Machine Learning Integration for Adaptive Building Envelopes

An Experimental Framework for Intelligent Adaptive Control

Shane Ida Smith, Ph.D.
Chris Lasch
University of Arizona

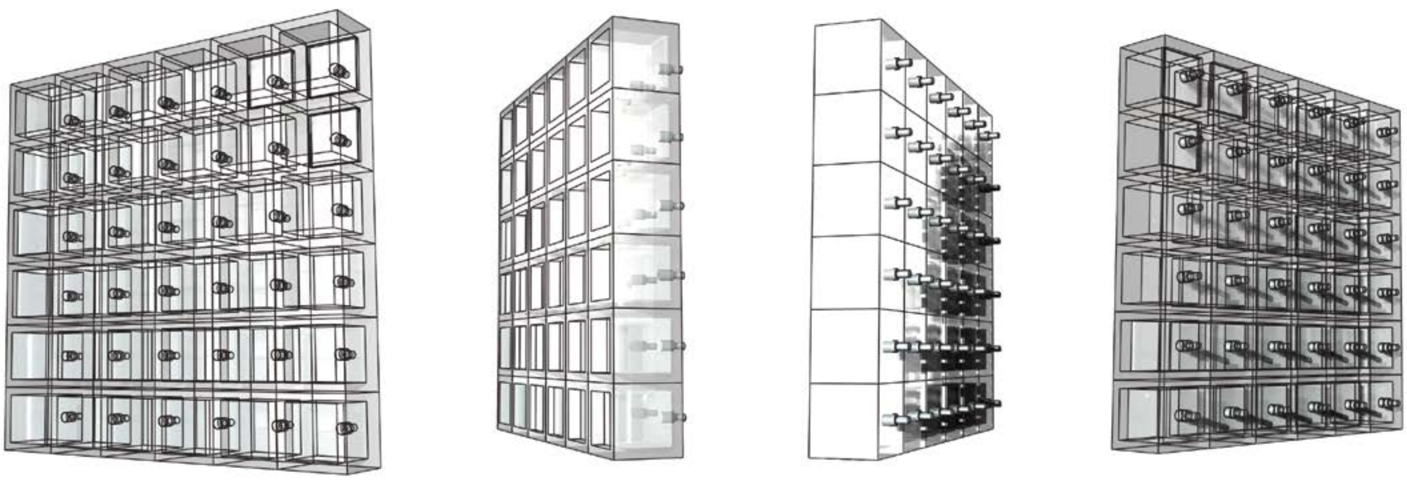

1

1 Responsive facade modules with distributed control actuation points.

ABSTRACT

This paper describes the development of an Intelligent Adaptive Control (IAC) framework that uses machine learning to integrate responsive passive conditioning at the envelope into a building's comprehensive conventional environmental control system. Initial results show that by leveraging adaptive computational control to orchestrate the building's mechanical and passive systems together, there exists a demonstrably greater potential to maximize energy efficiency than can be gained by focusing on either system individually, while the addition of more passive conditioning strategies significantly increases human comfort, health and wellness building-wide.

Implicitly, this project suggests that, given the development and ever increasing adoption of building automation systems, a significant new site for computational design in architecture is expanding within the post-occupancy operation of a building, in contrast to architects' traditional focus on the building's initial design. Through the development of an experimental framework that includes physical material testing linked to computational simulation, this project begins to describe a set of tools and procedures by which architects might better conceptualize, visualize, and experiment with the design of adaptive building envelopes. This process allows designers to ultimately engage in the opportunities presented by active systems that govern the daily interactions between a building, its inhabitants, and their environment long after construction is completed. Adaptive material assemblies at the envelope are given special attention since it is here that a building's performance and urban expression are most closely intertwined.

INTRODUCTION

Parallel advances are occurring in the fields of dynamic building facades and building automation control systems, exposing an increasingly complex terrain of dynamic systems' theory between exterior and interior built environments. Global trends in computational optimization strategies for automated control systems include the addition of intelligent control schemes, such as adaptive neuro-fuzzy inference systems, and optimization algorithms, such as multi-objective genetic algorithms, simulated annealing, meta-analysis, and others (Shaikh et al. 2014). In addition, efficiencies of conventional HVAC controllers are greatly improving, with emerging studies of applied reinforcement learning techniques indicating 4%–11% energy conservation over conventional control for heat-pumps (Ruelens et al. 2015). At the same time, emphasis on adaptive building envelope performance in response to dynamic environments is gaining heightened interest (Erickson 2013; Kolarevic and Parlac 2015; Zamella and Faraguna 2014). The ever-expanding portfolio of dynamic facade technologies exposes great promise to reduce a building's reliance on fossil-fuel based mechanical air conditioning in favor of natural, passive mechanisms that consume significantly less energy and simultaneously improve occupant well-being.[1] While each of these fields is receiving significant interest, there is not yet an explicit effort to link the two areas together for reciprocal benefits between proactive automated control systems and responsive envelope actuation functions.[2]

It is not possible to design in advance a system with a fixed control policy capable of anticipating dynamic outdoor and indoor conditions while also capitalizing on the qualitative and quantitative benefits that are possible in the synergistic interactions of these different socio-environmental control systems. In order to maximize the energy efficiency potential of these technologies, in addition to the qualitative potential for occupant experience and wellbeing, a building's environmental control operations must be considered holistically within an intelligent and adaptive framework. Such a framework shall be capable of orchestrating all of the building's systems in concert and adapt to simultaneous changes in internal and external conditions. The Intelligent Adaptive Control (IAC) architecture that we are developing is able to synthesize and adapt an integrated suite of control policies to coordinate building-wide active and passive environmental conditioning systems. IAC learns over time from sensors and history of control actions made during its operation. IAC policies constantly evolve so that its response becomes more finely tuned to the idiosyncrasies of each building's particular environmental landscape.

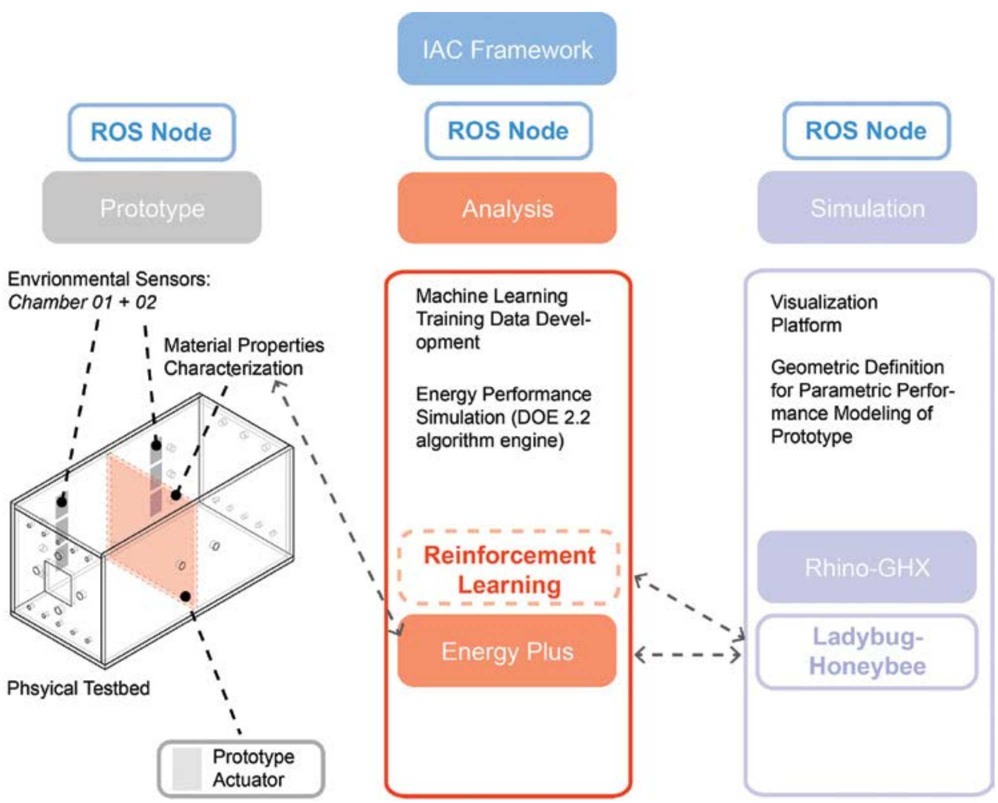

2

2 Intelligent Adaptive Control (IAC) experimental framework for building envelope integration.

In this project, our collaborators in Information Science and Electrical and Computer Engineering have identified two relevant control frameworks: Model Predictive Control and an area of machine learning known as Reinforcement Learning. Model Predictive Control (MPC) has become the dominant popular approach to HVAC control (Morari and Lee 1999; Maciejowski 2001; Ernst et al. 2009). Because MPC incorporates an accurate model of its task environment, it can anticipate future events and adjust accordingly based on decision point or fixed-horizon algorithms. MPC controllers require accurate knowledge of the operating environment conditions and become ineffective in unknown and changing operating environments. Reinforcement Learning (RL) is an area of machine learning concerned with how software agents learn to perform a series of sequential actions within an environment in order to maximize some notion of a long-term reward (Sutton and Barto 1998). A reinforcement learner does not rely on an a-priori model of its operating environment like an MPC does; it learns its optimal policy from its history of interaction with the environment. Reinforcement learning has been proposed as one approach to regulate controls within an environment as dynamic and complex as a building interior (Dalamagkidis et al. 2007).

For the conception of linking an intelligent automation system with the building envelope functions, an adaptive controller synthesis paradigm is preferred because the task environment dynamics are more uncertain. For this particular integration, our team has established a direction towards a hybrid approach to the computational control system, blending the benefits of RL with those of MPC (Peng and Morrison 2016). Comprehensively, the IAC framework engages concurrent development of physical dynamic envelope prototypes, simulation of digital design concepts, and analysis of building energy performance.

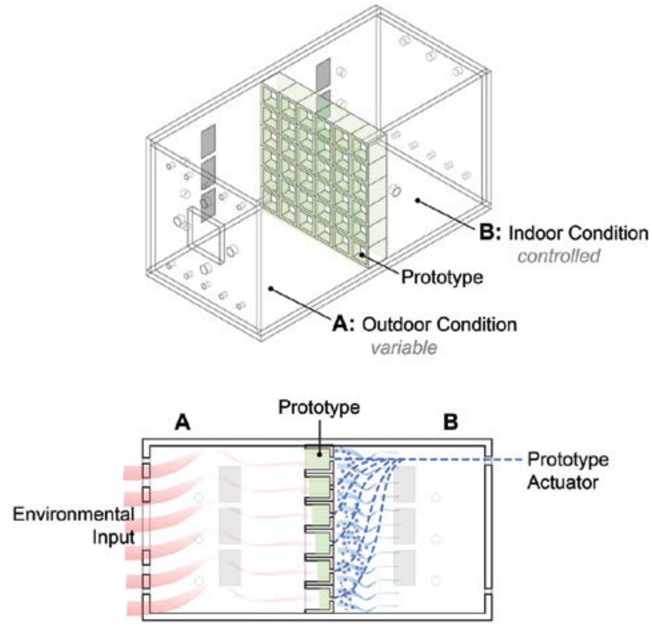

3

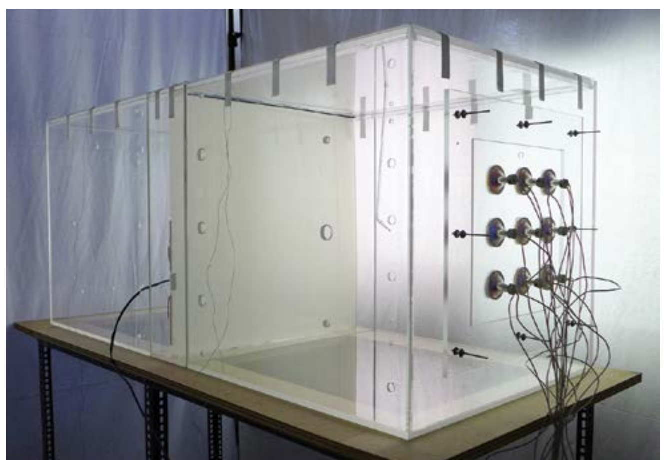

4

METHODS

Our experimental framework is an ecosystem consisting of a physical testing apparatus linked to both a digital simulation and analysis environment [Fig. 2]. A range of adaptive facade material assemblies can be inserted within the physical environmental test chamber. Digital configurations of these assemblies are simultaneously developed within a simulation environment for design purposes and in order to apply our experiments to the building scale for energy performance analysis. The bridge between these three environments is the IAC computational control framework. The IAC framework is an autonomous adaptive control architecture based on an adaptive machine-learning methodology. The IAC regulates the electronic controls within the physical testbed as well within the digital simulation. Over time, the data generated within these two experimental arenas train the IAC's control algorithms toward adaptive performance improvements.

Physical Testbed and Prototypes

The physical environmental test bed is an acrylic chamber designed as a modular kit of parts for testing a range of materials, control mechanisms and data processing models. The chamber is divided into two equal volumes separated by a slot for a removable prototype. Different active facade assembly prototypes can be inserted and tested between chamber A, representing exposure to an external environment, and chamber B, which represents a controlled internal environment [Fig. 3]. This test chamber permits experimental control of environmental conditions (humidity, temperature, light, heat flow) on each side of the testing facade and the monitoring of the response and adaptability of the apparatus to variations in conditions. The testbed is modular by design to enable experimental evaluation of the adequacy of various categories of adaptable materials and data-driven adaptive control policies within the IAC.

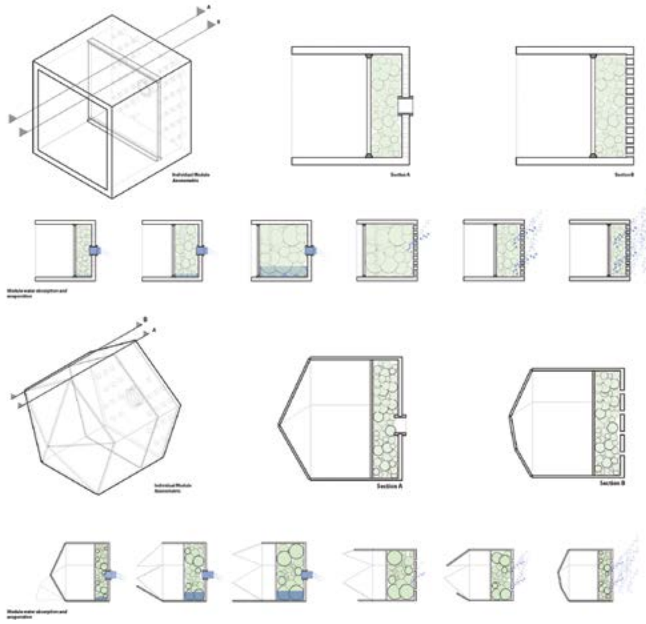

5

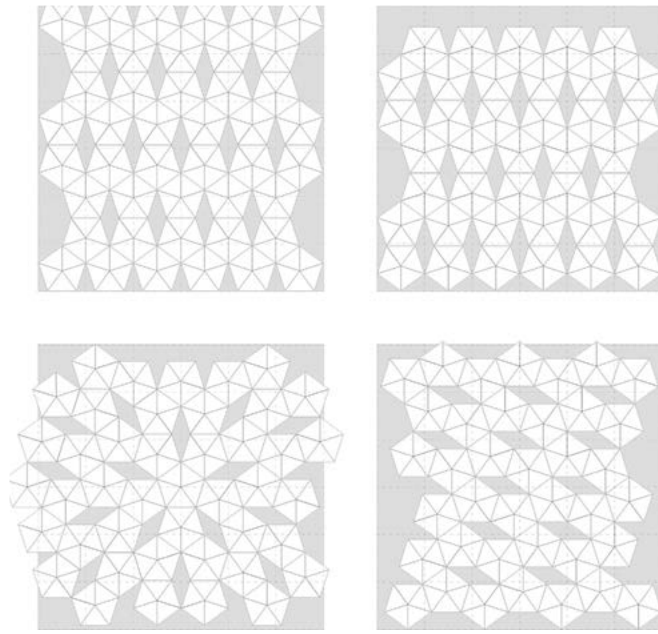

6

3 Test chamber for physical prototype insert dividing two equal chamber volumes of an environmental input (A) and a control environment (B).

4 Physical testbed construction with initial baseline dynamic glazing film and halogen dimmable lamp array input.

5 Digital design studies of dynamic responsive passive envelope conditioning hydrogel modules for evaporative cooling, heat capacitance, and natural daylighting functions.

6 Digital design studies of dynamic responsive building envelope modules in various pattern configurations.

Multi-sensory device arrays (thermocouples, photodiode, humidity sensors, infrared camera, etc.) embedded throughout the interior chamber, the facade itself, and the external space produce large-scale data flows used to generate responsive behaviors through adaptive learning. Initial physical prototype baseline studies are being prepared with electroactive photochromic dynamic glazing film technology, which responds to photometric measures in graduated increments of opacity and transparency based on a dimmable halogen lamp array input and photosensor data collection [Fig. 4]. The anticipated result of the combination of advanced facade materials with adaptive control is an autonomously responsive envelope that can maintain internal environmental conditions with appropriate performance levels compared to conventional methods. Specific targets in next-generation building energy management systems indicate the merging of sensor data and predictive statistical models to allow for more proactive modulation as signals are changing (Zavala et al. 2011). Future predictive modeling may also be linked with online sources provisioning communication from urban microclimate data from external sensory networks and utility providers (Pang, Hong, and Piette 2013).

Digital Simulation Environment

The work within the design development process of machine learning integration with adaptive building envelope and reciprocal building energy performance is conducted in the Rhino 3D – Grasshopper platform with Ladybug-Honeybee plug-ins and EnergyPlus simulations. In the current work, Python scripts access reinforcement learning algorithms that, along with weather data input and energy simulation output through base building analyses, inform dynamic changes in building envelope properties. Dynamic envelope design concepts developed with the parametric visualization tools [Figs. 5 and 6] can be correlated to dynamic properties for analysis engine input.

The framework developed for this project serves as a design process tool, in addition to informing potential building envelope technologies. There are three primary facets to the simulation process: a) defining the building envelope system and the influence of external environmental stimuli, b) determining the interior building environmental performance through dynamic envelope properties, and c) defining the learning algorithm to actuate change in properties or functionality of the building envelope system. Current building energy simulation models have two drawbacks in these areas—limitation on dynamic envelope analysis and limitation on reinforcement learning algorithm integration (Magoules and Zhao 2016; Sanyal et al. 2014).

The simulation process is more complex than current standards for energy performance models because real-time building performance results are continuously analyzed through algorithmic comparison with concurrent external stimuli to actuate

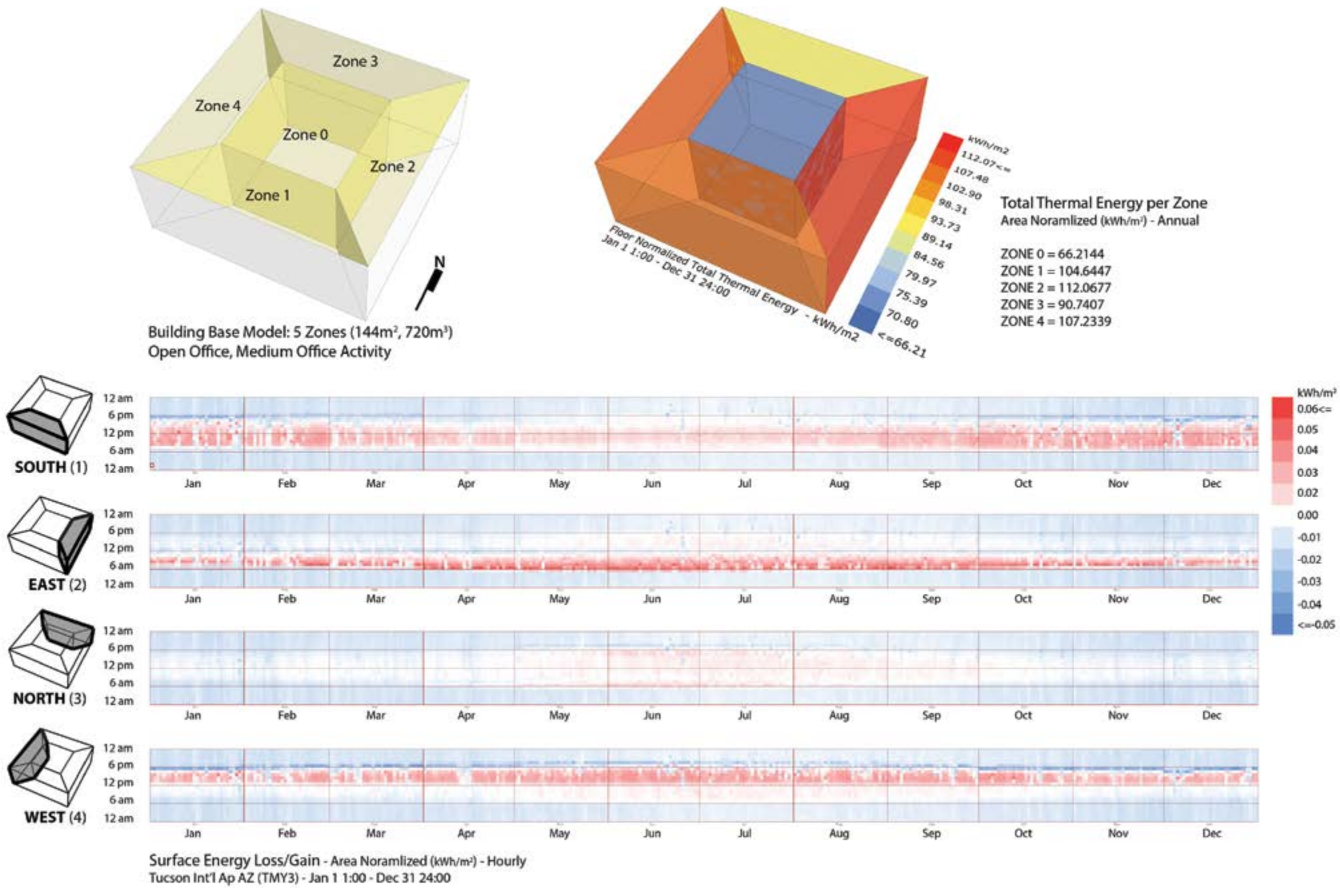

7 Base building energy analysis model establishing open office floor plan with four perimeter zones and specific analysis data for total thermal energy per zone (upper right) and surface energy losses and gains through envelope system (graphs) for each facade in cardinal orientations at hourly resolution for an annual timeframe.

change in the building envelope properties. The process is dynamic rather than static, and intends for adaptability of an envelope system beyond a two-state control process. Our current analysis framework includes a baseline building model with selective data processing for the surface energy losses and gains at the building envelope for each cardinal orientation [Fig. 7]. The data provides the MPC learning sets, which are utilized for initial building envelope property changes in response to the algorithm actuation. Further development is required to model the dynamic behaviors of envelope response stimulated through the IAC with the EnergyPlus interface for predictive environmental performance results.

Computational Control Framework

The multifaceted nature or our testing environment, containing physical as well as virtual components, supports our hybrid approach to the development of the IAC framework. The project methodology allows for two parallel sets of learning data to be developed - one in the test chamber with physical prototypes and sensors, and another in the simulation environment with analysis tools. Our project is also developing a compromise between planning and learning, where planning is represented by the framework of an MPC and learning is represented by a model-free reinforcement learning technique (Morari and Lee 1999; Maciejowski 2001; Afram and Janabi-Sharifi 2014). Both approaches seek to define a series of policies for state-dependent actions to maximize cumulative long term reward.

In planning, it is assumed that a complete model of the task environment is available and the planner induces a policy for choosing the action in each state that achieves optimal performance in terms of total long term reward. The RL approach, on the other hand, does not assume the environment is known ahead of time. Instead, the learning agent has to interact directly with the environment to gather data about the effects of its actions on the world and their reward value, and while doing so searches for an optimal policy for action.

The a-priori model of the task environment is both the strength and weakness of the planning approach. In an environment as complex as a building interior, these conditions are unlikely to be completely known in advance and may change over time.

The learning framework provides a general approach to solving sequential decision-making problems without relying on a pre-existing model of the task environment, but incurs the prohibitively high cost of real-time interaction with the environment that would require multiple parallel processors (Magoules and Zhao 2016). We can potentially get the advantages of both frameworks through a hybrid approach in which we use a suitable platform for offline training of an adaptive learning system through our complementary methodology of physical and simulation environments. In this case, the policy learned in simulation is used to initiate the learner with a reasonable performance that is then transferred to and fine-tuned in real-world interaction. This approach reduces the amount of costly real-world experience required to achieve high performance (Liu and Henze 2006; Cutler et al. 2015).

By constructing a feedback loop between actual and simulated environments, we streamline the development of the learner. At the same time, we iteratively increase the accuracy of our simulation by recalibrating it each cycle based on results recorded from the physical test chamber. The result is a prototyping environment where we can develop a novel environmental control framework that adapts to its ever changing context and continually improves its performance over time. This experimental setup is also designed to anticipate how our IAC might be employed in the field. While in operation as a building control system, a parallel simulation driven by real-time data collected from building sensors would provide an environment where alternative control policies may continuously be explored and evolve.

RESULTS

Initial work on the development of the IAC has focused on its integration with a conventional HVAC control system. Subsequent development will confront the more complex prospect of passive conditioning through an adaptive facade. In simulations of our initial model, the thermostat controller tuned via a reinforcement machine learning algorithm performed approximately 5% more efficiently than a simulated conventional automatic thermostat. We expect that this performance will improve over the course of our work towards single system performance improvements of 8–12%, and our target for accumulated efficiency between a passive facade system and a conventional HVAC system centrally controlled by the IAC technology is 25–35% (Jacobs 2003). The integration of the IAC with adaptive building envelope actuation could provide up to 50% reduction in energy demands.

CONCLUSION

Current work to date has demonstrated performance benefits from the application of an early version of the IAC. The time series of environmental conditions and the state of the facade will be analyzed to create models of system dynamics at multiple time scales. These dynamic models shall serve as the instrument for developing control algorithms to maintain desired chamber-internal environment states while optimizing for low energy consumption. This includes using adaptive learning techniques that explore control strategies under different optimizing constraints.

Our efforts have occurred through interdisciplinary collaborations with Information Sciences, Electrical and Computer Engineering, and Material Science Engineering. In order to develop robust IAC policies, further interdisciplinary connections are warranted that will enhance the possibilities for holistic evaluative frameworks in the reinforcement learning algorithms. The primary focus in near-term work required specific attention to linking the RL with the dynamic facade for parametric development of environmental data feedback informing actuation signals. The simplified physical model for the electrochromic dynamic glazing film and photometric analysis will be analyzed with concurrent development of dynamic simulation scripting in the EnergyPlus platform.

This work ultimately pursues the systematic evaluation of a range of adaptive envelope technologies with regard to environmental performance. When complete, the resulting comparative study will undoubtedly have value beyond the bounds of this project and be useful to architectural designers at large. Furthermore, our experimental testbed serves as a generalizable process and kit of parts. It may be adopted and improved upon by future designers for use as tool kit for the study and design of adaptive facades in all of their aspects.

ACKNOWLEDGEMENTS

The authors acknowledge a Renewable Energy Network (REN) Faculty Exploratory Grant award (#5825479) from the University of Arizona (UA) to make this research possible. We also thank our collaborators on the framework development: Clayton Morrison, Associate Professor in the School of Information Science, Technology, and Arts at the UA; Pierre Lucas, Professor in Material Science Engineering at the UA; and Kuo Peng, Doctoral Student in Electrical and Computer Engineering at the UA. We also thank all of the students involved on the project to date, including Alice Wilsey, Tyler Connel, Long Long, Nikki Hernandez, and Di Le.

NOTES

1. Design integration of high performance envelope systems has shown 30%–70% energy savings. "Federal R&D Agenda for Net-Zero Energy, High-Performance Green Buildings Report" National Science and Technology Council, Report of the Subcommittee on Buildings Technology Research and Development, Oct 2008, p.21.
2. The only identifiable emerging market technology in this area is the Provolta Energy OS patent pending machine learning platform by Rengen supported in part by R&D partner Lawrence Berkeley National Laboratory; however, this platform does not explicitly focus on integration with dynamic building envelopes, but rather on improving performance of HVAC control systems through building integration feedback.

REFERENCES

Afram, A. and F.Janabi-Shari. 2014. "Theory and applications of HVAC control systems-a review of model predictive control (MPC)." *Building and Environment* 72: 343–355.

Cutler, M., T.J. Walsh, and J.P. How. 2015. "Real-world reinforcement learning via multifidelity simulators." *IEEE Transactions on Robotics* 31 (3): 655–671.

Dalamagkidis, K., D. Kolokotsa, K. Kalaitzakis, and G.S. Stavrakakis. 2007. "Reinforcement learning for energy conservation and comfort in buildings." *Building and Environment* 42 (7): 2686–2698.

Erickson, James. 2013. *Envelope as Climate Negotiator: Evaluating adaptive building envelope's capacity to moderate indoor climate and energy.* Dissertation, Arizona State University.

Ernst, D., M. Glavic, F. Capitanescu, and L. Wehenkel. 2009. "Reinforcement learning versus model predictive control: A comparison on a power system problem." *IEEE Transactions on Systems, Man, and Cybernetics, PartB (Cybernetics)* 39 (2): 517–529.

Jacobs, Pete. 2003. "Small HVAC System Design Guide: Design Guidelines." *California Energy Commission* (October).

Kolarevic, Branko and Vera Parlac. 2015. *Building Dynamics: Exploring Architecture of Change*. United Kingdom: Taylor and Francis.

Liu, S. and G.P. Henze. 2006. "Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: Part1 – theoretical foundations." *Energy and Buildings* 38 (2): 142–147.

Maciejowski, J. 2001. *Predictive Control with Constraints*. Englewood Cliff, NJ: Prentice Hall.

Magoulès, Frederic and Hai-Xiang Zhao. 2016. *Data mining and machine learning in building energy analysis.* Hoboken, NJ: John Wiley & Sons.

Morari, M. and J.H. Lee. 1999. "Model predictive control: Past, present and future." *Computers & Chemical Engineering* 23 (4): 667–682.

Pang, Xiufeng, Tianzhen Hong, and Mary Ann Piette. 2013. "Improving Building Performance at Urban Scale with a Framework for Real-time Data Sharing" *Lawrence Berkeley National Laboratory* LBNL-6303E (June).

Peng, Kuo and Clayton Morrison. 2016. "Model Predictive Prior Reinforcement Leaning for a Heat Pump Thermostat." *IEEE International Conference on Automatic Computing: Feedback Computing* '16 (July).

Ruelens, F., Iacovella, S., Claessens, B., and Belmans, R. 2015. "Learning Agent for a Heat-Pump Thermostat with a Set-Back Strategy Using Model-Free Reinforcement Learning." *Energies* 8 (8): 8300–8318.

Shaikh, P. H., N.B.M. Nor, P. Nallagownden, I. Elamvazuthi, and T. Ibrahim. 2014. "A review on optimized control systems for building energy and comfort management of smart sustainable buildings." *Renewable and Sustainable Energy Reviews* 34: 409–429.

Sanyal, J., J. New, R.E. Edwards, and L. Parker. 2014. "Calibrating building energy models using supercomputer trained machine learning agents." *Concurrency and Computation: Practice and Experience* 26 (13): 2122–2133.

Sutton, R.S. and A.G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge: The MIT Press.

Zamella, Giovanni and Andrea Faraguna. 2014. *Evolutionary Optimisation of Facade Design: A New Approach for the Design of Building Envelopes*. New York, NY: Springer.

Zavala, V.M., C. Thomas, M. Zimmerman, and A. Ott. 2011. "Next-Generation Building Energy Management Systems and Implications for Electricity Markets" *Argonne National Laboratory* ANL/MCS-TM-315 (July).

IMAGE CREDITS

Figure 1: Responsive facade modules (Smith, 2016)
Figure 2: IAC Framework Methodology (Smith and Lasch, 2015)
Figure 3: Testbed diagram (Smith and Le, 2015)
Figure 4: Testbed photograph (Smith, 2016)
Figure 5: Adaptive module designs (Smith and Le, 2015)
Figure 6: Adaptive module patterns (Smith and Le, 2015)
Figure 7: Building energy analysis model results (Smith, 2016)

Shane Ida Smith is an architect, designer, educator, and philosopher. She has over fifteen years experience in professional architecture practice and is fluent with building performance analytics. She holds a Doctorate of Philosophy in Architectural Sciences from the Center for Architecture, Science, and Ecology at Rensselaer Polytechnic Institute. She is currently an Assistant Professor in Architecture at the University of Arizona conducting design research on emerging environmental building technologies.

Chris Lasch is a principal and co-founder of Aranda\Lasch, a design studio dedicated to experimental research and innovative building through a deep investigation of structure and materials. The studio was awarded the United States Artists Award and Young Architects + Designers Award in 2007, the Architectural Record Design Vanguard Award in 2014, the Architectural League Emerging Voices Award in 2015, and was named one of Architectural Digest's 2014 AD Innovators. Their early projects are the subject of the best-selling book, *Tooling*. Aranda\Lasch has exhibited its work internationally in galleries, museums, design fairs and architecture biennials.