IESO's Grid Innovation Fund

July 2025

Energy-efficient AI-powered Autonomous Greenhouses

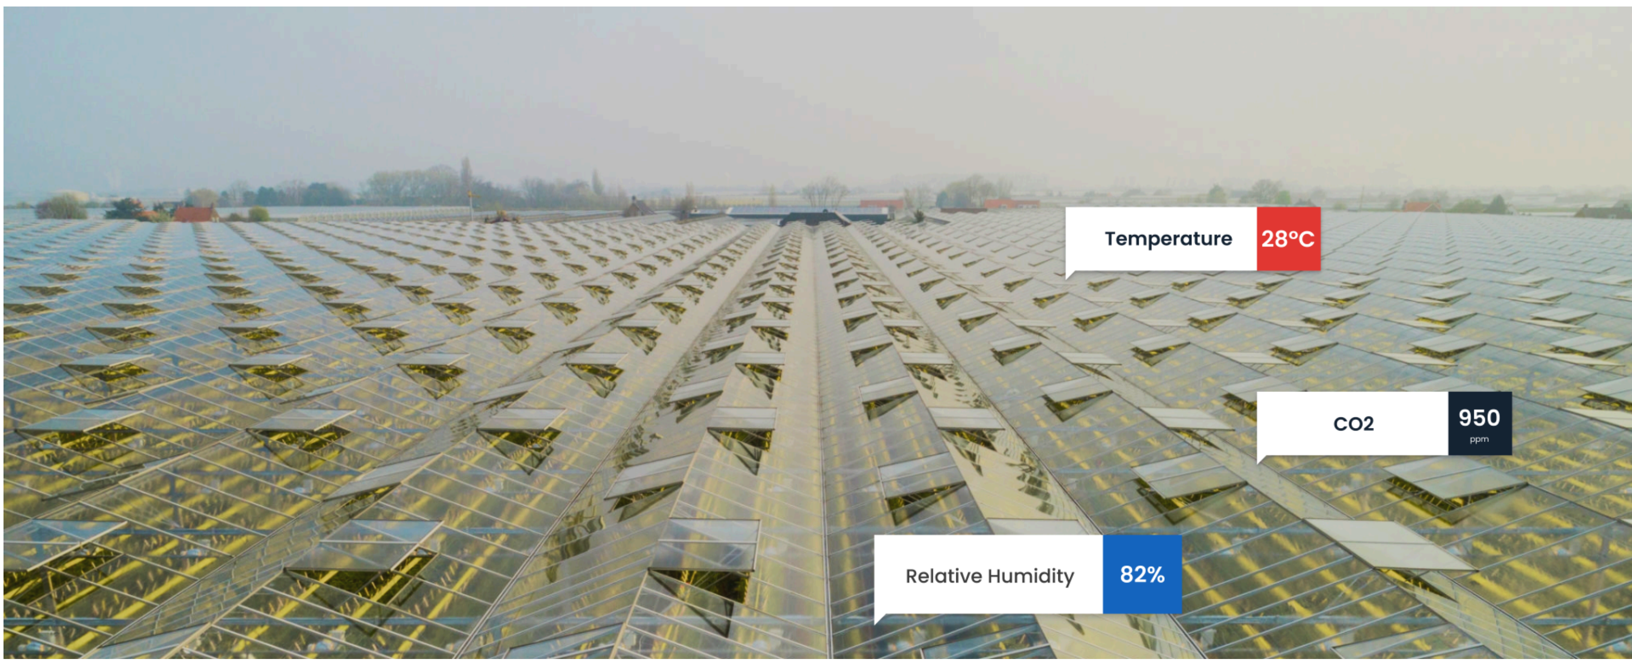

Grant Recipient: Great Lakes Greenhouses Inc

Technology Partner: Koidra Inc

Prepared by: Adele N. Tran (Koidra)

Connecting Today. Powering Tomorrow.

Disclaimer: This project is supported by the financial contribution of the Independent Electricity System Operator (IESO), through its Grid Innovation Fund. However, the views, opinions and learnings expressed in this report are solely those of Great Lakes Greenhouses Inc.

Table of Contents

Executive Summary

1. Project Motivation and Goals

1.1. Motivation

Modern commercial greenhouses are high-energy facilities facing increasing pressure to balance productivity with sustainability. Among the most pressing challenges are the escalating energy demands for lighting and climate control. In fully-lit greenhouses, lighting alone can account for over 90% of total electricity use. This issue is magnified in Ontario, where the greenhouse sector has been expanding rapidly to meet year-round food production demands in North America. With predictions of a 50% increase in greenhouse production and a corresponding 180% spike in electricity consumption over five years, the strain on the grid is already significant.

Despite technological progress, many greenhouse operations remain hampered by outdated control systems. Most growers are forced to rely on rule-based setpoints and fragmented sensor data, which cannot be integrated or interpreted cohesively due to proprietary systems with limited or no APIs. The result is an underutilization of available data and an over-reliance on human intuition.

At the same time, a generational shift is underway. As veteran growers retire and labor shortages persist, the sector risks losing crucial expertise. There is a clear need for intelligent, scalable tools that can preserve and replicate expert-level decision-making while improving overall operational efficiency.

Artificial Intelligence (AI), specifically in the form of model-based reinforcement learning (MBRL), presents a compelling solution. Previous research led by Koidra's founder, Dr. Kenenth Tran, demonstrated AI's potential to outperform even top-tier Dutch growers in yield and profitability during the Autonomous Greenhouse Challenge, organized by Wageningen University & Research. However, this success had not been validated in a large-scale, production-grade commercial greenhouse. This project was motivated by the need to bridge that gap and bring cutting-edge AI control systems into real-world horticulture.

1.2. Project Goals

The overarching goal of this project was to design, implement, and validate an AI-powered autonomous control system, hereafter coded as KoPilot, for a commercial greenhouse environment. KoPilot and its supporting software infrastructure would be specifically tested at Great Lakes Greenhouses (GLG), one of North America's largest cucumber growers.

More specifically, the project set out to

1. **Develop a scalable, interoperable software and hardware infrastructure** that could collect, clean, and analyze real-time sensor data from legacy systems (e.g. Hoogendoorn, Priva, Argus, etc.), regardless of vendor constraints.
2. **Create a grower-informed expert system** that formalized the decision logic of top horticulturalists, enabling a baseline AI that could operate with expert-level consistency.
3. **Build and train a reinforcement learning-based autonomous agent** that not only mimics but surpasses human control over time through continuous learning and feedback integration.
4. **Improve key performance metrics** such as yield, energy efficiency, and operational reliability compared to the status quo of manual grower control.
5. **Demonstrate safe, real-world AI deployment** in a high-stakes commercial greenhouse, with rigorous performance validation and safety mechanisms (manual override, recommendation mode, real-time monitoring).
6. **Evaluate the generalizability of the AI system** to other crops, environments, and industrial domains, including its adaptability to Multi-Unit Residential Buildings (MURBs) and other high-energy infrastructures.

This project was not about incremental improvement. It was a full-system redesign of how a greenhouse could be operated, replacing static rules with adaptive intelligence, and replacing manual control with scalable autonomy. The ultimate goal: create a blueprint for the next generation of energy-efficient and climate-smart agriculture, powered by AI.

2. Milestones Summary and Key Findings

2.1. Milestone 1: Software Infrastructure

Objective

Establish the software foundation required for AI-based greenhouse control by enabling automated data collection, integration with legacy systems, and remote monitoring.

Key Achievements

1. Successfully built an end-to-end IoT data infrastructure capable of interacting with both modern sensors and legacy control systems like Hoogendoorn's iSii.
2. Developed a custom solution using Robotic Process Automation (RPA) and Optical Character Recognition (OCR) to interface with Hoogendoorn, which lacks a programmatic API. This was essential for enabling automated data extraction and control setpoint adjustments.
3. Deployed an edge-computing-based IoT Hub for real-time data acquisition at 5-minute intervals. Data sources included:
    a. Hoogendoorn-linked devices
    b. Custom dataloggers (using a Revolution Pi) for new sensors
    c. Crop yield from an external database
4. Built the Control Center (previously called Krop Manager), a web and mobile dashboard for real-time and historical monitoring of climate, crop, and operational metrics.

Key Findings

1. Some legacy systems, especially Hoogendoorn's iSii, posed major interoperability challenges. These were successfully overcome through custom engineering, notably by building around UiPath and developing a proprietary OCR layer.
2. Real-time visibility and safety settings in the Control Center enhanced grower confidence and set the stage for AI deployment.
3. The IoT Suite (including the IoT Hub and the Control Center) developed in this milestone, and continuously improved over the years by Koidra, is now one of the most robust, scalable, and control-vendor-agnostic greenhouse data platforms globally.

Lesson Learned

1. Early assumptions about available APIs proved costly in terms of engineering effort.
2. Data infrastructure and integration with other systems tend to be underrated in an AI-centric project.

2.2. Milestone 2: Greenhouse Infrastructure

Objective

Prepare the physical greenhouse environment for A/B testing and AI deployment by establishing parallel test zones and equipping them with modern lighting and sensors.

Key Achievements

1. Established two isolated compartments at GLG for side-by-side testing: one AI-controlled (AI Zone) and one grower-controlled (Baseline Zone).
2. Installed new LED lighting, sensors, cameras, and environmental monitoring systems to enable precise measurement of outcomes.

Key Findings

1. Enabled a controlled environment to isolate and compare the effects of AI versus conventional growing practices.
2. Although it doesn't introduce innovations, its completion was critical in enabling the later milestones.

2.3. Milestone 3: Expert System

Objective

Codify experts' knowledge into an AI system that replicates optimal human decision-making as a baseline for future learning.

Key Activities

- Collaborated with expert growers and researchers from the Harrow Research Center (of Agriculture and Agri-Food Canada), to document and formalize decision rules.
- Developed an expert system that mirrors experienced growers' behavior for key environmental control decisions.
- Ran a comparative trial on organic eggplants, growing in two separated zones, with both crops started on April 19 and ended on Oct 25, 2022.
    - the Baseline zone: the climate was controlled by the on-site growers using the conventional practice
    - the AI zone: the climate was controlled by the newly-developed Expert system

Key Findings

1. The AI system improved yield by 28.5% compared to the Baseline zone, without additional energy input[1]. This validated its operational readiness for more advanced learning models.
2. The data system, and especially the Hoogendoor's climate computer integration, was able handle real-world conditions, including unpredictable data issues (e.g., missing sensors, GLG-imposed control limitations).
3. Due to the site limitations at GLG at the time, lighting was controlled by the onsite staff, limiting the ability of the AI system to fully control both the climate and lighting decisions.

Lessons Learned

The AI system was initially designed under the assumption of full control over lighting, climate, and irrigation. In practice, only partial control was available, which led to early mismatches, such as climate plans being based on AI-generated lighting schedules that differed from grower-set lighting. The system was updated to detect external control inputs and adjust its strategy accordingly. Future AI designs must account for shared or limited control environments from the outset.

2.4. Milestone 4: Autonomous Continuous Improvement

Objective

Enable the AI system to improve autonomously through reinforcement learning while maintaining alignment with grower expertise.

Key Achievements

- Deployed a physics-informed model-based reinforcement learning (MBRL) framework, combining domain knowledge with machine learning to optimize greenhouse control.
- Initialized the system using imitation learning from the expert system, ensuring baseline performance was comparable to expert growers.
- Developed a hierarchical control architecture with three tiers:
    - Tactical: daily decisions based on crop registration and weather forecast
    - Operational: 5-minute climate setpoint adjustments
    - Real-time: local actuator response via Hoogendoorn

[1] This milestone and trial results were later presented in a GrowOn seminar series, hosted by OMAFRA: https://www.youtube.com/watch?v=pmya3QPdf-E

- Integrated feedback loops across all levels of the control architecture to continually refine predictions and actions.
- Successfully ran two trials on mini-cucumbers to validate the approach.

Experimental Results

- **Trial #2:** 16.1% increase in cucumber yield using the same amount of electricity as the Baseline zone.
- **Trial #3:** 31.9% improvement in Normalized Yield and 24.7% reduction in heating energy use.

3. Discussion of Experimental Results and Rebuttal to the M&V Report

We conducted a total of 3 comparative trials, summarized as follows:

Trial	Crop	Growing Period	AI Method	Unique Characteristics
1	Organic Eggplants	Apr 19th to Oct 25th, 2022	Expert system	1/ Light control wasn't handled by AI. 2/ The trial was conducted in an older greenhouse.
2	Mini cucumbers	Jul 18th to Oct 18th, 2022	Physics-informed Reinforcement Learning	Light control wasn't yet handled by AI.
3	Mini cucumbers	Baseline zone: Oct 17th to Dec 16th, 2024 AI zone: Nov 8th, 2024 to Jan 9th, 2025	Physics-informed Reinforcement Learning	The two crops don't have the same starting and ending dates

Table 1: Summary of trials. Note that in trials #1 and #2, the two crops of comparison share the same crop cycles (same starting and ending dates) but in trial #3, the planting dates are shifted (as a business decision by GLG). Trials #2 and #3 share the same crop type, in the same greenhouse, but trial #1 was initially tested with eggplants in a different greenhouse.

Each trial has its own caveat which limits our ability to truly validate the electricity use efficiency of the AI-based algorithm. In particular:

- In trials #1 and #2, lighting was controlled manually by the growers. This is due to the fact that the lighting system was initially installed and the power generator and its schedule needed to be stabilized.
- In trial #3, we were able to control both lighting and climate control. However, the growing periods were shifted, leading to challenges for a rigorous comparison.

3.1. Trials #1 and #2 Results and Analysis

In these trials, lighting control was handled manually by the on-site growers and was the same in both zones. Therefore, **comparing electricity use efficiency is equivalent to comparing yield**, for which the results are as follows.

Trial	Crop Type	Yield of Baseline Zone	Yield of AI Zone	Yield and Electricity Use Improvement
1	Eggplants	17.06 kg/m^2	21.93 kg/m^2	28.5%
2	Mini Cucumbers	13.7 kg/m^2	15.9 kg/m^2	16.1%

Table 2: Yield and electricity use efficiency comparison after trials 1 and 2.

3.2. Trials #3 Results and Analysis

3.2.1. Background

Trial 3, conducted from Oct 17, 2024 to January 9, 2025, was designed to test the full autonomous capabilities of the AI system, including lighting control. Unlike the previous trials, this one sought to evaluate how well the system performed when it had more comprehensive control of climate parameters and the additional variable of artificial lighting.

3.2.2. The Seasonal Disparity

A major complication in Trial 3 arose because the **AI-controlled Autopilot zone started 24 days later** than the grower-controlled zone. This introduced significant **seasonal variation** between the two compartments:

- The later start meant the AI-controlled crop was grown deeper into the **winter period**, when **natural light levels are substantially lower**, and **heating demands are higher**.
- This led to the AI compartment receiving **only 54% of the total daily light integral (DLI)** that the Baseline zone received.
- The delay also overlapped more heavily with cold weather, increasing the complexity of climate control and energy efficiency performance evaluation.

3.2.3. Adjustment Metrics and Results

To ensure a fair and meaningful comparison between the AI-controlled (Autopilot) and grower-controlled zones, the evaluation relied on **normalized performance metrics**. These metrics account for differences in environmental inputs (e.g., light availability, heating demand) and are well-established in both agricultural research and energy systems engineering as appropriate for comparing system performance across non-identical operating conditions.

Normalized Yield: Light Use Efficiency (LUE)

Light Use Efficiency (LUE) is defined as the ratio of yield (kg/m²) to the cumulative **Daily Light Integral (DLI)**, expressed in mol/m², received by the crop over the growing period.

DLI quantifies the total amount of **photosynthetically active radiation (PAR)** available to plants, which is the primary driver of photosynthesis and biomass accumulation. In controlled environment agriculture, especially when comparing crops grown under **different lighting conditions** (natural vs. artificial or different seasonal light availability), LUE is a standard and physiologically sound metric for assessing **productivity per unit of light input**. It is widely used in indoor farming research and practice [1].

The theoretical justification for LUE as a normalized yield metric lies in the well-documented **linear relationship between PAR and photosynthetic rate** under non-saturating light levels [2], particularly relevant in winter conditions when light is often limiting.

The yield and LUE results are provided in Table 3. According to trial results, the AI-controlled zone achieved a **29.41% higher LUE** than the Baseline zone, indicating that it converted available light into yield more efficiently, despite operating under significantly lower natural light levels.

Metric	Baseline Zone	AI Zone	Difference
Cumulative Yield [kg/m^2]	8.69	7.55	-13%
Cumulative Light Integral [mol/m^2]	851.40	572.08	-32.81%
Solar PAR Light Integral [mol/m^2]	527.24	233.64	-55.69%
LED PAR Light Integral [mol/m^2]	324.16	338.44	+4,4%
LUE [g/mol]	10.2	13.20	+29.41%

Table 3: Yield, light integral, and Light Use Efficiency analysis

References

1. Singapore Food Agency. Focusing on light-use efficiency: A key factor for profitable indoor crop production. Link
2. Sun, Dongbao, and Qingsuo Wang. "Linear relationships between photosynthetic rate and photochemical energy expressed by PAR× Fv/Fm." *American Journal of Plant Sciences* 9.2 (2018): 125-138

Electricity Use Efficiency

During the trial, natural sunlight was scarce, averaging only **3.77 mol/m²/day**, compared to the **20+ mol/m²/day** typically ideal for cucumber cultivation. The LED systems used in both compartments had a relatively low intensity of **80 µmol/m²/s**, yielding a maximum of **6.9 mol/m²/day** if operated continuously for 24 hours.

In practice, both zones maximized their use of available artificial lighting:

- **Baseline zone average LED usage:** 18.76 hours/day
- **AI zone average LED usage:** 18.95 hours/day
  (*Source: Table 14, M&V Report*)

Given that both zones operated their LED systems near full capacity—with average daily runtimes of approximately 19 hours—and followed nearly identical lighting schedules, the **electricity consumption attributable to lighting was effectively equivalent** across treatments. Assuming that the **photosynthetic efficacy** of the LED grow lights is comparable to that of natural sunlight in terms of supporting plant growth, **Light Use Efficiency (LUE)** can be reasonably used as a **proxy for electricity use efficiency** in this context.

As such, the **29.41% improvement in LUE** also serves as a **reasonable estimate for the electricity use efficiency gain** in the AI-controlled zone.

Heat Energy Use Efficiency

To evaluate heating efficiency, we focused on the **overlapping period** during which both zones were simultaneously active: **November 8 to December 16**. This allowed for a fair comparison under similar ambient conditions.

According to Table 4 below:

- The **AI-controlled zone consumed 21% less gas heating energy** than the Baseline zone during this period.
- Importantly, this reduction occurred **despite the AI zone maintaining a slightly warmer average temperature**:
  - **Baseline zone:** 21.6 °C
  - **AI zone:** 21.8 °C

While the exact crop developmental stages may have differed between zones during this window, the **temperature profiles were nearly identical**, eliminating climate targets as a confounding factor. The ability of the AI system to deliver a comparable or warmer climate using significantly less heating input demonstrates a **higher thermal energy use efficiency**, attributable to more precise, data-informed control strategies.

Metric	Baseline Zone	AI Zone	Difference
Daily averaged temperature [°C]	21.6	21.8	+0.93%
Heating Energy [MWh]	108.14	85.35	-21.07%
CO2 concentration [ppm]	576	681	+18.23%
Vent opening [%]	3.44	3.05	-11.34%

Table 4: Climate and heat energy use comparison during the overlapping period

3.2.4. Analysis

We attribute the higher Light Use Efficiency (LUE) and greater energy use efficiency observed in the AI-controlled zone to its more effective venting strategy. Specifically, as shown in Table 4, the AI system maintained **reduced venting during the winter period**, which contributed to two key physiological and operational benefits:

1. **Higher CO_2 retention**: Minimizing ventilation loss helps maintain elevated CO_2 concentrations inside the greenhouse, which enhances photosynthesis and thereby supports increased biomass accumulation and higher yields.
2. **Improved thermal retention**: Reduced venting also allowed the greenhouse to retain more heat, decreasing reliance on active heating systems and lowering overall energy consumption.

It is generally accepted that excessive reduction in ventilation can increase relative humidity and reduce airflow, potentially promoting conditions favorable to pests and pathogens. However, in this trial, we **did not observe any pest or disease outbreaks** attributable to high humidity. We hypothesize this was due to two mitigating factors:

- **Low absolute humidity** during the winter, which facilitates rapid moisture exchange between the greenhouse interior and the colder, drier external environment.
- The **drying effect of heating**, which offsets internal moisture buildup during colder periods.

Together, these conditions enabled the AI system to adopt a lower-ventilation strategy without negative side effects, improving both energy and resource use efficiency.

3.3. Rebuttal to the M&V Report

3.3.1. Lack of Comprehensiveness

The report exclusively focuses on a single trial – the final cucumber crop grown from November 2024 to January 2025 – without accounting for two earlier trials where growing conditions between the AI and Baseline zones were concurrent and more directly comparable. This omission significantly undermines the report's evaluative scope. In particular:

- Trial 1 (eggplants, milestone 3) showed a 28.5% increase in yield;
- Trial 2 (cucumbers, milestone 4) demonstrated a 16.1% increase in yield;

both with the same light control and the same outdoor conditions.

3.3.2. Flawed Electricity Use Efficiency Analysis

The M&V report attempts to evaluate Electricity Use Efficiency (EUE) by comparing the supplemental lighting energy to the crop yield. However, the methodology is flawed for the following reason: the report uses only the **electrical energy consumed by LED lighting** in the denominator while failing to account for **natural solar radiation**, which also significantly contributes to biomass formation.

This leads to a fundamental mismatch between the energy input used in the denominator and the yield (which reflects total energy input from both sun and LEDs). This error distorts the EUE metric, especially because:

- The **Baseline Zone** received substantially more sunlight due to an earlier start in the season.
- The **AI Zone**, by contrast, operated deeper into winter with significantly **less DLI** (33% less total light according to the trial data).

Since the **sunlight contribution was not normalized**, and electricity was nearly equally applied in both zones (approximately **19 hours/day** LED usage), the reported EUE unfairly disadvantages the AI zone.

A more rigorous approach is to use **Light Use Efficiency (LUE)**, yield per unit of total photosynthetically active light energy, which properly integrates both natural and artificial lighting. The rationale was discussed in more detail in Section 3.3.3.

3.3.3. Mischaracterization of Lighting Utilization Potential

The report implies there is still an unexploited lighting efficiency gap that "neither the Baseline operator nor the autonomous grower system could exploit." This statement contradicts the operational data.

As shown in Table 14 of the M&V report, both zones utilized their available LED capacity to near-maximum:

- Baseline zone: 18.76 hours/day
- AI zone: 18.95 hours/day

This usage is close to the 24-hour maximum, especially considering plant photoperiod limitations. Given the **low LED capacity (80 μmol/m²/s)** and minimal sunlight during the trial period (avg. **3.77 mol/m²/day**), there was **no practical room for increased lighting**, and the AI system already operated at max lighting levels.

Therefore, the suggestion of untapped lighting potential is not supported by the system constraints or the trial's physical lighting deployment.

3.3.4. Understated Interpretation of Heating Energy Savings

The report acknowledges that the AI-controlled zone used **21% less heating energy** during the overlapping period (November 8 to December 16, 2024), as shown in Table 10 of the M&V report, but casts doubt on comparability due to possible differences in crop stages.

However, this skepticism is not justified by the data:

- The average temperature in the AI zone during the same period was **slightly higher** (21.8 °C vs 21.6 °C).
- Thus, **the AI zone achieved thermal savings despite maintaining a warmer environment**, indicating a clear gain in heating energy efficiency.

4. Main Lessons Learned

Throughout the course of this multi-phase project, the team encountered a diverse set of technical, operational, and methodological challenges. Each challenge provided critical learning opportunities that strengthened the development, deployment, and evaluation of the AI-powered autonomous greenhouse control system. This section summarizes the key difficulties encountered and the corresponding lessons that can inform future work in AI-driven controlled environment agriculture (CEA).

4.1. Asynchronous Crop Timelines and Seasonality

Challenge

Trial 3 (Milestone 4, Trial 2) involved crops planted at different times in the AI and Baseline zones, leading to **misaligned outdoor conditions**, particularly with respect to sunlight and outdoor temperature.

Lesson Learned

Trial design must carefully control for **seasonal and temporal variables** to enable valid comparisons. When synchronization is not possible, normalized performance metrics (e.g., Light Use Efficiency) must be used to ensure fair evaluation. Moreover, including **multiple trials across seasons** helps counterbalance individual trial biases.

4.2. Evaluation Methodology Misalignment

Challenge

The Measurement and Verification (M&V) report focused exclusively on a single, asynchronous trial and used flawed methodology to assess electricity use efficiency.

Lesson Learned

1. Fair and comprehensive evaluation of a new system requires the discussion of (normalized) performance metrics during the design of the trials. If the trial design is changed, the performance metrics should be discussed and finalized before the conclusion of the trial.
2. Future projects should involve M&V planning in parallel with trial design to align evaluation methods with scientific standards.

4.3. Integration with Legacy Infrastructure

Challenge

The initial lack of API access in the Hoogendoorn climate control system presented a major barrier to seamless data acquisition and actuation.

Lesson Learned

Legacy systems in commercial greenhouses are often closed and proprietary. Successful AI deployment requires creative integration solutions, such as **Robotic Process Automation (RPA)** and **custom OCR-based data extraction**. Future projects should allocate time and resources for system-level integration engineering.

4.4. Data Reliability and Sensor Issues

Challenge
During early testing (Milestone 3), the system encountered unreliable or missing sensor data (e.g., abnormally low leaf temperature due to disconnected sensors), which could undermine the control logic if left unhandled.

Lesson Learned
Robustness to sensor failures must be designed into the system from the outset. The team implemented **error filtering mechanisms** and **contingency logic** to detect and safely ignore outlier or missing data. Future AI deployments should include redundant sensing strategies and anomaly detection as core infrastructure components.

4.5. Hardware and Control Constraints

Challenge
Despite the AI system's ability to control climate metrics such as heating, CO_2, venting, and fogging, it was **not permitted to control lighting or irrigation** at various phases of the project due to technical and operational limitations on-site.

Lesson Learned
A flexible control architecture is essential. The AI model was successfully adapted to optimize the remaining controllable parameters in real time. Future deployments should clearly define control scope upfront and plan fallback strategies when full control is not feasible.

4.6. Multi-Scale Nature of Learning in Greenhouse Systems

Challenge
A greenhouse is a tightly coupled multi-timescale system, requiring the AI to learn and coordinate two fundamentally different digital twins:

- A **climate digital twin** that evolves on a fast timescale (minutes to hours), driven by control inputs like venting, heating, and lighting.

- A **crop digital twin** that evolves slowly (weeks to months), where biomass accumulation and yield outcomes are the result of long-term cumulative interactions.

Learning to control both dynamics simultaneously is **non-trivial** because:

- Fast dynamics require reactive, high-frequency decision-making.
- Slow dynamics provide sparse, delayed, and often noisy feedback signals.

Lesson Learned
We adopted a **hierarchical control architecture** to separate tactical (short-term) decisions from strategic (long-term) learning. Furthermore, initializing the agent with expert logic (via imitation learning) was essential to prevent instability caused by the long reward horizon of crop growth. Future work will benefit from further decomposing these learning loops and using curriculum learning techniques to stage training across temporal scales.

4.7. High Model Uncertainty and Limited Labeling

Challenge
Greenhouse systems exhibit high environmental variability and low observability of key internal states. As a result:

- The digital twins, especially the crop growth models, are learned under conditions of significant uncertainty and partial observability.

- Ground truth data for crop physiological state (e.g., canopy development, stress level, root biomass) is rarely available or consistently recorded.

This uncertainty makes it difficult for reinforcement learning agents to learn effective policies without overfitting or acting on spurious correlations.

Lesson Learned
To compensate, we embedded domain-specific constraints and grower heuristics directly into the control policy, creating hybrid AI systems that combine learning with rule-based logic. These constraints enforce operational safety, avoid actions that would violate biological or physical limits, and ensure the AI acts within agronomically acceptable bounds, even in the absence of direct crop state feedback. This design philosophy is essential when deploying AI in partially observed biological systems.

5. Generalization of the Model to Other Use Cases

The development of an AI-powered autonomous greenhouse controller in this project was intentionally designed with extensibility in mind: from the control architecture and software infrastructure to the learning algorithms and interface protocols. This section explores the broader applicability of the technology across industrial domains.

5.1. Applicability to Other Agricultural Crops

The AI control system developed in this project is not crop-specific. While the trial implementation focused on cucumbers and eggplants, the underlying control architecture – based on climate and crop digital twins, feedback loops, and reinforcement learning – can generalize to a wide range of crops under protected cultivation.

Key attributes that support generalization

- Modular climate control: The controller optimizes independently tunable variables such as temperature, humidity, CO_2, and lighting, which are relevant to virtually all horticultural crops.
- Adaptable learning framework: The reinforcement learning agent can be retrained or fine-tuned on new crops by updating the reward structure and constraints based on crop-specific agronomic targets (e.g., harvest timing, size, color, or stress thresholds).
- Sensor-agnostic data ingestion: The data pipeline supports integration of diverse sensor types, including SDI-12 and analog 4–20 mA devices, which makes the system flexible to retrofit across existing facilities with minimal reconfiguration.

That said, adaptation to different crops will require:

- Domain-specific calibration of the crop digital twin.
- Crop-specific constraints and heuristics to reflect growth stages, physiological limits, and market quality standards.
- A minimum baseline of historical crop and climate data to initiate effective model training.

5.2. Transferability to Buildings and Industrial Facilities

The AI control framework extends naturally to **Buildings** and **Industrial Process Automation**, all of which share three critical characteristics:

- Highly variable external conditions (weather, occupancy, process loads)
- Tight control over environmental conditions (e.g., temperature, ventilation, humidity)
- Strong incentives to minimize energy consumption, cost, and carbon footprint

Applications in Building Environments

In MURBs and commercial buildings, the control system can replace traditional rule-based HVAC automation with AI that learns to:

- Optimize HVAC setpoints based on occupancy, forecasted weather, and dynamic pricing.
- Reduce peak energy loads by pre-conditioning or load-shifting.
- Maintain thermal comfort while minimizing unnecessary energy use.

Integration with building standard protocols like BACnet or Modbus can allow the AI system to interact with existing building management systems, making it deployable without significant hardware retrofits.

Applications in Energy-Intensive Manufacturing Industries

In sectors such as:

- Food and beverage processing
- Pharmaceutical manufacturing
- Electronics fabrication
- Chemical and materials plants

…AI-powered control has massive potential to improve **energy use efficiency**, **process stability**, and **resource optimization**. These environments often involve:

- High thermal loads (e.g., ovens, dryers, cleanrooms)
- Multi-step production cycles
- Strict environmental controls (temperature, humidity, airflow, contamination risk)

With appropriate domain customization, the system could serve as an **intelligent process control layer** across many industrial verticals.

6. Conclusion

This project set out to evaluate the feasibility and performance of an AI-powered autonomous greenhouse control system, developed by Koidra, in partnership with Great Lakes Greenhouses (GLG) and Agriculture & Agri-Food Canada, with a central emphasis on improving electricity use efficiency in commercial greenhouse operations. Over the course of multiple trials spanning three years and several crop cycles, the project systematically tested the AI system's ability to manage climate variables, optimize resource use, and improve crop yield under real-world conditions.

The results indicate that the AI system can deliver measurable energy efficiency gains, particularly in the domains of light utilization and heating energy.

While direct measurement of electricity savings was complicated by seasonal shifts and limited lighting capacity in the trial infrastructure, these proxy metrics suggest that AI control has the potential to meaningfully reduce energy intensity per kilogram of yield, a core objective of IESO's funding mandate.

Beyond the numeric results, the project also surfaced valuable operational and methodological insights:

- The AI system is technically compatible with legacy greenhouse infrastructure, though initial integration requires custom workarounds when programmatic APIs are unavailable.
- Autonomous control can safely handle climate decisions, even in tightly constrained winter conditions, without adverse crop or plant health outcomes.
- Robust evaluation requires normalized metrics and seasonally synchronized trials. Future measurement frameworks should be co-designed with deployment teams to ensure fair attribution of energy gains.

Overall, this project demonstrates the viability of AI-driven climate control in achieving tangible energy efficiency outcomes in greenhouse agriculture. While the results should not be overgeneralized from the three trials, the evidence supports further exploration of this technology across more diverse environmental settings and crop types.

Finally, the architectural flexibility and learning-based control foundation developed here offer strong potential for cross-sector adaptation, including buildings and energy-intensive manufacturing environments. These extensions represent promising pathways for broadening the impact of this technology beyond agriculture and aligning with Ontario's long-term energy efficiency and emissions reduction goals.