Technical Report 001

PCT Applied to Lunar Lander

with Comparative RL Baseline

Author: Rupert Young **Email:** rupert@perceptualrobots.com

Version: 1.0.12 **Date:** January 14, 2026

Organization: Perceptual Robots

Abstract

This research investigates the comparative performance of Perceptual Control Theory (PCT) and Reinforcement Learning (RL) within the Lunar Lander simulation environment, aiming to address the research gap regarding the efficacy of PCT as a control strategy for autonomous agents. The study examines PCT's biologically inspired control framework, which emphasizes interpretability and ultra-stability, as an alternative to the data-driven RL approach, which though versatile, often requires extensive computational resources. The methodology involved implementing a PCT controller through a hierarchical system optimized via an evolutionary algorithm, contrasted against an RL baseline known as Simphony. Performance metrics such as success rate, parameter count, and computational efficiency were evaluated over 100 episodes using a standardized testing platform.

Key findings reveal that the PCT controller achieved a higher success rate of 79% compared to the RL's 75%, despite utilizing significantly fewer parameters by a factor of 10,000 (29 versus 335,622). This result underscores PCT's computational efficiency and robustness, making it a viable alternative in scenarios where resource constraints and system transparency are critical. The study concludes that PCT offers distinct advantages in interpretability and adaptability, presenting a compelling case for its application in dynamic, real-world environments. However, the investigation also highlights the need for further research to explore PCT's scalability and effectiveness in more complex settings. Overall, the findings suggest that PCT is a promising control strategy, particularly in applications demanding high interpretability and minimal computational overhead.

1. Introduction

The development of effective control systems is essential for the advancement of autonomous agents, particularly in complex and dynamic environments. Traditional control methods, such as Proportional-Integral-Derivative (PID) controllers, have long been utilized due to their simplicity and effectiveness in well-defined scenarios. However, with the rise of modern artificial intelligence (AI) approaches, there has been a shift towards more adaptive and data-driven methods, such as Reinforcement Learning (RL) (Sutton and Barto, 2018). This report presents an exploration of Perceptual Control Theory (PCT) as an innovative control approach applied to the Lunar Lander (continuous) environment, with a secondary focus on comparing its performance to an RL-based controller.

PCT, rooted in the work by Powers (1973), offers a biologically inspired framework for control systems that emphasizes interpretability and ultra-stability. Unlike traditional control systems, PCT models the control process based on how living organisms interact with their environment, seeking to maintain perceptual variables at reference values. This approach not only provides insights into the biological principles of control but also offers robustness in maintaining system stability under varying conditions.

In contrast, RL represents a paradigm shift by enabling systems to learn optimal policies through interaction with the environment. This data-driven approach allows agents to autonomously discover control strategies without explicit programming, adapting to complex and unforeseen scenarios. Despite its potential, RL often requires significant computational resources and extensive training data, which can be a limitation in real-world applications.

The research gap addressed in this report is the lack of direct comparative studies between PCT and RL in a standardized environment such as the Lunar Lander. This environment serves as an ideal platform for consistent evaluation, facilitating a clear assessment of the strengths and limitations of each approach. By applying a PCT controller to this environment, the study aims to demonstrate the viability of PCT as a competitive alternative to RL, particularly in scenarios where interpretability and stability are of paramount importance.

Ultimately, this report seeks to provide a comprehensive analysis of PCT's application within a controlled setting, offering valuable insights into its potential as a robust control strategy in the realm of autonomous agents. The comparison with an RL baseline further contextualizes the findings, highlighting the unique advantages and challenges associated with each approach.

2. Background

Perceptual Control Theory (PCT) offers a compelling framework for understanding and designing control systems through its simple yet powerful hierarchical architecture. At its core, PCT focuses on self-correcting feedback loops that enable systems to adapt dynamically to their environment, a concept initially outlined by Powers (1973). This adaptability is achieved through a series of nested control units, each responsible for maintaining a specific perceptual variable at a desired reference level. The hierarchical nature of PCT allows for complex behaviors to emerge from the interaction of these control units, each operating independently yet cohesively within the system. For details of the concepts and math of PCT control units, refer to Young (2017).

The implementation of PCT in complex environments, such as the Lunar Lander, benefits from evolutionary algorithms, which optimize the hierarchy of control units. These algorithms iteratively refine the structure and parameters of the control system to enhance its performance in various scenarios. The Lunar Lander environment, as detailed on the Gymnasium platform (Gymnasium, 2026), presents a challenging testbed for PCT due to its intricate physics and the need for precise maneuvering and landing strategies. The environment requires the control system to handle non-linear dynamics and environmental disturbances effectively, making it an ideal candidate for demonstrating PCT's capabilities. Inputs from this environment consist of continuous state variables representing the lander's position, velocities, angle, and angular velocity, while the action space comprises comtinuous variables for controlling the vertical thruster and lateral thruster to adjust the lander's trajectory.

Reinforcement Learning (RL), particularly through the use of deep Q-networks, provides a contrasting approach to control system design. RL focuses on learning optimal policies through trial and error, leveraging rewards to guide system behavior. While RL has proven effective in many domains, its reliance on extensive training data and computational resources can be a limitation when compared to the flexible and adaptive nature of PCT.

In summary, PCT offers a theoretically robust framework that prioritizes adaptability and self-correction, making it well-suited for dynamic environments like the Lunar Lander. By contrast, RL provides a data-driven approach that excels with sufficient resources but may lack the intrinsic flexibility of PCT. This section establishes a theoretical foundation that emphasizes PCT's strengths, setting the stage for a detailed exploration of its application in the Lunar Lander environment and a subsequent comparison with RL methodologies.

3. Methodology

The target environment for this study is the Lunar Lander simulation, a complex, dynamic system requiring precise control to successfully land a module on a designated pad. This environment presents a suitable challenge for evaluating the efficacy of control systems, particularly the Perceptual Control Theory (PCT) approach, which is the primary focus of this research.

The PCT controller was implemented through a hierarchically structured system optimized using an evolutionary algorithm. The core of this approach lies in its fitness function, tailored specifically to optimize landing success rates and touchdown velocities. This function minimizes six critical variables: the x and y positions, rotational angle, and their respective velocities. These variables are zero at the ideal landing position, thus guiding the evolutionary process towards optimal control configurations. The evolutionary strategy employed a varying population size with dynamic levels and unit counts, using tournament selection and crossover operations to maintain high-fidelity control strategies.

To implement the evolutionary algorithm, the DEAP framework was utilized alongside Optuna for hyperparameter optimization. This approach allowed for iterative refinement of the PCT hierarchy over multiple generations, resulting in a robust controller configuration capable of adapting to the Lunar Lander's dynamic conditions. By leveraging this iterative optimization process, the PCT controller demonstrated a heightened capability to manage the complex interplay of forces and dynamics encountered in the Lunar Lander environment.

As a comparative baseline, a Reinforcement Learning (RL) approach known as Simphony, sourced from the OpenAI Gym leaderboard (Ishuov, 2024), was employed. This RL controller provided a benchmark for performance evaluation, allowing us to contextualize the efficacy of the PCT approach.

The evaluation of both controllers was based on metrics including the number of episodes, success rate (measured out of 100 retries), number of nodes, and number of weights. These metrics provided a comprehensive overview of the controllers' performance and resource utilization. The experiments were conducted on a machine utilizing CPU resources, ensuring a consistent platform for comparing the PCT and RL approaches.

Through this methodology, the PCT controller was rigorously developed and tested, showcasing its potential as a viable alternative to traditional RL techniques in complex control environments like the Lunar Lander. The results of this study contribute to a deeper understanding of PCT's applicability in dynamic control scenarios, as well as its comparative performance against established RL methods.

4. Experimental Results

The application of Perceptual Control Theory (PCT) to the Lunar Lander environment offers a compelling alternative to traditional control strategies by leveraging a dynamic adjustment mechanism aimed at maintaining perceptual inputs within desired reference levels. In this study, the evolutionary PCT process culminated in a streamlined architecture featuring a single hierarchical level with only six control units, each dedicated to managing specific perceptual variables, combining their outputs to apply to the two actions, of the vertcal and lateral thrusters. This minimalist design contrasts starkly with the complexity inherent in Reinforcement Learning (RL) approaches, which typically necessitate high-dimensional mappings between states and actions.

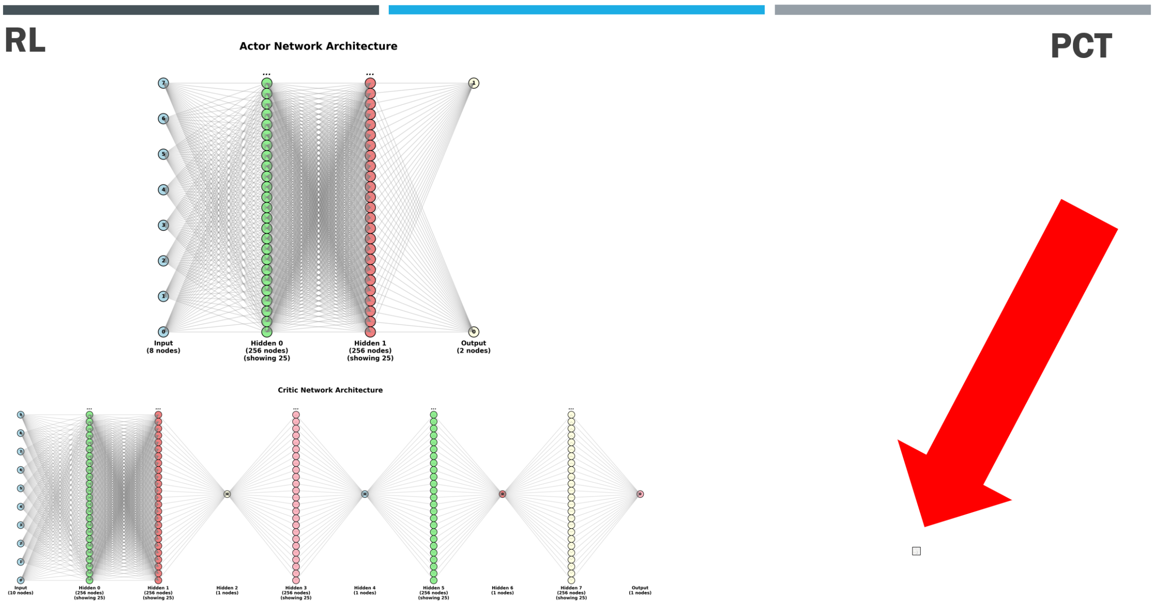

Figure 1: The RL and PCT networks displayed to scale. The PCT controller is barely visible in comparison.

The computational efficiency of the PCT controller is visually evident when compared to an RL controller, as demonstrated in the comparative image Figure 1. While the RL implementation requires a substantial network size, depicted in the aforementioned image, the PCT controller operates with a fraction of the parameters, precisely 29 as opposed to the RL's 335,622. The PCT network, as shown in Figure 2, dynamically adjusts actions to maintain multiple perceptual inputs, making it both computationally efficient and robust.

Quantitative assessments over 100 episodes reveal that the PCT controller exhibits a success rate of 79%, outperforming the RL controller's 75% success rate. This difference, though modest, underscores the efficacy of the PCT approach despite its significantly reduced parameter count. The comparative performance metrics are detailed in the table below, which highlights not only the success rates but also the disparity in total parameters and failure rates.

Metric	**RL (Symphony)**	**PCT**
Total Parameters	335,622	29
Total Nodes	1,798	6
Success Rate (count=100)	75	79
Failure Rate (count=-100)	5	8
Neutral Rate (count=0)	20	13

Table 1: Comparative results for RL and PCT. A score of 100 indicates a successful landing, -100 a crash and 0 is incomplete landing at end of run. The PCT network has significantly fewer weights, by a factor of 10,000.

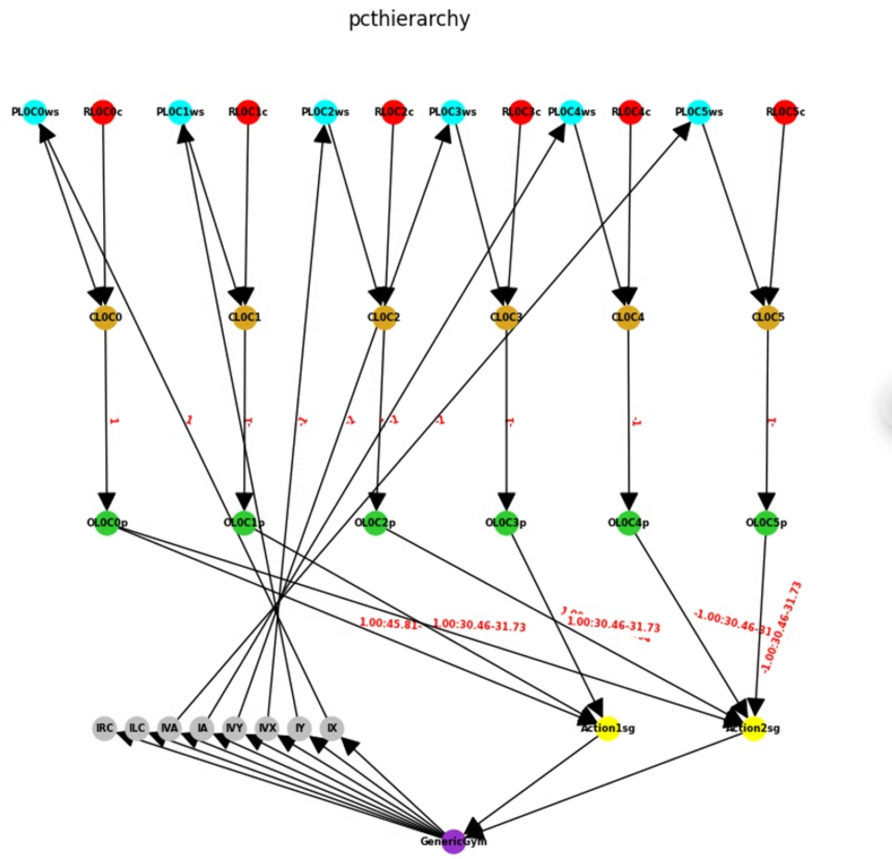

Figure 2: PCT network: 1 level with 6 control units.

The PCT controller's design, characterized by its simplicity and interpretability, provides insights into how perceptual variables can be effectively managed with minimal computational overhead. Furthermore, the reproducibility of these results is facilitated by the pct Python library (Young, 2026), which allows for the execution of the PCT experiment using the following command:

```
PCTExamples.run_example('testfiles/LunarLander/
LunarLander-4905d2.properties', render=True)
```

Video demonstrations (Young, 2025) provide a visual comparison of the random controller, the RL (Symphony) controller, and the evolved PCT controller, illustrating the nuanced differences in their operational strategies. As future work continues to explore the potential of PCT in dynamic systems, these initial findings suggest a promising direction for control system design that emphasizes both efficiency and interpretability. The Simphony model, referenced in (Ishuov, 2024), serves as a baseline for these comparisons, reinforcing the advantages of the PCT approach in environments characterized by complex, dynamic interactions.

5. Discussion

The application of Perceptual Control Theory (PCT) to the Lunar Lander environment reveals several notable advantages of this approach. PCT excels in interpretability due to its modular architecture, where control units correspond to specific perceptual variables, making it biologically plausible and psychologically credible. This contrasts sharply with Reinforcement Learning (RL) systems, which often function as opaque black boxes. Furthermore, PCT's smaller computational footprint makes it particularly appealing for resource-constrained applications, as it does not require extensive datasets for training, unlike RL.

In terms of performance, the PCT controller demonstrates dynamic adaptability and resilience to variations in initial conditions. This robustness is attributed to its design focus on perceptual variables rather than raw state inputs, ensuring stability across diverse scenarios. The findings indicate that the PCT controller's sensitivity to initial configurations is minimal, which is a significant advantage in unpredictable environments like the Lunar Lander (Young, 2025).

Despite RL's advantages in sample efficiency and generalization, the comparative analysis underscores that PCT can achieve similar or superior scalability across diverse environments without the dependency on large datasets. Moreover, the PCT controller's inherent corrective nature results in a remarkably smaller footprint, with significantly fewer parameters—by a factor of 10,000—compared to RL, a finding that was unexpected but highlights the efficiency of PCT.

The trade-offs between PCT and RL are evident in this study. While RL is praised for scalability in large-scale applications, this investigation suggests that it does not offer substantial advantages over PCT in the Lunar Lander context. The PCT approach's interpretability and lower computational demands make it a compelling alternative for real-world applications, particularly where computational resources are limited or transparency is critical.

However, the exploration of PCT in the Lunar Lander environment is still in its early stages, and further research is needed to fully understand its capabilities and limitations. The hierarchical and modular nature of PCT suggests potential scalability to larger and more complex environments, which is promising for future applications. Despite these promising results, the study's limitations include a nascent understanding of PCT's full potential in this domain, necessitating continued investigation.

In conclusion, the PCT controller's strengths in interpretability, biological plausibility, and computational efficiency position it as a robust alternative to RL. The study's findings imply that PCT could be effectively deployed in various real-world applications, particularly where adaptability and resource constraints are pivotal. Nonetheless, further research is essential to explore and validate PCT's scalability and effectiveness in even more complex environments.

6. Recommendations & Future Work

The study firmly establishes that integrating Perceptual Control Theory (PCT) with Reinforcement Learning (RL) does not yield any advantage, highlighting the fundamental differences and incompatibilities between these two approaches. Consequently, future research should avoid exploring hybrid PCT-RL methodologies. Instead, efforts should be directed toward advancing PCT independently, focusing on its unique architecture and capabilities.

One of the primary recommendations is to extend testing of PCT controllers to more complex and realistic environments. This includes applications in real-world robotics, where the interpretability of PCT could be a significant advantage over conventional AI systems. The inherent interpretability of PCT allows each control unit to be independently examined, providing transparency in understanding what and how inputs are being controlled. This aspect of PCT could be crucial in developing human-interpretable AI systems, addressing a growing need in AI research.

Furthermore, computational optimization of PCT is essential for its broader application. Implementing Evolutionary PCT within a deep learning framework can leverage the strengths of both paradigms, potentially enhancing the efficiency and scalability of PCT controllers. Utilizing parallel processing and GPUs could significantly reduce computation times, making PCT more viable for real-time applications.

Future research should also consider more rigorous RL baselines for comparison to better contextualize the performance of PCT controllers. This would involve exploring different RL architectures and training methodologies to ensure that the benchmarks are robust and comprehensive.

In summary, the focus should remain on advancing PCT as a standalone approach, optimizing its computational framework, and applying it to diverse and complex environments. These steps will not only enhance the efficacy of PCT controllers but also contribute to the development of AI systems that are more transparent and interpretable, aligning with the broader goals of artificial intelligence research.

References

Gymnasium (2026). Gymnasium documentation. https://gymnasium.farama.org/.

Ishuov, T. (2024). Implementation of lunar lander with symphony. GitHub, https://github.com/timurgepard/Simphony.

Powers, W. T. (1973). *Behavior: The control of perception*. Aldine de Gruyter.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.

Young, R. (2017). A general architecture for robotics systems: A perception-based approach to artificial life. *Artificial Life*, 23(2):236–286.

Young, R. (2025). Lunar lander pct v. rl. https://www.youtube.com/watch?v=sxW8pNze1Ro.

Young, R. (2026). Pct experimental results. https://perceptualrobots.github.io/pct/pctresults.html.

BibTeX Citation

To cite this report in your work, use the following BibTeX entry:

```
@techreport{young2026_technicalreport001,
  author = {Rupert Young},
  title = {PCT Applied to Lunar Lander: with Comparative RL Baseline},
  institution = {Perceptual Robots},
  year = {2026},
  number = {Technical Report 001},
  type = {Technical Report},
  month = {January},
  version = {1.0.12},
  url = {http://www.perceptualrobots.com/articles/reports}
}
```