ClimAIte Control

Sebastian Fischer Stripp

P5 Presentation

Supervisors: Dr. Michela Turrin & Dr. Regina Bokel

External Examiner: Dr. Olindo Caso

Climate Change

Literature

RQ & Method

Case Study

EnergyPlus API

DRL Agents

Results

Climate Change

Building Sector

EU Greenhouse Gas

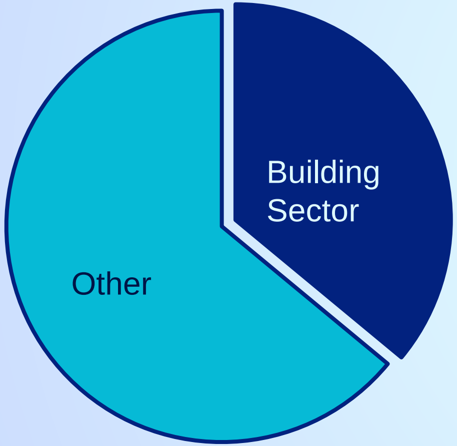

UK Greenhouse Gas

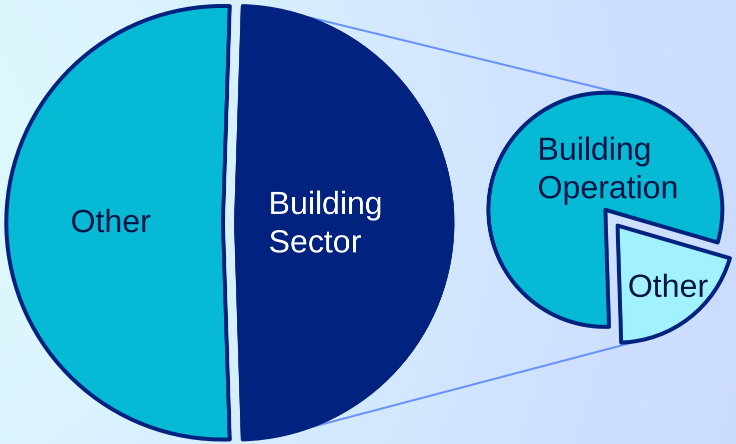

1%

Improvement

1% Change = Impact x Reach

1% Change = Impact x Reach

1% Change = Impact x Reach

Impact x Reach

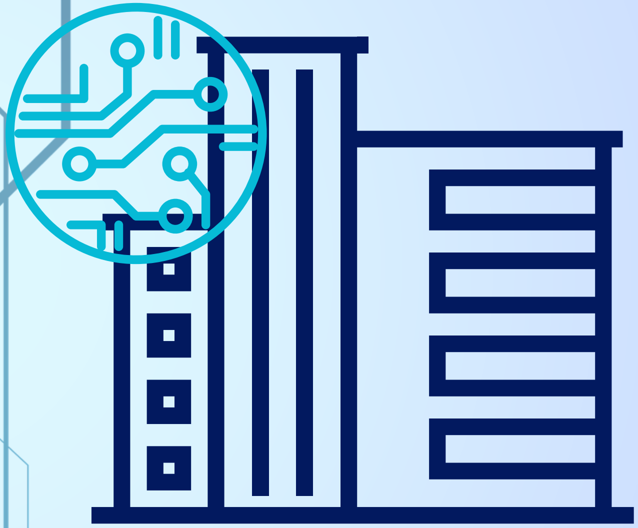

Impact x Reach

Background & Literature

Building as Energy Balance System

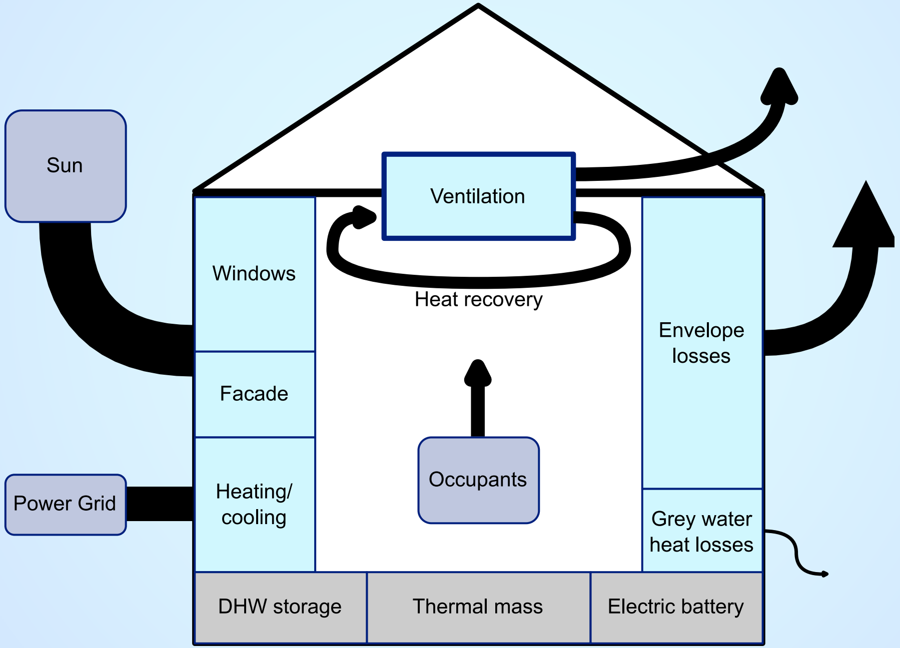

Building as Energy Balance System

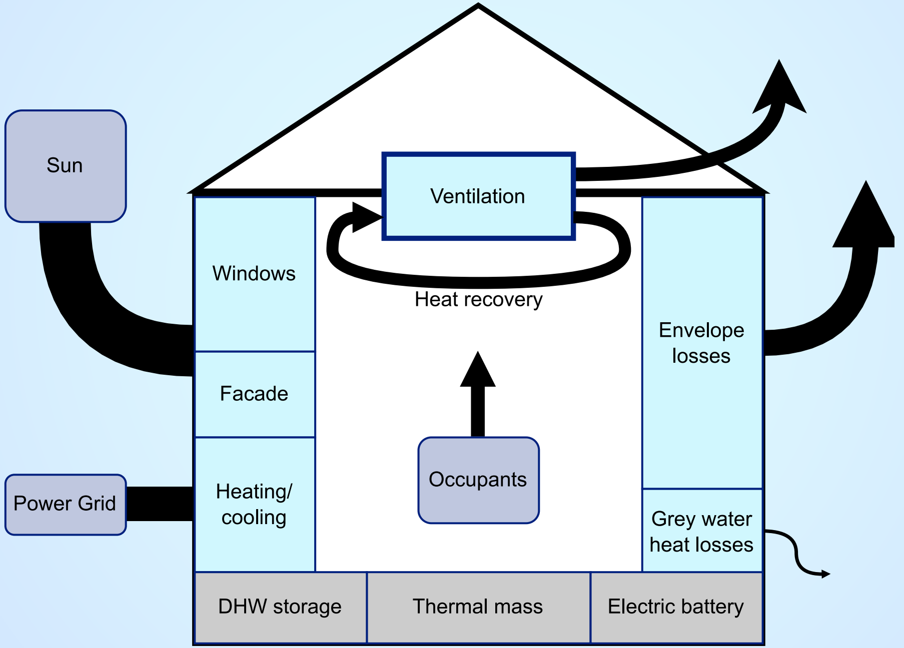

AI & Machine Learning

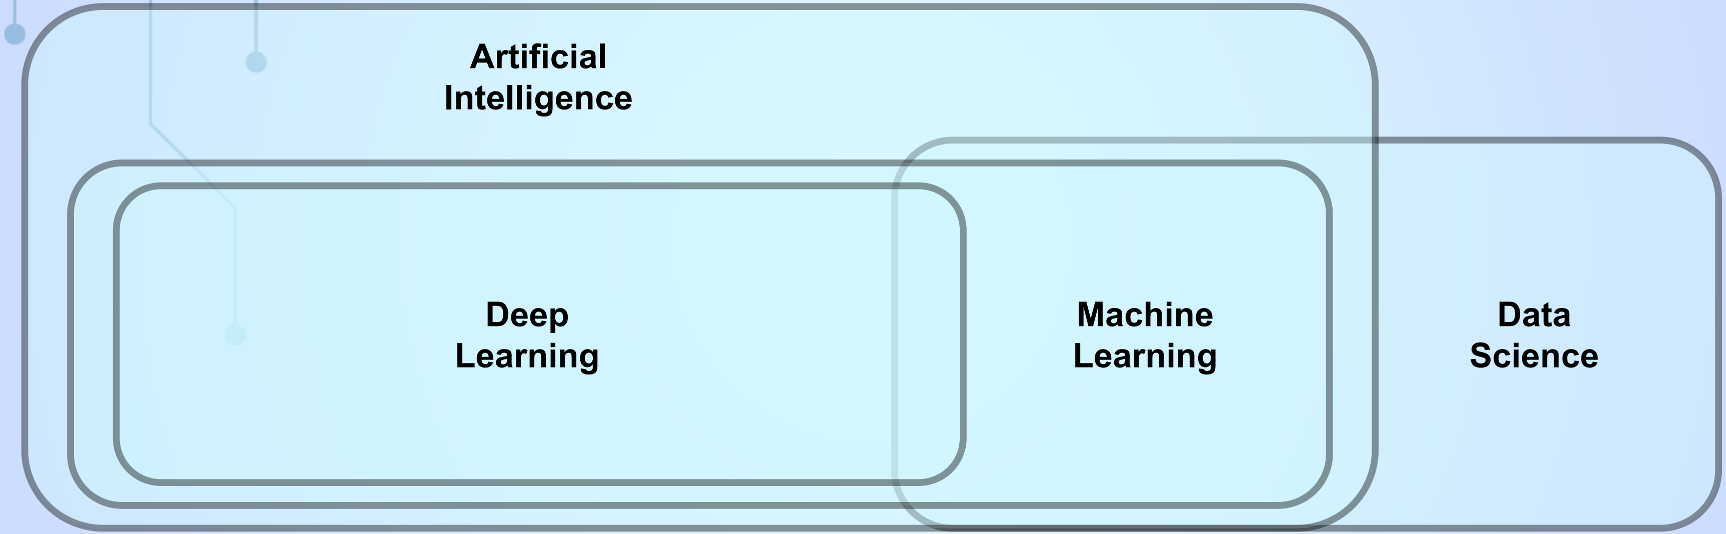

Reinforcement Learning

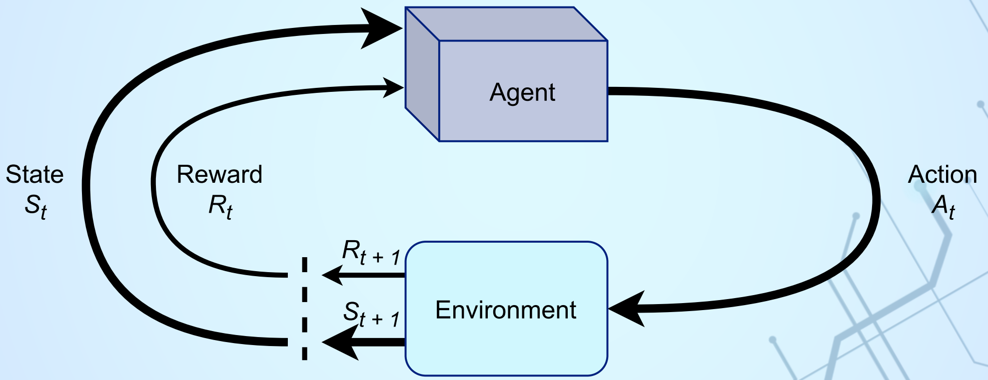

Climate Change | **Literature** | RQ & Method | Case Study | EnergyPlus API | DRL Agents | Results

Deep Q-Learning

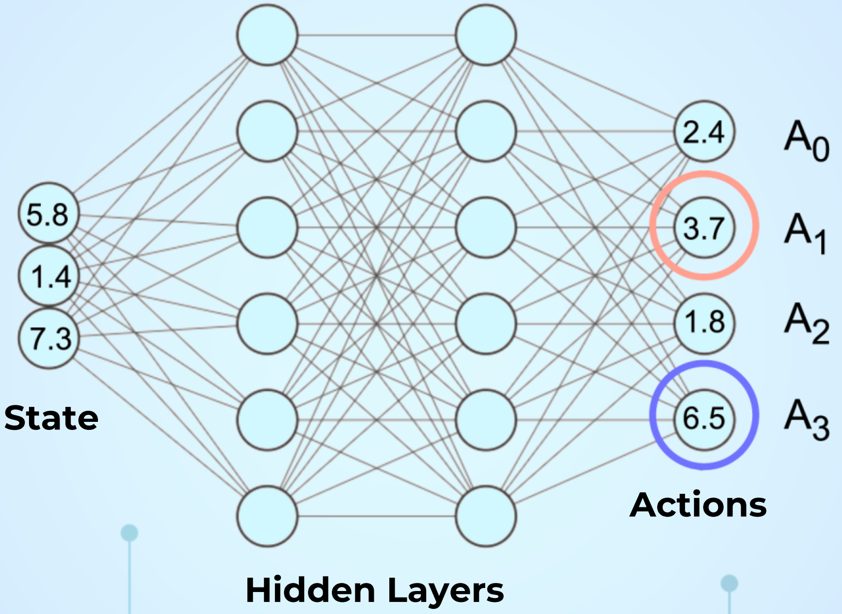

Model Predictive Control

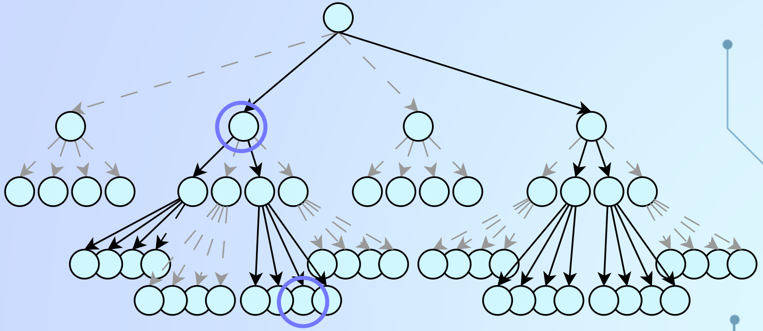

Explicitly programmed model

Searches for optimum action

Expertise-based

Deep Reinforcement Learning

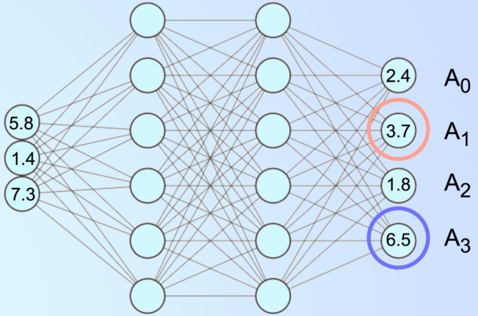

Generalist model

Learns optimum action

Data-based

Literature – Smart Operation

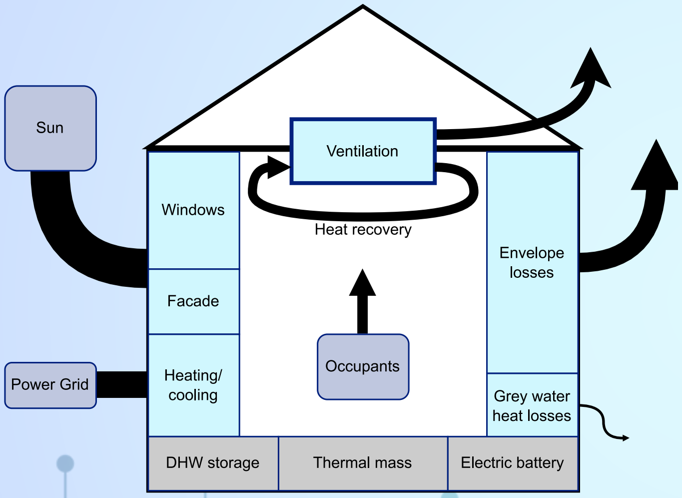

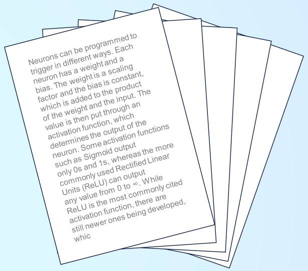

Energy Storage

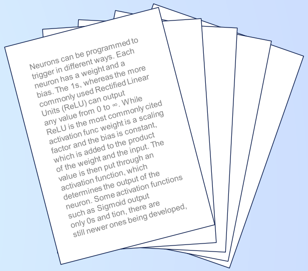

Energy Control

Literature – Storage & Control

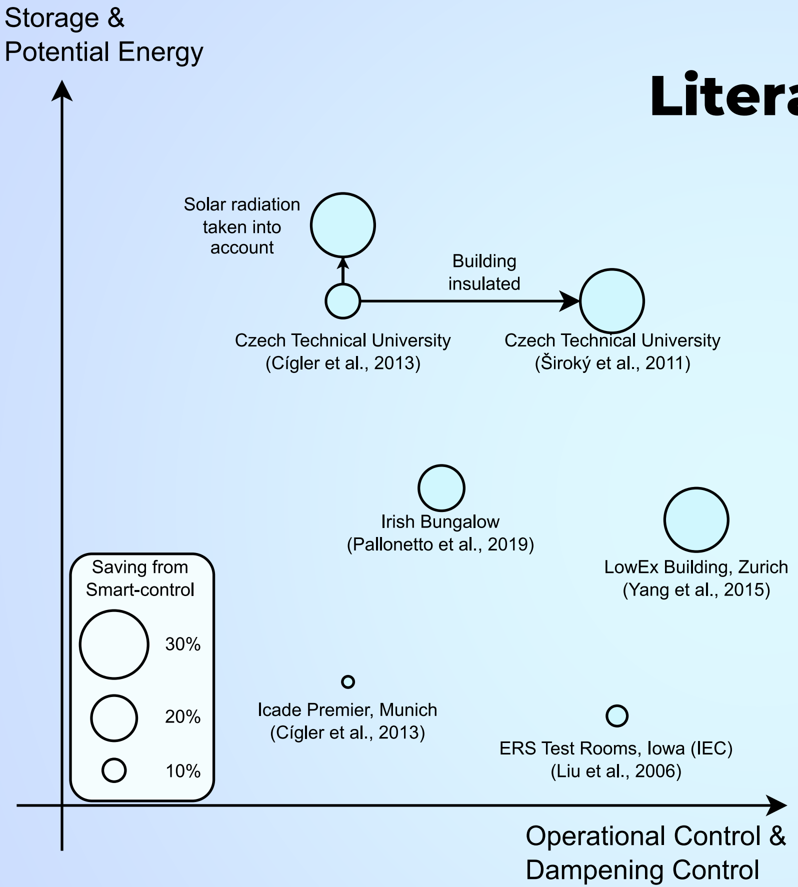

Research Question & Methodology

*“To what extent do **Reinforcement Learning control** algorithms compared with **building fabric factors** influence a building’s **energy-saving** potential?“*

Methodology

Building, Digital Twin

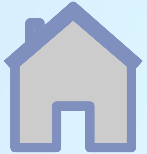

Control Algorithm

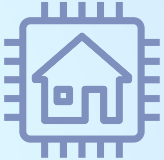

Methodology

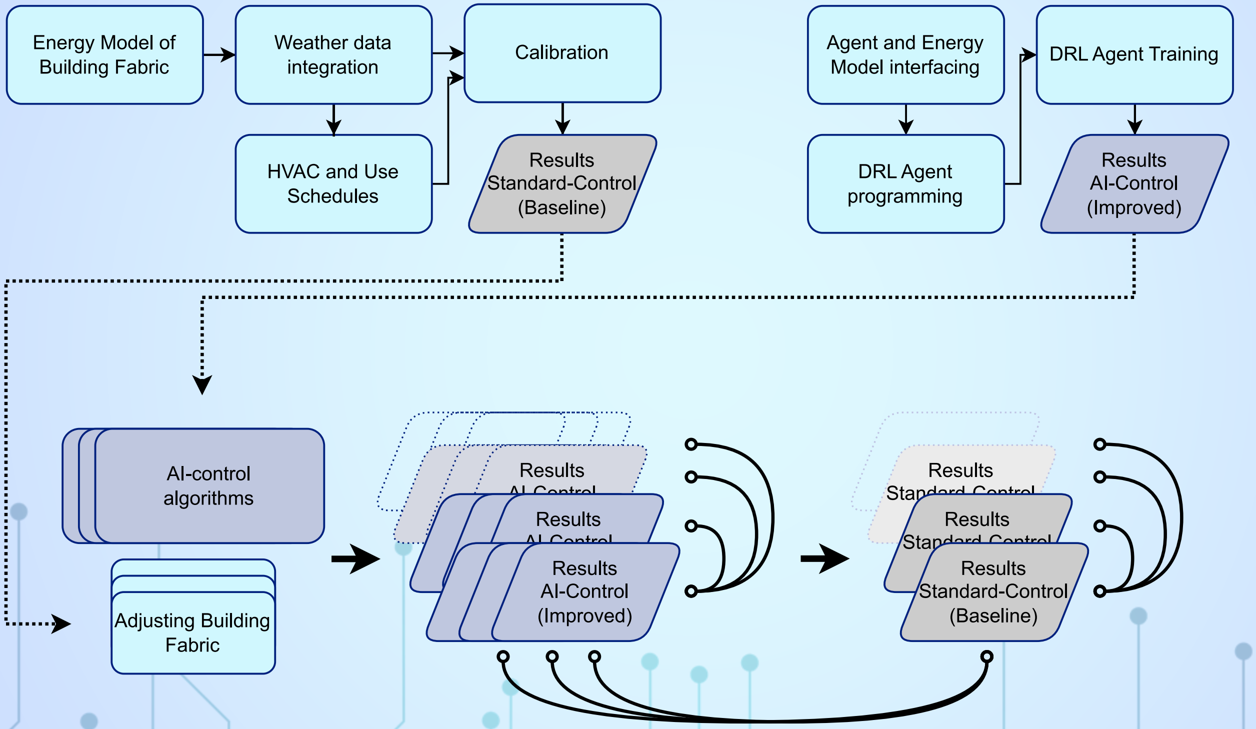

Methodology

	Algorithm 1	Algorithm 2	Algorithm 3
Building 1			
Building 2			
Building 3			

Case Study

Primary School
Wales, UK

Welsh School

Gross area: **2,480 m^2**
Treated floor area: **1,992 m^2**

Passivhaus
heating demand: **13 kWh $m^{-1}a^{-1}$** (TFA)

Airtightness: **0.25 ACH^{-1}** N50
U-value: **0.1 $Wm^{-2}K^{-1}$**

14 classrooms
Occupancy: **455**
Ages 4 - 11

Aerial photo by Powys County Council

Welsh School

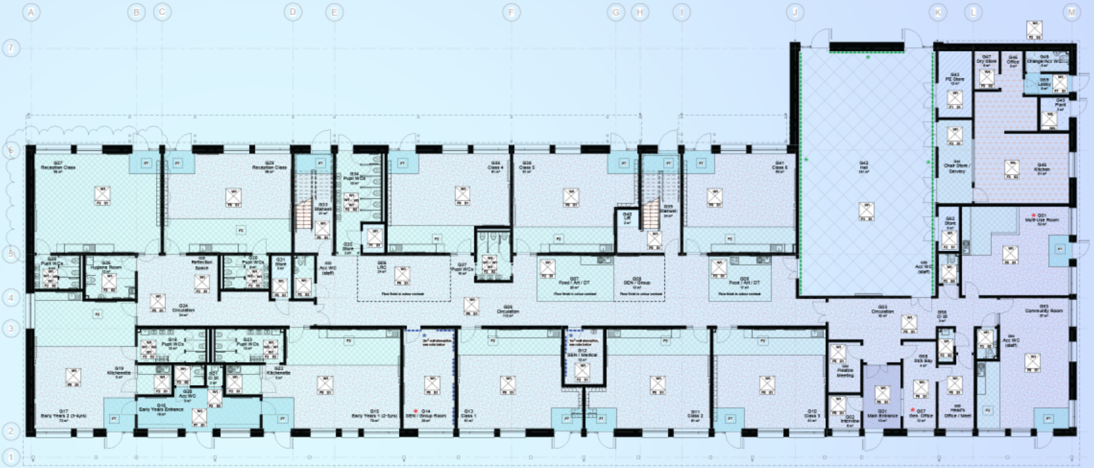

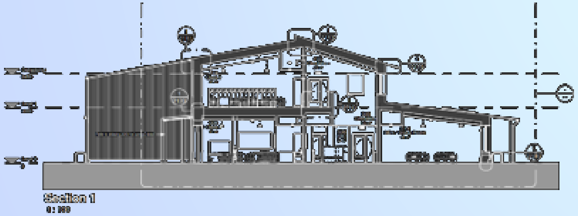

Information Flow

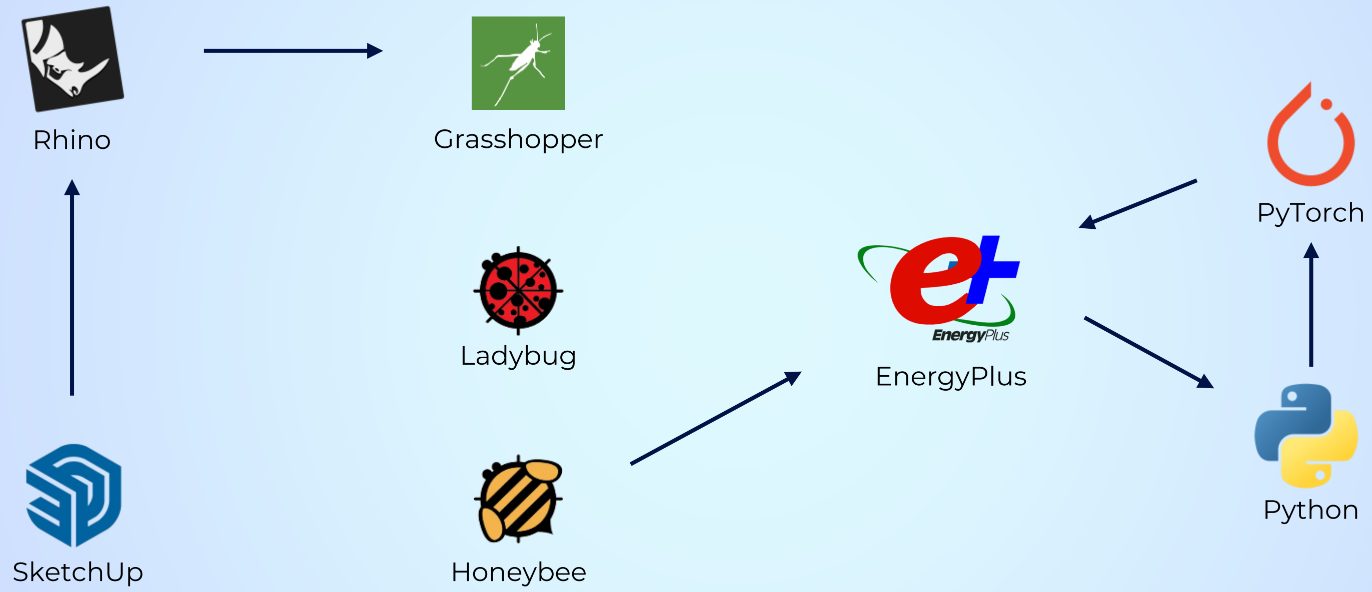

3D Modelling **Digital Twin – Energy Model** **Reinforcement Learning**

Fabric - Digital Twin

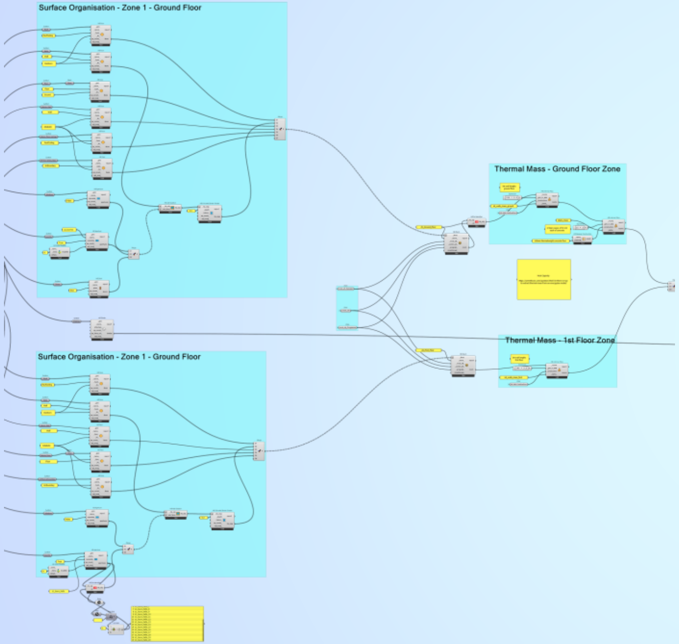

Building fabric definition

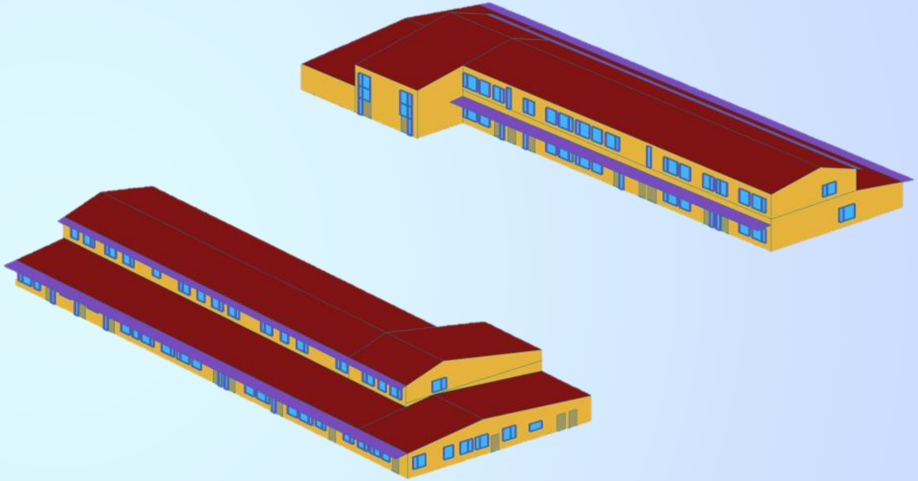

Building fabric visualisation

Programme - Digital Twin

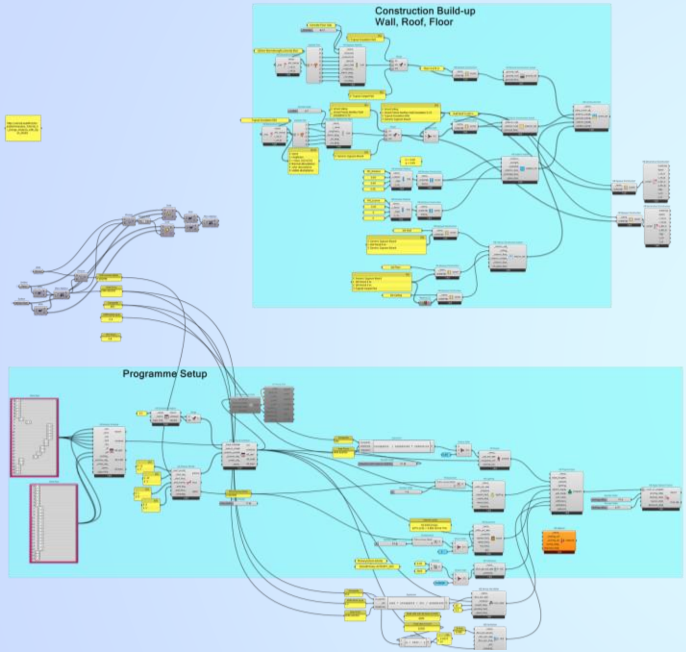

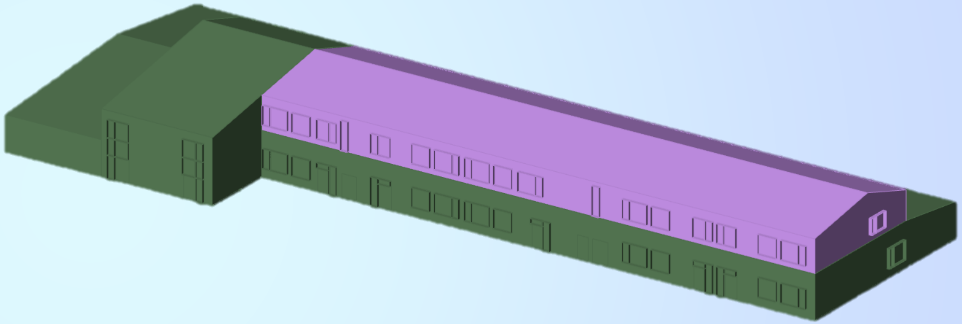

Programme definition

Building Zoning

Temperature Monitoring Data

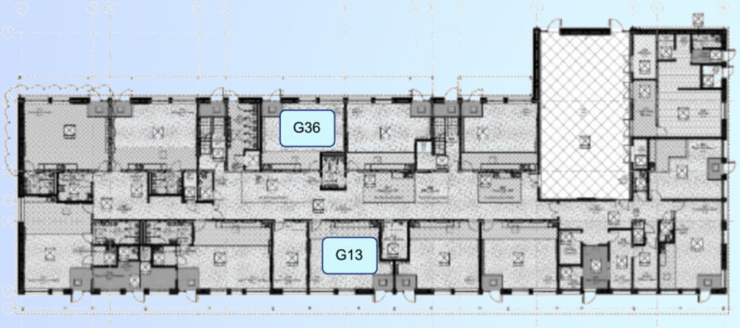

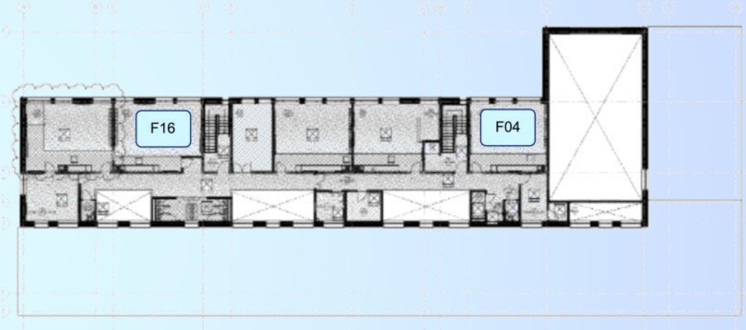

Validation & Calibration

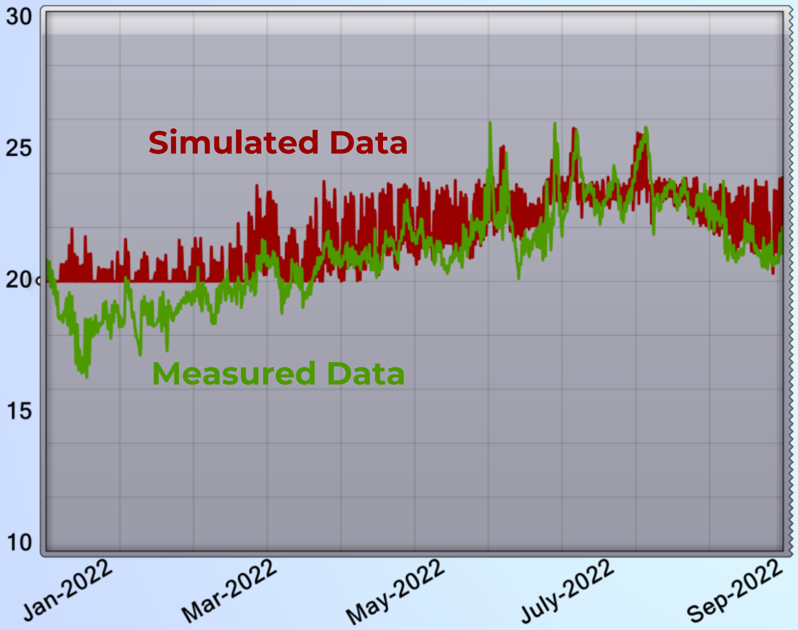

ASHRAE Standard	Hourly Calibration Criterion	Results (varying periods)
NMBE	<10%	-8% to 10.45%
CV(RMSE)	<30%	2% to 11.32%

EnergyPlus Runtime API

EnergyPlus

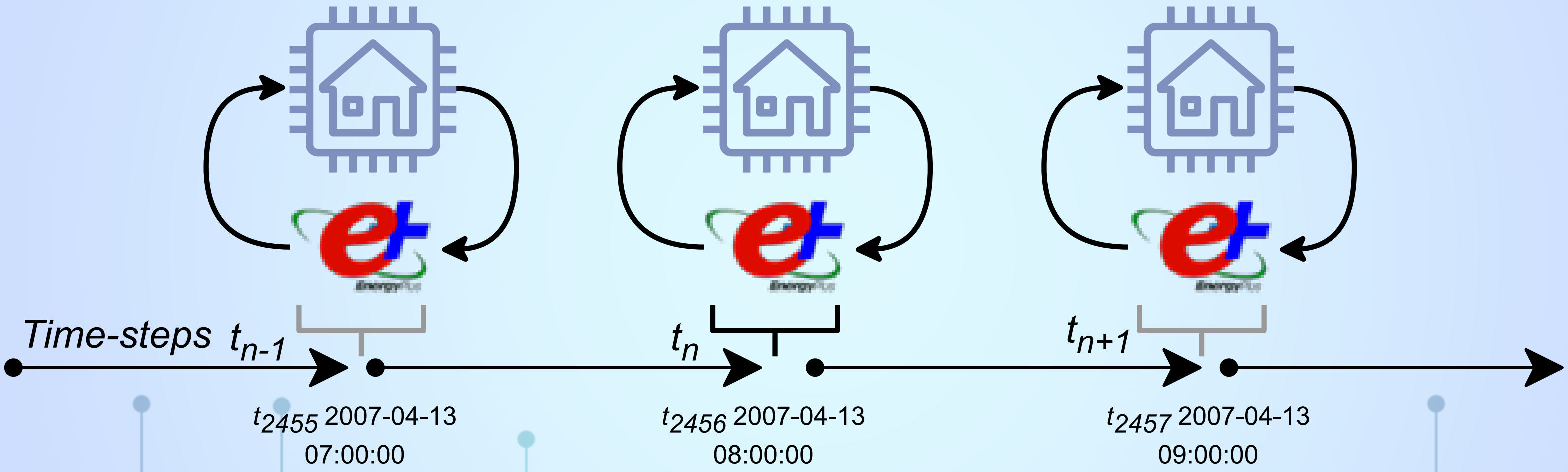

Time-steps
t_{n-1}
t_n
t_{n+1}
t_{2455} 2007-04-13 07:00:00
t_{2456} 2007-04-13 08:00:00
t_{2457} 2007-04-13 09:00:00

EnergyPlus API

EMS and Python API, since 2020

Example uses

IBM research team

AirBoxLab

MechyAI

DRL

Deep
Reinforcement
Learning

General RL Setup

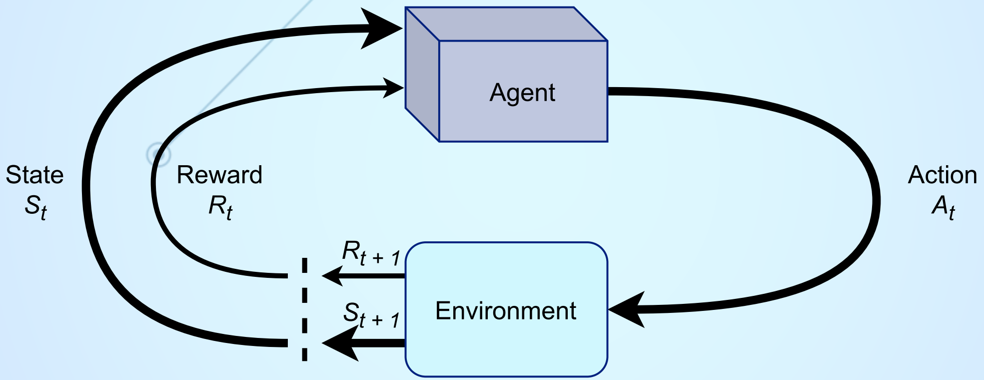

Example Building Control RL Setup

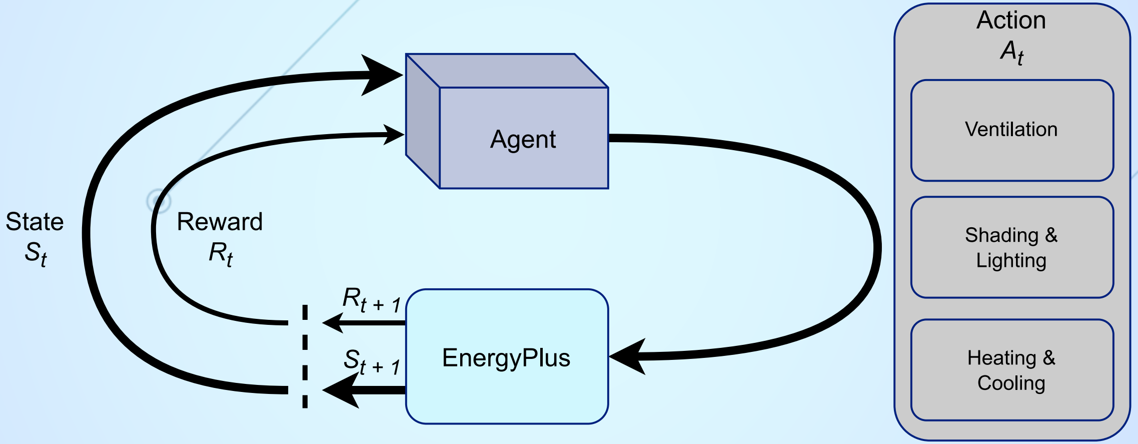

Current RL Setup

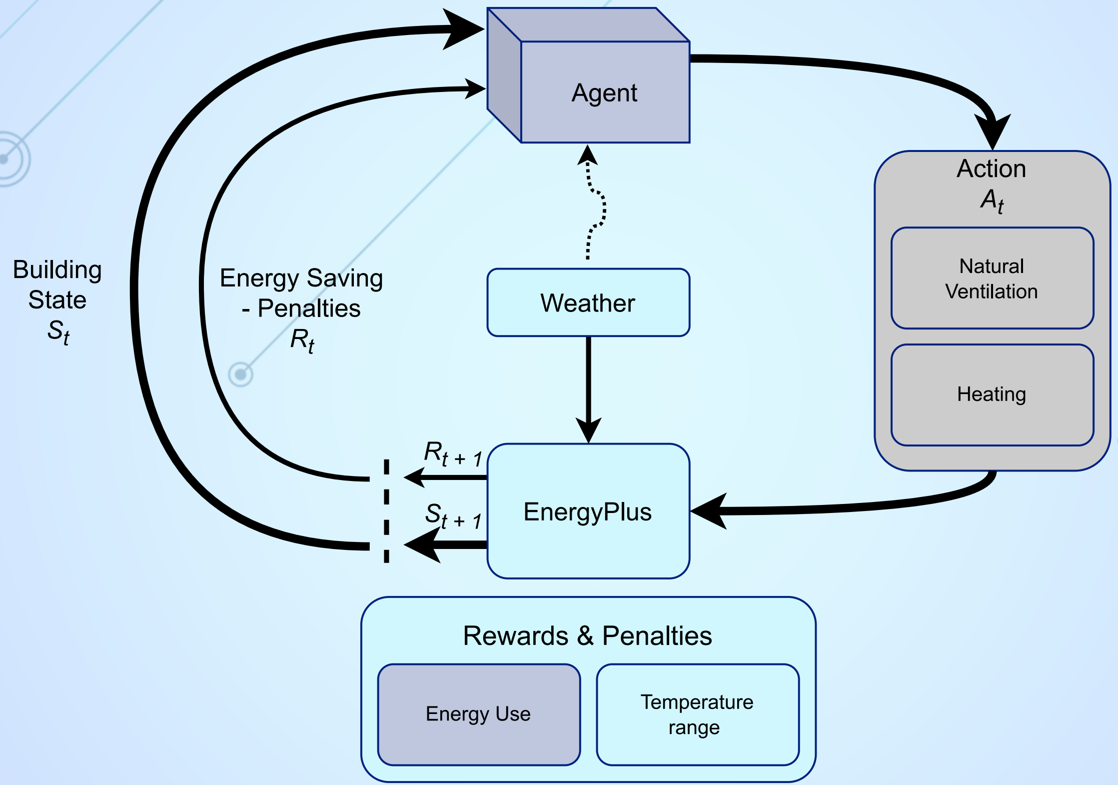

Actions

1. 18.0 °C - louvres **closed**
2. 18.5 °C - louvres **closed**
3. 19.0 °C - louvres **closed**
4. 19.5 °C - louvres **closed**
5. 20.0 °C - louvres **closed**

6. 18.0 °C - louvres **open**
7. 18.5 °C - louvres **open**
8. 19.0 °C - louvres **open**
9. 19.5 °C - louvres **open**
10. 20.0 °C - louvres **open**

11. 5.0 °C - louvres **closed**

Rewards & Penalties

Comfort Penalty

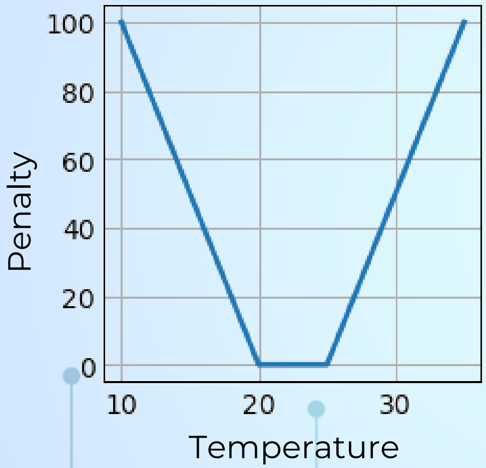

Energy Use Penalty

0.6 per hour per 10 kWh
or
1 per 60 MJ

Runtime Interaction

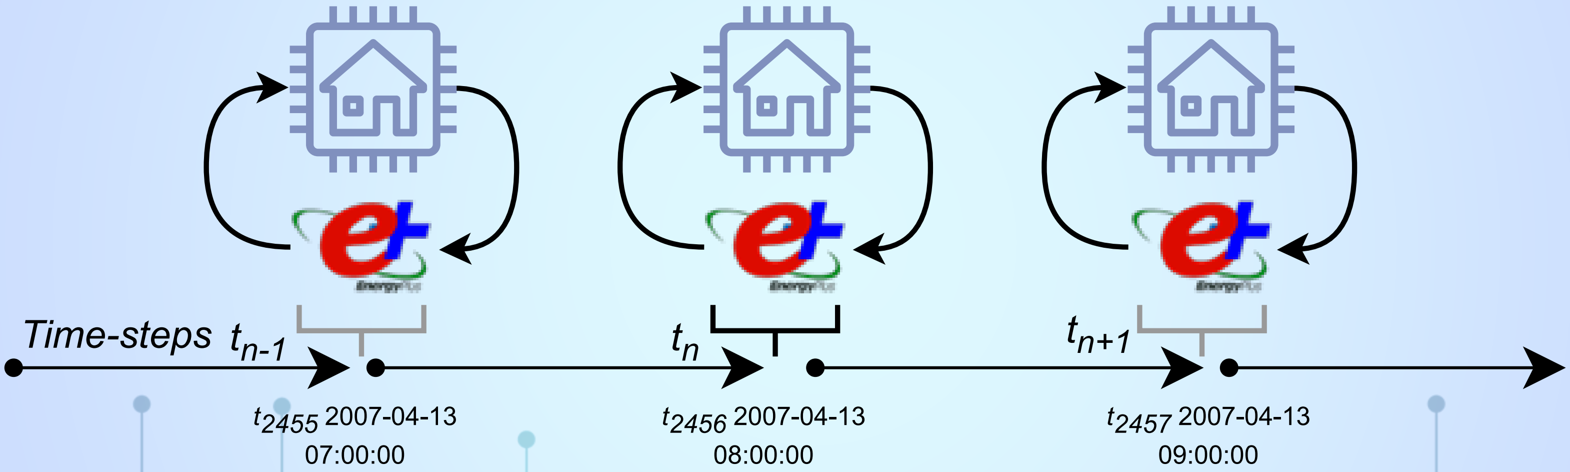

Runtime Step

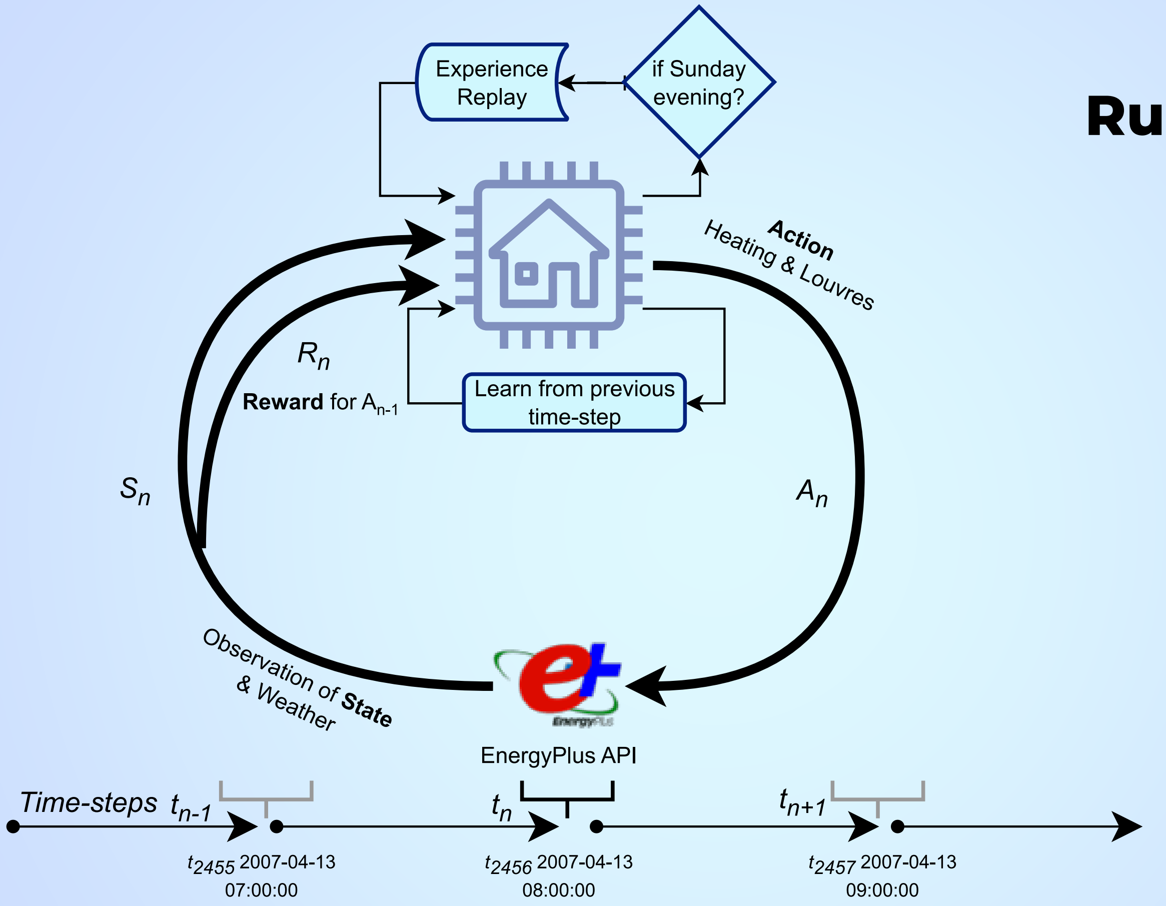

Agent Neural Network

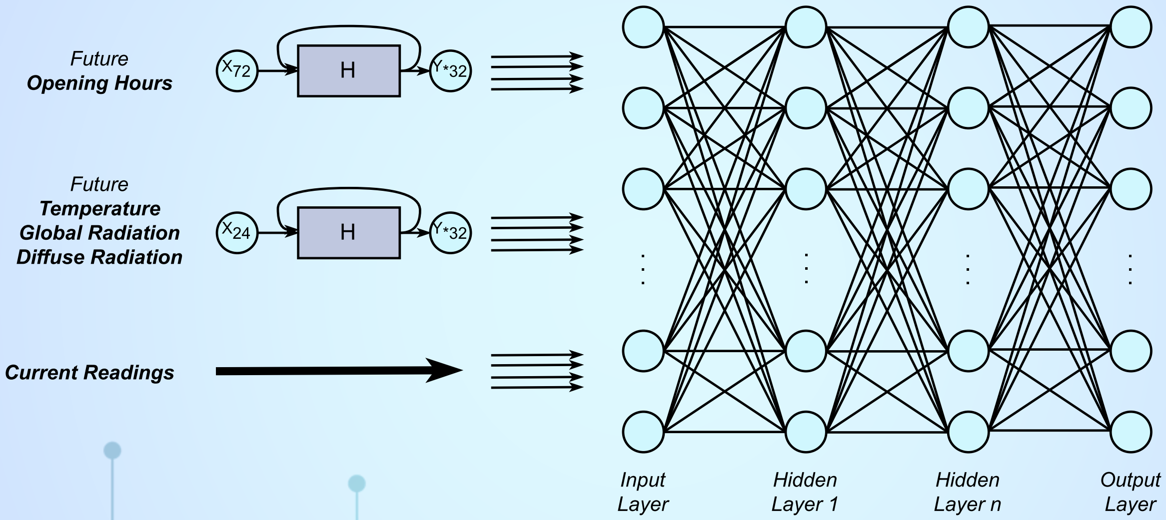

State Example

-1. 1. 1. 1. 1. 1. 1. 1. 1. -1. -1.
-1. -1. -1. -1. -1. -1. -1. -1. -1.
-1. -1. -1. -1. -1. 1. 1. 1. 1. 1.
1. 1. 1. -1. -1. -1. -1. -1. -1. -1.
-1. -1. -1. -1. -1. -1. -1. -1. -1.
1. 1. 1. 1. 1. 1. 1. 1. -1. -1. -1.
-1. -1. -1. -1. -1. -1. -1. -1. -1.
-1. -1. -1. 0.148 0.148 0.036
0.052 0.108 0.104 0.04 0.076
0.028 -0.028 -0.092 -0.152 -
0.104 -0.108 -0.108 -0.08 -
0.072 -0.04 -0.028 -0.004
0.048 0.132 0.112 0.12 -
0.654 -0.866 -0.94 -0.914 -
0.912 -0.962 -0.94 -0.944 -
0.934 -0.332 -0.568 -1. -1. -1.
-1. -1. -1. -1. -1. -1. -1. -1. -1.
-0.98 -0.346 -0.496 -0.328 -
0.27 -0.236 -0.208 -0.29 -
0.39 -0.534 -0.676 -0.788 -
0.958 -1. -1. -1. -1. -1. -1. -1. -
1. -1. -0.982 -0.886 -0.742
0.35902 0.36312 -1. 0. 0.12
0.12 0.68 0.12 -0.0609 1. 0. -
0.11111 -0.58667 0.19616
0.20669 0.35549 0.35969
0.36254 0.36655 0.00606
0.0185 0.00594 0.01441
0.00107 0.00559 1.

72 - Open/closed, Boolean

-1. 1. 1. 1. 1. 1. 1. 1. 1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. 1. 1. 1. 1. 1. 1. 1. 1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. 1. 1. 1. 1. 1. 1. 1. 1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1.

24 – Ext. Temperature

0.148 0.148 0.036 0.052 0.108 0.104 0.04 0.076 0.028 -0.028 -0.092 -0.152 -0.104 -0.108 -0.108 -0.08 -0.072 -0.04 -0.028 -0.004 0.048 0.132 0.112 0.12 -0.654

24 – Global Radiation

-0.866 -0.94 -0.914 -0.912 -0.962 -0.94 -0.944 -0.934 -0.332 -0.568 -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -1. -0.98 -0.346

24 – Diffuse Radiation

-0.496 -0.328 -0.27 -0.236 -0.208 -0.29 -0.39 -0.534 -0.676 -0.788 -0.958 -1. -1. -1. -1. -1. -1. -1. -1. -1. -0.982 -0.886 -0.742 0.35902

26 – Current Readings

0.36312 -1. 0. 0.12 0.12 0.68 0.12 -0.0609 1. 0. -0.11111 -0.58667 0.19616 0.20669 0.35549 0.35969 0.36254 0.36655 0.00606 0.0185 0.00594 0.01441 0.00107 0.00559 1.

Agent Types

DRQN	DRQN	DRQN	DRQN	DRQN	*DQN*
24 hours	**4 hours**	**24 hours**	**4 hours**	**0 hours**	**4 hours**
Temperature	Temperature	Temperature	Temperature		Temperature
Global Radiation	Global Radiation				Global Radiation
Diffuse Radiation	Diffuse Radiation				Diffuse Radiation

Agent Neural Network

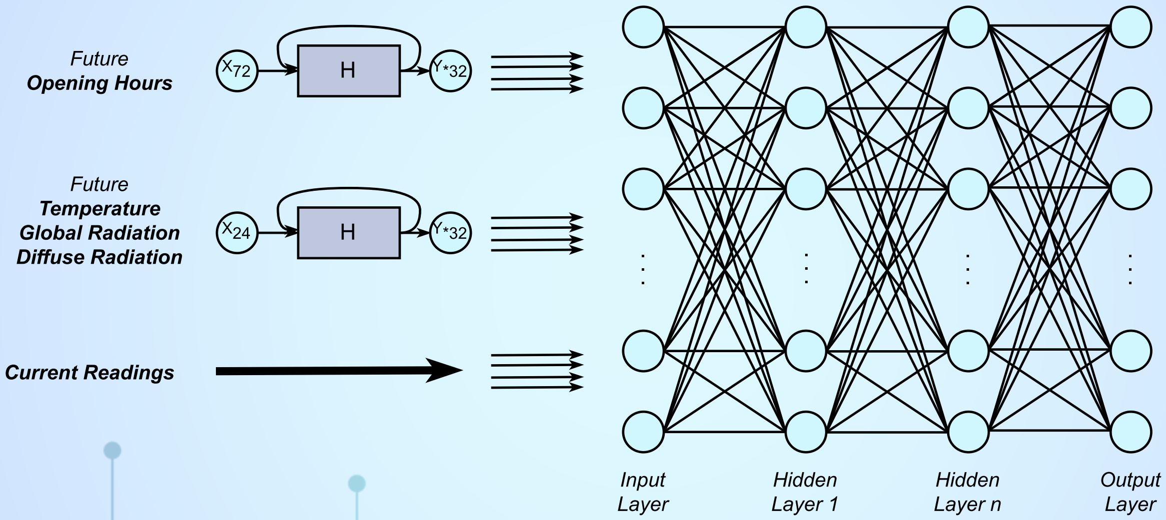

Building Variants

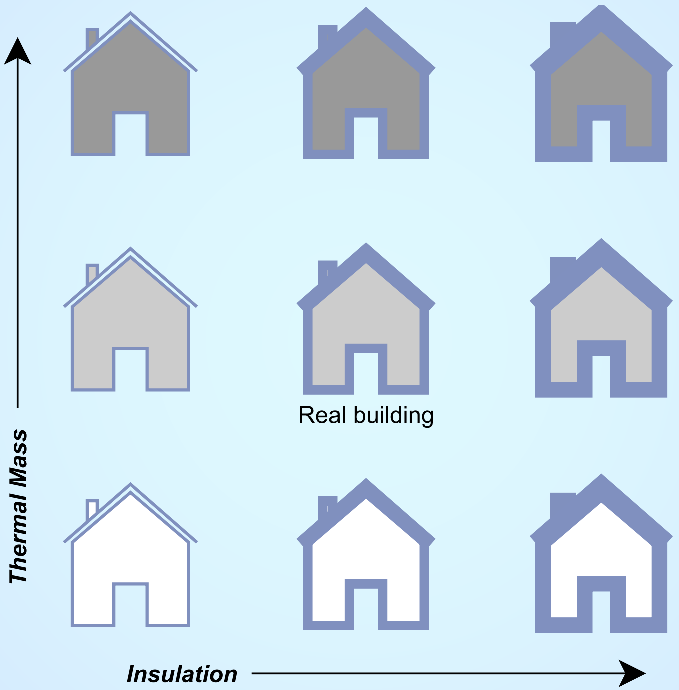

Experiments Setup

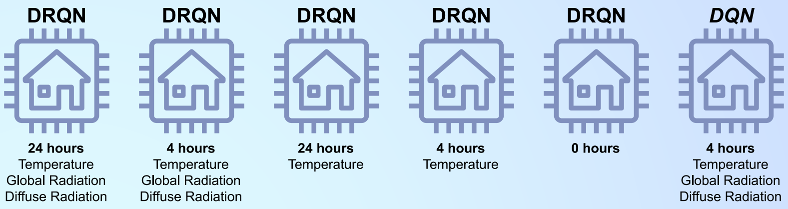

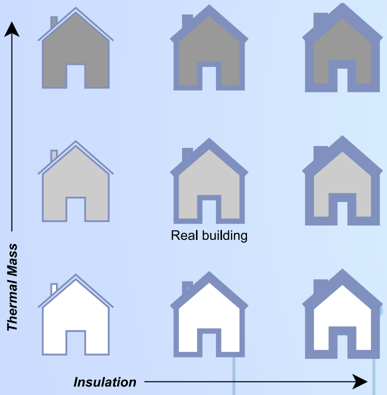

	Algorithm 1	Algorithm 2	Algorithm 3
Building 1			
Building 2			
Building 3			

DRL Agent Training

- **35 simulated years** of training – 306,600 hourly timesteps
- **8 hours** – real-time to train 1 agent on 1 building
- **7 years of local weather data**

Speed

- EnergyPlus: **1,000** steps/min
- PyTorch: >**60,000** steps/min

- **Parallelisation** can improve training speed (see appendix)

Results & Conclusion

Specific Agents

Building type name	Insulation	Thermal Mass	Baseline EnergyPlus Heating [kWh/m2/a]	RL 24h Foresight RNN Heating [kWh/m2/a]	RL 4h Foresight RNN Heating [kWh/m2/a]	RL No Foresight Heating [kWh/m2/a]	RL 24h Foresight RNN w/o Solar Heating [kWh/m2/a]	RL 4h Foresight RNN w/o Solar Heating [kWh/m2/a]	RL 4h Foresight flat inputs Heating [kWh/m2/a]
Building-InsuBASE-MassBASE	Baseline	Baseline	8.20	5.83	6.36	8.20	6.55	7.35	8.08
Building-InsuBASE-MassDW	Baseline	Decreased 50%	8.40	5.68					
Building-InsuBASE-MassUP	Baseline	Increased 2x	7.89	5.07					
Building-InsuDW-MassBASE	Decreased 50%	Baseline	25.02	21.02					
Building-InsuDW-MassDW	Decreased 50%	Decreased 50%	25.29	24.04					
Building-InsuDW-MassUP	Decreased 50%	Increased 2x	24.63	21.67					
Building-InsuUP-MassBASE	Increased 2x	Baseline	3.98	2.74					
Building-InsuUP-MassDW	Increased 2x	Decreased 50%	4.14	3.39					
Building-InsuUP-MassUP	Increased 2x	Increased 2x	3.62	2.44					

Base Agents Across All Buildings

Building type name	Insulation	Thermal Mass	Baseline EnergyPlus Heating [kWh/m2/a]	Baseline EnergyPlus deg-hours < 20C	Baseline EnergyPlus deg-hours time > 25C	RL 24h Foresight RNN Heating [kWh/m2/a]	RL 24h Foresight RNN deg-hours < 20C	RL 24h Foresight RNN deg-hours time > 25C	RL 4h Foresight RNN Heating [kWh/m2/a]	RL 4h Foresight RNN deg-hours < 20C	RL 4h Foresight RNN deg-hours time > 25C	RL No Foresight Heating [kWh/m2/a]	RL No Foresight deg-hours < 20C	RL No Foresight deg-hours > 25C	RL 24h Foresight RNN w/o Solar Heating [kWh/m2/a]	RL 24h Foresight RNN w/o Solar deg-hours < 20C	RL 24h Foresight RNN w/o Solar deg-hours > 25C	RL 4h Foresight RNN w/o Solar Heating [kWh/m2/a]	RL 4h Foresight RNN w/o Solar deg-hours < 20C	RL 4h Foresight RNN w/o Solar deg-hours > 25C	RL 4h Foresight flat inputs Heating [kWh/m2/a]	RL 4h Foresight flat inputs deg-hours < 20C	RL 4h Foresight flat inputs deg-hours > 25C
Building-InsuBASE-MassBASE	Baseline	Baseline	8.20	6.22	10.16	5.83	153.68	10.08	6.36	40.67	10.13	8.20	6.22	10.16	6.55	67.60	10.13	7.35	11.87	10.15	8.08	6.14	10.16
Building-InsuBASE-MassDW	Baseline	Decreased 50%	8.40	7.43	13.59	6.09	150.59	13.44	6.68	38.07	13.55	8.40	7.43	13.59	6.79	64.87	13.48	7.56	12.45	13.57	8.28	7.69	13.59
Building-InsuBASE-MassUP	Baseline	Increased 2x	7.89	4.95	7.31	5.31	158.76	7.21	5.86	42.77	7.29	7.89	4.95	7.31	6.09	69.93	7.23	7.02	10.03	7.30	7.77	4.87	7.31
Building-InsuDW-MassBASE	Decreased 50%	Baseline	25.02	242.39	15.69	20.85	598.83	14.84	21.89	429.95	15.30	25.02	242.39	15.69	22.16	479.79	15.11	22.91	323.52	15.43	24.64	251.01	15.66
Building-InsuDW-MassDW	Decreased 50%	Decreased 50%	25.29	252.08	21.36	21.21	600.76	19.65	22.30	423.69	20.02	25.29	252.08	21.36	22.48	477.29	20.03	23.32	326.90	20.16	24.95	263.86	21.31
Building-InsuDW-MassUP	Decreased 50%	Increased 2x	24.63	216.35	9.71	20.36	610.84	9.01	21.31	438.42	9.31	24.63	216.35	9.71	21.72	489.77	9.26	22.28	318.70	9.55	24.22	235.37	9.69
Building-InsuUP-MassBASE	Increased 2x	Baseline	3.98	0.09	8.77	2.37	68.18	8.68	2.82	6.19	8.76	3.98	0.09	8.77	2.93	14.97	8.69	3.60	0.31	8.77	3.92	0.10	8.77
Building-InsuUP-MassDW	Increased 2x	Decreased 50%	4.14	0.07	11.98	2.56	61.80	12.00	2.99	5.67	11.96	4.14	0.07	11.98	3.07	13.95	12.01	3.71	0.22	11.98	4.07	0.10	11.98
Building-InsuUP-MassUP	Increased 2x	Increased 2x	3.62	0.08	6.18	1.92	63.17	6.40	2.44	6.33	6.16	3.62	0.08	6.18	2.57	13.55	6.42	3.26	0.19	6.15	3.57	0.04	6.18

Literature Comparison

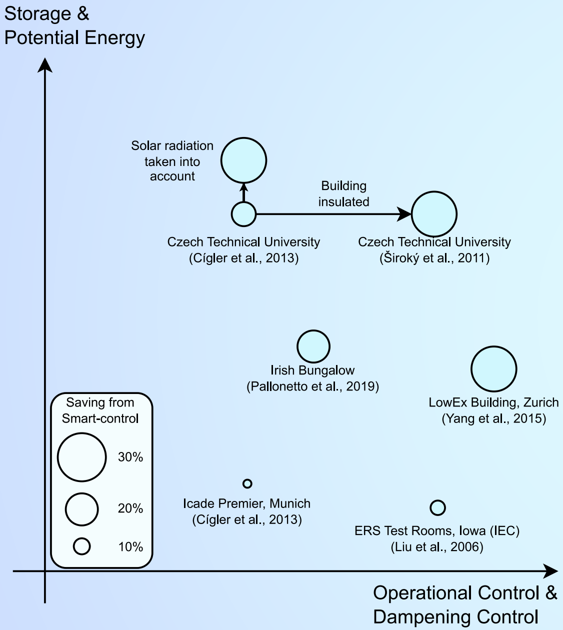

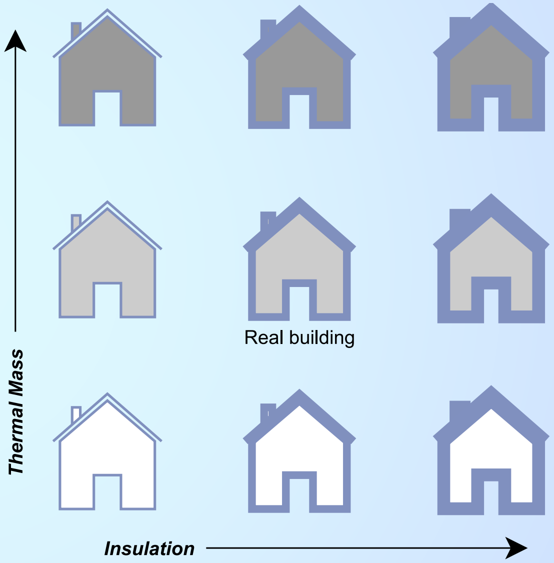

Literature Comparison

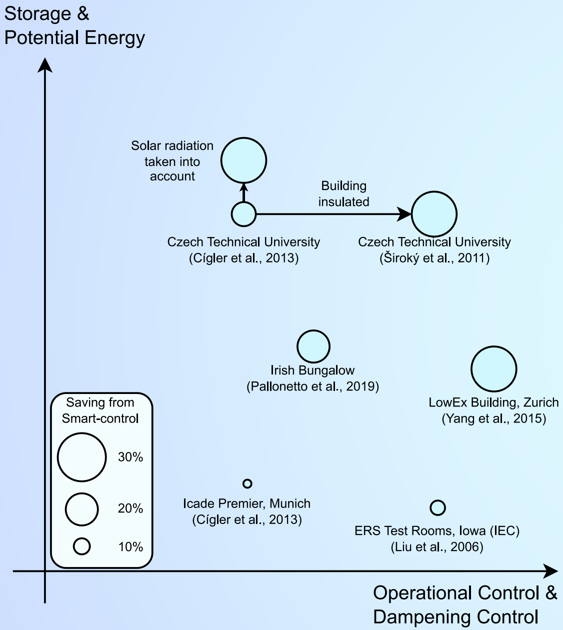

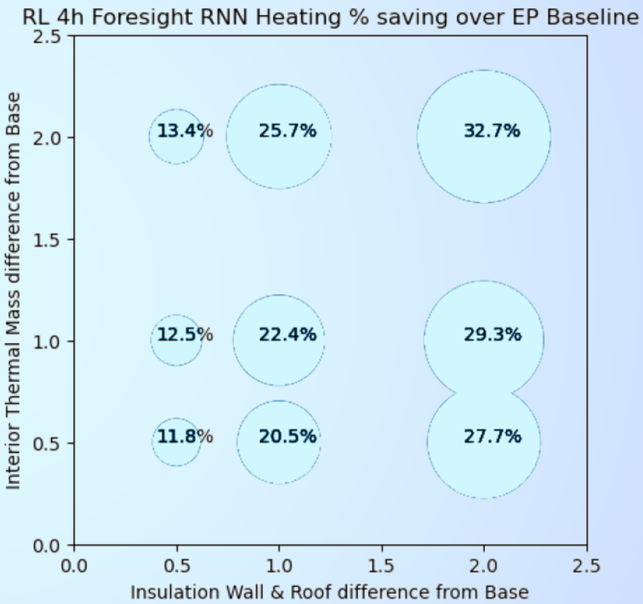

24h Foresight Agent

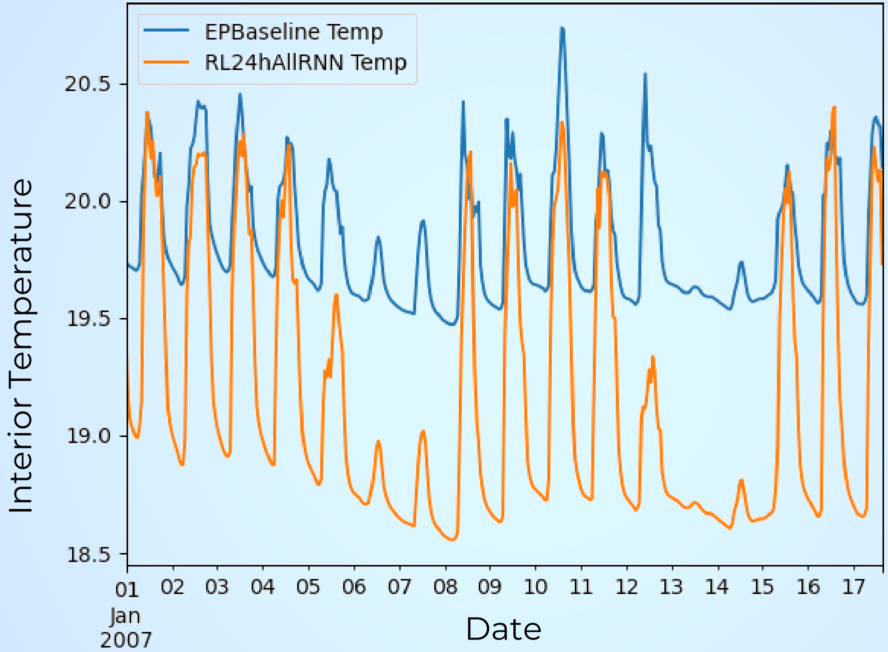

4h Foresight Agent

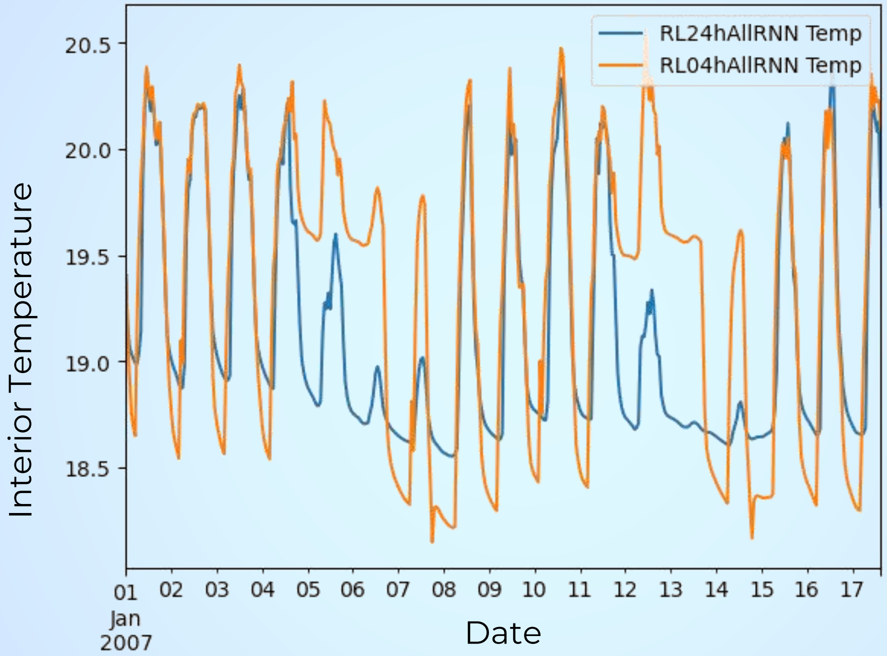

Energy Saving vs Comfort

$$\int_{p_0}^{p_i} \frac{Q_b - Q_a}{d}$$

Linear Ratio

$$\int_{p_0}^{p_i} \frac{(Q_b - Q_a)(hf - d)}{f}$$

Multiplication Score

$$\int_{p_0}^{p_i} \frac{k^{Q_b - Q_a}}{k^d}$$

Power Ratio

p is the period being investigated
Q_b energy use by the baseline
Q_a energy use by the agent
d degree-hours outside comfort
h open-hours per period step p
f optional factor, prevents **(hf – d)** being negative
k power base factor

Multiplication Score - Example

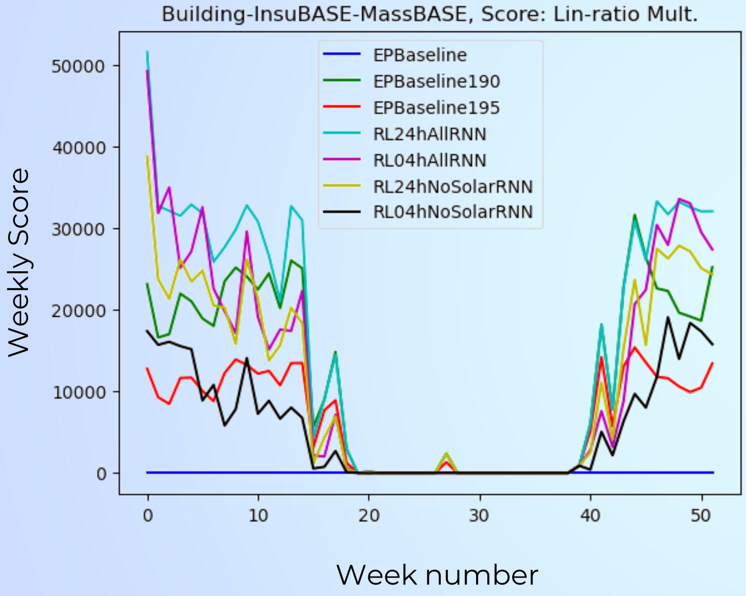

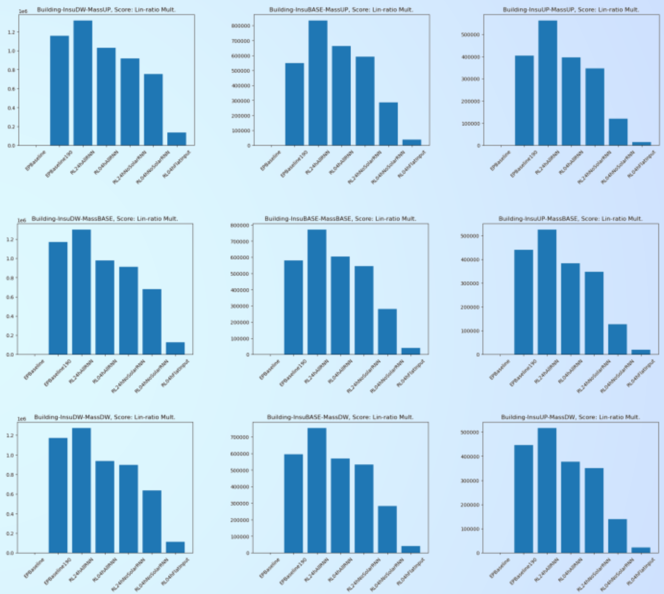

Interior vs Exterior Temperatures

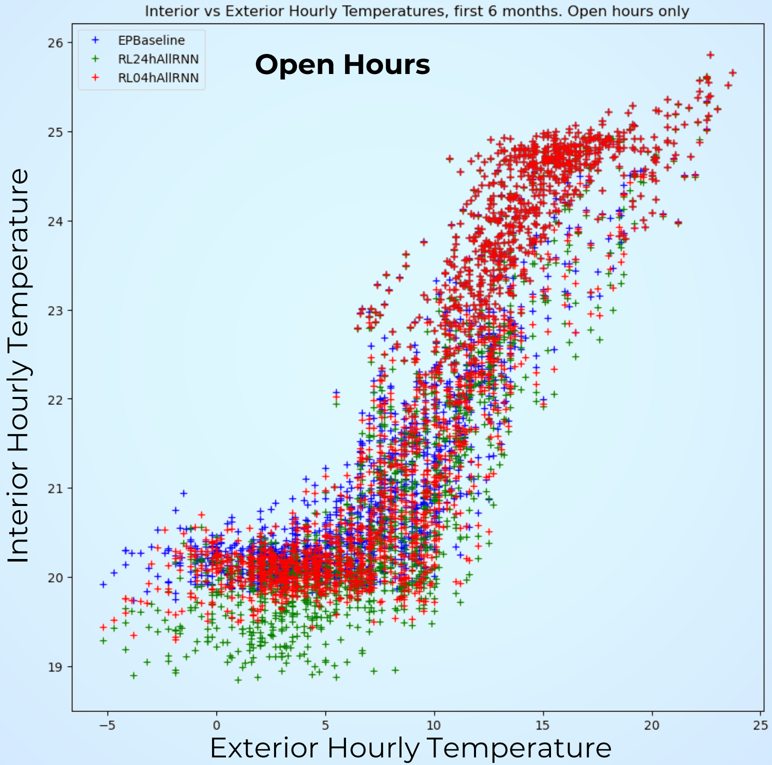

Interior vs Exterior Temperatures

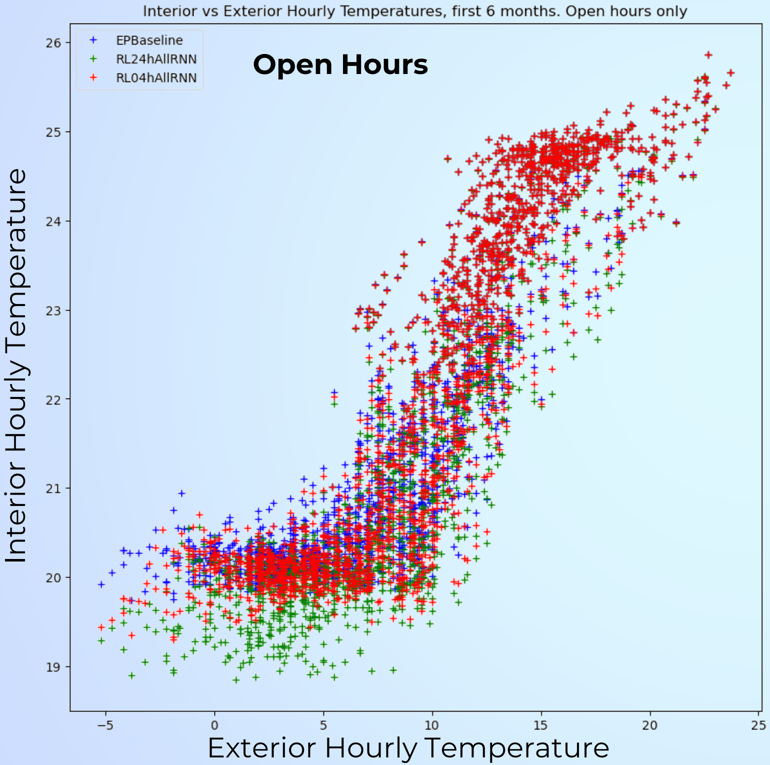

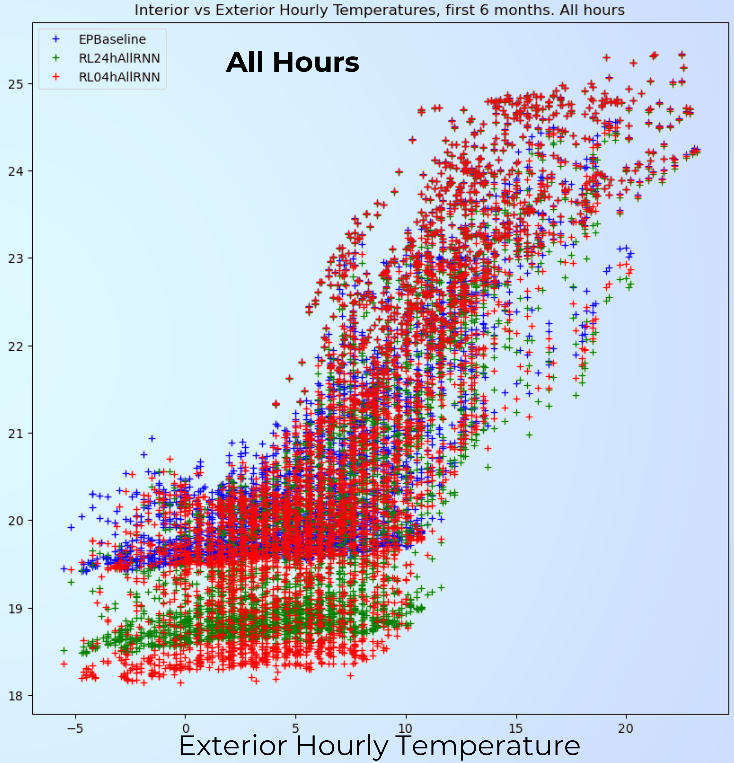

Climate Change | Literature | RQ & Method | Case Study | EnergyPlus API | DRL Agents | **Results**

Degree-hours vs Energy Use

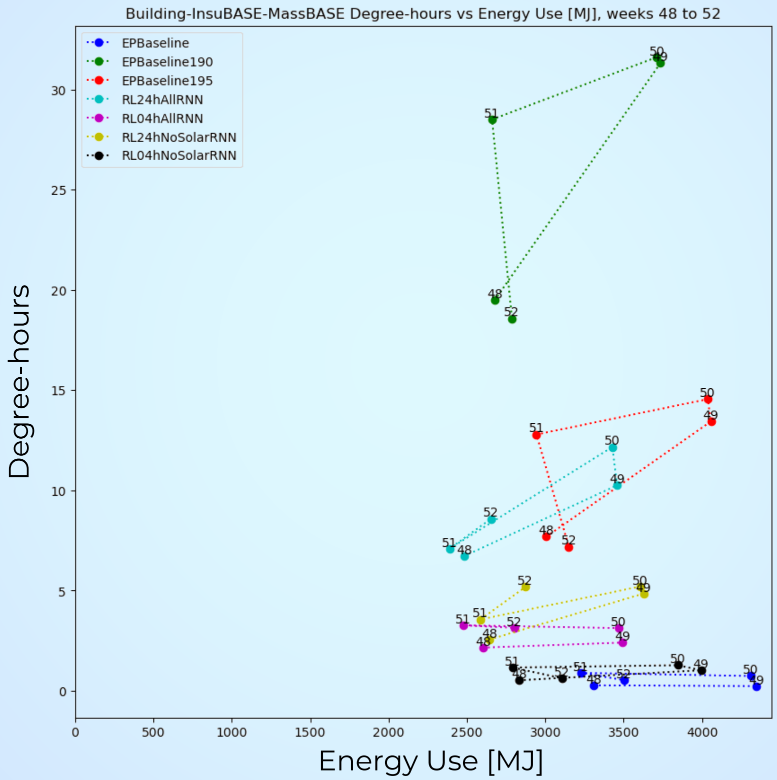

RL Training Metrics

Agent loss during training

Mean Square Error loss

Rolling mean of 17,520 steps (2 years)

Performance Conclusion

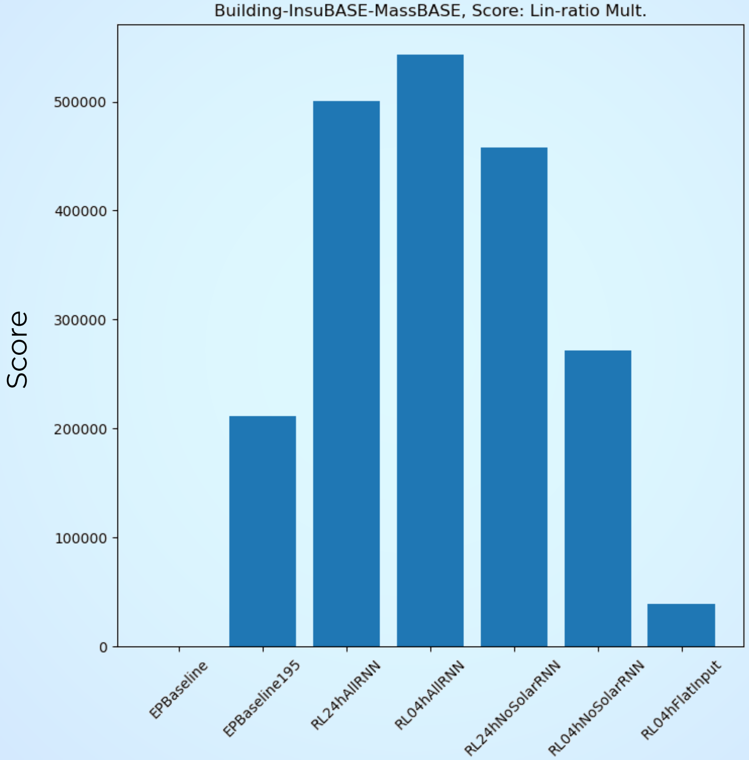

Multiplication Score metric

Weekly performance
over 52 weeks

f = 0.4

Rule Based Controllers

A new baseline?

Rule Based Controllers

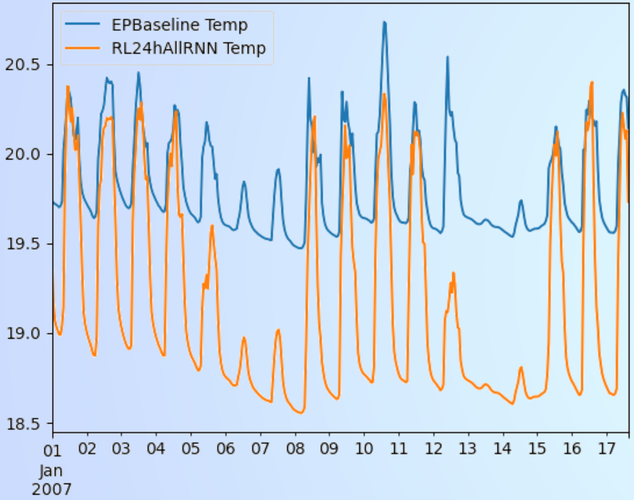

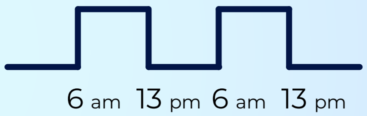

20 °C at day

18-20 °C at night

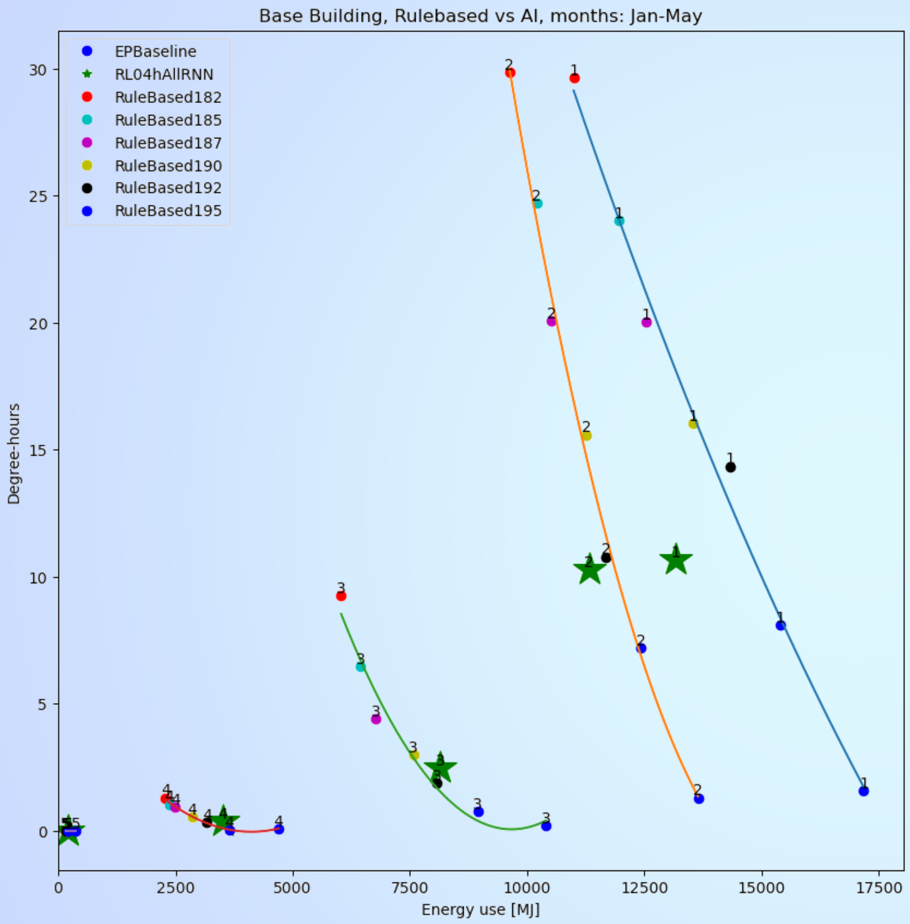

Rule Based Controllers

Monthly results

Trendline for rule based controllers

RL04hAllRNN outperforms RBCs

Climate Change | Literature | RQ & Method | Case Study | EnergyPlus API | DRL Agents | **Results**

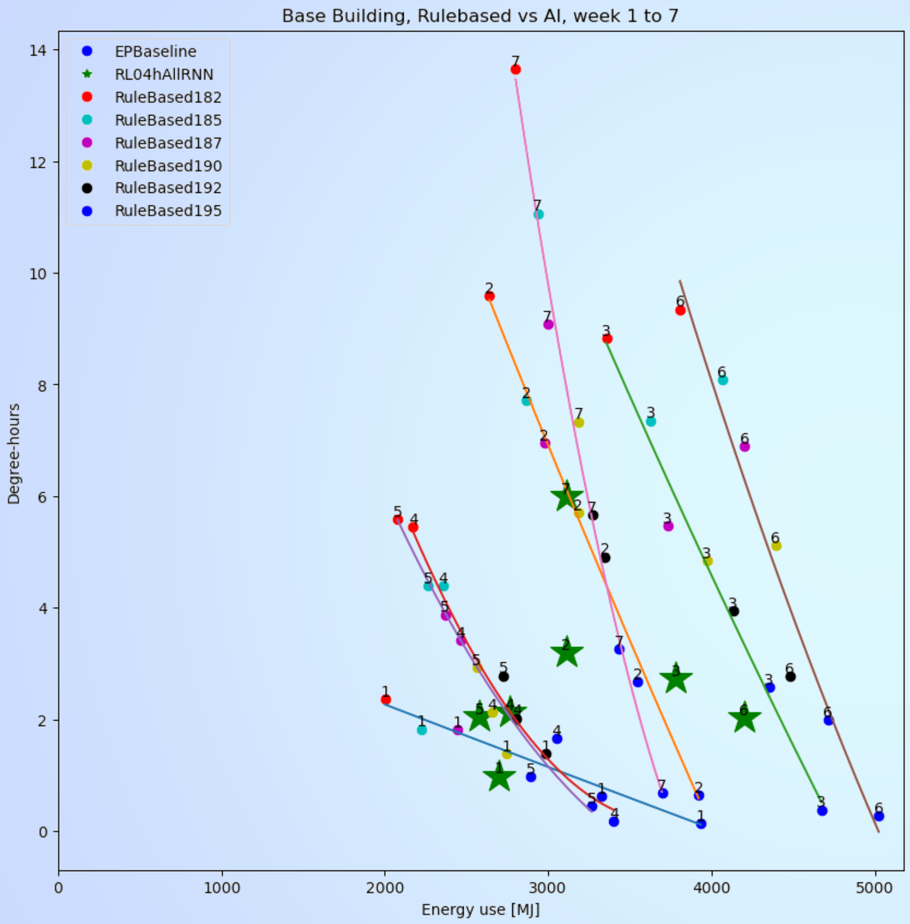

Rule Based Controllers

Weekly results

Climate Change | Literature | RQ & Method | Case Study | EnergyPlus API | DRL Agents | **Results**

1%

CO_2 Reduction*

From a 3% energy reduction

*Based on UK building CO_2 emissions

Key Findings

- **Theory confirmation**: Increase in energy storage and control leads to greater relative energy saving potential
- **RL outperforms rule based controller**
- Largest absolute savings in low-performing buildings
- Additional information (solar) and longer horizon improves performance, but can lead to instability in learning
- RL algorithm design has large influence on performance

Industry & Societal Impact

- Greater potential saving if combined with secondary measures; energy storage, variable tariffs, additional control surfaces etc.
- Building equipment may require new APIs or control access
- Potential for low-capital technological retrofit for old buildings, if full retrofit is not suitable or affordable

Future Research & Improvements

- Smart-control for low-insulated, high thermal mass building
- Metrics for assessing energy savings against comfort compromise
- Algorithms' robustness for weather prediction uncertainty
- Physics based RL
- Multi-objective RL
- Benchmarking of controllers on standardised buildings

Thank You

Questions?

RNN – Recurrent Neural Network

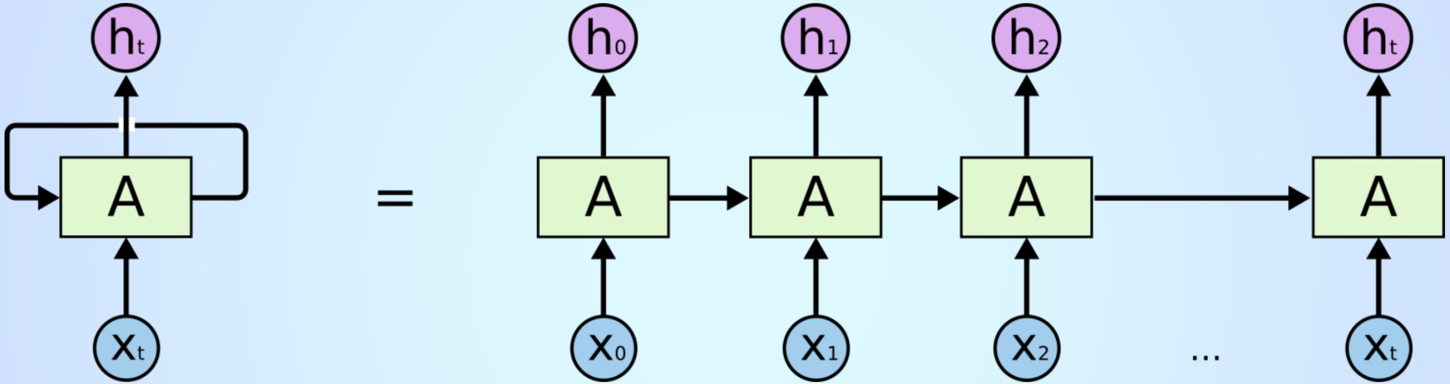

Diagrams by Christopher Olah

RNN – Recurrent Neural Network

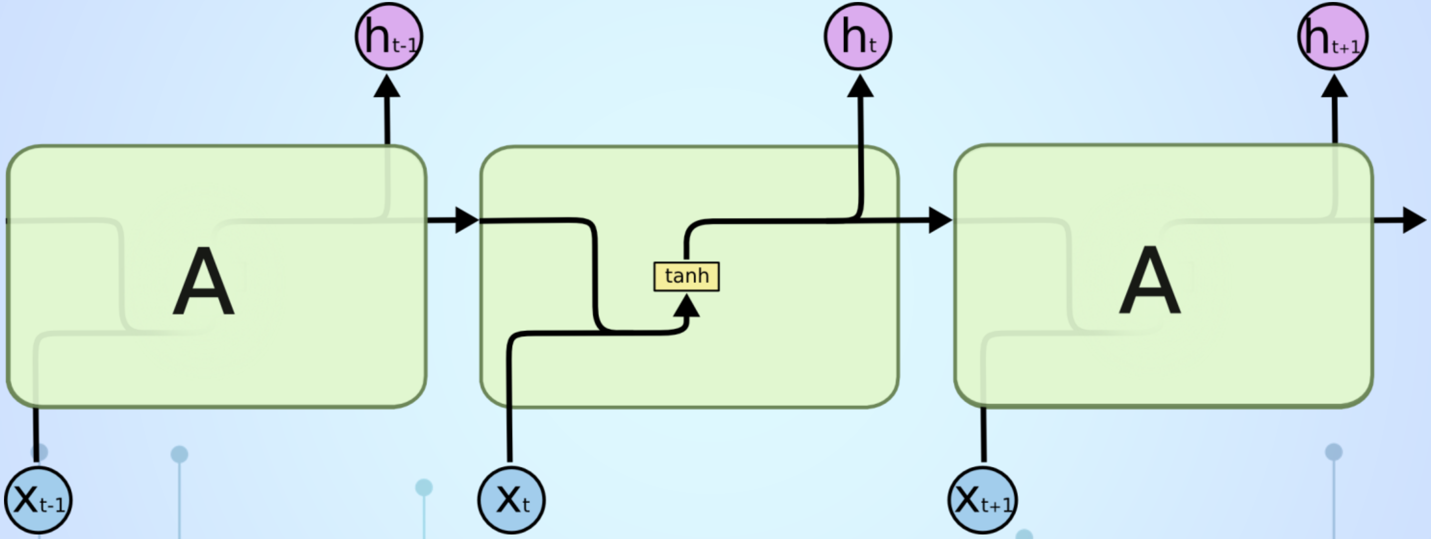

Diagrams by Christopher Olah

LSTM – Long Short Term Memory

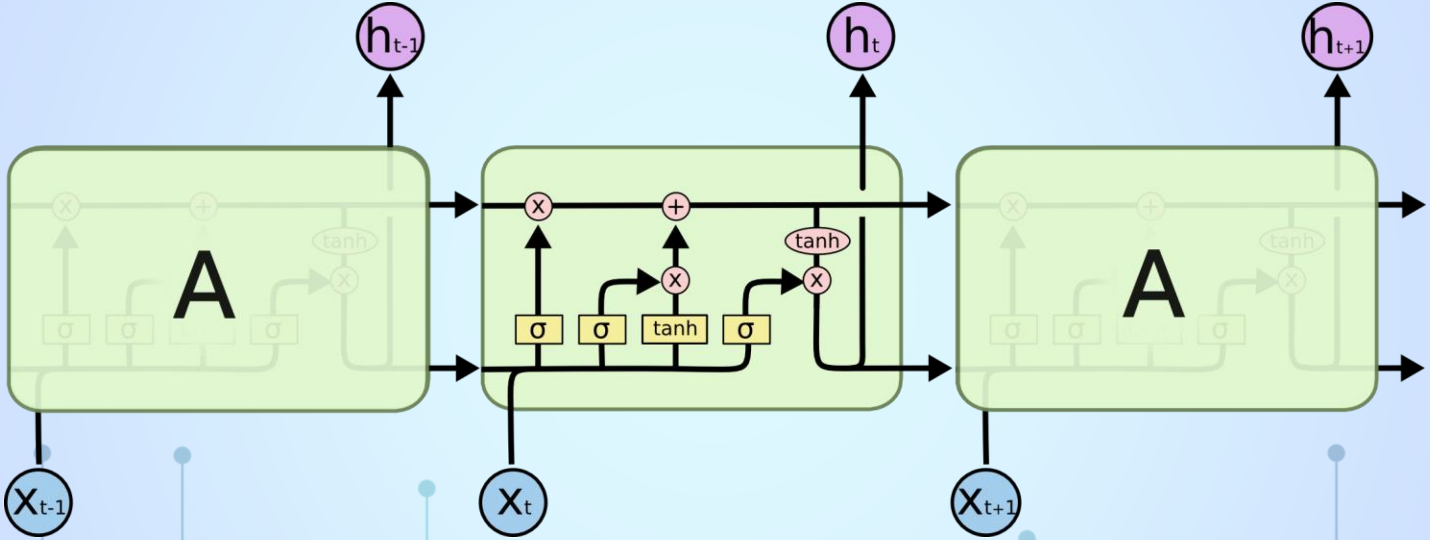

Diagrams by Christopher Olah

Bibliography

For full list of sources, see the accompanying report

Thanks!

CREDITS: This presentation template was created by Slidesgo, including icons by Flaticon, and infographics & images by Freepik.

Please keep this slide for attribution.