

AI-Based Solar Thermal Cooling Optimization for Large-Scale Data Centers: A Sustainable Approach

Robinson García^{1*}, Carlos Garzón², Karen Betancout³, Luis Caicedo⁴

¹Universidad Internacional de Investigación México, Postgraduate Program, Post grado, 77500, Cancún-México

* **Corresponding Author Email:** robinson.garcia@uiimex.edu.mx- **ORCID:** 0009-0005-0929-3844

²Universidad Internacional de Investigación México, Postgraduate Program, Post grado, 77500, Cancún-México

Email: carlo2s@gmail.com- **ORCID:** 0009-0005-1414-5547

³Universidad Internacional de Investigación México, Postgraduate Program, Post grado, 77500, Cancún-México

Email: kare2n@gmail.com - **ORCID:** 0000-0002-1886-032X

⁴Universidad Internacional de Investigación México, Postgraduate Program, Post grado, 77500, Cancún-México

Email: lui2s@gmail.com - **ORCID:** 0000-0003-1904-7303

Article History:

DOI: 10.22399/ijasrar.57

Received: Feb. 23, 2026

Revised: Apr. 10, 2026

Accepted: Apr. 13, 2026

Keywords:

Artificial Intelligence,
Data Center,
Solar thermal Cooling,
Reinforcement Learning,
Sustainability

Abstract: This study investigates the performance of three cooling configurations for a large-scale (1 MW) data center in a high solar irradiance region. The configurations include: (i) a conventional electric chiller system, (ii) a solar thermal cooling system with PID control, and (iii) an AI-optimized solar thermal cooling system using Deep Reinforcement Learning (DRL). A dynamic co-simulation environment was developed integrating TRNSYS, EnergyPlus, and a Python-based PPO agent to evaluate the trade-off between energy efficiency and water sustainability. The results demonstrate that the AI-based system reduces grid electricity consumption by 79% and water usage by 92% compared to the conventional baseline. Furthermore, it significantly improves thermal stability, achieving a Root Mean Square Error (RMSE) of 0.24°C relative to the setpoint, and shortens the response time to load disturbances to 11 seconds. This work validates that AI-enhanced solar cooling is a scalable and sustainable solution for data centers in water-stressed regions..

1. Introduction

The exponential growth in demand for digital services driven by technologies such as artificial intelligence (AI), cloud computing, and advanced analytics has led to an unprecedented expansion of data centers worldwide [8].

This evolution comes with significant energy consumption and a concerning water footprint, particularly in regions with limited natural resources. Large data centers can consume over 25 million liters of water per megawatt annually for cooling needs, posing a critical environmental challenge [6].

Recent studies underscore the urgency for sustainable cooling solutions [13, 15]. While conventional mechanical chillers are effective, they are energy-intensive and rely heavily on evaporative cooling towers, which consume vast amounts of water. Alkrush provide a critical review noting that purely mechanical solutions are reaching efficiency limits, prompting the exploration of renewable-energy-assisted cooling [1].

Similarly, Azarifar highlight that while liquid cooling offers promise, retrofitting existing air-cooled facilities remains a substantial barrier [5]. Other works have explored waste heat utilization [3] and customized cooling for fluctuating temperatures [7].

Solar thermal cooling has emerged as a viable alternative, capable of reducing electrical load by harnessing solar radiation [14]. However, its application is hindered by the intermittency of solar input [12, 11] and the complex, non-linear dynamics required to integrate it with 24/7 data center operations.




Figure 1. Comparison of typical water and energy consumption between traditional data center cooling and sustainable solar-assisted models (Illustrative comparison based on literature data).

Traditional control methods, such as Proportional-Integral-Derivative (PID) loops, often struggle to manage these hybrid systems efficiently. They typically maintain setpoints via reactive control but cannot proactively optimize for efficiency or anticipate thermal load changes. Emerging research suggests that AI-based controllers [10], specifically Deep Reinforcement Learning (DRL), can outperform classical controls by learning optimal policies through interaction. For instance, Hosseinionari and Seethaler [9] demonstrated that DRL integrated with predictive control yields superior thermal management in complex industrial processes compared to static controllers. Despite these advances, there is a gap in the literature regarding the holistic optimization of solar thermal cooling specifically for water sustainability in data centers. Most studies, such as those by Almaraz Almaraz et al. [2], focus on industrial applications or purely on energy savings. This study addresses this gap by developing a high-fidelity co-simulation to rigorously quantify the benefits of an AI-driven approach. Objective: The primary objective of this paper is to evaluate and validate a Deep Reinforcement Learning control strategy that optimizes the trade-off between energy efficiency, water sustainability, and thermal stability in a solar-assisted data center cooling system compared to conventional baselines

2. Methodology

2.1 Case Study Description

A representative 1 MW data center model located in Phoenix, Arizona, USA, was developed for this study. This location was selected to evaluate the system under challenging environmental conditions, characterized by high annual solar irradiance ($\sim 6.5 \text{ kWh/m}^2/\text{day}$) and significant water scarcity constraints. The simulation parameters were defined based on standard industry specifications: the IT equipment generates a constant thermal load of 1000 kW, and the server room thermal environment is maintained at 24°C , strictly adhering to ASHRAE TC 9.9 thermal guidelines (allowable range $18\text{--}27^\circ\text{C}$). Weather data was sourced from the Typical Meteorological Year (TMY3) database to ensuring statistically representative climatic conditions.

2.2 System Architectures and Technical Specifications

Three cooling system architectures are evaluated. Key physical parameters were selected based on standard engineering practices and manufacturer data.

- **System A (Conventional Baseline):** A standard water-cooled centrifugal chiller system (1 MW capacity) with cooling towers, regulated by classical PI control. A nominal COP of 5.0 was assumed for the baseline electric chiller.
- **System B (Solar Thermal + PID):** Integrates 300 m² of flat-plate solar collectors tilted at 30° (latitude optimization), a 30 m³ stratified thermal storage tank, and a single-effect LiBr/H₂O absorption chiller (COP ~0.7). The collector loop utilizes a 50/50 water–propylene glycol mixture to prevent freezing and corrosion. The collectors are modeled with an optical efficiency (n_0) of 0.70. Control is handled by conventional PID loops².
- **System C (Solar Thermal + AI):** Utilizes the same physical hardware and fluid properties as System B but replaces the control logic with a Deep Reinforcement Learning (DRL) agent. The agent manages valves, pumps, and backup chillers to optimize performance.




Figure 2. Conceptual architecture of the proposed system integrating solar thermal collectors, a stratified thermal storage tank, an absorption chiller, and an AI-based intelligent controller.




Figure 3. Schematic architectures of the three cooling systems compared. Top: System A (Baseline); Middle: System B (Solar + PID); Bottom: System C (Solar + AI).

2.3 AI Model Reproducibility and Experimental Setup

To ensure the reproducibility of the proposed Deep Reinforcement Learning (DRL) approach, this section provides the complete specification of the dataset, model architecture, and training hyperparameters. The control logic relies on the Proximal Policy Optimization (PPO) algorithm implemented using the Stable-Baselines3 library with a PyTorch backend.

2.3.1 Dataset and Simulation Environment

To ensure the reproducibility of the proposed Deep Reinforcement Learning (DRL) approach, this section provides the complete specification of the dataset, model architecture, and training hyperparameters. The control logic relies on the Proximal Policy Optimization (PPO) algorithm implemented using the Stable-Baselines3 library with a PyTorch backend.

Table 1. Dataset and Environment Specifications

Feature Category	Parameter Description	Range / Characteristics	Source
State Space (Inputs) (Normalized 0-1)	Outdoor Dry Bulb Temp.	28°C – 42°C (Summer profile)	TMY3 (Phoenix, AZ)
	Solar Irradiance (GHI)	0 – 1000 $\frac{W}{m^2}$	TMY3 (Phoenix, AZ)
	Rack Inlet Temperature	18°C – 28°C	EnergyPlus Sensor
	Thermal Storage Temp.	60°C – 95°C	TRNSYS Sensor
	Current Valve Position	0 – 100%	System Feedback
Action Space (Outputs) (Normalized -1 to 1)	Valve Mixing Ratio	Continuous [0, 1]	Agent Decision
	Abs. Chiller Setpoint	Continuous Offset $\pm 2^\circ\text{C}$	Agent Decision
	Backup Chiller Threshold	Continuous [24, 26] $^\circ\text{C}$	Agent Decision
Data Split	Training Set	3,000 Episodes	Randomized Weather
	Validation Set	500 Episodes	Unseen Weather Data

2.3.2 Dataset and Simulation Environment

The DRL agent employs an Actor-Critic architecture based on Multi-Layer Perceptrons (MLP). The network configuration is detailed in Table 2.

Table 2. Deep Neural Network Architecture (PPO Policy)

Component	Layer Type	Units / Neurons	Activation Function
Shared Feature Extractor	Input Layer	5 (State Vector)	--
	Dense (Fully Connected)	64	Tanh
	Dense (Fully Connected)	64	Tanh
Actor Head (Policy)	Dense (Output)	3 (Gaussian Dist.)	Linear
Critic Head (Value)	Dense (Output)	1 (Value Estimate)	Linear
Total Parameters	--	4,800	--

2.3.3 Training Procedure and Hyperparameters

The agent was trained to maximize a reward function that penalizes energy consumption, water usage, and thermal deviations. To ensure result replicability, a fixed random seed (42) was set. The hardware environment consisted of a workstation with an NVIDIA RTX A5000 GPU. The specific hyperparameters used are listed in Table 3.

Table 3. Training Hyperparameters

Hyperparameter	Value	Description
Algorithm	PPO	Proximal Policy Optimization
Optimizer	Adam	Adaptive Moment Estimation
Learning Rate (α)	3×10^{-4}	Linear decay schedule
Discount Factor (γ)	0.99	Emphasis on long-term rewards
Batch Size	64	Minibatch size
n_steps	2048	Steps per update
Clip Range (ϵ)	0.2	Policy update constraint
Entropy Coefficient	0.01	Encourages exploration

2.3.4 Reward Function Formulation

The core of the DRL agent's learning process is governed by the reward function, designed to balance competing objectives: minimizing energy consumption, eliminating water waste, and maintaining thermal comfort. The reward r_t at each time step t is defined as follows (**Eq. 1**):

$$r_t = -(w_E \cdot E_{grid,t} + w_W \cdot W_{water,t} + w_T \cdot \phi(T_{rack,t}) + w_\Delta \cdot \Delta u_t) \quad (1)$$

Where:

$E_{grid,t}$ is the electrical energy consumed (kWh).

$W_{water,t}$ is the water consumption (m^3).

$\phi(T_{rack,t})$ is a penalty function for temperature deviations outside the ASHRAE allowable range (18-27 $^\circ\text{C}$).

Δu_t represents the change in control actions (smoothness penalty) to prevent actuator wear.

w_E, w_W, w_T, w_Δ are the weighting factors derived from the ablation study.

3. Results and discussion

The simulation results were analyzed to quantify the performance of the three systems. To address the requirement for rigorous validation, quantitative metrics including Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Standard Deviation (SD) were calculated.

3.1 Thermal Stability and Error Analysis

Figure 4 illustrates the rack inlet temperature profile over a typical hot day. System A (Conventional) exhibits oscillations of ± 1.5 due to the reactive nature of PI control. System B (Solar+PID) improves stability to ± 1.0 °C. However, System C (AI) maintains the temperature within a tight band of ± 0.5 °C.




Figure 4. Data center return air temperature (rack inlet temperature) over a typical hot day for each system. System C (Green) demonstrates superior stability with minimal variance.

Quantitative error analysis confirms these visual observations. As shown in Table 4, System C achieves an RMSE of 0.24°C, representing an 83.4% improvement over the baseline. The response time to load disturbances was reduced from 42 seconds (PID) to 11 seconds (AI).

Table 4. Quantitative Performance Metrics (Mean $\pm$ SD over 5 runs)

Metric	System A (Base-line)	System B (So-lar+PID)	System C (So-lar+AI)	Improvement (C vs A)
Thermal Control				
MAE °C	1.12 $\pm$ 0.45	0.65 $\pm$ 0.20	0.18 $\pm$ 0.05	83.9%
RMSE °C	1.45	0.88	0.24	83.4%
Response Time (s)	180 $\pm$ 30	42 $\pm$ 10	11 $\pm$ 2	93.8%
Sustainability				
Avg. COP	3.4 $\pm$ 0.2	4.7 $\pm$ 0.5	6.1 $\pm$ 0.3	79.4%
Water Use (m ³ /day)	17.8 $\pm$ 1.2	3.9 $\pm$ 0.8	1.4 $\pm$ 0.3	92.1%




Figure 5. Daily coefficient of performance (COP) comparison with error bars representing standard deviation.

3.2 Energy and Water Sustainability

The Coefficient of Performance (COP) analysis is presented in Figure 5. The AI-optimized System C achieves an average COP of 6.1, compared to 3.4 for the baseline. This corresponds to a 79% reduction in grid electricity usage for cooling.

Regarding water sustainability, Figure 6 highlights a massive reduction. System C consumes only ~ 1.4 m³/day versus 17.8 m³/day for System A. This 92% reduction is achieved because the AI agent intelligently schedules heat rejection during cooler nighttime hours and maximizes dry cooling modes.




Figure 6. Daily water consumption for cooling. System C achieves a 92% reduction relative to the conventional baseline.

3.3 Model Convergence and Robustness

The training convergence of the DRL agent is shown in Figure 7. The agent reaches a stable policy after approximately 2,500 episodes, maximizing the reward function which balances energy, water, and thermal penalties.




Figure 7. Reinforcement Learning training convergence. The shaded area represents ± 1 Standard Deviation across training runs.




Figure 8. COP sensitivity to solar irradiance variation (Robustness test).

Sensitivity analysis (Figure 8) further demonstrates that System C maintains high efficiency (COP > 5.0) even when solar irradiance drops to 80% of nominal values, whereas the PID-controlled System B suffers significant performance degradation.

3.4 Ablation Study

An ablation study was conducted to validate the reward function components. When the "smoothness penalty" (W_{Δ}) was removed from the reward function (Eq. 1), the system's energy efficiency remained high, but the thermal RMSE degraded by 40% (from 0.24°C to 0.34°C). This confirms that the multi-objective formulation is essential for achieving both stability and efficiency simultaneously.

4. Conclusions

This study presented a comprehensive evaluation of an AI-optimized solar thermal cooling system for data centers. By integrating Deep Reinforcement Learning with a high-fidelity co-simulation model, the research rigorously validated the following contributions:

1. **Precision Control:** The DRL agent achieved a thermal RMSE of 0.24°C, reducing temperature variability by over 83% compared to the conventional baseline and shortening response times to 11 seconds.
2. **Sustainability Impact:** The proposed solution demonstrated a transformative reduction in resource consumption, lowering water usage by 92% (1.4 m³/day vs. 17.8 m³/day) and increasing energy efficiency (COP) to 6.1.
3. **Robustness:** The AI controller exhibited superior adaptability to environmental fluctuations compared to standard PID loops.

Future Work: Future research will focus on: (1) Experimental validation through a pilot implementation in a physical data center; (2) Integrating IT workload prediction into the state space to enable proactive cooling strategies; and (3) Evaluating the generalization of the model across different climatic zones (e.g., tropical and temperate).

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript

References

- [1] Alkrush, A., Mohamed, S., Abdelrehim, O., & Hegazi, A. A. (2024). Data centers cooling: A critical review of techniques, challenges, and energy saving solutions. *International Journal of Refrigeration*, 160, 246-262. <https://doi.org/10.1016/j.ijrefrig.2024.02.007>
- [2] Almaraz Almaraz, V., Méndez Lagunas, L. L., & Rodríguez Ramírez, J. (2024). Potencial de la energía solar para aplicaciones en la industria alimentaria en los PODEBIS. *Tendencias en Energías Renovables y Sustentabilidad*, 3(1), 226-230. <https://doi.org/10.56845/terys.v3i1.219>
- [3] Amiri, L., Madadian, E., Bahrani, N., & Ghoreishi-Madiseh, S. (2021). Techno-Economic Analysis of Waste Heat Utilization in Data Centers: Application of Absorption Chiller Systems. *Energies*, 14(9), 2433. <https://doi.org/10.3390/en14092433>

- [4] Arroyo, J., Manna, C., Spiessens, F., & Helsen, L. (2022). Reinforced model predictive control (RL-MPC) for building energy management. *Applied Energy*, 309, 118346. <https://doi.org/10.1016/j.apenergy.2021.118346>
- [5] Azarifar, M., Arik, M., & Chang, J.-Y. (2024). Liquid cooling of data centers: A necessity facing challenges. *Applied Thermal Engineering*, 247, 123112. <https://doi.org/10.1016/j.applthermaleng.2024.123112>
- [6] Bakar Siddik, M., Shehabi, A., & Marston, L. (2021). The environmental footprint of data centers in the United States. *Environmental Research Letters*, 16. <https://doi.org/10.1088/1748-9326/abfba1>
- [7] Borkowski, M., & Piłat, A. K. (2022). Customized data center cooling system operating at significant outdoor temperature fluctuations. *Applied Energy*, 306, 117975. <https://doi.org/10.1016/j.apenergy.2021.117975>
- [8] Colomina Saló, C., & Galceran-Vercher, M. (2024). Las otras geopolíticas de la inteligencia artificial. *Revista CIDOB d'Afers Internacionals*, 138, 27-50. <https://doi.org/10.24241/rcai.2024.138.3.27>
- [9] Hosseinionari, H., & Seethaler, R. (2024). The integration of Model Predictive Control and deep Reinforcement Learning for efficient thermal control in thermoforming processes. *Journal of Manufacturing Processes*, 115, 82-93. <https://doi.org/10.1016/j.jmapro.2024.01.085>
- [10] Kwame Nti, E., Cobbina, S. J., Attafuah, E. E., Opoku, E., & Amoah Gyan, M. (2022). Environmental sustainability technologies in biodiversity, energy, transportation and water management using artificial intelligence: A systematic review. *Sustainable Futures*, 4, 100068. <https://doi.org/10.1016/j.sfr.2022.100068>
- [11] Milidonis, K., Blanco, M. J., Grigoriev, V., Panagiotou, C. F., Bonanos, A. M., Constantinou, M., ... & Asselineau, C. A. (2021). Review of application of AI techniques to Solar Tower Systems. *Solar Energy*, 224, 500-515. <https://doi.org/10.1016/j.solener.2021.06.009>
- [12] Reza Maghami, M., & Mutambara, A. G. (2023). Challenges associated with Hybrid Energy Systems: An artificial intelligence solution. *Energy Reports*, 9, 924-940. <https://doi.org/10.1016/j.egyr.2022.11.195>
- [13] Sartori, V., Vidwans, A., & Bakhar, T. (2021). Data center cooling efficiency and water sustainability. *Journal of Thermal Science*, 30, 112-124.
- [14] Temiz, M., & Dincer, I. (2022). A newly developed solar-based cogeneration system with energy storage and heat recovery for sustainable data centers: Energy and exergy analyses. *Sustainable Energy Technologies and Assessments*, 52, 102145. <https://doi.org/10.1016/j.seta.2022.102145>
- [15] Zhang, Y., Zhao, Y., Dai, S., Nie, B., Ma, H., Li, J., ... & Ding, Y. (2022). Cooling technologies for data centres and telecommunication base stations – A comprehensive review. *Journal of Cleaner Production*, 334, 130280. <https://doi.org/10.1016/j.jclepro.2021.130280>