Towards self-adaptive building energy control in smart grids

Juan Gómez-Romero
Universidad de Granada (Spain)
jgomez@decsai.ugr.es

Miguel Molina-Solana
Universidad de Granada (Spain), Imperial College (UK)
miguelmolina@ugr.es

Buildings are the largest energy-demanding sector in the world, representing over one third of the total worldwide consumption and a similarly important source of CO_2 emissions [15]. More than a half of the energy consumption during the building life-cycle is due to the operation of the HVAC systems (heating, ventilation and air conditioning) [27, 31]. Renovation works and retrofitting are essential to progress towards highly-efficient buildings, and to be effective, they must be accompanied by suitable operation protocols [22]. As a matter of fact, selecting optimal setpoints and deadbands for HVAC control could alone lead to energy savings up to 35% while maintaining occupants' comfort [12]. Nevertheless, most state-of-the-art technologies for automatic control cannot achieve these figures [24], and those successful require a considerable effort to be adapted to different scenarios [4].

The overarching objective of this research initiative —namely IA4SG (Intelligent Agents for the Smart Grid)— is to create the technologies supporting the smart energy control of the future building ecosystem. Machine Learning has a tremendous potential to achieve great energy savings, reduce contaminant emissions and make the best use of renewable resources at an affordable cost for the building owners and residents [29]. We envision a future energy system in which building control will be performed by autonomous self-adaptive agents that, with minimal configuration, will learn how to operate the HVAC equipment more efficiently and how to collaborate with other actors of the grid. To this aim, we pursue to develop new Deep Learning and Reinforcement Learning methods, algorithms and tools to address three key issues: (A) generation of optimal control instructions for HVAC to save energy while guaranteeing comfort; (B) simulation of buildings under different operations and contexts; (C) coordination between components of the energy system to achieve an overall reduction of the contaminant emissions. The concept behind IA4SG is shown in Figure 1:

A. **Control optimization by Deep Reinforcement Learning (DRL): From the game board to the building**. An IA4SG agent teaches itself to efficiently operate the building from multiple "trial and error" episodes, as AlphaZero did [30]. Since the agent requires a considerable number of episodes to refine its skills, collection of such experience is not only performed on the real building, but also on a simulation model (see B). This process is more complex if the agent has capabilities for energy production and storage, which at the same time offers more possibilities for improvement. Once the agent is trained, it can be deployed to operate the building and updated as the scenario evolves. Preliminary works —-using hypothetical simple building models with a reduced action space [2, 34, 38]— suggest that HVAC control with DRL is feasible and has potential to dramatically transform the area, since it would overcome scalability limitations of similar control approaches based on data-driven simulation [36].

B. **Deep neural simulation models (DNSM): Learning and predicting the building behaviour from data**. Instead of using a manually-crafted physical simulation model, the IA4SG agent automatically learns a digital twin of the real building from historical sensor data. The simulation model can be re-calibrated from live data. Learning is not performed from scratch, but by reusing simulation models created for other scenarios (i.e. by performing transfer learning). Shallow neural networks have been previously used to create data-driven simulation models, but only to estimate consumption (not the whole environment state) at a microscale (not for a large section of the building) [35]. Recently deep neural networks have proved to be useful to model single HVAC components and simulate short-term thermal behavior [6, 10, 21].

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

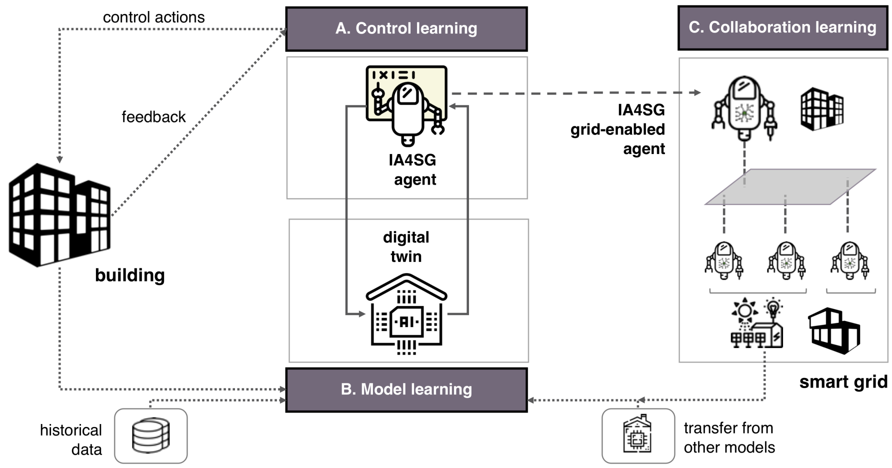

Figure 1: IA4SG concept

C. **Cooperation among multiple energy system actors: Leveraging agent cliques to federations**. An IA4SG agent cooperates with other agents in the energy system to trade energy in order to maximize the overall use of clean energies in the whole grid. Each agent describes and publishes its capabilities, resources and objectives, and after been trained to interact with other agents, it is ready to negotiate energy acquisition and sharing by following established protocols. Multi-Agent Systems (MAS) have been proposed before as a suitable technology to manage different aspects of power systems [25] and buildings [20], given their capabilities for autonomous organization, semantic interoperability, and system scalability. Analyses of research on MAS applications for smart and micro grids have shown that MAS solutions are feasible [18], but also that the increasing complexity and decentralization of the grid require novel negotiation protocols, expressive knowledge models and advanced validation methods [7, 33]. Traditionally, agents' interaction has been implemented by means of a formal specification of their operational cycle [8]. More recently, optimization methods based on dynamic programming, games theory, and heuristic search have shown better use of energy, although most of them have still a centralized supervisory control which reduces the flexibility of the grid [39]. Recent advances on multi-agent learning [32, 26, 16] will be put forward to train IA4SG agents.

This concept is different from the current approaches in the area. The prevalent paradigm for optimal control in energy systems is Model Predictive Control (MPC) [19], in which a simulation model of the building is used to estimate its response to alternative control sequences and situations. Physical models —which characterize the thermal behavior of the building by using differential equations encoding the principles of mass, energy and momentum transfer— have been traditionally used in MPC, since they are interpretable and their accuracy exceeded that of data-driven models [3]. However, creating physical simulation models require a very significant human effort and running them is very time-consuming. Hence, most works tend to simplify the models [28] —e.g. by reducing the differential equations to linear combinations—, even though the accuracy and the coverage of the simulation is a crucial aspect to avoid the generation of control instructions under wrong assumptions [11]. Model simplification results in short-scope, limited-extensibility and reduced-performance solutions involving a great deal of manual work [1]. Additionally, implementations for joint management of groups of buildings in the grid are still scarce and not flexible enough for open environments [37], since until recently most of them have assumed a centralized manually-designed control [14]. Here we propose a decentralized self-learning system, taking into account the particular features of each component while assuming a shared objective of reducing energy cost and increasing the use of renewable sources.

We are confident that our proposal could create in 5 years the technologies to reduce a 30% the energy consumed by buildings and increment a 30% the use of clean energies in grids, in line with current regulations [9] and furthering other works with a narrower scope [5, 23, 17, 13]. Not only do we expect a reduction on energy consumption and an increment on the use of renewable sources, but also a reduction on the cost of controlling energy in buildings, which is crucial to achieve a widespread adoption of climate mitigation technologies.

Acknowledgments

This research work has been partially funded by the European Union's Horizon 2020 Research and Innovation Programme (g.a. 754446) and the Spanish Ministry of Science, Innovation and Universities (TIN2017-91223-EXP).

Icons made by *SmashIcons* from `https://www.flaticon.com`.

References

[1] Abdul Afram and Farrokh Janabi-Sharifi. Theory and applications of HVAC control systems - A review of model predictive control (MPC). *Building and Environment*, 72:343–355, 2014.

[2] Abdul Afram, Farrokh Janabi-Sharifi, Alan S. Fung, and Kaamran Raahemifar. Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy and Buildings*, 141:96–113, apr 2017.

[3] Zakia Afroz, GM Shafiullah, Tania Urmee, and Gary Higgins. Modeling techniques used in building HVAC control systems: A review. *Renewable and Sustainable Energy Reviews*, 83:64–84, 2018.

[4] Xiaodong Cao, Xilei Dai, and Junjie Liu. Building energy-consumption status worldwide and the state-of-the-art technologies for zero-energy buildings during the past decade. *Energy and Buildings*, 128:198–213, 2016.

[5] Miquel Casals, Marta Gangolells, Núria Forcada, Marcel Macarulla, Alberto Giretti, and Massimo Vaccarini. SEAM4US: An intelligent energy management system for underground stations. *Applied Energy*, 166:150–164, 2016.

[6] Young Tae Chae, Raya Horesh, Youngdeok Hwang, and Young M. Lee. Artificial neural network model for forecasting sub-hourly electricity usage in commercial buildings. *Energy and Buildings*, 111:184–194, 2016.

[7] Vitor N. Coelho, Miri Weiss Cohen, Igor M. Coelho, Nian Liu, and Frederico Gadelha Guimarães. Multi-agent systems applied for energy systems integration: State-of-the-art applications and trends in microgrids. *Applied Energy*, 187:820–832, 2017.

[8] Aris L. Dimeas and Nikos D. Hatziargyriou. Operation of a multiagent system for microgrid control. *IEEE Transactions on Power Systems*, 20(3):1447–1455, 2005.

[9] European Parliament. Directive (EU) 2018/844 of the European Parliament and of the Council of 30 May 2018 amending Directive 2010/31/EU on the energy performance of buildings and Directive 2012/27/EU on energy efficiency, 2018.

[10] Francesco Ferracuti, Alessandro Fonti, Lucio Ciabattoni, Stefano Pizzuti, Alessia Arteconi, Lieve Helsen, and Gabriele Comodi. Data-driven models for short-term thermal behaviour prediction in real buildings. *Applied Energy*, 204:1375–1387, 2017.

[11] Aurélie Foucquier, Sylvain Robert, Frédéric Suard, Louis Stéphan, and Arnaud Jay. State of the art in building modelling and energy performances prediction: A review. *Renewable and Sustainable Energy Reviews*, 23:272–288, 2013.

[12] Ali Ghahramani, Kenan Zhang, Kanu Dutta, Zheng Yang, and Burcin Becerik-Gerber. Energy savings from temperature setpoints and deadband: Quantifying the influence of building and system properties on savings. *Applied Energy*, 165:930–942, 2016.

[13] Juan Gomez-Romero, Carlos J. Fernandez-Basso, M. Victoria Cambronero, Miguel Molina-Solana, Jesus R. Campana, M. Dolores Ruiz, and Maria J. Martin-Bautista. A probabilistic algorithm for predictive control with full-complexity models in non-residential buildings. *IEEE Access*, 7:38748–38765, 2019.

[14] IEEE Smart Grid R&D. The role of control systems research in smart grids. Technical report, 2018.

[15] International Energy Agency. Transition to sustainable buildings: Strategies and opportunities to 2050. Technical report, 2013.

[16] Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castañeda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, Nicolas Sonnerat, Tim Green, Louise Deason, Joel Z Leibo, David Silver, Demis Hassabis, Koray Kavukcuoglu, and Thore Graepel. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.

[17] Hussain Kazmi, Fahad Mehmood, Stefan Lodeweyckx, and Johan Driesen. Gigawatt-hour scale savings on a budget of zero: Deep reinforcement learning based optimal control of hot water systems. *Energy*, 144:159–168, 2018.

[18] Muhammad Waseem Khan and Jie Wang. The research on multi-agent system for microgrid control and optimization. *Renewable and Sustainable Energy Reviews*, 80:1399–1411, 2017.

[19] Michaela Killian and Martin Kozek. Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment*, 105:403–412, 2016.

[20] Timilehin Labeodan, Kennedy Aduda, Gert Boxem, and Wim Zeiler. On the application of multi-agent systems in buildings for improved building operations, performance and smart grid interaction – A survey. *Renewable and Sustainable Energy Reviews*, 50:1405–1414, 2015.

[21] Yang Liu, Nam Dinh, Yohei Sato, and Bojan Niceno. Data-driven modeling for boiling heat transfer: Using deep neural networks and high-fidelity simulation results. *Applied Thermal Engineering*, 144:305–320, 2018.

[22] Yuehong Lu, Shengwei Wang, and Kui Shan. Design optimization and optimal control of grid-connected and standalone nearly/net zero energy buildings. *Applied Energy*, 155:463–477, 2015.

[23] Diana Manjarres, Ana Mera, Eugenio Perea, Adelaida Lejarazu, and Sergio Gil-Lopez. An energy-efficient predictive control for HVAC systems applied to tertiary buildings based on regression techniques. *Energy and Buildings*, 152:409–417, 2017.

[24] Miguel Molina-Solana, María Ros, M. Dolores Ruiz, Juan Gómez-Romero, and M.J. Martin-Bautista. Data science for building energy management: A review. *Renewable and Sustainable Energy Reviews*, 70:598–609, 2017.

[25] Mohammad H. Moradi, Saleh Razini, and S. Mahdi Hosseinian. State of art of multiagent systems in power engineering: A review. *Renewable and Sustainable Energy Reviews*, 58:814–824, 2016.

[26] OpenAI. OpenAI Five. `https://blog.openai.com/openai-five/`.

[27] Luis Pérez-Lombard, José Ortiz, and Christine Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.

[28] Peter Rockett and Elizabeth Abigail Hathway. Model-predictive control for non-domestic buildings: a critical review and prospects. *Building Research & Information*, 45(5):556–571, 2017.

[29] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. Tackling climate change with Machine Learning. `http://arxiv.org/abs/1906.05433`, 2019.

[30] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.

[31] U.S. Energy Information Administration. What's new in how we use energy at home: results from EIA's 2015 residential energy consumption survey (RECS), 2018.

[32] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[33] Pavel Vrba, Vladimir Marik, Pierluigi Siano, Paulo Leitao, Gulnara Zhabelova, Valeriy Vyatkin, and Thomas Strasser. A review of agent and service-oriented concepts applied to intelligent energy systems. *IEEE Transactions on Industrial Informatics*, 10(3):1890–1903, 2014.

[34] Tianshu Wei, Yanzhi Wang, and Qi Zhu. Deep Reinforcement Learning for building HVAC control. In *Proceedings of the 54th Annual Design Automation Conference (DAC '17)*, pages 1–6, New York, New York, USA, 2017. ACM Press.

[35] Yixuan Wei, Xingxing Zhang, Yong Shi, Liang Xia, Song Pan, Jinshun Wu, Mengjie Han, and Xiaoyun Zhao. A review of data-driven approaches for prediction and classification of building energy consumption. *Renewable and Sustainable Energy Reviews*, 82:1027–1047, 2018.

[36] Baris Yuce, Yacine Rezgui, and Monjur Mourshed. ANN–GA smart appliance scheduling for optimised energy management in the domestic sector. *Energy and Buildings*, 111:311–325, 2016.

[37] Rehman Zafar, Anzar Mahmood, Sohail Razzaq, Wamiq Ali, Usman Naeem, and Khurram Shehzad. Prosumer based energy management and sharing in smart grid. *Renewable and Sustainable Energy Reviews*, 82:1675–1684, 2018.

[38] Zhiang Zhang, Adrian Chong, Yuqi Pan, Chenlu Zhan, Silian Lu, Khee Poh Lam, and D Ph. A Deep Reinforcement learning approach to using whole building energy model for HVAC optimal control. In *Proceedings of the 2018 ASHRAE/IBPSA-USA Building Performance Modeling Conference and SimBuild*, pages 1–9, 2018.

[39] Muhammad Fahad Zia, Elhoussin Elbouchikhi, and Mohamed Benbouzid. Microgrids energy management systems: A critical review on methods, solutions, and prospects. *Applied Energy*, 222:1033–1055, 2018.

Supplementary material

See Figures 2–4.

More information

`https://jgromero.github.io/ia4sg/`

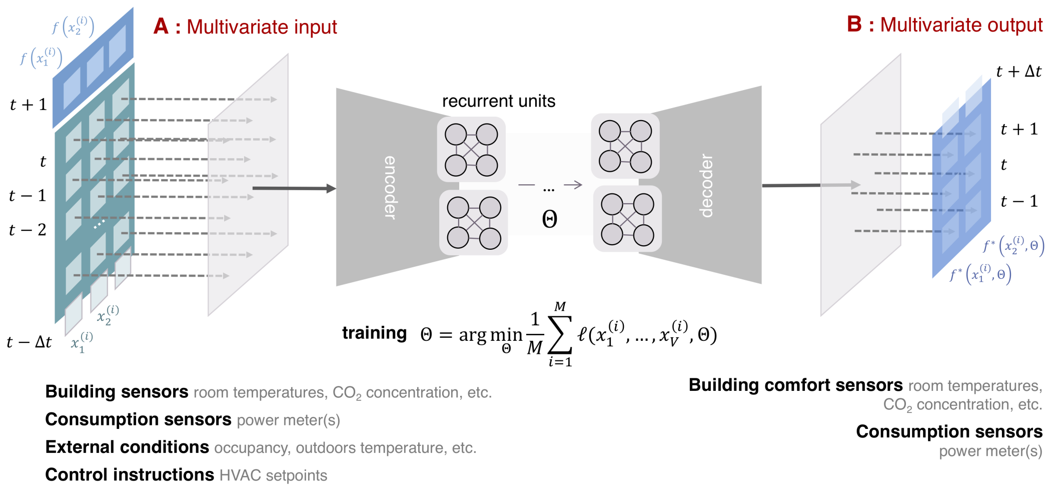

Figure 2: **Deep neural simulation models (DNSM)**: Encoder-decoder architecture for a multivariate deep recurrent neural network implementing a DNSM for HVAC. M is the number of training samples and v is the number of variables.

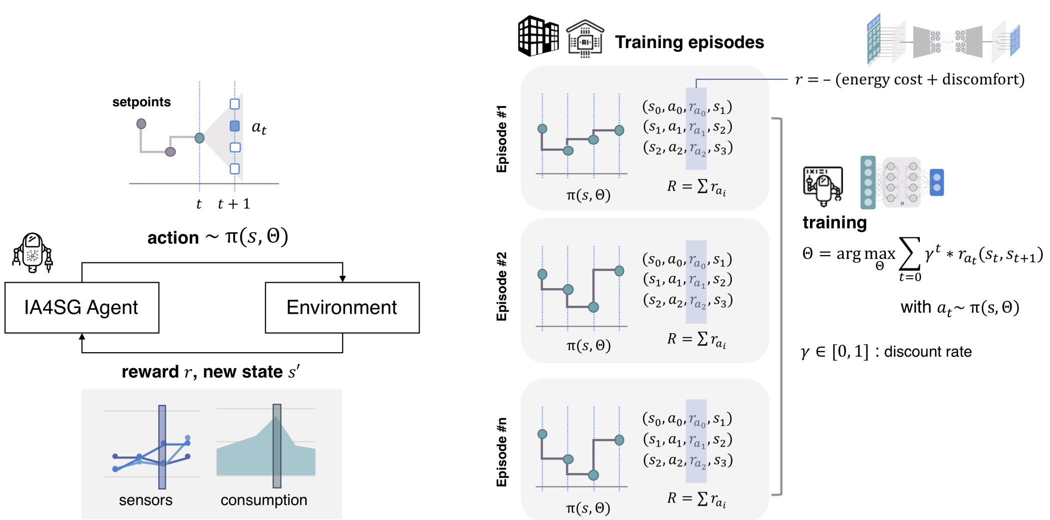

Figure 3: **Control optimization by Deep Reinforcement Learning (DRL)**: (left) building control as a DRL problem, π gives preferred action probabilities for the current state s; (right) training from multiple episodes collected from a real building and from a simulation. A DNSM is used to estimate the agent learning rewards in terms of energy cost and discomfort.

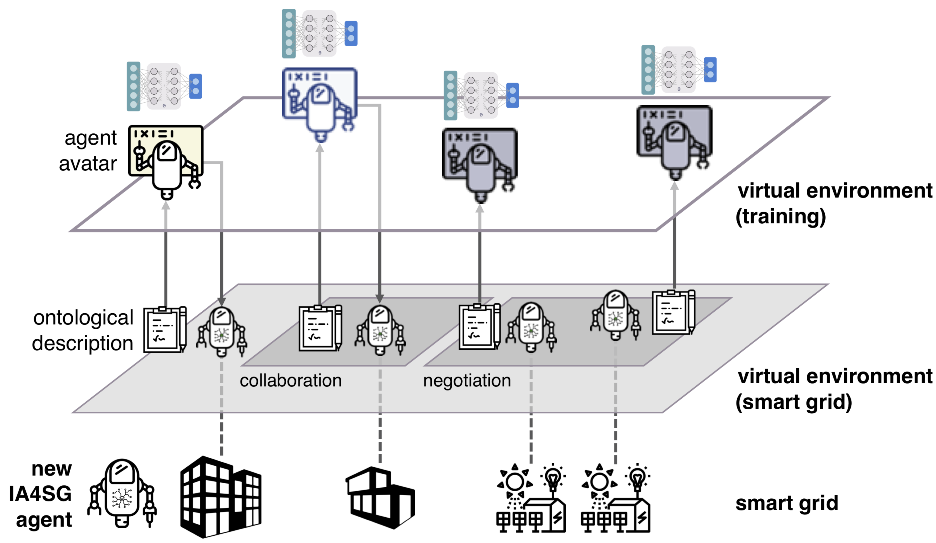

Figure 4: **Cooperation among multiple energy system actors**: An agent is trained on the virtual training environment along with the avatars of other actors in the grid by repeated simulation. After training, the new agent and the fully-collaborative agents are updated.