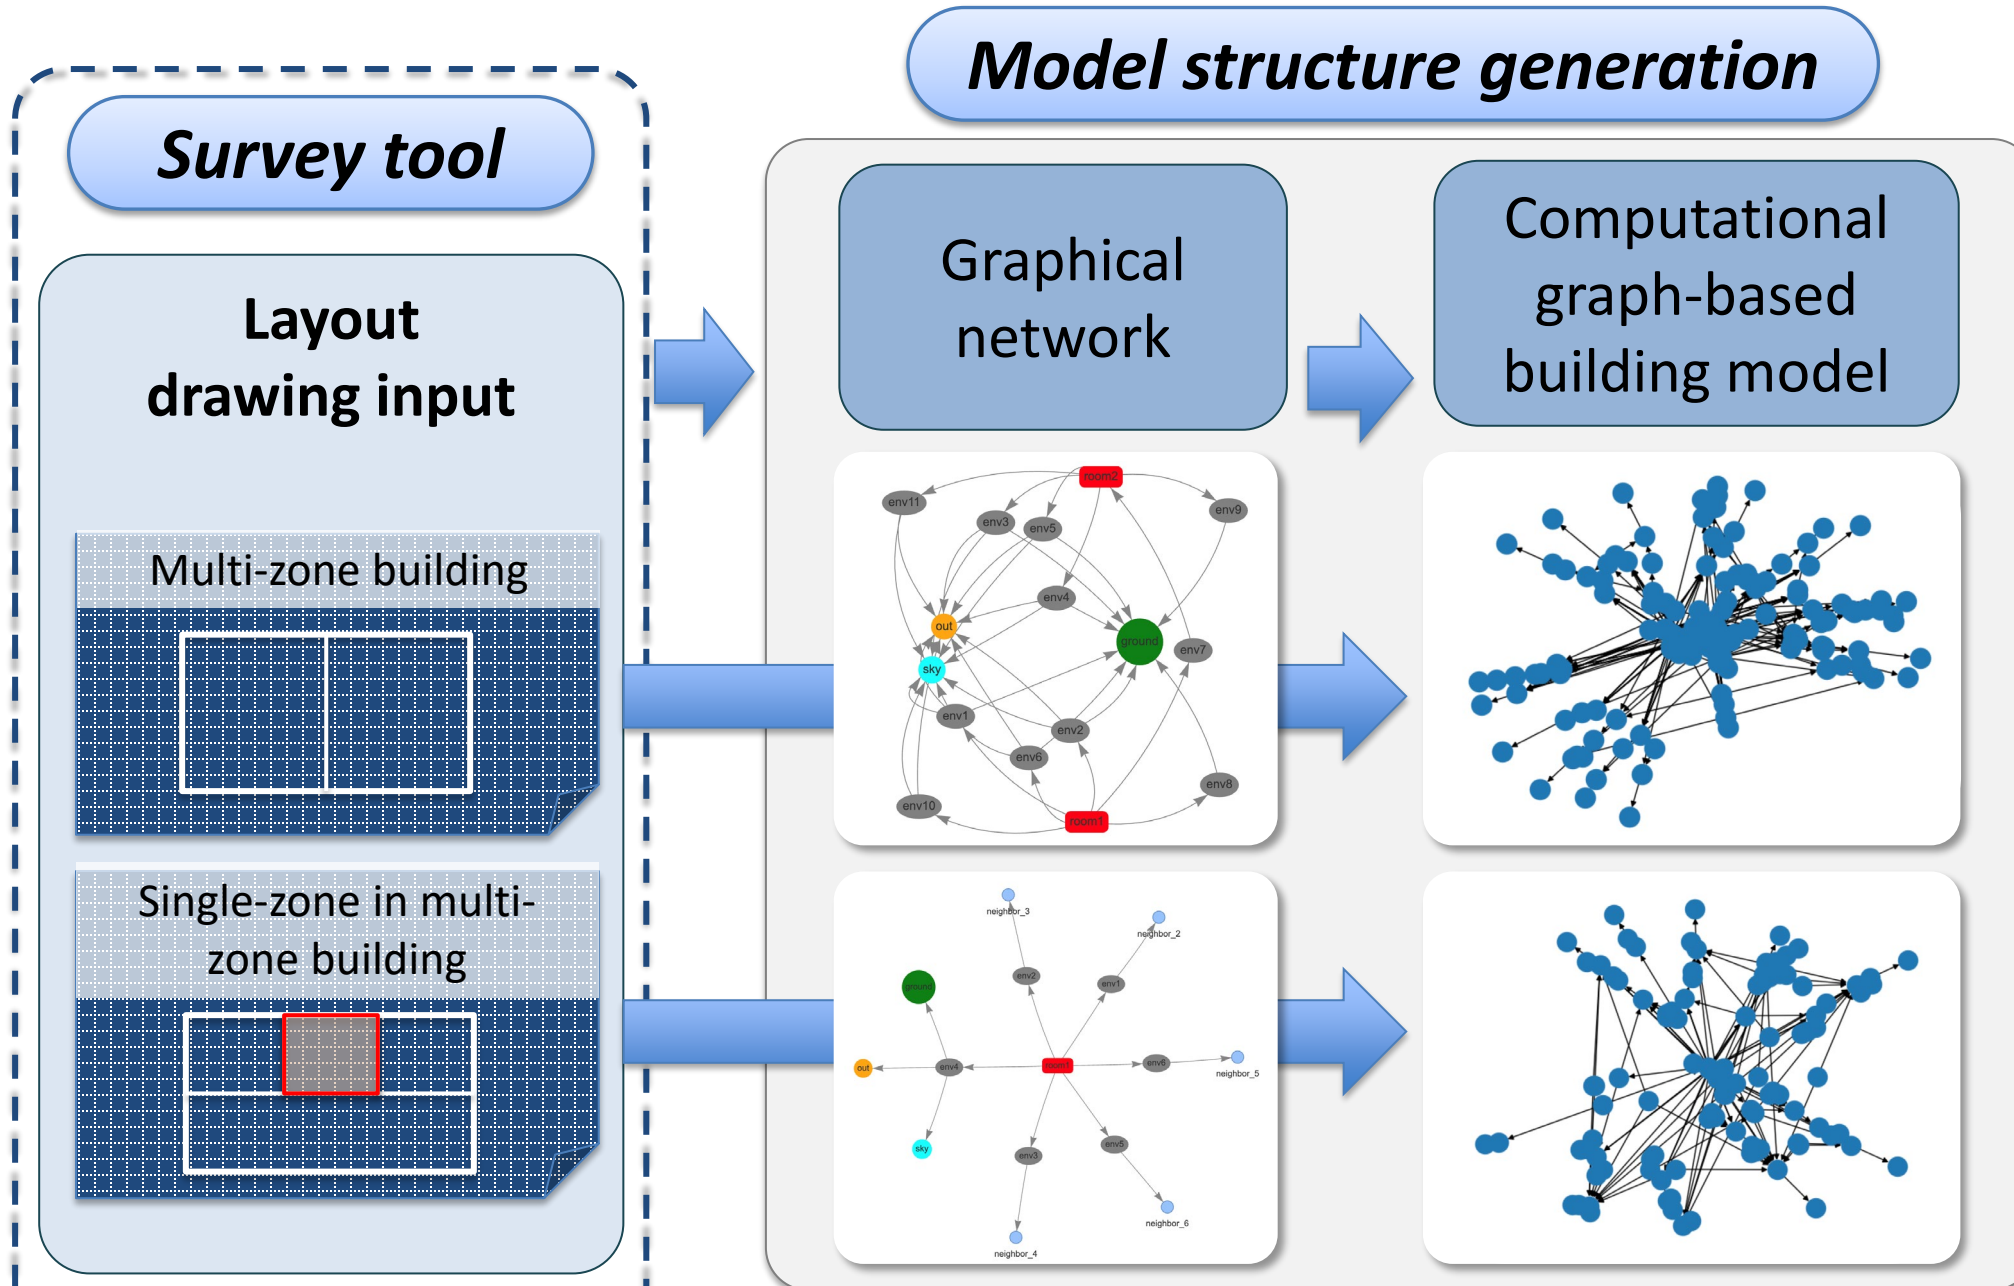
- CPS solution for autonomous buildings enabling deployment of asset-specific smart control policies by non-experts
- Field deployment and outreach through Purdue's Industry Consortium on Center for High Performance Buildings

Key Innovations

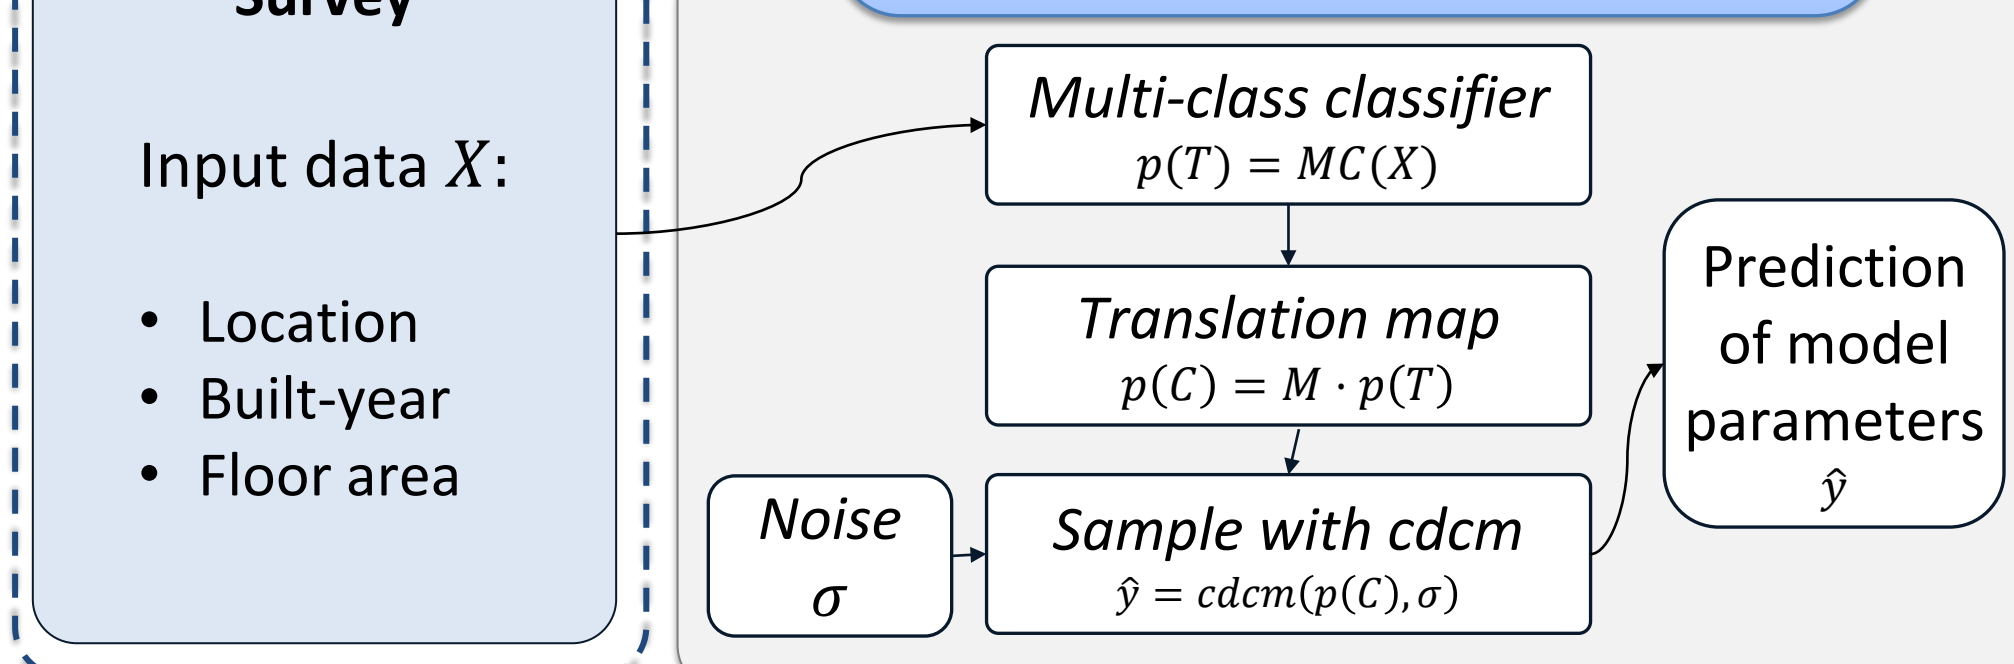
AI-BEE Generate Model Universe

Scalability in Model Structure

Generate model structure for any building layout



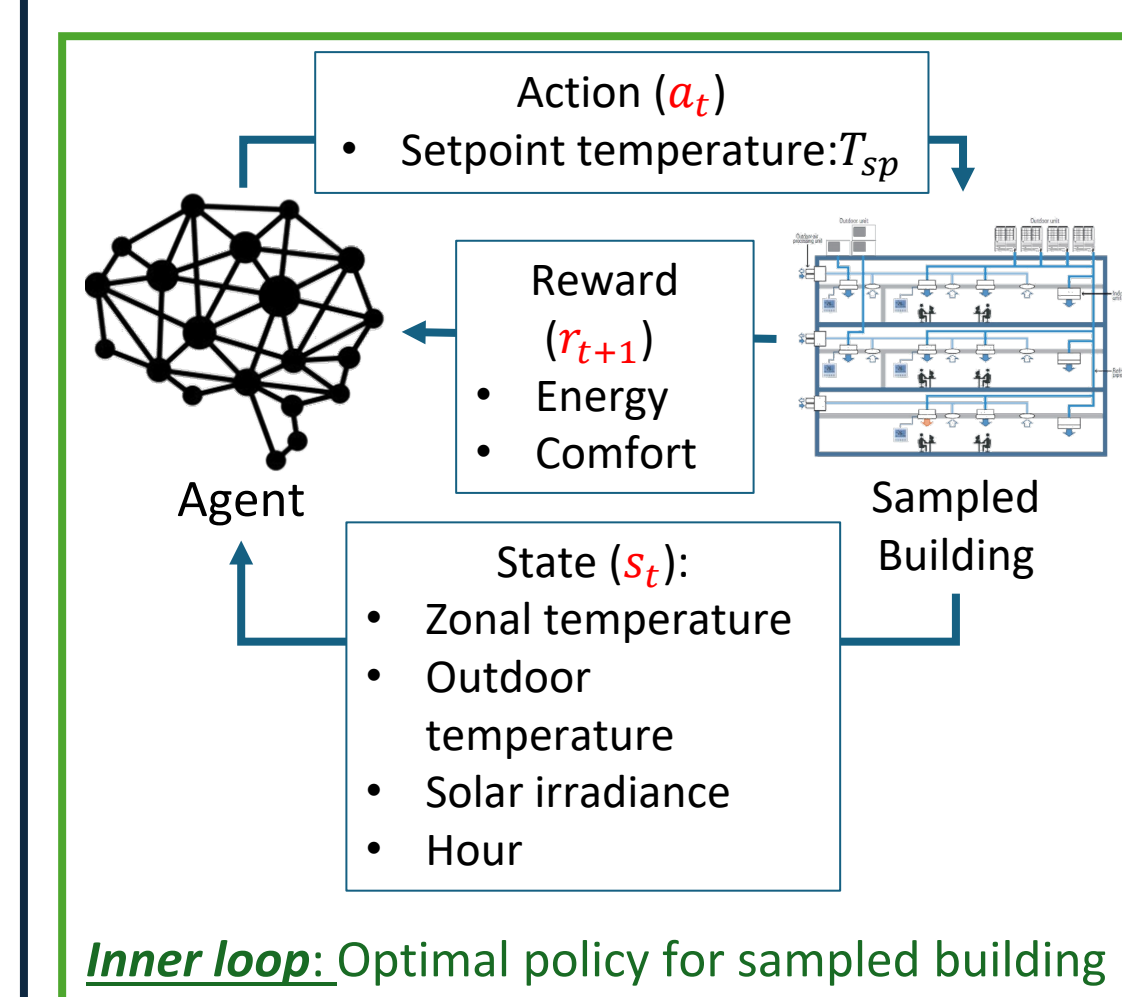
Quantification of model parameter uncertainty



Meta-Reinforcement Learning

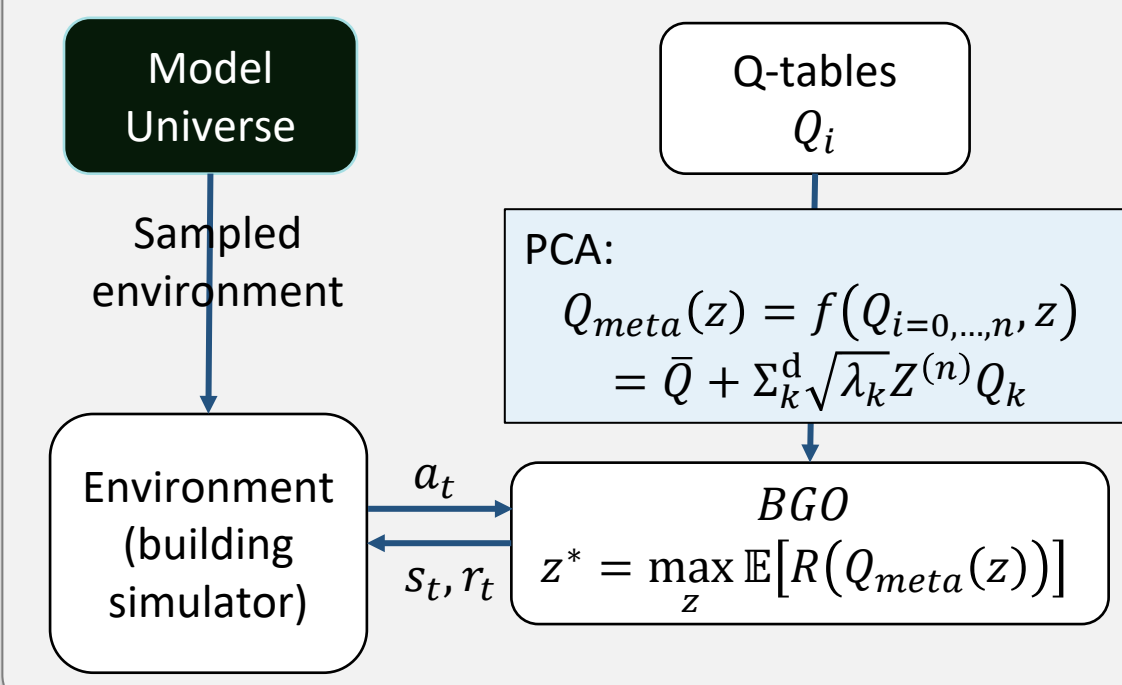
Building Control Task Formulation

Outer loop: Learns a RL-algorithm optimized for the model universe



Algorithm 1: TD-based Meta-RL

- Inner loop: tabular Q-learning
- More computationally efficient
- Low-dimensional control problems



Algorithm 2: Gradient-based Meta-RL

- Inner loop: policy gradient approach
- High-dimensional control problems
- More computationally expensive

Algorithm 2: Adapted version of MAML for Reinforcement Learning

Setup $p(\mathcal{M})$: distribution over tasks (i.e. Model Universe)

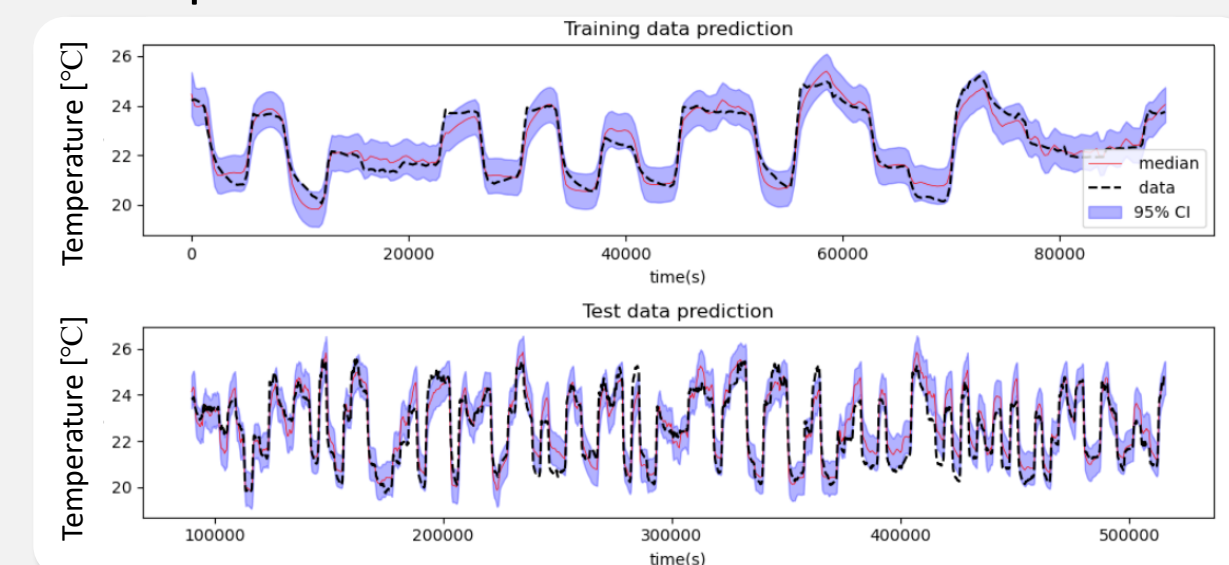
Setup hyperparameters: $\omega = \theta_0$;

external reward function $R(\mathcal{M}_i, f_{\theta}, \omega)$

- 1: **External Loop:**
- 2: **while** not done **do**
- 3: Randomly initialize $\theta = \theta_0$
- 4: **Internal Loop:**
- 5: **while** not done **do**
- 6: Sample batch of tasks $\mathcal{M}_i \sim p(\mathcal{M})$
- 7: **for all** \mathcal{M}_i **do**
- 8: Sample K trajectories $\mathcal{D} = \{(x_1, a_1, \dots, x_H)\}$ using f_{θ} in \mathcal{M}_i
- 9: Optimize f_{θ} with gradient descent $\theta'_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{M}_i}(f_{\theta})$
- 10: **end for**
- 11: **Update** $\omega' \leftarrow \omega + \alpha \frac{\partial \mathbb{E}_{\mathcal{M}_i \sim p(\mathcal{M})} [R_{\mathcal{M}_i}]}{\partial \omega}$

Differentiable building models

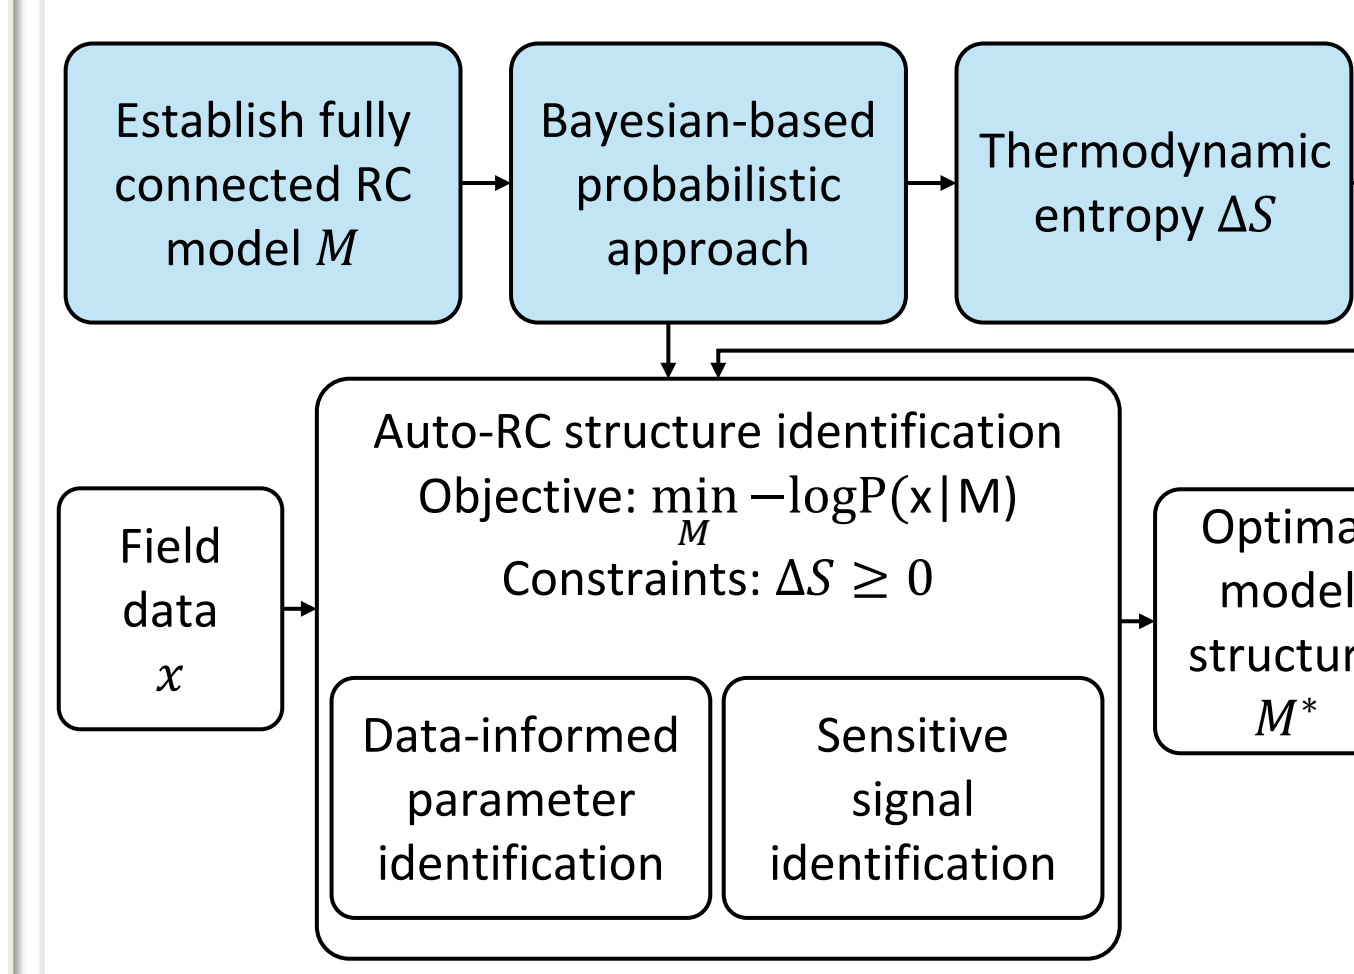
- Gradient-based meta-RL requires a differentiable reward function
- JAX-based implementation of building model
- Automatic differentiation enables Bayesian updates



Field Study

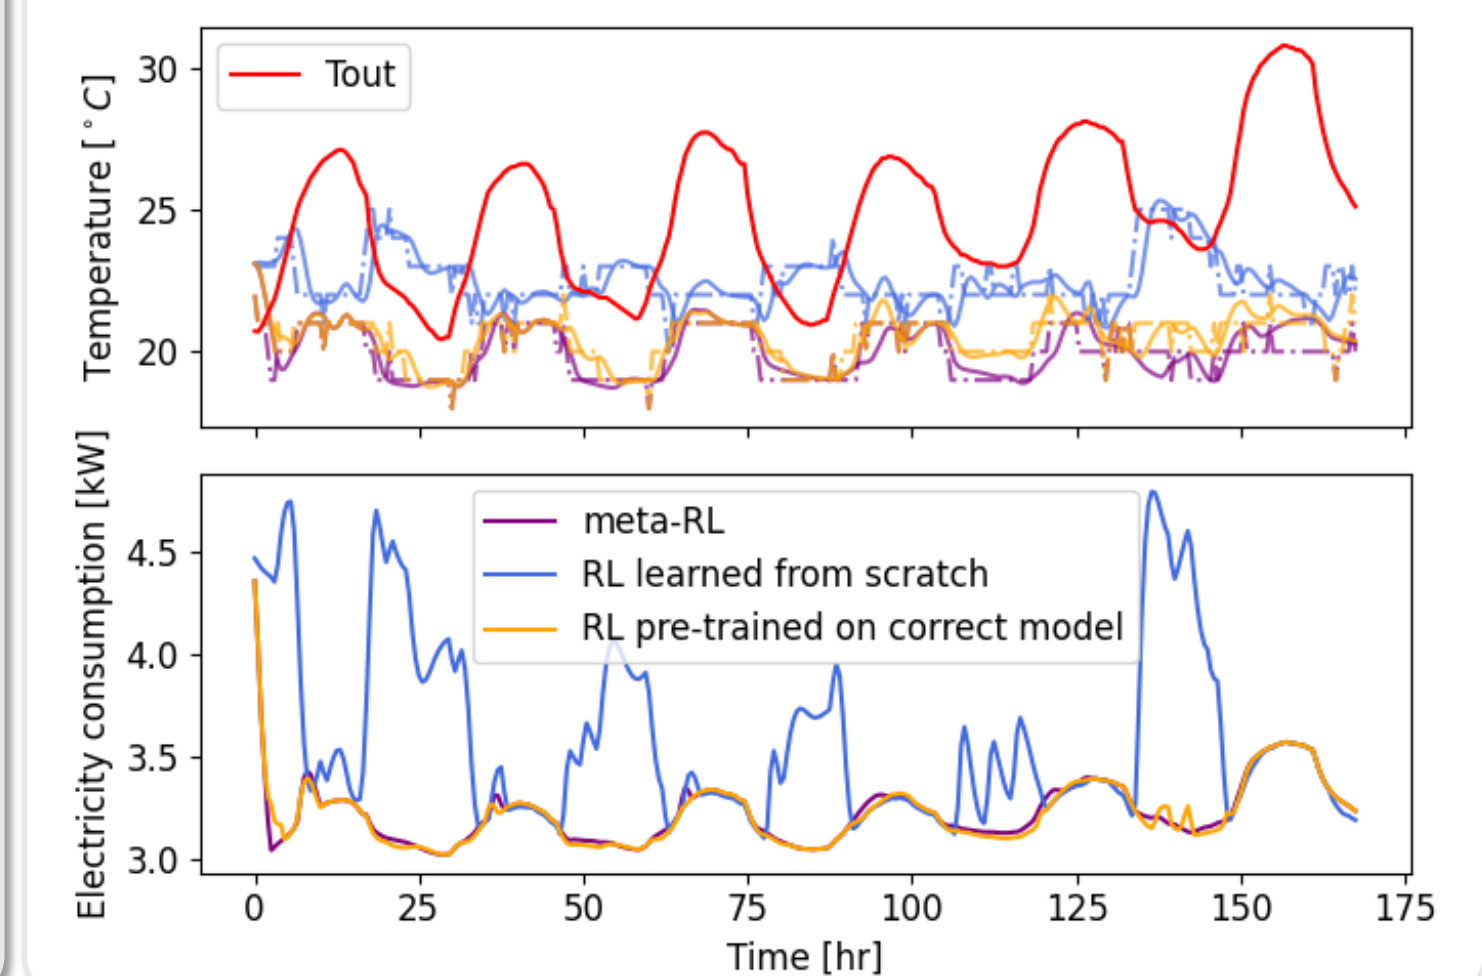
Model Update

- Model structure selection via field data



Simulation results

- **Meta-RL** outperforms **RL learned from scratch** and performs close to **RL pre-trained on correct model**



Future Goal

Field implementation of autonomous control



- Deploy and validate meta-RL on a real building
- Verification of Bayesian update using synthetic building data
- Validation of Bayesian updates in a real building
- Couple Bayesian updates of the model universe with meta-RL (referred to as BUMRL)
- Verify BUMRL using synthetic building data
- Deploy and validate BUMRL