

Biosphere 2 Ocean: Heat Exchange Optimization Through Machine Learning

Project Description

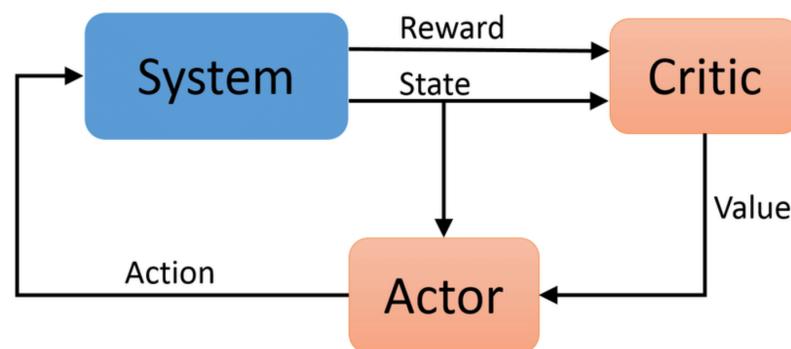
- The Biosphere 2 Ocean (B2O) uses a heat exchanger (HX-1) to maintain the setpoint temperature ($75^{\circ}\text{F} \pm 0.20^{\circ}\text{F}$).
- Amount of energy and money required to heat/cool the B2O water is very large.
- The Biosphere 2 uses a heat exchanger to regulate ocean temperature within a set point. There are currently no models relating the heating and cooling loop with optimal energy usage.
- Machine learning techniques can be used to develop a controller that adjusts temperature to optimize cost via energy consumption.

Scientific Challenges

- One prominent challenge was that of creating an effective thermodynamic model of the B2O and its surrounding thermal zones.
- Another challenge was creating a model or algorithm that can account for the large variability in Biosphere 2 Ocean thermal zone conditions.
- Access to large amount of data: finding what's important to energy optimization.

Potential Applications

- A controller utilizing machine learning techniques could be run year-round to control the temperature and flow rate of the B2O heat exchanger.
- Application of an RL controller like the one described here could be extended to similar thermal zones which require temperature set point maintenance.



Flow chart diagram describing the architecture of the reinforcement learning control used in the 2017 paper by Wang et al [1].

Team Members (by work for this method):

Carl Solway
Sasha Sepulveda
Lyra Troy

Methodology

1. A model-free actor-critic Reinforcement Learning (RL) controller is designed using a variant of artificial recurrent neural networks – Long Short-Term Memory (LSTM) networks based on [1].
2. The RL algorithm aims to maximize a total accumulated sum of rewards over time, given a previously unknown environment through a trial-and-error learning process. This occurs in two phases.
3. At each time step, the actor chooses an action, which is taken on the system. The system then posts a state value, and a reward value, depending on how closely the system matches some optimal state.
4. The critic LSTM network computes a state action value, which in conjunction with the state of the system, gives the actor LSTM network the information required to take a new action at the next time step. The entire process then repeats.
5. The LSTM's "learn" by adjusting the weights of the various nodes in the LSTM's hidden layer. To update these weights, backpropagation of the gradient vector is used.

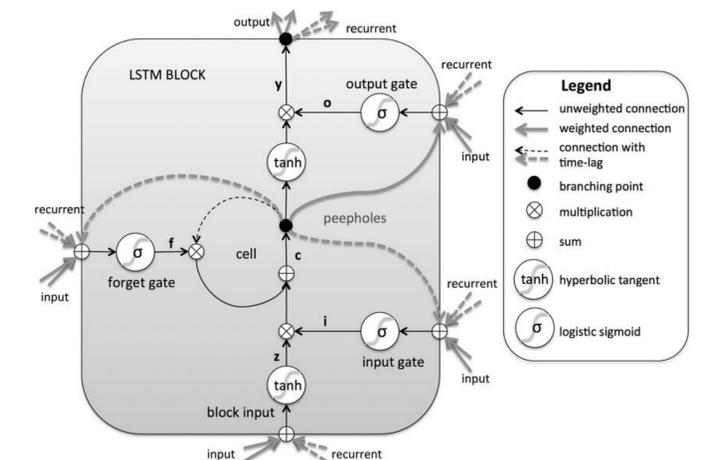
Results

1. Ultimately, this approach to constructing a control system for the Biosphere 2 Ocean heat exchanger did not pass the theoretical stage. This was due to a lack of a workable simulation of the system for an RL controller to optimize and learn from.
2. A basic LSTM Network was constructed using MATLAB's deep learning toolbox. As discussed previously, training the network proved untenable.
3. It is important to note Lyra Troy's thermodynamic modelling of the system culminated in the creation of a MATLAB GUI capable of approximating system input output pairings. It is possible that this GUI could be paired with the LSTM RL architecture described above as a starting point to train a rudimentary controller.

Glossary of Technical Terms

Long-Short-Term Memory (LSTM) Network: A variation on a vanilla Recurrent Neural Network that includes a set number of hidden layers, with input, output, and forget gates that modulate the flow of information into and out of the hidden layer cell state.

Recurrent Neural Network: A multi-layer series of interconnected, weighted nodes. Information is input in the form of an array, and then fed through the nodes. High activations in certain node sequences allow for the approximation of complex, nonlinear functions without the use of models.



General architecture diagram of an LSTM network.

References

1. Wang, Y.; Velswamy, K.; Huang, B. A Long-Short Term Memory Recurrent Neural Network Based Reinforcement Learning Controller for Office Heating Ventilation and Air Conditioning Systems. *Processes* 2017, 5, 46.
2. Singh, Satinder, et al. "Advances in Neural Information Processing Systems 12." MIT Press, Policy Gradient Methods for Reinforcement Learning with Function Approximation, 2000, pp. 1057–1063.
3. Yang, K. (July 9, 2008). "Artificial Neural Networks (ANNs): A New Paradigm for Thermal Science and Engineering." *ASME. J. Heat Transfer*. September 2008; 130(9): 093001. <https://doi.org/10.1115/1.2944238>

Acknowledgments

This project was mentored by Sam McLaren, whose help is acknowledged with great appreciation. Additionally we would like to thank all the individuals at the Biosphere 2 who took time out of their schedules to show us the premises and provide the data necessary to conduct subsequent analysis.