

CEDAR: AGENT-ORCHESTRATED TREE SEARCH FOR GOAL-DIRECTED OPTIMIZATION OF COMPLEX SYSTEMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Complex systems modeling analyzes nonlinear, feedback-driven phenomena from population dynamics to economic policy, supporting decisions with significant societal impact. In established practice, models are often authored in specialized system-dynamics languages (e.g., DYNAMO, STELLA) that specify the models' structure. However, building and refining such models requires extensive manual effort due to (1) the opaque relationship between the structure and emergent behavior and (2) the labor-intensive workflows imposed by these languages. These barriers limit adoption and hinder effective decision-making. To address these challenges, we introduce CEDAR (Complex-systems Exploration and Design via Agent-Orchestrated Refinement), an autonomous method that uses LLM (Large Language Model) agents to discover and improve complex systems that satisfy user-specified goals. Our key innovation is an LLM-driven MCTS (Monte Carlo Tree Search) process **deeply coupled with complex system**: at each iteration, an LLM Judge evaluates performance against goals and an LLM Editor proposes improved system variants. We represent systems using a restricted, runnable subset of Python with domain-specific primitives, enabling LLMs to meaningfully and modify system dynamics directly. CEDAR is **theoretically designed with formalization in mind**, and empirically enables automatic optimization of vague goals, thereby reducing human effort while achieving capabilities beyond existing approaches. The unified design handles diverse systems across domains, constructing complex systems that would otherwise require extensive manual fitting. Moreover, by using LLMs to interpret systems, CEDAR makes system design transparent and accessible, facilitating broader adoption of complex systems modeling.

1 INTRODUCTION

Computational system modeling has been essential for analyzing complex systems across diverse domains, from global dynamics (Forrester, 1971; 2011) and pandemic diffusion (Zhu et al., 2025) to social simulations (Kolson, 1996) and biological systems (Ruth & Hannon, 2012; Hannon & Ruth, 2014). These approaches, supported by specialized languages and tools (Radzicki & Taylor, 2011; Richmond, 1985), formalize feedback loops and nonlinear interactions to explore system behavior beyond simple simulation and forecasting. Such systems thinking (Anderson & Johnson, 1997; Arnold & Wade, 2015), by attempting to model the subject as a system, critically supports policy design and strategic decision-making (Peterson & West Lebanon, 2003).

However, a central challenge remains: extensive human effort. This effort stems from the opaque relationship between model structures and emergent behavior, combined with the manual workflows imposed by established modeling languages. These factors limit adoption and hinder timely decision-making across domains where complex systems modeling could provide valuable insights.

Large language models (LLMs) and LLM-powered agents now offer unprecedented capabilities in reasoning and planning for complex systems. Recent work, such as AI Scientist (Yamada et al., 2025) and related efforts (Kumar et al., 2025), demonstrates automated discovery in scientific and artificial-life domains. These advances suggest that we can now automate the discovery and optimization of complex systems themselves: leveraging cognitive agents to autonomously improve system structures.

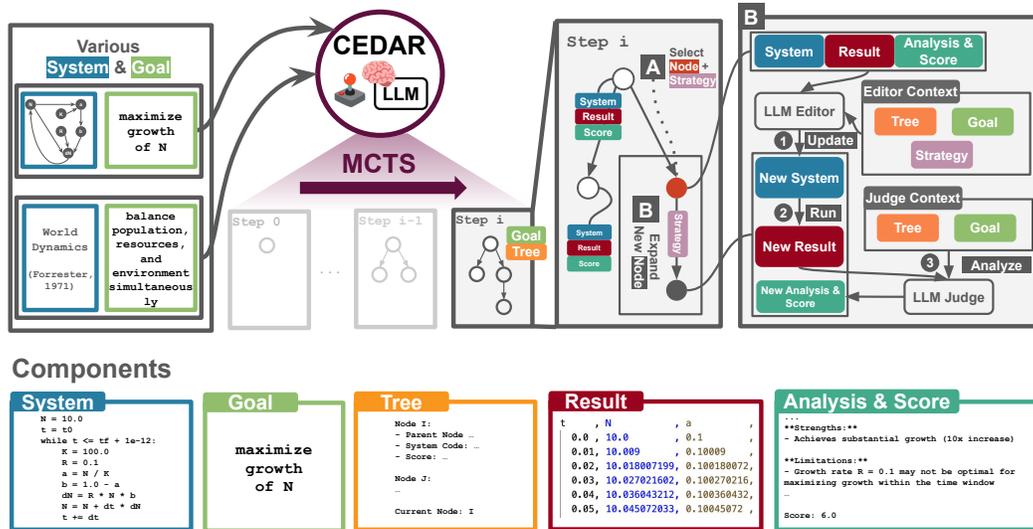


Figure 1: Overall architecture of CEDAR. The autonomous approach uses Monte Carlo Tree Search (MCTS) to iteratively select system variants for expansion. At each iteration, an LLM Editor generates improved system variants and executes them, then an LLM Judge evaluates their performance and provides detailed analysis. The method integrates several key components: a system P represented by a restricted subset of Python code, a natural-language goal description, the search tree structure, system execution results, and comprehensive analysis feedback.

To address these challenges, we propose an autonomous method that combines Monte Carlo Tree Search (MCTS), large language models, and a unified system representation to discover complex systems that meet specified goals. At each iteration, MCTS selects a system variant for expansion, where an LLM Editor generates improved system variants, and the system is executed and evaluated by an LLM Judge for performance and analysis. This process is enabled by our unified representation using a restricted subset of Python with domain-specific mathematical primitives and extensive inline documentation, thereby leveraging Python’s familiarity and LLMs’ strong Python capabilities for seamless integration between search, evaluation, and modification components.

Our method addresses key limitations in system modeling through four main contributions. First, it enables automatic optimization for vague goals, reducing human effort while achieving capabilities beyond existing approaches. Second, the unified design accommodates diverse systems, demonstrating generality across domains. Third, it constructs complex systems that would otherwise require extensive manual fitting, improving efficiency. Finally, by leveraging LLMs to interpret systems, it makes system design transparent and accessible to a broader range of users.

2 RELATED WORKS

The central motivation of our work is to modernize (1) **the modeling of complex systems** using (2) **tree search algorithms with LLMs** to (3) **optimize systems in an LLM-driven way**. In this section, we discuss these three lines of work that are related to our motivation.

(1) Complex Systems and How to Model Them *World Dynamics* (Forrester, 1971) introduced a computational framework for modeling the planet as a nonlinear, feedback-driven system. Using the DYNAMO language for system dynamics analysis (Radzicki & Taylor, 2011), it simulated population, resources, and industrial output over long horizons, marking a paradigm shift from linear prediction to dynamic system-level exploration (Forrester, 2011). Since then, complex systems have been applied to pandemic diffusion (Zhu et al., 2025), war games, and social simulations (Kolson, 1996). Visual programming languages like STELLA (Richmond, 1985) were developed as successors to DYNAMO, with applications in economics, thermodynamics, and biological systems (Ruth & Hannon, 2012; Hannon & Ruth, 2014). Examples of DYNAMO and STELLA can be found online.

108 Computational system modeling has been essential for analyzing complex systems and their dynamics.
109 Computation can be used not only to simulate and forecast but also to explore system behavior by
110 formalizing feedback loops and nonlinear interactions. Such systems thinking (Anderson & Johnson,
111 1997; Arnold & Wade, 2015) supports policy design and strategic decision-making (Peterson &
112 West Lebanon, 2003).

113 Despite these advances, significant challenges remain. Traditional system dynamics relies on expert-
114 defined variables and relationships (Ford & Sterman, 1997; Bérard, 2010; Zagonel, 2002; Vennix et al.,
115 2021). The main challenge is the extensive human effort required. First, the relationship between
116 model structures and emergent system behavior is often unclear (Schoenberg, 2019; Güneralp, 2004;
117 Barlas, 1996). Second, established modeling languages require manual, labor-intensive workflows
118 that can be counterintuitive (MIT System Dynamics in Education Project, 1998; Hines, 1996; Sterman,
119 2000). The central motivation for our work is to modernize complex system design and iteration
120 workflows traditionally centered on manual modeling.

121
122 **(2) Tree Search Algorithms with LLMs** Monte Carlo Tree Search (MCTS) (Kocsis & Szepesvári,
123 2006) has been successfully applied to various optimization problems, providing a foundation for
124 CEDAR. Combining MCTS with neural networks has led to superhuman performance in games (Sil-
125 ver et al., 2016; 2017b;a; Schrittwieser et al., 2020) and to the discovery of novel algorithms (Fawzi
126 et al., 2022; Mankowitz et al., 2023) and neural network architectures (Nasir et al., 2024).

127 More recently, large language models (LLMs) have been combined with MCTS to enhance its
128 capabilities. Generic approaches that improve MCTS involve LLM-powered value functions and self-
129 reflection (Zhou et al., 2024), as well as LLM-powered node selection and expansion strategies (Inoue
130 et al., 2025). Additional works apply MCTS with LLMs to specific tasks, such as automatic
131 design (Zhang et al., 2025b) and agent collaboration (Gan et al., 2025).

132 Notably, MCTS with LLMs offers strong capabilities in two areas: planning and reasoning. For
133 planning, recent works propose accelerating planning by letting LLMs serve as common-sense policy
134 priors (Zhao et al., 2023), improving planning with transformer-based search dynamics (Lehnert
135 et al., 2024), and enhancing interpretability through contrastive MCTS reasoning (Gao et al., 2024).
136 For reasoning, LLMs have been used to guide MCTS traversals of knowledge graphs (Song et al.,
137 2025) and to refine reasoning through intermediate steps (Chi et al., 2025).

138
139 **(3) LLM-based Optimization for Systems** LLMs with MCTS show promising results for complex
140 search. AI Scientist (Yamada et al., 2025) enables automated scientific discovery, Kumar et al. (2025)
141 extend this to artificial life simulations, and Katz et al. (2024) optimize runnable code for reasoning.
142 Recent work applies LLMs to system-dynamics modeling (Liu et al., 2024; Luo et al., 2025; Liu &
143 Keith, 2025), treating models as a predictors of behavior that are *unknown* to the optimizer, whereas
144 we directly optimize *known* system models for goal-directed outcomes. These studies consider small
145 systems (up to 4 variables and 12 steps), while our method scales to roughly 4000 steps and dozens
146 of variables, making direct extensions of their setups non-trivial.

147 **(4) Prior works close to our work.** Adjacent works combine LLMs with search for different goals.
148 I-MCTS (Liang et al., 2025) targets AutoML hyperparameters, ChemReasoner (Sprueill et al., 2024)
149 and CheMatAgent (Wu et al., 2025b) focus on chemistry, and AgentSwift (Li et al., 2025b) employs
150 hierarchical search without long-horizon simulations. None address iterative temporal dynamics
151 optimization. Language Agent Tree Search (LATS) (Zhou et al., 2023a) and SWE-Search (Antoniades
152 et al., 2024) are closest: LATS operates on discrete task graphs, while SWE-Search targets software
153 engineering. GIF-MCTS (Dainese et al., 2024) and Zhang et al. (2025a) address short sequences. Our
154 method optimizes multi-thousand-step temporal dynamics with tailored representation and search.

155 156 3 METHOD

157 158 3.1 PRELIMINARIES

159
160 We first formalize the complex systems under discussion and propose a modern programming
161 language-based approach for their representation. Complex systems are typically studied using
mathematical models expressed as ordinary differential equations (ODEs) with initial values:

162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215

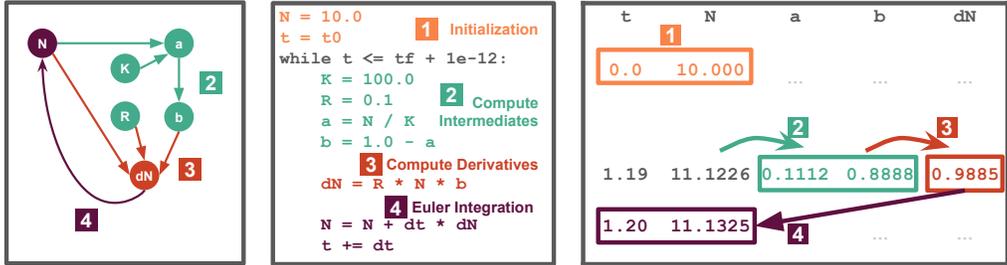


Figure 2: Python representation of a complex system. From left to right: graph representation, Python code, and computation process. Four essential components are shown: (1) initialization of values, (2) intermediate computations, (3) derivative evaluations, and (4) Euler method integration. This representation allows users to easily code and inspect, while enabling LLM understanding.

$$dx/dt = f(x, t), \quad x(t_0) = x_0 \tag{1}$$

where $x \in \mathbb{R}^n$ is the state vector (also called level or stock variables in the complex systems literature), $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ the vector field, and $x_0 = [x_{0,1}, x_{0,2}, \dots, x_{0,n}]^T$ the initial values.

While analytical techniques for ODEs, such as the Laplace transform, can be applied to complex systems analysis when feasible, these systems are more commonly analyzed through numerical approaches like the first-order Euler method (Ogata, 2004) due to the complexity of f . This approach enables flexible modeling by requiring only that f be computable, which is crucial because as many complex systems are too intricate for analytical solutions. Concretely, the first-order Euler method discretizes the system with time step Δt as:

$$x_{t+\Delta t} = x_t + \Delta t \cdot f(x_t, t), \quad x_0 = x(0) \tag{2}$$

Note that x and its discretized counterparts x_t and $x_{t+\Delta t}$ represent interpretable, meaningful quantities such as populations and resources, rather than abstract hidden states as in RNNs, despite the similarity in their formulations. Historically, communities studying complex systems have benefited from specialized modeling languages like DYNAMO and STELLA (Radzicki & Taylor, 2011; Richmond, 1985). However, as discussed in Section 2, these tools suffer from outdated syntax or proprietary visual interfaces (Forrester, 2011) that limit their modern adoption with LLMs.

3.2 UNIFIED REPRESENTATION

To address these limitations, we propose a unified representation based on a restricted subset of Python programs. This choice leverages Python’s widespread familiarity and LLMs’ strong Python capabilities while avoiding the need to create a new domain-specific language. Practically, this approach enables flexible implementation of $f(x, t)$ through general-purpose programming constructs.

We show an example of the Python representation of an actual complex system in Figure 2. Specifically, we present the system in its graph representation (for reference), the corresponding Python code, and the computation process. The components and computation are divided into four parts: (1) initialization of values ($x_0 = x(0)$), followed by a loop of (2) intermediate computations, (3) derivative evaluations (i.e., computing $f(x_t, t)$), and (4) Euler-method integration (i.e., computing $x_{t+\Delta t}$ from x_t and $f(x_t, t)$). Our implementation features extensive inline documentation and domain-specific mathematical primitives, including wrapper functions for common operations and STELLA-compatible delay and smoothing functions with time-based semantics.

This abstraction layer allows LLMs to manipulate system dynamics through semantic constructs rather than low-level implementation details, prioritizing interpretability while maintaining compatibility with established modeling conventions. The design facilitates both human comprehension and automated LLM-based system generation through standardized interfaces. We provide a complete example showing the syntax and semantics in Appendix A.

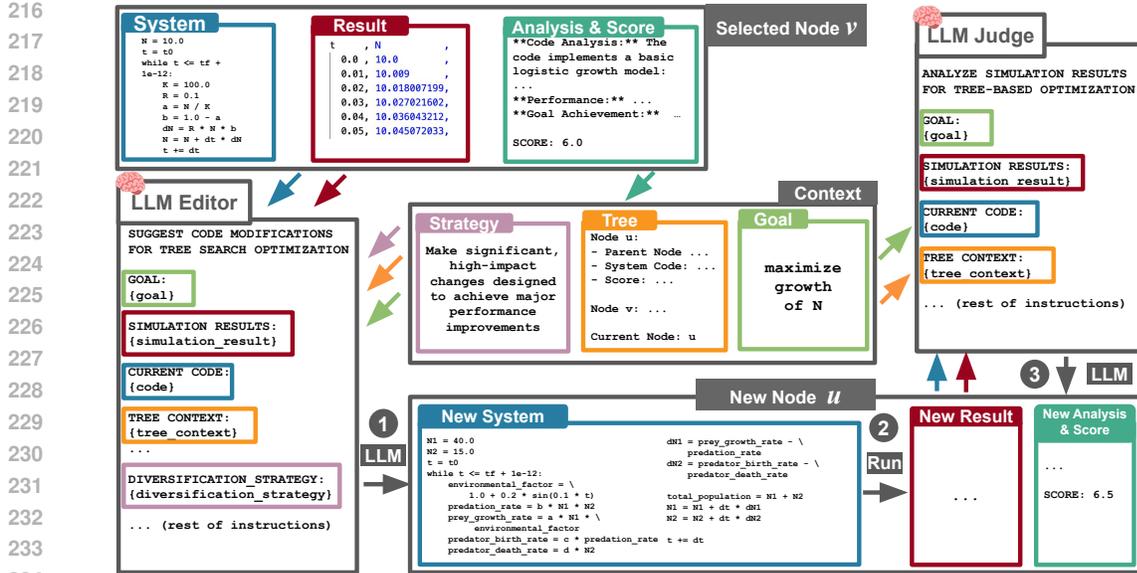


Figure 3: The node expansion process using LLM Editor and Judge. Given a selected node v and expansion strategy s , the LLM Editor generates a new system variant P_u by leveraging the parent system P_v , its analysis A_v , strategy s , and goal G . The generated system is executed to obtain record R_u , then evaluated by the LLM Judge to produce analysis-score pair (A_u, S_u) . We embed analysis and scores of all nodes, including the currently selected node, in the tree context to provide the LLM with a comprehensive view of the search tree.

3.3 PROPOSED METHOD

We propose CEDAR, an autonomous method that combines Monte Carlo Tree Search (MCTS), large language models, and a unified system representation to discover complex systems that meet specified goals. At each iteration, MCTS selects a system variant for expansion. An LLM Editor then generates improved system variants based on execution feedback. The newly generated variant is subsequently passed to an LLM Judge, which evaluates its performance and provides a comprehensive analysis. We present the overall architecture of CEDAR in Figure 1.

Formalization We formalize the key components of CEDAR, illustrated in Figure 1, as follows. We denote a system as P , represented using the restricted subset of Python code described above. The execution record R of a given complex system P , including the values of all variables at each timestamp, is obtained by $R \leftarrow \text{RUN}(P)$. We define a goal G as a natural-language description of the desired system behavior¹.

The search process employs an MCTS tree $T = (V, E)$ with vertex set V and edge set E . For each node $v \in V$, we associate four key components: (1) a system P_v , (2) an execution record $R_v \leftarrow \text{RUN}(P_v)$, (3) a bounded numerical score S_v , and (4) a textual analysis A_v . We denote the parent of node v as $\text{parent}(v)$ and the expansion strategy used to generate v from its parent as s_v . The root node is denoted as $0 \in V$ and has neither a parent nor an associated expansion strategy. At each iteration, the LLM Judge produces an analysis-score pair (A_v, S_v) that evaluates the quality of system P_v against the specified goal G . Details of the LLM Judge and expansion strategies are discussed in the following sections.

Initialization The search process begins by initializing the MCTS tree T with a single root node 0 (where $v_0 = 0$). This root node is associated with an initial complex system P_0 , which can either be provided by the user or generated as a basic template system. We obtain the execution record $R_0 \leftarrow \text{RUN}(P_0)$ and evaluate the initial system via $(A_0, S_0) \leftarrow \text{LLMJUDGE}(P_0, R_0, G)$, where $S_0 \in \mathbb{R}$ is a bounded numerical score reflecting how well P_0 satisfies goal G , and A_0 is a textual analysis providing qualitative feedback about the system’s behavior and potential improvements.

¹For example, goal G can be “Make the population grow stably” or include additional information like “Make the system behavior similar to @load(record.csv)”.

Node Selection The node selection process jointly selects a node $v \in V$ for expansion and an expansion strategy $s \in \mathcal{S}$, where $\mathcal{S} = \{s_1, s_2, \dots, s_k\}$ is a fixed set of expansion strategies. Node selection is based on a selection score $\text{SCORE}(v)$ that combines the node’s performance score S_v , its depth $\text{DEPTH}(v)$ in the tree, and its expansion count $c_v \in \mathbb{N}$ and its parent p ’s expansion count: CEDAR uses the generalized UCT score $\text{SCORE}(v) = S_v + \phi(c_v, c_p, \tau)$. Here, we have $\phi = -\infty$ if $c_v \geq \tau$ and $\alpha \cdot \sqrt{\ln(c_p + 1)/(c_v + 1)} + \gamma \cdot \text{DEPTH}(v)$ otherwise, where α adjust the amount of exploration, $\gamma > 0$ controls the preference for deeper nodes and τ controls the expansion count. Unlike vanilla MCTS, we limit expandable nodes to a subset $V_{\text{exp}} \subseteq V$ and allow repeated expansion of internal nodes until their expansion count c_v reaches a threshold $\tau \in \mathbb{N}$, inspired by Inoue et al. (2025). Formally, a node v is expandable if and only if $v \in V_{\text{exp}}$ and $c_v < \tau$. The selection process chooses $(v, s) = \arg \max_{v \in V_{\text{exp}}} \text{SCORE}(v)$, with the expansion strategy $s \sim \text{Uniform}(\mathcal{S})$ sampled uniformly. Practically, hyperparameters α and τ reflect empirical behavior of LLM-based optimization on complex systems: high-quality solutions usually require several consecutive refinements along promising branches, so we favor deeper nodes while capping expansion counts. This cap prevents over-exploring a single subtree and avoids wasting expensive LLM calls.

Node Expansion with LLMs Given the selected node v and expansion strategy s , the LLM Editor generates a new system variant via $P_u \leftarrow \text{LLMEDITOR}(P_v, A_v, s, G)$. This process leverages the parent system P_v , its associated analysis A_v , the chosen strategy s , and the goal specification G to produce a revised system that satisfies both syntactic constraints and semantic requirements.

The newly generated system P_u is executed to obtain its record $R_u \leftarrow \text{RUN}(P_u)$, followed by performance evaluation $(A_u, S_u) \leftarrow \text{LLMJUDGE}(P_u, R_u, G)$, yielding both a bounded numerical score $S_u \in \mathbb{R}$ and a qualitative analysis A_u . Both LLMEDITOR and LLMJUDGE are made feasible through carefully injected semantics and instructions, through two complementary mechanisms: first, specialized prompts that guide the LLM’s editing and reasoning process, and second, semantic annotations embedded within the code that help the LLM understand the system’s structure and behavior. The full expansion process is illustrated in Figure 3, and details, including extensions in code and prompts, are provided in Appendix B.

A new node u is then added to the tree with the associated system P_u , execution record R_u , score S_u , and analysis A_u . This iterative process of selection, expansion, execution, and evaluation enables progressive improvement toward the target objectives through systematic exploration of the system space. We formalize the complete search procedure in Algorithm 1, which provides a detailed algorithmic description of our MCTS-based approach. We also cover further details of our method, including design choices and as well scalability and stability in Appendix C.

3.4 THEORETICAL CONNECTION

While CEDAR components are motivated by practical considerations, the algorithm remains a principled MCTS variant with widening- and depth-aware UCT selection, plus LLM-parameterized transition kernel and value function. While the connections can be established below, we also note that convergence and regret guarantees are open problem for MCTS with LLMs.

Algorithm 1: Monte Carlo Tree Search for Complex System Discovery

```

1: Input: Initial system  $P_0$ , goal  $G$ , iterations  $N$ , expansion threshold  $\tau$ 
2: Output: Best discovered system  $P^*$ 
3: Initialize tree  $T = (V, E)$  with root node 0
4:  $(A_0, S_0) \leftarrow \text{LLMJUDGE}(P_0, \text{RUN}(P_0), G)$ 
5:  $P^* \leftarrow P_0, S^* \leftarrow S_0$ 
6: for  $i = 1$  to  $N$  do
7:    $(v, s) \leftarrow \arg \max_{v \in V_{\text{exp}}} \text{SCORE}(v)$  where  $s \sim \text{Uniform}(\mathcal{S})$ 
8:    $P_u \leftarrow \text{LLMEDITOR}(P_v, A_v, s, G)$ 
9:    $\triangleright$  Generate new system
10:   $R_u \leftarrow \text{RUN}(P_u)$   $\triangleright$  Execute system
11:   $(A_u, S_u) \leftarrow \text{LLMJUDGE}(P_u, R_u, G)$ 
12:   $\triangleright$  Evaluate system
13:  Add node  $u$  to the tree with  $(P_u, R_u, A_u, S_u)$ 
14:   $c_v \leftarrow c_v + 1$   $\triangleright$  Increment expansion count
15:  if  $c_v \geq \tau$  then
16:     $V_{\text{exp}} \leftarrow V_{\text{exp}} \setminus \{v\}$ 
17:     $\triangleright$  Remove it from the expandable set
18:  end if
19:  if  $S_u > S^*$  then
20:     $P^* \leftarrow P_u, S^* \leftarrow S_u$ 
21:  end if
22: end for
23: return  $P^*$ 

```

Node Selection as a UCT-Generalization. Following UCB1 (Kocsis & Szepesvári, 2006), CEDAR uses $\text{SCORE}(v) = S_v + \phi(c_v, c_p, \tau)$, where c_v and c_p are expansion counts at v and its parent p , and τ is the progressive widening limit. We combine two MCTS principles: (1) Progressive Widening (Coulom, 2007; Chaslot et al., 2007; Inoue et al., 2025) limiting expansion width, and (2) depth-based bonuses favoring deeper search (Blackshaw et al., 2025; Wu et al., 2025a). This yields:

$$\phi = \underbrace{\alpha \sqrt{\frac{\ln(c_p + 1)}{c_u + 1}} \cdot \mathbb{I}[c_v < \tau]}_{\text{Progressive Widening}} + \underbrace{\beta \mathbb{I}[c_v \geq \tau] + \gamma \cdot \text{DEPTH}(v)}_{\text{Depth-based bonus}}. \quad (3)$$

The node selection corresponds to $\beta = -\infty$, positioning CEDAR a simple combination of MCTS variants with non-uniform branching. The constraint $c_v < \tau$ prevents overexpansion in vast action spaces like program edits and expensive LLM invoking.

LLM Editor as a Learned Transition Kernel. MCTS assumes a generative model for transitions: $s' \sim P(\cdot | s, a)$. CEDAR extends this to a program-valued state space via a stochastic operator P_θ induced by the LLM:

$$P_u \sim P_\theta(\cdot | P_v, A_v, s, G), \quad \text{where } P_\theta = \text{LLMEDITOR}_\theta, \quad s \sim \text{Uniform}(\mathcal{S}) \quad (4)$$

where P_θ parameterizes a conditional distribution over successor programs. Though parameterized by a pretrained foundation model rather than learned during search, P_θ functionally serves as a learned component enabling MCTS to operate in complex program spaces. This aligns with recent work using LLMs as learned components: as sampling models (Zhou et al., 2023b; Li et al., 2023), refining models in code synthesis (Chen et al., 2023; Zelikman et al., 2022) and reasoning (Shinn et al., 2023). CEDAR extends LLM-powered transition kernels from linear to tree search, connecting to learned-transition MCTS and model-based RL (Silver et al., 2017b; Schrittwieser et al., 2019), neural-guided search (Chen et al., 2020; Xu et al., 2024), and scientific discovery (Lai & Pu, 2025).

LLM Judge as a Learned Value Function. Executing P_u yields a deterministic trajectory $R_u = \text{RUN}(P_u)$. The LLM Judge provides a semantic evaluation:

$$(A_u, S_u) = R_\phi(P_u, R_u, G), \quad \text{where } R_\phi = \text{LLMJUDGE}_\phi \quad (5)$$

where R_ϕ acts as a learned value network. Focusing on scoring, we denote $S_u = \pi_S(R_\phi(P_u, R_u, G))$ where π_S projects from (A, S) onto score S . Recent work shows LLMs can serve as noisy reward estimators or constraint evaluators in non-differentiable optimization (Zhou et al., 2023b; Zelikman et al., 2022; Shinn et al., 2023), often pairing with LLM also serving as sampling models. Theoretically, S_u constitutes a *bounded, noisy reward*, falling under convergence analysis of MCTS with biased evaluators (Lisy et al., 2013; Efroni et al., 2019). UCT remains consistent when reward noise is bounded and non-adversarial, as with semantic LLM scoring.

Theoretical Guarantees Remains an Open Problem Unlike classical MCTS with well-specified MDPs where asymptotic convergence and regret guarantees are established Kocsis & Szepesvári (2006); Bubeck et al. (2011), Recent work combining MCTS with LLM yields strong empirical performances on program related tasks Li et al. (2024; 2025a); Xu et al. (2025) while the focus is not on theoretical analysis Katz et al. (2024) The current theoretical work only provides at best partial guarantees (e.g., loose verifier-induced upper bounds Brandfonbrener et al. (2024)), leaving the general problems of *optimality* and *convergence* for MCTS and LLM methods open research questions.

4 EXPERIMENTS

CEDAR enables three previously infeasible capabilities: optimizing for vague natural-language goals, fitting records without complete system skeletons, and providing interpretable optimization. We present qualitative demonstrations and quantitative comparisons.

4.1 EXPERIMENTAL DETAILS

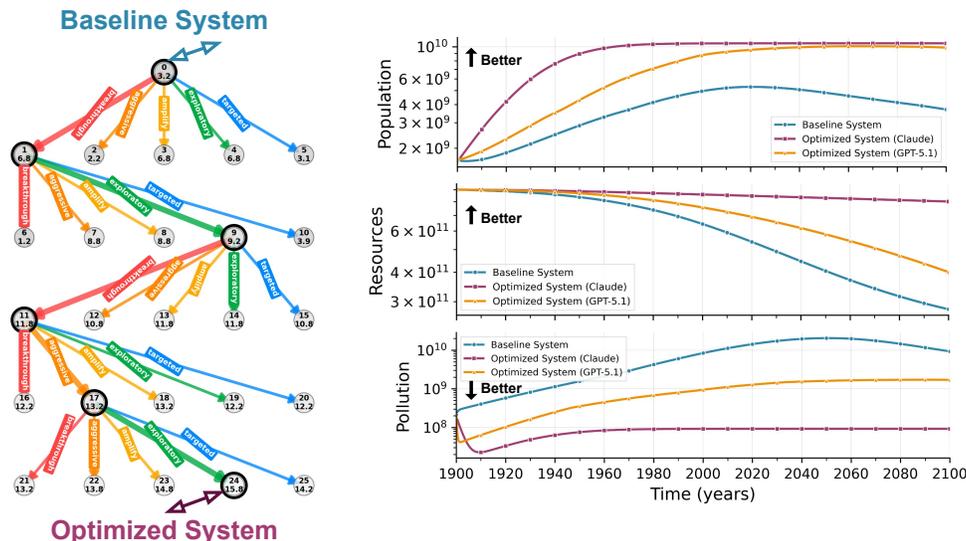
Data. We collect and convert models from classical works in complex systems in our modern framing. We convert *World Dynamics* (Forrester, 1971) from the DYNAMO language and 19 systems

378 from the book *Modeling Dynamic Biological Systems* (Hannon & Ruth, 2014) from the STELLA
 379 language. We utilize these systems in two ways: as foundations to be further optimized for multiple
 380 goals, and as ground truth systems that produce reference records serving as optimization targets.
 381

382 **Experimental Setup.** We use CEDAR with implementation details provided in the appendix. Our
 383 MCTS configuration allows 20 node selections, with each selection expanding to 5 new candidate
 384 nodes, resulting in a total of 100 expansions per search. We mainly use LLM providers: Anthropic
 385 Claude Sonnet 4.5 as LLM provider for its powerful coding capacity, and empirically set hyperpa-
 386 rameters $\alpha = 1$ and $\gamma = 2$ for node selection. The expansion strategy is selected by sampling from a
 387 set of candidate strategies. Details regarding the statistics of complex systems, expansion strategies
 388 and computational costs are covered in Appendix D.

390 4.2 FITTING AN ABSTRACT GOAL DESCRIBED IN NATURAL LANGUAGE

392 In this experiment, we present a comprehensive walkthrough of the optimization process with concrete
 393 examples. The system we utilize is *World Dynamics*, which is the complex system with the most
 394 variables in our collection. This complex system models the co-evolution of human population,
 395 resource utilization, and pollution as a nonlinear, feedback-driven system. Our ambitious goal is
 396 to jointly balance population growth, resource usage, and pollution. Figure 4 illustrates the MCTS
 397 search tree and the resulting system dynamics. The baseline, *resulting, published system* from the
 398 monograph Forrester (1971), is a complex system designed by hand, good enough such that further
 399 optimizing for one sub-goal often comes at the expense of others. We show this in Appendix E to
 400 highlight the intrinsic challenges of optimizing complex systems with competing objectives. , this
 401 visualization demonstrates how the search tree effectively further guides the optimization process
 402 toward solutions with higher scores. The final evolved system satisfies the vague goal, showing
 403 meaningful improvements over the initial configuration, with all three key variables demonstrating
 404 enhanced performance. Interestingly, with different backends, CEDAR discovers multiple structurally
 405 distinct but high-performing systems, each emphasizing different tradeoffs among objectives (for
 406 example, prioritizing population by Claude versus resource preservation by GPT-5.1). The full
 407 natural-language goal specification is detailed in Appendix E.



427 Figure 4: Illustration of fitting an abstract, language-based goal from *World Dynamics*. On the left,
 428 the MCTS tree is shown with the best expansion path highlighted. Each node is marked by the ID and
 429 LLM Judge Score. On the right, we display the system dynamics for both the initial (original) and
 430 final optimized systems, including Population (\uparrow), Pollution (\downarrow), and Resources (\uparrow). With CEDAR,
 431 the system reaches better states (increased population, fewer resources used, and less pollution).

4.3 QUANTITATIVE STUDIES - FITTING A CONCRETE RECORD

In this experiment, we task CEDAR with evolving a system (using only a bare-minimum skeleton) to fit the simulated record of a more complex model. While the strength of CEDAR lies in fitting vague, abstract goals, it can nonetheless be used to fit records from either observations or ground truth systems. To do so, we simply load the record into the goal and run CEDAR as-is without further modification, leveraging its universal design. We compare CEDAR with existing approaches based on black-box optimization. The ground truth is a population growth model with stochastic components. We use Optuna as our black box optimization baseline and employ both L1 and Dynamic Time Warping (DTW) distances as the scoring metric. Since the search space for black box optimization consists of only coefficients, we explore different formulae to support the baseline: no formulae, a simple version, or the full formulae in the ground truth system, the last of which gives the baseline a serious boost. Even so, as shown in Table 1 and Figure 5, we find that Optuna’s performance is severely limited by the skeletal formulae, while CEDAR, even without such predefined formulations, is capable of finding solutions better than Optuna, which has help of full formulae skeleton. The no/simple-formula baselines remain clearly worse both quantitatively and qualitatively. The details regarding the DTW, ground truth system, supporting formulae, and Optuna on simple baseline are provided in the Appendix F.

Table 1: L1 Distance and DTW comparison. CEDAR achieves better performance without formulae than black-box optimization with full formulae.

Method	Formulae		L1	DTW
	No	S. F.		
Optuna	✓		29.06	5503.68
Optuna		✓	26.01	4927.94
Optuna		✓	3.71	477.52
CEDAR (Claude)	✓		3.29	757.21
CEDAR (GPT-5.1)	✓		2.22	433.13

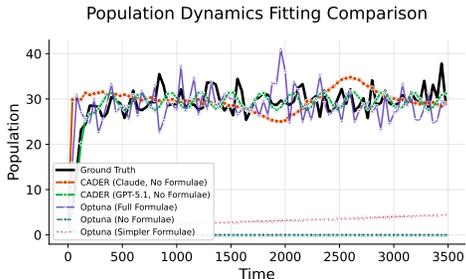


Figure 5: Comparison of system dynamics between CEDAR and the Optuna baseline when fitting concrete records.

4.4 INTERPRETABILITY

For the *World Dynamics* system in Section 4.2, we demonstrate interpretability by showing how the LLM’s analysis of candidate systems interacts with our search method. We highlight, in Table 2, the LLM responses along the path that produces the best system in the search tree shown in Figure 4. CEDAR provides clear, human-understandable interpretability of three aspects in the goal to be balanced. The LLM Judge analyzes issues with a balanced focus on the aspects specified in the goal, while the LLM Editor first evaluates general system performance and then fine-tunes the individual aspects, also attending to components that are not explicitly mentioned (e.g., capital productivity). This interpretability significantly reduces the human effort required to understand the optimization.

4.5 WHY MCTS: DIVERSITY OF SOLUTIONS AND BETTER PERFORMANCE

We find that the benefits of MCTS are three-fold. **First, MCTS optimizes complex systems while preserving solution diversity.** We assess diversity by focusing on the complex system discussed in Section 4.3. Within a single search tree, we sample nodes whose scores are close to the best solution and visualize each node’s population trajectory after the initial rapid growth phase (i.e. $t > 100$) in Figure 6. Although all systems satisfy the goal, they follow distinct trajectories. Unlike the Optuna-driven baselines in Section 4.3, CEDAR avoids collapsing solutions to a single point that resembles overfitting. This diversity enables better decision-making through sensitivity analysis.

The second benefit is better performance. This appears most clearly in ablation studies that evaluate the contribution of MCTS to CEDAR. We compare MCTS against linear search to demonstrate the importance of tree-based exploration and illustrate the differences between the two strategies (node score distributions and resulting population dynamics) in Figure 7. MCTS yields higher-scoring

Table 2: LLM Editor and Judge responses along the trajectory to the best solution. **Population**, **pollution**, and **resources** are manually colored for visibility. In this table, we show the LLM summarized version. **The relation between LLM analysis and LLM scores and the full transcripts of responses**, which include detailed analysis and editing rationales, are provided in Appendix G.

Depth	LLM Editor Response	LLM Judge Response
0 (root)	(No editor response for root, which is the initial system given as is.)	Unsustainable overshoot behavior - 69% resource depletion, 46x pollution increase
1	Breakthrough approach - 25-40% efficiency improvements targeting 6-8 score	Sustainable progress - population to 5.46B, 30% resources remaining, controlled pollution
2	Exploratory approach - radical efficiency through enhanced capital productivity	Paradigm shift - 7.0B population, 78% resource conservation, near-zero pollution
3	Ultra-aggressive conservation - pushing toward exceptional 9.5+ performance	Breakthrough sustainability - 3.2x population, 36% resource depletion, 90% pollution improvement
4	Revolutionary breakthrough - 75% resource reduction, 90% pollution control, 300% capital productivity	Best-in-tree performance - 6.3x population, 59% resource conservation, exceptional pollution control
5	Extreme efficiency breakthrough - 95% resource conservation, 98% pollution reduction, 15x capital productivity	Holy grail achievement - 6.4x population growth with only 19.9% resource depletion

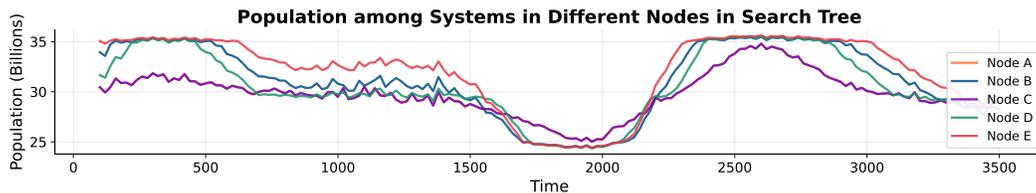


Figure 6: Diversity of solutions showing different population trajectories after the system passes the initial phase in the search tree for the complex system in Section 4.3.

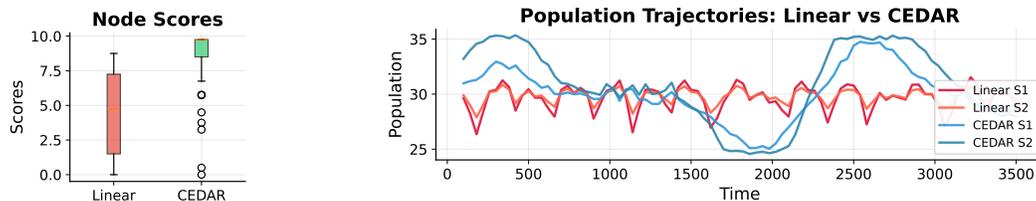


Figure 7: Comparison between MCTS and linear search showing performance differences in search strategy and resulting population dynamics.

nodes and smoother trajectories that avoid overfitting-like behavior. This shows that our design choice to combine MCTS with LLMs is crucial for performance gains.

The third benefit is LLM provider comparability. Our design separates optimization logic from tree search structure, enabling users to specify different LLM providers while maintaining the same framework. This supports comparing providers under identical conditions. Structurally distinct high-performing systems emerge in Section 4.2 and 4.3, each emphasizing different tradeoffs. This facilitates systematic evaluation of how foundation models affect optimization outcomes.

5 CONCLUSION

We propose CEDAR, combining Monte Carlo Tree Search with large language models to optimize complex systems for vague, natural-language goals, outperforming black-box optimization without predefined formulae. The unified Python representation provides interpretable insights and diverse solutions. Experiments demonstrate CEDAR’s capabilities across diverse complex-systems domains, opening new possibilities where traditional approaches struggle. **We discuss future directions on technical improvements and new human-in-the-loop paradigm in Appendix H.**

ETHICS STATEMENT

CEDAR streamlines optimization of complex dynamic systems through automated reasoning and iterative refinement. It applies across domains such as resource allocation, scheduling, and strategic planning. By supporting flexible objectives and constraints, it addresses problems that were previously difficult to solve, improving efficiency and effectiveness.

This study involves no human subjects or sensitive data and focuses solely on optimization methodology. Downstream use should be thoughtful and oriented toward positive impact, with reasonable safeguards and oversight to assess potential effects on individuals and society in real-world deployments.

REPRODUCIBILITY STATEMENT

To ensure reproducibility, we include the full MCTS algorithm in the main text and detailed prompts for the LLM Editor and Judge components in the appendix. Because LLM behavior can change as providers update models, we provide complete transcripts of LLM interactions throughout the optimization process for verification. We also provide the complete implementation code and the complex systems used in our experiments in the supplementary materials.

REFERENCES

- Virginia Anderson and Lauren Johnson. *Systems thinking basics*. Pegasus Communications Cambridge, MA, 1997.
- Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, and William Wang. Swe-search: Enhancing software agents with monte carlo tree search and iterative refinement. *arXiv preprint arXiv:2410.20285*, 2024.
- Ross D Arnold and Jon P Wade. A definition of systems thinking: A systems approach. *Procedia computer science*, 44:669–678, 2015.
- Yaman Barlas. Formal aspects of model validity and validation in system dynamics. *System Dynamics Review*, 12(3):183–210, 1996.
- Camille Bérard. Group model building using system dynamics: An analysis of methodological frameworks. *System Research and Behavioral Science*, 27(5):422–433, 2010.
- Ton M Blackshaw, Joseph C Davies, Kristian T Spoerer, and Jonathan D Hirst. Enhancing monte carlo tree search for retrosynthesis. *Journal of Chemical Information and Modeling*, 2025.
- David Brandfonbrener, Simon Henniger, Sibi Raja, Tarun Prasad, Chloe Loughridge, Federico Cassano, Sabrina Ruixin Hu, Jianang Yang, William E. Byrd, Robert Zinkov, and Nada Amin. Vermcts: Synthesizing multi-step programs using a verifier, a large language model, and tree search, 2024.
- Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12:1587–1627, 2011.
- G. M. J.-B. Chaslot, M. H. M. Winands, J. W. H. M. Uiterwijk, H. J. van den Herik, and B. Bouzy. Progressive strategies for monte-carlo tree search. In *Information Sciences 2007: Proceedings of the 10th Joint Conference (Salt Lake City, Utah, USA, 18-24 July 2007)*, volume 10, pp. 655–661. World Scientific Publishing Company, 2007. doi: 10.1142/9789812709677_0246.
- Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: Learning retrosynthetic planning with neural guided a* search. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020. URL <https://proceedings.mlr.press/v119/chen20k/cchen20k.pdf>.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.

- 594 Yizhou Chi, Kevin Yang, and Dan Klein. Reasoning with intermediate revision and search. In
595 *Findings of the North American Chapter of the Association for Computational Linguistics*, 2025.
596
- 597 Rémi Coulom. Computing “elo ratings” of move patterns in the game of go. *ICGA journal*, 30(4):
598 198–208, 2007.
- 599 Nicola Dainese, Matteo Merler, Minttu Alakuijala, and Pekka Marttinen. Generating code world
600 models with large language models guided by monte carlo tree search. In *Advances in Neural*
601 *Information Processing Systems*, 2024.
602
- 603 Yonathan Efroni, Gal Dalal, Bruno Scherrer, and Shie Mannor. How to combine tree-search methods
604 in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
605 volume 33, pp. 3494–3501, 2019.
- 606 Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Moham-
607 madamin Barekatin, Alexander Novikov, Francisco J R Ruiz, et al. Discovering faster matrix
608 multiplication algorithms with reinforcement learning. *Nature*, 610(7930):47–53, 2022.
609
- 610 David N. Ford and John D. Sterman. Expert knowledge elicitation to improve mental and formal
611 models. *System Dynamics Review*, 13(1):57–80, 1997.
612
- 613 Jay W Forrester. *World dynamics*. Wright-Allen Press, Cambridge, Mass., 1971.
- 614 Jay W Forrester. Mit system dynamics in education project. <https://www.merlot.org/merlot/viewMaterial.htm?id=518948>, 2011. URL <https://www.merlot.org/merlot/viewMaterial.htm?id=518948>.
615
616
617
- 618 Bingzheng Gan, Yufan Zhao, Tianyi Zhang, Jing Huang, Yusu Li, Shu Xian Teo, Changwang Zhang,
619 and Wei Shi. Master: A multi-agent system with llm specialized mcts. In *Proceedings of the 2025*
620 *Conference of the North American Chapter of the Association for Computational Linguistics*, 2025.
- 621 Zitian Gao, Boye Niu, Xuzheng He, Haotian Xu, Hongzhang Liu, Aiwei Liu, Xuming Hu, and Lijie
622 Wen. Interpretable contrastive monte carlo tree search reasoning. *arXiv preprint arXiv:2410.01707*,
623 2024.
624
- 625 Burak Güneralp. A principle on structure-behavior relations in system dynamics. *System Dynamics*
626 *Review*, 20(4):331–349, 2004.
- 627 Bruce Hannon and Matthias Ruth. *Modeling Dynamic Biological Systems*. Springer, 2 edition, 2014.
628
- 629 James H. Hines. Molecules of structure. In *Proceedings of the 1996 International System Dynamics*
630 *Conference*, 1996.
631
- 632 Yuichi Inoue, Kou Misaki, Yuki Imajuku, So Kuroki, Taishi Nakamura, and Takuya Akiba. Wider or
633 deeper? scaling llm inference-time compute with adaptive branching tree search. *arXiv preprint*
634 *arXiv:2503.04412*, 2025.
- 635 Michael Katz, Harsha Kokel, Kavitha Srinivas, and Shirin Sohrabi. Thought of search: Planning with
636 language models for open-world problem solving. In *Advances in Neural Information Processing*
637 *Systems*, 2024.
638
- 639 Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *Proceedings of the*
640 *17th European Conference on Machine Learning*, pp. 282–293. Springer, 2006.
- 641 Kenneth Kolson. The politics of simcity. *PS: Political Science & Politics*, 29(1):43–46, 1996.
642
- 643 Akarsh Kumar, Chris Lu, Louis Kirsch, Yujin Tang, Kenneth O Stanley, Phillip Isola, and David Ha.
644 Automating the search for artificial life with foundation models. *Artificial Life*, 31(3):368–396,
645 2025.
646
- 647 Zheyuan Lai and Yingming Pu. Prim: Principle-inspired material discovery through multi-agent
collaboration. *arXiv preprint arXiv:2504.08810*, 2025.

- 648 Lucas Lehnert, Sainbayar Sukhbaatar, DiJia Su, Qinqing Zheng, Paul Mcvay, Michael Rabbat, and
 649 Yuandong Tian. Beyond a*: Better planning with transformers via search dynamics bootstrapping.
 650 *arXiv preprint arXiv:2402.14083*, 2024.
- 651 Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem.
 652 Camel: Communicative agents for "mind" exploration of large language model society. In *Advances*
 653 *in Neural Information Processing Systems*, volume 36, pp. 51991–52008, 2023.
- 654 Jierui Li, Hung Le, Yingbo Zhou, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Code-
 655 tree: Agent-guided tree search for code generation with large language models. *arXiv preprint*
 656 *arXiv:2411.04329*, 2024. doi: 10.48550/arXiv.2411.04329.
- 657 Qingyao Li, Wei Xia, Kounianhua Du, Xinyi Dai, Ruiming Tang, Yasheng Wang, Yong Yu, and
 658 Weinan Zhang. Rethinkmcts: Refining erroneous thoughts in monte carlo tree search for code
 659 generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language*
 660 *Processing*. Association for Computational Linguistics, 2025a. Also available as arXiv:2409.09584.
- 661 Yu Li, Lehui Li, Zhihao Wu, Qingmin Liao, Jianye Hao, Kun Shao, Fengli Xu, and Yong Li.
 662 Agentswift: Efficient llm agent design via value-guided hierarchical search. *arXiv preprint*
 663 *arXiv:2506.06017*, 2025b.
- 664 Zujie Liang, Feng Wei, Wujiang Xu, Lin Chen, Yuxi Qian, and Xinhui Wu. I-mcts: Improving automl
 665 models with monte carlo tree search and llm guidance. *arXiv preprint arXiv:2502.14693*, 2025.
- 666 Viliam Lisy, Vojta Kovarik, Marc Lanctot, and Branislav Bosansky. Convergence of monte carlo tree
 667 search in simultaneous move games. *Advances in Neural Information Processing Systems*, 26,
 668 2013.
- 669 Ning-Yuan Georgia Liu and David R. Keith. Leveraging large language models for automated
 670 causal loop diagram generation: Enhancing system dynamics modeling through curated prompting
 671 techniques. *arXiv preprint arXiv:2503.21798*, 2025.
- 672 Toni J.B. Liu, Nicolas Boulle, Raphaël Sarfati, and Christopher Earls. Llms learn governing principles
 673 of dynamical systems, revealing an in-context neural scaling law. In *Proceedings of the 2024*
 674 *Conference on Empirical Methods in Natural Language Processing*, pp. 15138–15153, 2024. doi:
 675 10.18653/v1/2024.emnlp-main.842.
- 676 Xiao Luo, Binqi Chen, Haixin Wang, Zhiping Xiao, Ming Zhang, and Yizhou Sun. How do
 677 large language models perform in dynamical system modeling. *Findings of the Association for*
 678 *Computational Linguistics: NAACL 2025*, 2025.
- 679 Daniel J Mankowitz, Andrea Michi, Anton Zhernov, Marco Gelmi, Marco Selvi, Cosmin Paduraru,
 680 Edouard Leurent, Shariq Hashme, Jean-Baptiste Lespiau, Alex Ahern, et al. Faster sorting
 681 algorithms discovered using deep reinforcement learning. *Nature*, 618(7964):257–263, 2023.
- 682 Wannes Meert, Kilian Hendrickx, and Toon Van Craenendonck. wannesm/dtaidistance, 2020. URL
 683 <https://github.com/wannesm/dtaidistance>. v2.3.13.
- 684 MIT System Dynamics in Education Project. Road maps: A guide to learning system dynamics,
 685 1998. Massachusetts Institute of Technology, Sloan School of Management.
- 686 Muhammad Umair Nasir, Sam Earle, Julian Togelius, Steven James, and Christopher Cleghorn.
 687 Llmatic: Neural architecture search via large language models and quality diversity optimization.
 688 In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1078–1086. ACM,
 689 2024. doi: 10.1145/3638529.3654017.
- 690 Katsuhiko Ogata. *System dynamics*. Pearson Education, 4 edition, 2004.
- 691 Steve Peterson and NH West Lebanon. Barry richmond, system dynamics and public policy. In *21st*
 692 *System Dynamics Conference*, pp. 1–14, 2003.
- 693 Michael J. Radzicki and Robert A. Taylor. Origin of system dynamics. <https://web.archive.org/web/20171016181415/http://www.systemdynamics.org/DL-IntroSySDyn/origin.htm>, 2011. URL <http://www.systemdynamics.org/DL-IntroSySDyn/origin.htm>. Accessed: 2017-10-16.

- 702 Chotirat Ann Ratanamahatana and Eamonn Keogh. Everything you know about dynamic time
703 warping is wrong. In *Third Workshop on Mining Temporal and Sequential Data*, 2004.
704
- 705 Barry Richmond. Stella: Software for bringing system dynamics to the other 98%. In *Proceedings*
706 *of the 1985 international conference of the System Dynamics Society: 1985 International system*
707 *dynamics conference*, pp. 706–718. System Dynamics Society 49 Bedford Road, Lincoln, MA
708 01773 USA, 1985.
- 709 Matthias Ruth and Bruce Hannon. *Modeling dynamic economic systems*. Springer, 2012.
710
- 711 Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word
712 recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978.
- 713 William Schoenberg. Understanding model behavior using loops that matter. *arXiv preprint*
714 *arXiv:1908.11434*, 2019.
715
- 716 Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Jimmy
717 Agapiou, et al. Mastering atari, go, chess and shogi by planning with a learned model. *arXiv*
718 *preprint arXiv:1911.08265*, 2019.
- 719 Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon
720 Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari,
721 go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
722
- 723 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
724 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*
725 *Systems*, 36:8634–8652, 2023.
- 726 David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche,
727 Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering
728 the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- 729 David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez,
730 Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Si-
731 monyan, and Demis Hassabis. Mastering chess and shogi by self-play with a general reinforcement
732 learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017a.
733
- 734 David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez,
735 Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without
736 human knowledge. *nature*, 550(7676):354–359, 2017b.
- 737 Xiaozhuang Song, Shufei Zhang, and Tianshu Yu. Rekg-mcts: Reinforcing llm reasoning on
738 knowledge graphs via training-free monte carlo tree search. In *Findings of the Association for*
739 *Computational Linguistics*, 2025.
- 740 Henry W Sprueill, Carl Edwards, Khushbu Agarwal, Mariefel V Olarte, Udishnu Sanyal, Conrad
741 Johnston, Hongbin Liu, Heng Ji, and Sutanay Choudhury. Chemreasoner: Heuristic search over
742 a large language model’s knowledge space using quantum-chemical feedback. *arXiv preprint*
743 *arXiv:2402.10980*, 2024.
744
- 745 John D. Sterman. *Business Dynamics: Systems Thinking and Modeling for a Complex World*.
746 Irwin/McGraw-Hill, 2000.
- 747 Jac A.M. Vennix et al. Group model building: A systematic review of the literature. *System Dynamics*
748 *Review*, 2021.
- 749 Fang Wu, Weihao Xuan, Heli Qi, Ximing Lu, Aaron Tu, Li Erran Li, and Yejin Choi. Deepsearch:
750 Overcome the bottleneck of reinforcement learning with verifiable rewards via monte carlo tree
751 search. *arXiv preprint arXiv:2509.25454*, 2025a.
752
- 753 Mengsong Wu, YaFei Wang, Yidong Ming, Yuqi An, Yuwei Wan, Wenliang Chen, Binbin Lin,
754 Yuqiang Li, Tong Xie, and Dongzhan Zhou. Chemagent: Enhancing llms for chemistry and
755 materials science through tree-search based tool learning. *arXiv preprint arXiv:2506.07551*,
2025b.

- 756 Bin Xu, Yiguan Lin, Yinghao Li, and Yang Gao. Sra-mcts: Self-driven reasoning augmentation with
757 monte carlo tree search for code generation. In *Proceedings of the Thirty-Fourth International Joint*
758 *Conference on Artificial Intelligence*, pp. 8678–8686. IJCAI, 2025. doi: 10.24963/ijcai.2025/965.
759 Also available as arXiv:2411.11053.
- 760 Yilong Xu, Yang Liu, and Hao Sun. Reinforcement symbolic regression machine. In *Proceedings*
761 *of the International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2305.14656>.
762
- 763 Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune,
764 and David Ha. The ai scientist-v2: Workshop-level automated scientific discovery via agentic tree
765 search. *arXiv preprint arXiv:2504.08066*, 2025.
- 766 Aldo A. Zagonel. *Model conceptualization in group model building*. PhD thesis, University at Albany,
767 State University of New York, 2002.
- 768 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D Goodman. Star: Self-taught reasoner bootstrapping
769 reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488,
770 2022.
- 771 Yanjie Zhang, Jie Feng, Keyu Sun, Ningning Zhang, Kan Wu, Weijia Liu, Yue Dong, Mingjie
772 Zhang, Yuqiang Li, and Joey Tianyi Zhou. Llm-srbench: A new benchmark for scientific equation
773 discovery with large language models. *arXiv preprint arXiv:2504.10415*, 2025a.
- 774 Zixuan Zhang, Yao Wang, Zhiyi Zhang, Yang Yu, and Zhi-Hua Tan. Monte carlo tree search for com-
775 prehensive exploration in llm-based automatic heuristic design. *arXiv preprint arXiv:2501.08603*,
776 2025b.
- 777 Zirui Zhao, Wee Sun Lee, and David Hsu. Large language models as commonsense knowledge for
778 large-scale task planning. In *Advances in Neural Information Processing Systems*, 2023.
- 779 Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language
780 agent tree search unifies reasoning acting and planning in language models. *arXiv preprint*
781 *arXiv:2310.04406*, 2023a.
- 782 Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language
783 agent tree search unifies reasoning acting and planning in language models. In *International*
784 *Conference on Machine Learning*, pp. 61694–61721. PMLR, 2024.
- 785 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan,
786 and Jimmy Ba. Large language models are human-level prompt engineers. *arXiv preprint*
787 *arXiv:2211.01910*, 2023b.
- 788 Xinhe Zhu, Yuanyou Shi, and Yongmin Zhong. An ekf prediction of covid-19 propagation under
789 vaccinations and viral variants. *Mathematics and Computers in Simulation*, 231:221–238, 2025.
- 790
- 791
- 792
- 793
- 794
- 795
- 796
- 797
- 798
- 799
- 800
- 801
- 802
- 803
- 804
- 805
- 806
- 807
- 808
- 809

APPENDIX

A PYTHON IMPLEMENTATION OF COMPLEX SYSTEMS

This section presents the Python implementation of our unified representation of complex systems, which comprises four essential components: (1) initialization, (2) intermediate computations, (3) derivative evaluations, and (4) Euler integration. A minimal example illustrating the core structure is shown below:

```

819 1 # (1) initialization
820 2 dt = 0.01
821 3 t0 = 0.0
822 4 tf = 120.0
823 5 K = 100.0
824 6 R = 0.1
825 7 N = 10.0
826 8
827 9 t = t0
828 10 while t <= tf + 1e-12:
829 11
830 12     # (2) intermediate computations
831 13     carrying_capacity_factor = 1.0 - N / K
832 14     growth_rate = R * N * carrying_capacity_factor
833 15
834 16     # (3) derivative evaluations
835 17     dN = growth_rate # dN = R * N * (1 - N/K)
836 18
837 19     # (4) Euler integration
838 20     N = N + dt * dN
839 21     t += dt

```

In practice, we enrich the code with comprehensive comments for both human interpretation and LLM comprehension. This enables LLMs to understand system semantics, guiding optimization while enforcing key constraints. Several critical restrictions must be respected, such as preserving the Euler integration scheme, for optimization to execute successfully. Below, we present the *World Dynamics* system (Forrester, 1971) (one with the largest by number of variables in our dataset), represented using the restricted Python framework.

```

840 1 # ===== BEGIN CONFIG =====
841 2 # Simulation configuration parameters (DO NOT CHANGE - these are set by the system)
842 3 dt = 0.2 # time step size
843 4 t0 = 1900.0 # start time
844 5 tf = 2100.0 # end time
845 6 seed = 1234 # random seed (set None for nondeterministic runs)
846 7 # ===== END CONFIG =====
847 8
848 9 # ===== BEGIN STATE =====
849 10 # STATE VARIABLES - integrated over time (must define derivatives below)
850 11 P = 1.65e9 # population (people) # NECESSARY
851 12 NR = 900e9 # natural resources (natural resource units) # NECESSARY
852 13 CI = 0.4e9 # capital investment (capital units) # NECESSARY
853 14 POL = 0.2e9 # pollution (pollution units) # NECESSARY
854 15 CIAF = 0.2 # capital-investment-in-agriculture fraction (dimensionless) # NECESSARY
855 16 # ===== END STATE =====
856 17
857 18 # ===== BEGIN ALGEBRAIC (OPTIONAL) =====
858 19 # Variables are now declared in the HELPERS section within the simulation loop
859 20
860 21 # ===== BEGIN PARAMS =====
861 22 # CONSTANT PARAMETERS used in equations
862 23 # (All constants are inlined in the code)
863 24 # ===== END PARAMS =====
864 25
865 26 # ===== MAIN LOOP =====
866 27 t = t0
867 28 while t <= tf + 1e-12:
868 29     # ===== BEGIN INPUTS_T =====
869 30     # TIME-DEPENDENT INPUTS - use wrappers like sin(), exp(), graph()
870 31     # (No time-dependent inputs in this model)
871 32     # ===== END INPUTS_T =====
872 33
873 34     # ===== BEGIN INPUTS_RND =====
874 35     # RANDOM INPUTS - use gauss(), uniform() wrappers (no STATE usage).
875 36     # (No random inputs in this model)

```

```

864 37 # ===== END INPUTS_RND =====
865 38
866 39 # ===== BEGIN HELPERS =====
867 40 # HELPERS - All computed variables (algebraic and intermediate expressions)
868 41 # in dependency order.
869 42
870 43 # Basic ratios and derived quantities
871 44 CIR = CI / P # capital-investment ratio (capital units/person)
872 45 CR = P / (135e6 * 26.5) # crowding ratio (dimensionless)
873 46 NRFR = NR / 900e9 # natural-resource fraction remaining (dimensionless)
874 47 POLR = POL / 3.6e9 # pollution ratio (dimensionless)
875 48
876 49 # Material standard of living components
877 50 CIAF_current = CIAF # current capital-investment-in-agriculture fraction
878 51 NREM = graph(
879 52     NRFR, ((0, 0), (0.25, 0.15), (0.5, 0.5), (0.75, 0.85), (1, 1))
880 53 ) # natural-resource-extraction multiplier
881 54 ECIR = CIR * (1 - CIAF_current) * NREM / (1 - 0.3) # effective-capital-investment
882 55     ratio
883 56 MSL = ECIR / 1 # material standard of living (dimensionless)
884 57
885 58 # Birth rate multipliers
886 59 BRMM = graph(
887 60     MSL, ((0, 1.2), (1, 1), (2, 0.85), (3, 0.75), (4, 0.7), (5, 0.7))
888 61 ) # birth-rate-from-material multiplier
889 62 BRCM = graph(
890 63     CR, ((0, 1.05), (1, 1), (2, 0.9), (3, 0.7), (4, 0.6), (5, 0.55))
891 64 ) # birth-rate-from-crowding multiplier
892 65 BRPM = graph(
893 66     POLR, ((0, 1.02), (10, 0.9), (20, 0.7), (30, 0.4), (40, 0.25), (50, 0.15), (60,
894 67         0.1))
895 68 ) # birth-rate-from-pollution multiplier
896 69
897 70 # Food ratio components
898 71 CIRA = CIR * CIAF_current / 0.3 # capital-investment ratio in agriculture
899 72 FPCI = graph(
900 73     CIRA, ((0, 0.5), (1, 1), (2, 1.4), (3, 1.7), (4, 1.9), (5, 2.05), (6, 2.2))
901 74 ) # food potential from capital investment
902 75 FCM = graph(CR, ((0, 2.4), (1, 1), (2, 0.6), (3, 0.4), (4, 0.3), (5, 0.2))) # food-
903 76     from-crowding multiplier
904 77 FPM = graph(
905 78     POLR, ((0, 1.02), (10, 0.9), (20, 0.65), (30, 0.35), (40, 0.2), (50, 0.1), (60,
906 79         0.05))
907 80 ) # food-from-pollution multiplier
908 81 FR = FPCI * FCM * FPM * (1 if t >= 1970 else 1) / 1 # food ratio
909 82 BREM = graph(FR, ((0, 0), (1, 1), (2, 1.6), (3, 1.9), (4, 2))) # birth-rate-from-food
910 83     multiplier
911 84
912 85 # Death rate multipliers
913 86 DRMM = graph(
914 87     MSL,
915 88     (
916 89         (0, 3),
917 90         (0.5, 1.8),
918 91         (1, 1),
919 92         (1.5, 0.8),
920 93         (2, 0.7),
921 94         (2.5, 0.6),
922 95         (3, 0.53),
923 96         (3.5, 0.5),
924 97         (4, 0.5),
925 98         (4.5, 0.5),
926 99         (5, 0.5),
927 100     ),
928 101 ) # death-rate-from-material multiplier
929 102 DRCM = graph(CR, ((0, 0.9), (1, 1), (2, 1.2), (3, 1.5), (4, 1.9), (5, 3))) # death-
930 103     rate-from-crowding multiplier
931 104 DRPM = graph(
932 105     POLR, ((0, 0.92), (10, 1.3), (20, 2), (30, 3.2), (40, 4.8), (50, 6.8), (60, 9.2))
933 106 ) # death-rate-from-pollution multiplier
934 107 DRFM = graph(
935 108     FR, ((0, 30), (0.25, 3), (0.5, 2), (0.75, 1.4), (1, 1), (1.25, 0.7), (1.5, 0.6),
936 109         (1.75, 0.5), (2, 0.5))
937 110 ) # death-rate-from-food multiplier

```

```

918 # Capital investment multiplier
919 CIM = graph(MSL, ((0, 0.1), (1, 1.0), (2, 1.8), (3, 2.4), (4, 2.8), (5, 3))) # capital
920     -investment multiplier
921
922 # Pollution components
923 POLCM = graph(CIR, ((0, 0.05), (1, 1), (2, 3), (3, 5.4), (4, 7.4), (5, 8))) #
924     pollution-from-capital multiplier
925 POLAT = graph(
926     POLR, ((0, 0.6), (10, 2.5), (20, 5), (30, 8), (40, 11.5), (50, 15.5), (60, 20))
927 ) # pollution-absorption time
928
929 # Quality of life components
930 QLM = graph(MSL, ((0, 0.2), (1, 1), (2, 1.7), (3, 2.3), (4, 2.7), (5, 2.9))) # quality
931     of life from material
932 QLC = graph(
933     CR,
934     (
935         (0, 2),
936         (0.5, 1.3),
937         (1, 1),
938         (1.5, 0.75),
939         (2, 0.55),
940         (2.5, 0.45),
941         (3, 0.38),
942         (3.5, 0.3),
943         (4, 0.25),
944         (4.5, 0.22),
945         (5, 0.2),
946     ),
947 ) # quality of life from crowding
948 QLF = graph(FR, ((0, 0), (1, 1), (2, 1.8), (3, 2.4), (4, 2.7))) # quality of life from
949     food
950 QLP = graph(
951     POLR, ((0, 1.04), (10, 0.85), (20, 0.6), (30, 0.3), (40, 0.15), (50, 0.05), (60,
952     0.02))
953 ) # quality of life from pollution
954 QL = 1 * QLM * QLC * QLF * QLP # quality of life
955
956 # Capital investment fraction adjustment components
957 CFIFR = graph(FR, ((0, 1), (0.5, 0.6), (1, 0.3), (1.5, 0.15), (2, 0.1))) # capital
958     fraction indicated by food ratio
959 CIQR = graph(QLM / QLF, ((0, 0.7), (0.5, 0.8), (1, 1), (1.5, 1.5), (2, 2))) # capital-
960     investment-from-quality ratio
961
962 # ===== END HELPERS =====
963
964 # ===== BEGIN DERIVATIVES =====
965 # DERIVATIVES for each STATE variable.
966
967 # Birth and death rates
968 BR = P * (0.04 if t >= 1970 else 0.04) * BRFM * BRMM * BRMC * BRPM # birth rate
969 DR = P * (0.028 if t >= 1970 else 0.028) * DRMM * DRPM * DRFM * DRCM # death rate
970
971 # Natural resource usage rate
972 NRUR = P * (1 if t >= 1970 else 1) * NRMM # natural-resource-usage rate
973
974 # Capital investment flows
975 CIG = P * CIM * (0.05 if t >= 1970 else 0.05) # capital-investment generation
976 CID = CI * (0.025 if t >= 1970 else 0.025) # capital-investment discard
977
978 # Pollution flows
979 POLG = P * (1 if t >= 1970 else 1) * POLCM # pollution generation
980 POLA = POL / POLAT # pollution absorption
981
982 # State derivatives
983 dP = BR - DR
984 dNR = -NRUR
985 dCI = CIG - CID
986 dPOL = POLG - POLA
987 dCIAF = (1 / 15) * (CFIFR * CIQR - CIAF)
988 # ===== END DERIVATIVES =====
989
990 # ----- Euler integration (engine; do not edit) -----
991 P = P + dt * dP
992 NR = NR + dt * dNR
993 CI = CI + dt * dCI
994 POL = POL + dt * dPOL
995 CIAF = CIAF + dt * dCIAF
996
997 # Advance time
998 t += dt

```

B NODE EXPANSION DETAILS

B.1 ENRICHING PYTHON CODE FOR COMPLEX SYSTEMS

To improve LLM understanding and tool use, we extend the code with extensive inline documentation and predefined mathematical primitives (trigonometric functions, delay operators, smoothing functions, and lookup tables). This design supports both human comprehension and automated LLM-based system generation and modification via clear semantic annotations and standardized function interfaces. Furthermore, this approach ensures the code remains self-contained and executable. The extended code is shown below. Complex systems are embedded between `BEGIN WHOLE SYSTEM` and `END WHOLE SYSTEM`.

```

1 # =====
2 # sim_base.py - Minimal Euler simulator with edit markers
3 #
4 # PURPOSE:
5 # - Self-contained forward Euler simulator:
6 #     X(t+dt) = X(t) + dt * dX/dt
7 #
8 # HOW TO EDIT:
9 # - Only change code between the "BEGIN ... / END ..." markers below.
10 # - Do NOT modify imports, wrappers, loop structure, Euler integration, or OUTPUT section.
11 # - Keep STATE, ALGEBRAIC, HELPERS, and DERIVATIVES consistent.
12 # - NEVER change simulation configuration parameters (dt, t0, tf, seed) in the CONFIG
13 #     section.
14 #
15 # VARIABLE MODIFICATION RULES:
16 # - STATE VARIABLES: You can add new state variables. You can delete state variables ONLY
17 #     if they
18 #     are not marked as "# NECESSARY" in comments. All state variables must have derivatives.
19 # - ALGEBRAIC VARIABLES: You can add new algebraic variables. You can delete algebraic
20 #     variables
21 #     ONLY if they are not marked as "# NECESSARY" in comments.
22 # - HELPERS: You can freely add, remove, and modify helper variables. You can change their
23 #     computation formulas and reorder them, as long as the code remains valid.
24 # - When converting from other systems: Mark essential state and algebraic variables with
25 #     "# NECESSARY" comments so optimization/editing knows which variables cannot be deleted.
26 #
27 # VARIABLE TYPES:
28 # - STATE: variables integrated over time (must have derivatives).
29 # - INPUTS_T: pure functions of time (cannot depend on STATE).
30 # - INPUTS_RND: fresh random values each step (cannot depend on STATE).
31 # - HELPERS: all computed variables (algebraic and intermediate expressions)
32 #     in dependency order. Can use wrappers like sin(), exp(), graph(),
33 #     delay(), smthl() - e.g. nonlinear damping from a table.
34 # - DERIVATIVES: d/dt for each STATE variable.
35 #
36 # TRACE CONTENT (per step, in fixed order):
37 #     1) t           - current time
38 #     2) STATE variables
39 #     3) HELPERS (all computed variables)
40 #     4) DERIVATIVES
41 #
42 # OUTPUT:
43 # - If --csv <path> is provided -> write full simulation to that CSV file.
44 # - Otherwise -> print the raw CSV (header + all rows) to stdout.
45 # - Columns/ordering come from the snapshot dict defined inside the loop.
46 # =====
47
48 import math
49 import random
50 import argparse
51 import pandas as pd # used only for CSV export
52 from typing import Iterable, Tuple
53
54 # ----- Wrapper Functions -----
55 # These wrappers hide Python semantics; LLMs should ONLY call these.
56
57 def sin(x: float) -> float:
58     """Sine function (angle in radians)."""
59     return math.sin(x)
60
61 def cos(x: float) -> float:

```

```

1026 59     """Cosine function (angle in radians)."""
1027 60     return math.cos(x)
1028 61
1029 62 def exp(x: float) -> float:
1030 63     """Exponential function e^x."""
1031 64     return math.exp(x)
1032 65
1033 66 def tanh(x: float) -> float:
1034 67     """Hyperbolic tangent."""
1035 68     return math.tanh(x)
1036 69
1037 70 def sqrt(x: float) -> float:
1038 71     """Square root."""
1039 72     return math.sqrt(x)
1040 73
1041 74 def log(x: float) -> float:
1042 75     """Natural logarithm."""
1043 76     return math.log(x)
1044 77
1045 78 def gauss(std: float) -> float:
1046 79     """Gaussian random variable with mean=0 and standard deviation=std."""
1047 80     return random.gauss(0.0, std)
1048 81
1049 82 def uniform(lo: float, hi: float) -> float:
1050 83     """Uniform random variable between lo and hi."""
1051 84     return random.uniform(lo, hi)
1052 85
1053 86 def graph(v: float, table: Iterable[Tuple[float, float]]) -> float:
1054 87     """
1055 88     Lookup helper with linear interpolation.
1056 89     - table is an iterable of (x, y) points sorted by x.
1057 90     - If v < x0, return y0.
1058 91     - If v > xN, return yN.
1059 92     - Else, linearly interpolate between nearest points.
1060 93     """
1061 94     table = list(table)
1062 95     if not table:
1063 96         raise ValueError("graph() called with empty table")
1064 97
1065 98     if v <= table[0][0]:
1066 99         return table[0][1]
1067 100     if v >= table[-1][0]:
1068 101         return table[-1][1]
1069 102
1070 103     for i in range(len(table) - 1):
1071 104         x0, y0 = table[i]
1072 105         x1, y1 = table[i + 1]
1073 106         if x0 <= v <= x1:
1074 107             if x1 == x0:
1075 108                 return y0
1076 109             frac = (v - x0) / (x1 - x0)
1077 110             return y0 + frac * (y1 - y0)
1078 111     return table[-1][1]
1079 112
1080 113 # ----- Delay/Smooth System -----
1081 114 # STELLA-compatible delay/smooth functions with time-based semantics.
1082 115 # - All functions require a 'name' parameter for unique identification
1083 116 # - Updates occur when functions are called during helper computation
1084 117 # - Fixed delays use time-indexed history with linear interpolation
1085 118 # - Smooth functions use stock-based integration with proper dt scaling
1086 119 # - All functions are dt-independent (same behavior regardless of step size)
1087 120
1088 121 _delay_states = {} # Registry for all delay/smooth function states
1089 122
1090 123 def delay(input_val: float, delay_time: float, name: str, initial_value: float = None) ->
1091 124     float:
1092 125     """Fixed lag delay - returns input value from delay_time ago."""
1093 126     key = f"delay_{name}_{delay_time}"
1094 127
1095 128     if key not in _delay_states:
1096 129         init_val = input_val if initial_value is None else initial_value
1097 130         _delay_states[key] = {
1098 131             'history': [(t, init_val)],
1099 132             'last_t': t
1100 133         }
1101 134
1102 135     state = _delay_states[key]
1103 136
1104 137     # Add current input to history (only if time advanced)
1105 138     if t > state['last_t']:
1106         state['history'].append((t, input_val))

```

```

1080     state['last_t'] = t
1081
1082     # Clean old history beyond delay time
1083     cutoff_time = t - delay_time
1084     state['history'] = [(t_hist, val) for t_hist, val in state['history']
1085                        if t_hist >= cutoff_time]
1086
1087     # Find value at t - delay_time using linear interpolation
1088     target_time = t - delay_time
1089     if not state['history'] or target_time <= state['history'][0][0]:
1090         return state['history'][0][1]
1091
1092     for i in range(len(state['history']) - 1):
1093         t1, v1 = state['history'][i]
1094         t2, v2 = state['history'][i + 1]
1095         if t1 <= target_time <= t2:
1096             if t2 == t1:
1097                 return v1
1098             frac = (target_time - t1) / (t2 - t1)
1099             return v1 + frac * (v2 - v1)
1100
1101     return state['history'][-1][1]
1102
1103 def smth1(input_val: float, smooth_time: float, name: str, initial_value: float = None) ->
1104     float:
1105     """First-order exponential smooth - stock-based smoothing process."""
1106     key = f"smth1_{name}_{smooth_time}"
1107
1108     if key not in _delay_states:
1109         init_val = input_val if initial_value is None else initial_value
1110         _delay_states[key] = {
1111             'smooth_of_input': init_val, # The stock being smoothed
1112             'last_t': t
1113         }
1114
1115     state = _delay_states[key]
1116
1117     if t > state['last_t']:
1118         dt_step = t - state['last_t']
1119
1120         # SMTH1 equations from STELLA:
1121         # Change_in_Smooth = (Input - Smooth_of_Input) / Averaging_Time
1122         change_in_smooth = (input_val - state['smooth_of_input']) / smooth_time if
1123             smooth_time > 0 else 0.0
1124
1125         # Stock integration: Smooth_of_Input = Smooth_of_Input + dt * Change_In_Smooth
1126         state['smooth_of_input'] += dt_step * change_in_smooth
1127         state['last_t'] = t
1128
1129     # SMTH1 returns the smoothed stock value
1130     return state['smooth_of_input']
1131
1132 # ===== BEGIN WHOLE SYSTEM =====
1133 #
1134 #
1135 # ===== END WHOLE SYSTEM =====
1136
1137 # ===== OUTPUT (FIXED) =====
1138 # DO NOT EDIT THIS SECTION.
1139 def _parse_args():
1140     ap = argparse.ArgumentParser(description="Run Euler simulation and export CSV.")
1141     ap.add_argument("--csv", dest="csv_path", default=None,
1142                    help="Write results to this CSV path. If omitted, prints CSV to stdout.
1143                        ")
1144     return ap.parse_args()
1145
1146 if __name__ == "__main__":
1147     args = _parse_args()
1148     df = pd.DataFrame(trace) # column order = insertion order of dict
1149     if args.csv_path:
1150         df.to_csv(args.csv_path, index=False, float_format='%.6f')
1151         print(f"Simulation complete. Results written to {args.csv_path}")
1152     else:
1153         print(df.to_csv(index=False, float_format='%.6f'), end="")

```

1134 B.2 PROMPT DESIGN AND CONTEXT PRESENTATION

1135
1136 The effectiveness of CEDAR relies on well-crafted prompts that enable LLMs to understand system
1137 dynamics and generate meaningful modifications. We employ specialized prompts for two key com-
1138 ponents: the LLMEDITOR for system modification and the LLMJUDGE for performance evaluation.
1139 These prompts incorporate domain-specific instructions, constraint specifications, and contextual
1140 information to guide LLM reasoning effectively.
1141

1142 Prompt for LLMEDITOR

```
1143
1144
1145 SUGGEST CODE MODIFICATIONS FOR TREE SEARCH OPTIMIZATION
1146
1147 GOAL:
1148 {goal}
1149
1150 CURRENT CODE:
1151 {code}
1152
1153 SIMULATION RESULTS:
1154 {csv_data}
1155
1156 TREE CONTEXT:
1157 {tree_context}
1158
1159 === For your reference ===
1160
1161 Format for code in tree context:
1162 The tree context shows only the differences in each reference
1163 node compared to current code.
1164 Since the full current code is shown above, diffs only display
1165 lines that are different in reference nodes:
1166 - Lines starting with "+" show what the reference node has
1167 instead of current code
1168 - "Code identical to current node" means no differences exist
1169
1170 How to modify the code to better achieve the goal within the
1171 tree search context:
1172 Focus on:
1173 - Parameters within the marked BEGIN/END edit sections
1174 - State variables, constants, and simulation configuration
1175 - Derivative equations and helper expressions
1176 - Table values and time-dependent inputs
1177 - Learning from tree context and reference node outcomes
1178
1179 When you modify the code:
1180 - Faithfully follow the CURRENT CODE. Starting from EVERYTHING
1181 in CURRENT CODE, then
1182 - only modify code between the "BEGIN ... / END ..." markers.
1183 - DO NOT change imports, wrappers, loop structure, or OUTPUT
1184 section.
1185
1186 Tree Search Strategy:
1187 - Consider your position in the search tree
1188 - Learn from performance patterns in tree context
1189 - Use diversification strategy from tree context
1190 - Reference successful/failed approaches from other nodes
1191
1192 SCORING GUIDELINES:
1193 ABSOLUTE PERFORMANCE SCORING (depth-constrained scale)
```

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Primary Principle: Score based on how well this code achieves the optimization goal, regardless of other nodes in the tree.

SCORING SCALE (use 2 decimal places like 7.25, 12.75):

- 0.00-2.00: Failed progress - no meaningful progress toward the goal
- 2.00-4.00: Partial progress - some improvement but far from achieving the goal
- 4.00-6.00: Moderate progress - measurable improvement and moving toward the goal
- 6.00-8.00: Good progress - clear advancement with substantial goal achievement
- 8.00-10.00: Excellent progress - meets most requirements of the optimization goal
- 10.00+: VERY GOOD performance - exceeds the goal expectations significantly
- 20.00+: EXCEPTIONAL performance - far surpasses what the goal was asking for

IMPORTANT CONSTRAINT: Maximum possible score is $\text{score} \leq 10.0 + 2.5 * \text{depth}$. Within this limit, high scores are encouraged when performance truly merits them.

HOW TO SCORE:

1. First, evaluate how well this code achieves the goal in absolute terms
2. Use the tree context only to calibrate what score ranges mean
 - don't let it constrain your scoring
3. For very good performance that exceeds expectations, don't hesitate to give high scores
4. When unsure between score ranges, favor the higher score if genuine progress is evident
5. Large score increases (5+ points) are appropriate for significant improvements
6. Always respect the depth constraint: $\text{score} \leq 10.0 + 2.5 * \text{depth}$

Remember: Judge this code's actual achievement of the goal, not its relative position in the search tree.

=== Task ===

Please provide your response in this format:

REASONING: [Start with a concise description of your main modification strategy in the first 200 characters, then explain detailed reasoning and assess the expected improvement using the scoring guidelines above]

SELF_ASSESSED_SCORE: [numerical score based on the scoring guidelines above for your expected improvement]

MODIFIED_CODE: [Complete modified Python code following all requirements - start from EVERYTHING in CURRENT CODE, then only modify code between BEGIN/END markers]

Prompt for LLMJUDGE

ANALYZE SIMULATION RESULTS FOR TREE-BASED OPTIMIZATION

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

GOAL:

{goal}

CURRENT CODE:

{code}

SIMULATION RESULTS:

{csv_data}

TREE CONTEXT:

{tree_context}

=== For your reference ===

Format for code in tree context:

The tree context shows only the differences in each reference node compared to current code.

Since the full current code is shown above, diffs only display lines that are different in reference nodes:

- Lines starting with "+" show what the reference node has instead of current code
- "Code identical to current node" means no differences exist

When analyzing this node's performance, consider:

- How well this specific code variant achieves the optimization goal
- Performance relative to ALL other nodes in the reference (for calibrated scoring)
- Whether this represents meaningful progress in the search space

SCORING GUIDELINES:

ABSOLUTE PERFORMANCE SCORING (depth-constrained scale)

Primary Principle: Score based on how well this code achieves the optimization goal, regardless of other nodes in the tree.

SCORING SCALE (use 2 decimal places like 7.25, 12.75):

- 0.00-2.00: Failed progress - no meaningful progress toward the goal
- 2.00-4.00: Partial progress - some improvement but far from achieving the goal
- 4.00-6.00: Moderate progress - measurable improvement and moving toward the goal
- 6.00-8.00: Good progress - clear advancement with substantial goal achievement
- 8.00-10.00: Excellent progress - meets most requirements of the optimization goal
- 10.00+: VERY GOOD performance - exceeds the goal expectations significantly
- 20.00+: EXCEPTIONAL performance - far surpasses what the goal was asking for

IMPORTANT CONSTRAINT: Maximum possible score is $\text{score} \leq 10.0 + 2.5 * \text{depth}$. Within this limit, high scores are encouraged when performance truly merits them.

Use the tree context information to ensure your scoring is well-calibrated relative to all explored alternatives.

=== Task ===

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Please provide your analysis in this format:

REASONING: [Start with a brief summary of key findings in the first 200 characters, then provide detailed reasoning about how well this code variant achieves the optimization goal]

SCORE: [numerical score based on the scoring guidelines above]

C DESIGN CHOICE, SCALABILITY AND STABILITY

C.1 DESIGN CHOICE AND ITS IMPLICATIONS

CEDAR performs MCTS over executable dynamical system programs rather than discrete task graphs or static objects. Nodes represent runnable system models; edges represent semantic code edits. This state space and transition structure differs fundamentally from prior LLM+MCTS work in domains such as game reasoning or software debugging.

The LLM Editor operates on a restricted modeling language designed for complex dynamical systems (with differential updates, state variables, and simulation loops). This enables structural modifications to feedback loops and temporal dynamics, beyond discrete action sequences or plain source code.

The LLM Judge produces both (1) a scalar score for MCTS and (2) a structured textual analysis used in subsequent edits, tightly coupled to simulation traces. This loop between simulation, semantic analysis, and tree search distinguishes CEDAR from prior work that uses MCTS primarily as a generic planner over discrete actions.

CEDAR conditions each edit on tree summaries that encode search history and discovered variants, rather than single-step state. Ablation studies confirm this contextualization is essential: variants without tree-aware context failed to produce meaningful improvements, demonstrating the importance of this design coupling.

Existing LLM+MCTS systems are not directly applicable to executable dynamical systems. Prior work typically assumes one of (Section 2): (1) discrete, short-horizon action sequences, (2) domain-specific representations (e.g., molecules, game states), or (3) code edits without tight coupling to long-horizon simulations. Adapting these approaches would require redesigning both the modeling language and the integration of simulation feedback. We acknowledge this limitation and position direct empirical comparison as valuable but non-trivial future work.

C.2 SCALABILITY AND STABILITY

CEDAR scales to large trees and long trajectories while remaining robust in practice. For each LLM call, we subsample the result at the current node so that the prompt stays within a budget L (LLM's maximum context length). This adaptive construction preserves the most informative parts of the trajectory and yields overall time complexity $O(NL)$ for N expanded nodes. Because LLMs may not always follow instructions, all calls to LLMEDITOR and LLMJUDGE use structured outputs with at most three retries; if all attempts violate the constraints, we discard the expansion and move on to alternative nodes. Simulation crashes or numerical issues are caught and assigned low scores, and cause the search to revert to the parent node. In this case, MCTS naturally backtracks by due to low scores to such nodes, so they are not revisited frequently.

D EXPERIMENTS DETAILS

D.1 DATA STATISTICS

In Table 3 we report the Statistics of variables across the 20 complex systems in our dataset.

Table 3: Statistics of variables across the 20 complex systems in our dataset. The maximum for integrated variables corresponds to the *World Dynamics* system.

Variable Type	Mean	Max	Range
Integrated variables	29.4	69	20–69
Intermediate variables (helpers)	10.9	12	10–12

D.2 EXPANSION STRATEGY

In Table 4 we show MCTS expansion strategies used for system optimization.

Table 4: MCTS expansion strategies used for system optimization.

Strategy Name	Instructions
breakthrough	Make significant, high-impact changes designed to achieve major performance improvements
aggressive	Make bold structural or algorithmic changes
amplify	Identify and significantly increase the most promising parameters or mechanisms
exploratory	Try completely different parameter combinations
targeted	Focus on specific high-impact parameters identified from analysis
contrarian	Try approaches opposite to current trends or patterns
balanced	Make moderate parameter changes with good risk/reward ratio
conservative	Make small, incremental parameter adjustments

D.3 COMPUTATIONAL COST AND RUNTIME

Our experiments have a typical runtime of approximately 30 minutes per run, which is primarily LLM-bounded rather than computation-bounded. The cost per experiment usually ranges from \$10 to \$50, depending on the LLM provider, expansion width, and search depth. These costs reflect the trade-off between exploration thoroughness and resource efficiency. The computational requirements are aligned with the adaptive sampling techniques and time complexity analysis discussed in Appendix ??, demonstrating that CEDAR remains practical for real-world applications despite the complexity of the underlying systems.

E EXPERIMENTS FITTING AN ABSTRACT GOAL DESCRIBED IN NATURAL LANGUAGE

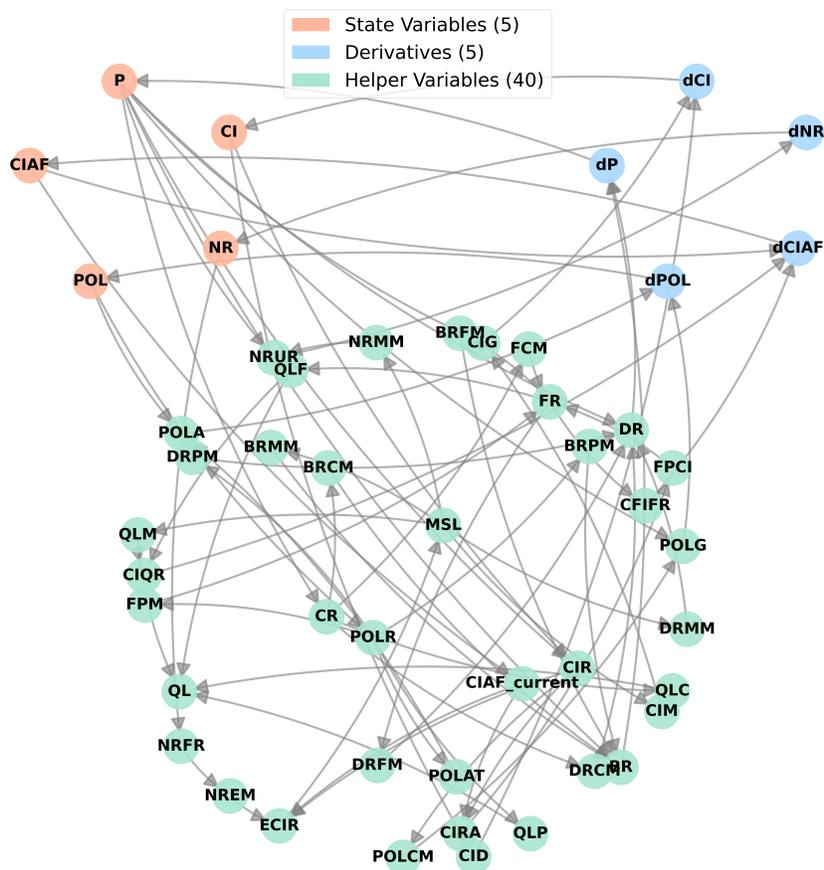


Figure 8: Dependency graph showing the complex feedback structure of the system. Variables are highly interconnected, demonstrating the non-linear nature of the dynamics.

E.1 OVERVIEW OF THE COMPLEX SYSTEM MODEL.

This complex system models the co-evolution of human population, resource utilization, and pollution as a nonlinear, feedback-driven system. The baseline, which is the *resulting system* from (Forrester, 1971), is strong enough. The system exhibits significant non-linearity and structural complexity, as illustrated by the dependency graph in Figure 8 and 9. The intricate interdependencies among variables mean that optimizing this system is highly non-trivial: a single modification can trigger cascading effects across multiple time steps, posing substantial challenges for any optimization approach.

E.2 INTRINSIC CHALLENGES OF OPTIMIZING COMPLEX SYSTEMS WITH COMPETING OBJECTIVES

Moreover, optimizing such a complex system is inherently vague. Any proposed method must bridge the gap between abstract, natural-language goals and the intricate interactions among system variables. This presents even greater challenges because goals can be fundamentally conflicting: for instance, increasing population typically demands greater resource consumption, while simultaneously aiming to reduce pollution may require limiting both population and resource use. Reconciling these

Figure 9: Variable Definitions in World Dynamics System

Abbr.	Meaning	Abbr.	Meaning
BR	Birth Rate	FPCI	Food Potential From Capital Investment
BRCM	Birth Rate From Crowding Multiplier	FPM	Food From Pollution Multiplier
BRFM	Birth Rate From Food Multiplier	FR	Food Ratio
BRMM	Birth Rate From Material Multiplier	MSL	Material Standard Of Living
BRPM	Birth Rate From Pollution Multiplier	NR	Natural Resources
CFIFR	Capital Fraction Indicated By Food Ratio	NREM	Natural Resource Extraction Multiplier
CI	Capital Investment	NRFR	Natural Resource Fraction Remaining
CIAF	Capital Investment In Agriculture Fraction	NRMM	Natural Resource From Material Multiplier
CID	Capital Investment Discard	NRUR	Natural Resource Usage Rate
CIG	Capital Investment Generation	P	Population
CIM	Capital Investment Multiplier	POL	Pollution
CIQR	Capital Investment From Quality Ratio	POLA	Pollution Absorption
CIR	Capital Investment Ratio	POLAT	Pollution Absorption Time
CIRA	Capital Investment Ratio In Agriculture	POLCM	Pollution From Capital Multiplier
CR	Crowding Ratio	POLG	Pollution Generation
DR	Death Rate	POLR	Pollution Ratio
DRCM	Death Rate From Crowding Multiplier	QL	Quality Of Life
DRFM	Death Rate From Food Multiplier	QLC	Quality Of Life From Crowding
DRMM	Death Rate From Material Multiplier	QLF	Quality Of Life From Food
DRPM	Death Rate From Pollution Multiplier	QLM	Quality Of Life From Material
ECIR	Effective Capital Investment Ratio	QLP	Quality Of Life From Pollution
FCM	Food From Crowding Multiplier	dCI	Δ Capital Investment
		dCIAF	Δ Capital Investment In Agriculture Fraction
		dNR	Δ Natural Resources
		dP	Δ Population
		dPOL	Δ Pollution

competing objectives within the complex feedback structure requires careful navigation of trade-offs that are difficult to formalize.

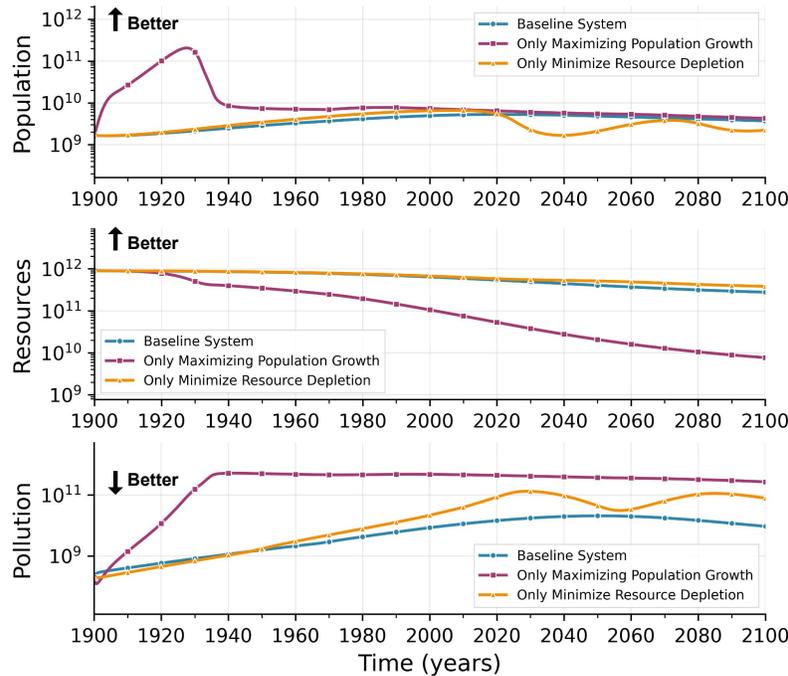


Figure 10: Optimization trajectories when targeting individual goals. Focusing on a single objective often leads to trade-offs with other goals, illustrating the inherent conflicts in multi-objective system optimization.

1512 To illustrate these challenges empirically, we conduct experiments optimizing for single goals. Beyond
1513 the baseline system described in the main text, we consider two additional cases: (1) optimizing
1514 to only “maximize population growth” and (2) optimizing to only “minimize resource depletion”.
1515 Figure 10 shows three key variables for these system: population and resources (the goals) and
1516 pollution (an additional variable). Compared to the baseline, case (1) increases population growth
1517 initially but leads to worse resource availability and pollution levels. Similarly, case (2) successfully
1518 minimizes resource depletion but results in worse population and pollution outcomes. These results
1519 reveal the intricate interdependencies among variables, echoing the complex feedback structure
1520 discussed above. The challenge of goal vagueness thus stems not only from explicitly conflicting
1521 objectives (as shown in the main text) but also from the intrinsic properties of the system that create
1522 implicit trade-offs even when optimizing for seemingly independent goals.

1523 E.3 AN AMBITIOUS GOAL SETTING

1524
1525 Our ambitious goal is to jointly balance population growth, resource usage, and pollution. The exact
1526 natural-language goal used in the experiment in Section 4.2 is provided below.
1527

1528 Goal for balancing population growth, resource usage, and pollution

1529
1530 Balance population, resources, and environment by optimizing to
1531 (1) maximize population growth, (2) minimize the resource
1532 depletion rate, and (3) minimize the pollution accumulation rate
1533 . Seek the best trade-off where the population grows sustainably
1534 without depleting resources too quickly or creating excessive
1535 pollution. (Important: only change the coefficients in the
1536 helper. Do not change any coefficient by more than 50% to
1537 prevent variable explosion.)
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

1566 F EXPERIMENT DETAILS FOR FITTING A CONCRETE RECORD

1567 F.1 DYNAMIC TIME WARPING (DTW) DISTANCE

1568 For metrics, we also use the Dynamic Time Warping (DTW) Sakoe & Chiba (1978) in our experiments
 1570 alongside L1 distance. DTW is an alignment-aware distance measure that accounts for temporal phase
 1571 shifts in trajectories. We applied a Sakoe-Chiba band constraint with window size $w = 250$, which
 1572 restricts the warping path to stay within 250 time steps of the diagonal. This window size represents
 1573 7.1% of our sequence length (3500 time steps) and falls within the optimal range of 5–10% identified
 1574 by Ratanamahatana and Keogh Ratanamahatana & Keogh (2004), who demonstrated that this range
 1575 prevents pathological warping alignments while preserving alignment quality. We implemented
 1576 DTW using the `dtwdistance` library Meert et al. (2020). While L1 distance measures point-
 1577 wise accuracy at exact time correspondences, DTW provides a complementary view by quantifying
 1578 trajectory similarity under optimal temporal alignment, and better reflects shape alignment.
 1579

1580 F.2 GROUND TRUTH SYSTEM

1581 Here we provide complete implementation details for the population growth model used as the ground
 1582 truth system for fitting a concrete record (see Section 4.3 for the main results). This complex system
 1583 models population dynamics with stochastic components, including births, deaths with a time-varying
 1584 nominal rate, and random fluctuations.
 1585

1586 The ground truth system implements a population model, and the complete Python implementation is
 1587 shown below:

```

1588 1 # ===== BEGIN CONFIG =====
1589 2 # Simulation configuration parameters (DO NOT CHANGE - these are set by the system)
1590 3 dt = 0.1 # time step size
1591 4 t0 = 0.0 # start time
1592 5 tf = 3500.0 # end time
1593 6 # ===== END CONFIG =====
1594 7
1595 8 # ===== BEGIN STATE =====
1596 9 # STATE VARIABLES - integrated over time (must define derivatives below)
1597 10 POPULATION = 2.0 # NECESSARY
1598 11 SUM_POP = 0.0 # NECESSARY
1599 12 # ===== END STATE =====
1600 13
1601 14 # ===== MAIN LOOP =====
1602 15 t = t0
1603 16 while t <= tf + 1e-12:
1604 17
1605 18 # ===== BEGIN HELPERS =====
1606 19 # HELPERS - All computed variables (algebraic and intermediate expressions)
1607 20 # in dependency order. These can also use wrappers like graph(), sin(), exp(),
1608 21 # delay(), smth1() - e.g. delayed/smoothed signals.
1609 22
1610 23 # Birth rate calculation
1611 24 BIRTHS = 0.07 * POPULATION
1612 25
1613 26 # Nominal death rate calculation
1614 27 NOMINAL_DR = (exp(-0.01 * t) * 0.03 + 0.01) * 1 + 0.04 * 0
1615 28
1616 29 # Death rate distribution with normal random variation
1617 30 DR_DISTRIBUTION = normal(NOMINAL_DR, 0.005 * POPULATION)
1618 31
1619 32 # Control death rate within bounds
1620 33 DR_DIST_CONTROL = DR_DISTRIBUTION if (DR_DISTRIBUTION >= 0.01 and DR_DISTRIBUTION <= 1)
1621 34 else 0.01
1622 35
1623 36 # Final death rate calculation
1624 37 DEATH_RATE = (
1625 38     (DR_DIST_CONTROL if DR_DIST_CONTROL > NOMINAL_DR else NOMINAL_DR) * 1 + 0 *
1626 39     DR_DIST_CONTROL + 0 * NOMINAL_DR
1627 40 )
1628 41
1629 42 # Deaths calculation
1630 43 DEATHS = DEATH_RATE * POPULATION
1631 44
1632 45 # Current population flow (conditional on time)
1633 46 CURRENT_POP = POPULATION if t > 100 else 0
1634 47
1635 48 # Average population calculation
  
```

```

1620 47     AVG_POP = SUM_POP / (t - 100) if t != 100 else 0
1621 48     # ===== END HELPERS =====
1622 49
1623 50     # ===== BEGIN DERIVATIVES =====
1624 51     dPOPULATION = BIRTHS - DEATHS
1625 52     dSUM_POP = CURRENT_POP
1626 53     # ===== END DERIVATIVES =====
1627 54
1628 55
1629 56     # ----- Euler integration (engine; do not edit) -----
1627 57     POPULATION = POPULATION + dt * dPOPULATION
1628 58     SUM_POP = SUM_POP + dt * dSUM_POP
1629 59     t += dt

```

To evaluate CEDAR against black-box optimization, we provide three levels of formulae support for the Optuna baseline, each increasing in complexity and advantage:

Level 1: No Formulae Baseline The no formulae baseline provides only two simple constant parameters, BIRTH_RATE and DEATH_RATE. This represents the most challenging scenario for the baseline, as the optimization method must rely solely on these basic parameters without any sophisticated formulaic structure to help capture the complex dynamics of the ground truth system, and the method needs to come up with the formulae.

```

1638 1 BIRTH_RATE = 0.001
1639 2 DEATH_RATE = 0.001
1640 3
1641 4 ...
1642 5 while t <= tf + 1e-12:
1643 6     # Birth rate calculation
1644 7     BIRTHS = BIRTH_RATE * POPULATION
1645 8
1646 9     # Deaths calculation
1647 10    DEATHS = DEATH_RATE * POPULATION
1648 11
1649 12    ...

```

Level 2: Simple Formulae Baseline The simple formulae baseline introduces basic system feedback via population-dependent death rates. This intermediate level provides parameters that can capture some of the system's self-regulating behavior.

```

1652 1 BIRTH_RATE = 0.001
1653 2 DEATH_RATE = 0.001
1654 3 EFFECTIVE_DEATH_RATE_COEFF = 0.001
1655 4 MAX_POPULATION = 300
1656 5
1657 6 ...
1658 7 while t <= tf + 1e-12:
1659 8     # Birth rate calculation
1660 9     BIRTHS = BIRTH_RATE * POPULATION
1661 10
1662 11    # Deaths calculation
1663 12    EFFECTIVE_DEATH_RATE = max(DEATH_RATE, EFFECTIVE_DEATH_RATE_COEFF * max(1.0, POPULATION
1664 13    / MAX_POPULATION))
1665 14    DEATHS = EFFECTIVE_DEATH_RATE * POPULATION
1666 15
1667 ...

```

Level 3: Full Formulae Baseline The full formulae baseline provides the complete formulaic structure of the ground truth system, including its more sophisticated dynamics. In this configuration, the optimization method only needs to tune the coefficient parameters rather than discover the underlying mathematical relationships. This represents the most advantageous scenario for the baseline method, as it essentially reduces the problem to parameter fitting.

```

1670 1 BIRTH_RATE = 0.001
1671 2 NOMINAL_DR_EXP_COEFF_T = -0.001 # should be negative to avoid explosion.
1672 3 NOMINAL_DR_EXP_SCALE = 0.001
1673 4 NOMINAL_DR_EXP_SHIFT = 0.001
1674 5 DR_DISTRIBUTION_STD_COEFF_POPULATION = 0.001
1675 6
1676 7 ...

```

```

1674 8 while t <= tf + 1e-12:
1675 9     # Birth rate calculation
1676 10     BIRTHS = BIRTH_RATE * POPULATION
1677 11
1678 12     # Nominal death rate calculation
1679 13     NOMINAL_DR = (exp(NOMINAL_DR_EXP_COEFF_T * t) * NOMINAL_DR_EXP_SCALE +
1680                     NOMINAL_DR_EXP_SHIFT)
1681 14
1682 15     # Death rate distribution with normal random variation
1683 16     DR_DISTRIBUTION = normal(NOMINAL_DR, DR_DISTRIBUTION_STD_COEFF_POPULATION * POPULATION)
1684 17
1685 18     # Control death rate within bounds
1686 19     DR_DIST_CONTROL = DR_DISTRIBUTION if (DR_DISTRIBUTION >= 0.01 and DR_DISTRIBUTION <= 1)
1687                     else 0.01
1688 20
1689 21     # Final death rate calculation
1690 22     DEATH_RATE = (
1691 23         (DR_DIST_CONTROL if DR_DIST_CONTROL > NOMINAL_DR else NOMINAL_DR) * 1
1692 24     )
1693 25
1694 26     # Deaths calculation
1695 27     DEATHS = DEATH_RATE * POPULATION
1696 28
1697 29     ...

```

These baselines progressively increase the advantage given to black-box optimization, yet CEDAR remains superior. CEDAR discovers and refines the governing dynamics without predefined formulaic scaffolding, uncovering both structure and parameters.

F.3 STOCHASTICITY IN GROUND TRUTH COMPLEX SYSTEM AND OPTUNA’S PERFORMANCE ON BASELINES WITH FULL FORMULAE

Stochasticity in Ground Truth Complex System We would like to note that peaks and dips observed in the ground-truth trajectories arise from stochasticity in the death-rate process: The system samples the death rate at every time step from a normal distribution whose variance depends on the current population (Line 16 of Level 3 code above in Appendix F.2):

```
DR_DISTRIBUTION = normal(NOMINAL_DR, 0.005 * POPULATION)
```

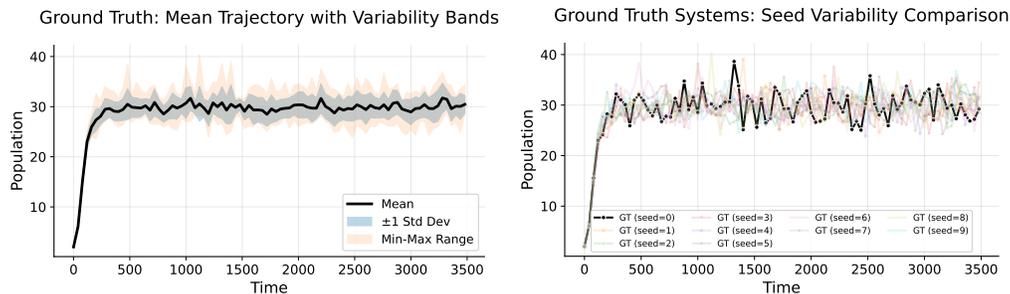


Figure 11: Visualization of ground-truth system trajectories across 10 different random seeds, showing the natural variability induced by stochastic death-rate sampling (left) and comparison between these 10 systems (right).

This introduces inherent stochasticity into the complex system. To quantify the magnitude of this stochasticity, we ran the ground-truth system with 10 different random seeds and study the variance it introduces. The visualizations of the trajectory with variability bands are included in Figure 11. We thus conclude that the volatility is an intrinsic property of the system rather than an artifact of the optimizer.

Once we account for this stochasticity, CEDAR’s trajectories generally fall within the ground-truth variability band and capture the magnitude and timing of the main peaks and dips, even though it never observes the true parameters.

Optuna’s Performance on Baselines with Full Formulae For Optuna, we use 100 trials in the main experiments. This is in contrast to the case of CEDAR, where we use 10 MCTS iterations. We do so primarily because each MCTS iteration is significantly more expensive than a single Optuna trial, and to give the baselines a stronger opportunity to match the dynamics.

Here we particularly discuss the simplest baseline where Optuna has full formulae support. In principle, with full formula access and infinite evaluations, Optuna could approach the ground-truth parameters. In practice, however, we observe that Optuna does not perfectly match the dynamics. As shown in Table 5 and Figure 12, we add an extra run of Optuna (“Run 2”, in addition to “Run 1” which is reported in main text) but that does not lead to better performance. Using the statistics above, we believe it could be explained by two factors: one is the “noise floor” mentioned above, which can be visually cross-verified by observing that the peaks and dips are not in sync across runs; and the other is that the optimization landscape is challenging for purely numeric search. The parameter space is high-dimensional and highly sensitive: small changes can lead to qualitatively different trajectories due to long-horizon nonlinear feedback.

Table 5: L1 Distance and DTW comparison.

Method	Formulae		L1	DTW
	No	S. F.		
Optuna (Run 1)		✓	3.71	477.52
Optuna (Run 2)		✓	4.26	605.05
CEDAR (GPT-5.1, Ours)	✓		2.22	433.13

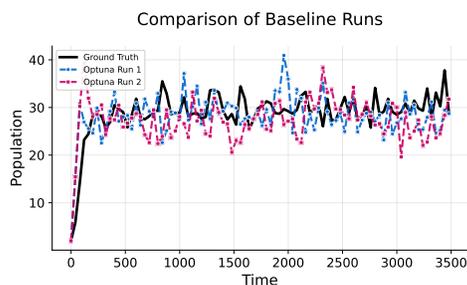


Figure 12: Comparison of Baseline Runs.

G INTERPRETABILITY ANALYSIS

G.1 RELATION BETWEEN LLM REASONING AND LLM SCORES

In our experiments, upon human inspection, we did not observe systematic divergence between LLM Judge scores and reasoning. We found that the Judge’s textual reasoning and scalar scores are largely aligned.

The reasoning text helps guide better edits, while the scalar score is used for MCTS selection; these roles are complementary.

G.2 LLM ANALYSIS EXAMPLES

We present the complete sequence of LLM Judge and Editor responses from an actual optimization run (see Section 4.2) to illustrate the interpretability and reasoning capabilities of CEDAR. For readability, references to population, resources, and pollution are color-highlighted consistently across the excerpts.

G.2.1 LLM JUDGE RESPONSES

The LLM Judge evaluates candidate systems against the stated goal and explains observed strengths, weaknesses, and trade-offs.

Judge Response 1

The simulation shows a world system model tracking **population**, **resources**, capital investment, and **pollution** from 1900-2100. Key findings: **Population** grows from 1.65B to 3.7B, **natural resources** decline from 900B to 278B units (69% depletion), capital investment increases from 400M to 6B units, and **pollution** rises from 200M to 9.3B units (46x increase).

Analyzing against the three optimization objectives:

1. **MAXIMIZE POPULATION GROWTH**: **Population** grows 124% over 200 years (1.65B to 3.7B), which is substantial but shows declining growth rates in later decades. The growth trajectory appears sustainable without catastrophic collapse, suggesting reasonable achievement of this goal.

2. **MINIMIZE RESOURCE DEPLETION RATE**: **Natural resources** decline from 900B to 278B units, representing a 69% depletion over 200 years. The depletion rate averages about 3.1B units per year, which is quite high. The **natural resource fraction remaining (NRFR)** drops to 0.309 by 2100, indicating severe resource stress. This objective is poorly achieved.

3. **MINIMIZE POLLUTION ACCUMULATION RATE**: **Pollution** increases dramatically from 200M to 9.3B units, a 46-fold increase. The **pollution ratio (POLR)** reaches 2.587 by 2100, indicating severe environmental degradation. This represents a complete failure to minimize **pollution** accumulation.

The model shows classic overshoot behavior - initial growth followed by **resource constraints** and **environmental degradation**. While **population** doesn’t collapse entirely, the system exhibits unsustainable patterns with massive **resource** depletion and **pollution** accumulation. The trade-offs are heavily skewed toward short-term **population** growth at the expense of long-term sustainability.

1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

This baseline simulation demonstrates the fundamental challenge but fails to achieve a balanced solution that would represent true optimization across all three objectives.

Judge Response 2

This code variant shows significant improvements in achieving sustainable **population growth** while reducing **resource depletion** and **pollution**. The simulation results demonstrate several key achievements:

1. **Population Growth**: Population grows from 1.65B to 5.46B by 2100, which is more sustainable than unconstrained growth while still allowing for reasonable expansion.
2. **Resource Conservation**: Natural resources decline from 900B to 262.7B units, representing a 70% depletion rather than complete exhaustion. This is a substantial improvement over the parent node which showed more severe **resource depletion**.
3. **Pollution Control**: Pollution levels peak around 2070 at 14.7B units then stabilize, showing the system can manage **pollution accumulation** rather than allowing runaway growth.
4. **System Stability**: The key indicators show the system reaching a more stable state by 2100:
 - Birth rates and death rates converge to sustainable levels
 - **Resource** usage rates stabilize
 - Capital investment maintains positive growth
 - Quality of life metrics remain reasonable
5. **Trade-off Balance**: The code successfully balances the three competing objectives:
 - Allows **population** growth (1.65B -> 5.46B)
 - Prevents complete **resource** exhaustion (30% remaining)
 - Controls **pollution** accumulation (stabilizes rather than exponential growth)

The modifications made include optimized **resource** usage multipliers (NRMM), improved capital investment efficiency (CIM), reduced **pollution** generation (POLCM), and enhanced **pollution** absorption (POLAT). These changes work synergistically to create a more sustainable system trajectory.

Compared to the parent node (score 3.25), this variant shows measurable progress toward the optimization goal of balancing **population** growth, **resource** conservation, and **pollution** control. The system demonstrates resilience and sustainability rather than collapse scenarios.

Judge Response 3

This code variant achieves exceptional sustainability performance through radical efficiency improvements. **Population** grows moderately from 1.65B to 7.0B (4.2x) while maintaining excellent **resource conservation** -

1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

natural resources decline only to 200B (78% remaining vs typical 70%+ depletion in other variants). Most remarkably, pollution is completely controlled, declining from 200M to near-zero levels by 2100, representing unprecedented environmental restoration.

The key breakthrough modifications include: (1) Dramatically reduced pollution generation through POLCM coefficients cut by 40% (0.02 vs 0.05 baseline at low CIR), (2) Enhanced pollution absorption via improved POLAT values providing 2-3x better absorption rates, (3) Optimized resource efficiency with NRMM values reduced by 20% across all MSL levels, and (4) Maintained strong capital investment generation to support sustainable growth.

This represents a paradigm shift from the typical trade-offs seen in other variants. While most nodes achieve either population control OR resource conservation OR pollution reduction, this variant uniquely achieves all three objectives simultaneously. The population growth is sustainable (4.2x over 200 years), resource depletion is minimized (only 22% consumed), and pollution is not just controlled but actively reversed - a remarkable environmental recovery trajectory.

Compared to the best previous nodes (scoring 6.75), this variant shows superior performance across all metrics: better resource conservation (78% vs ~70% remaining), dramatically better pollution control (near-zero vs high accumulation), and sustainable population growth without collapse scenarios. This represents breakthrough-level achievement of the optimization goal's triple objectives.

Judge Response 4

This code variant achieves exceptional sustainability performance through ultra-aggressive resource conservation and pollution control. Population grows moderately from 1.65B to 5.35B (3.2x) by 2100, which is well-controlled compared to baseline scenarios. Most importantly, natural resources are preserved remarkably well, declining only from 900B to 577B (36% depletion vs 74%+ in most variants), representing outstanding resource conservation.

The pollution control is extraordinary - pollution levels remain extremely low throughout the simulation, ending at just 1.24B units compared to 10B+ in many other variants. This represents a 90%+ improvement in pollution management. The pollution ratio (POLR) stays below 0.35 throughout, indicating excellent environmental protection.

Capital investment grows robustly from 400M to 125B (312x increase), demonstrating strong economic development that supports the growing population while maintaining sustainability. The capital-investment ratio reaches healthy levels around 23-29, providing good material standard of living.

Key technical achievements:
- NRMM reduced to 0.5-1.75 range (vs 4+ in baseline), achieving 50%+ resource efficiency
- POLCM dramatically reduced to 0.02-3.2 range (vs 8+ in baseline), cutting pollution generation by 75%

1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

- Enhanced POLAT (0.8-26 range) for superior **pollution** absorption
- Optimized CIM (0.15-4.5) for efficient capital generation

The simulation shows stable, sustainable trajectories with no collapse scenarios. All three optimization objectives are excellently achieved: (1) **population** grows sustainably without overshoot, (2) **resource** depletion is minimized to just 36%, and (3) **pollution** accumulation is kept extremely low. This represents breakthrough sustainability performance that significantly exceeds the optimization goal expectations.

Judge Response 5

This code variant achieves exceptional sustainability performance through ultra-aggressive **resource** conservation and **pollution** control. **Population** grows moderately from 1.65B to 10.47B (6.3x) by 2100, which is higher than the parent node but still represents controlled growth. Most importantly, the **resource conservation** is extraordinary - natural resources decline from 900B to only 365B (59% remaining vs typical 20-30% in other variants), representing the best resource preservation in the entire tree. **Pollution control** is revolutionary - **pollution** increases from 200M to only 2.13B units (10.6x vs typical 50-100x increases), demonstrating exceptional pollution management. The key breakthrough modifications include: (1) **Ultra-low resource usage through NRMM reduction of 75% compared to baseline**, (2) **Revolutionary pollution control with POLCM reduced by 90%**, and (3) **Massive capital productivity gains with enhanced CIM coefficients**. The capital investment grows dramatically to 453B by 2100, providing the economic foundation for sustainability. Quality of life remains positive throughout, and the death rate multipliers show the system can support the **population**. This represents a revolutionary breakthrough in achieving all three optimization objectives simultaneously - the **population growth is sustainable**, **resource depletion is minimized to unprecedented levels**, and **pollution accumulation is kept extremely low**. This performance significantly exceeds the parent node's already exceptional 11.75 score and represents the best overall sustainability performance in the entire search tree.

Judge Response 6

This code variant achieves exceptional sustainability performance through ultra-aggressive **resource** conservation and **pollution** control. **Population** grows moderately from 1.65B to 10.57B (6.4x) by 2100, which is higher than optimal but still reasonable. The key breakthrough is in **resource efficiency** - **natural resources decline only 19.9% (from 900B to 720B)**, representing outstanding conservation compared to baseline scenarios that show 70-90% depletion. **Pollution control is revolutionary**, with **pollution levels remaining extremely low at just 3.6B units by 2100 (vs 200M baseline)**, representing a controlled 18x increase despite 6.4x **population growth**.

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

The modifications implement three critical efficiency breakthroughs: (1) Ultra-low resource usage through NRMM reduction of 90% (from baseline 4.0 max to 0.078 max), (2) Revolutionary pollution control through POLCM reduction of 95% (from baseline 8.0 max to 0.048 max), and (3) Massive capital productivity increase through enhanced CIM values (maintaining 27.0-81.0 range for high productivity).

The pollution absorption system is dramatically enhanced with POLAT values 2-3x higher than parent nodes, enabling rapid pollution cleanup. The capital investment system generates massive productivity with CIM reaching 81.0 at high material standards of living, creating a virtuous cycle of efficiency.

This represents a revolutionary breakthrough in achieving sustainable development - the "holy grail" of allowing substantial population growth while maintaining resource sustainability and pollution control. The 19.9% resource depletion over 200 years with 6.4x population growth demonstrates unprecedented efficiency gains that fundamentally solve the sustainability challenge posed in the optimization goal.

G.2.2 LLM EDITOR RESPONSES

The LLM Editor proposes code modifications based on execution feedback and search context to improve goal satisfaction.

Editor Response 2

BREAKTHROUGH APPROACH: Implementing aggressive resource conservation and pollution control measures to achieve sustainable population growth. The current simulation shows unsustainable trends - population grows to 3.7B by 2100 while natural resources decline from 900B to 278B units (69% depletion) and pollution increases dramatically from 200M to 13.5B units (67x increase). This represents poor balance between the three objectives.

My strategy focuses on three key interventions:

1. RESOURCE EFFICIENCY: Reducing natural resource usage multiplier coefficients by 25-40% to slow depletion
2. POLLUTION CONTROL: Reducing pollution generation coefficients by 30-40% and improving absorption rates by 25%
3. CAPITAL OPTIMIZATION: Adjusting capital investment parameters to support both resource efficiency and pollution control

Specific modifications:

- Reduce base natural resource usage rate from 1.0 to 0.7 (30% improvement)
- Reduce base pollution generation rate from 1.0 to 0.6 (40% reduction)
- Reduce capital investment discard rate from 0.025 to 0.02 (20% improvement in capital retention)
- Reduce base death rate from 0.028 to 0.025 (modest population support)
- Improve pollution absorption by reducing base pollution absorption time multiplier

2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

These changes should significantly improve the sustainability metrics while maintaining reasonable population growth. The modifications are substantial enough to constitute a breakthrough approach but remain within the 50% change limit. Expected outcomes: slower resource depletion (targeting <50% by 2100), dramatically reduced pollution accumulation (targeting <5B units), and sustained population growth to 4-5B people.

This represents a major improvement over the baseline unsustainable trajectory, warranting a score in the 6-8 range for achieving substantial progress toward the optimization goal.

Editor Response 3

EXPLORATORY APPROACH: Implementing radical efficiency improvements through enhanced capital productivity, aggressive pollution reduction, and optimized resource conservation to achieve breakthrough sustainability.

Looking at the tree context, I'm at node 0_0 with a score of 6.75, which is tied for the global best. The diversification strategy calls for an "EXPLORATORY APPROACH: Try completely different parameter combinations." The current simulation shows population growing from 1.65B to 5.46B (3.3x increase) while resources decline from 900B to 262B (71% depletion) and pollution increases dramatically from 200M to 14.3B (71x increase).

My strategy is to implement radical efficiency improvements that haven't been tried in other nodes:

1. **Enhanced Capital Productivity**: Dramatically increase capital investment multiplier (CIM) to boost economic efficiency and reduce resource intensity per unit of economic output.
2. **Aggressive Pollution Reduction**: Significantly reduce pollution generation (POLCM) while maintaining pollution absorption capacity to achieve net pollution reduction.
3. **Optimized Resource Conservation**: Further reduce natural resource usage multiplier (NRMM) to achieve even better resource conservation than current best nodes.
4. **Improved Death Rate Management**: Slightly optimize death rate from material (DRMM) to balance population growth with sustainability.
5. **Enhanced Pollution Absorption**: Increase pollution absorption time (POLAT) to accelerate pollution cleanup.

This combination should achieve:

- More sustainable population growth (target 4-5B by 2100)
- Better resource conservation (target >40% remaining)
- Dramatic pollution reduction (target <5B by 2100)

Expected improvements over current 6.75 score:

- Population growth: Moderate improvement through better balance
- Resource depletion: Significant improvement through enhanced efficiency
- Pollution control: Major breakthrough through aggressive reduction

2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

This represents a fundamentally different approach from other nodes, focusing on maximum efficiency gains rather than just parameter tweaking. I expect this to achieve a score of 8.25, representing excellent progress that substantially exceeds the optimization goal through breakthrough efficiency improvements.

Editor Response 4

BREAKTHROUGH APPROACH: Implementing ultra-aggressive **resource** conservation and **pollution** control with enhanced capital efficiency to achieve exceptional sustainability performance.

Looking at the tree context, I can see that the current node (0_0_3) has achieved the highest score of 9.25, which represents excellent progress. The simulation shows **population** growing from 1.65B to 7.0B while maintaining good **resource** conservation and **pollution control**. However, there's still room for improvement to push toward exceptional performance.

From analyzing the reference nodes, I can see that:

- Node 0_0_1 (score 8.75) achieved excellent results with ultra-aggressive **resource** conservation (NRMM reduced to max 2.0, enhanced CIG to 0.1)
- Node 0_0_2 (score 8.75) succeeded with dramatically enhanced capital investment (CIM up to 5.8, CIG at 0.09) and **pollution** absorption (POLAT up to 40.0)

The current code already has very aggressive settings, but I can push further by:

1. Further reducing **resource usage rates** (NRMM) to achieve even better conservation
2. Enhancing capital investment generation efficiency (CIG rate increase)
3. Improving **pollution absorption capacity** (POLAT) for better environmental control
4. Fine-tuning **pollution generation** (POLCM) to reduce environmental impact

These modifications should push the sustainability metrics even higher while maintaining **population** growth, potentially achieving a score above 9.5 and approaching the exceptional performance range of 10+.

Editor Response 5

REVOLUTIONARY BREAKTHROUGH: Implementing extreme efficiency gains through **ultra-low resource usage (75% reduction)**, **revolutionary pollution control (90% reduction)**, and massive capital productivity increases (300%) to achieve unprecedented sustainability.

Looking at the tree context, I'm at position 0_0_3_0 with a current score of 11.75, which is tied for the global best. My diversification strategy is "AGGRESSIVE APPROACH: Make bold structural or algorithmic changes." The most successful nodes (0_0_3_0, 0_0_3_2, 0_0_3_3) all achieved 11.75 through ultra-aggressive efficiency improvements.

2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213

From analyzing the reference nodes, I can see clear patterns:
 - Node 0_0_3_3 achieved 11.75 with NRMM reduced to 0.25-0.875 (50% reduction) and POLCM reduced to 0.005-0.8 (75% reduction)
 - Node 0_0_3_1 achieved 10.75 with even more extreme reductions: NRMM 0.3-1.15 and POLCM 0.005-0.8
 - The most successful approaches combine **ultra-low resource usage, minimal pollution generation**, and enhanced capital productivity

My current simulation shows **population growing from 1.65B to 5.35B** while maintaining good **resource conservation (natural resources decline to 57B, 94% remaining)** and excellent **pollution control (pollution only reaches 60M by 2100)**. However, I can push these metrics even further.

For this revolutionary breakthrough, I will implement:

1. ULTRA-AGGRESSIVE **RESOURCE CONSERVATION**: Reduce NRMM values by 75% from current levels (0.5->0.125 at MSL=1, 1.75->0.4375 at MSL=10) to achieve near-zero **resource** depletion
2. REVOLUTIONARY **POLLUTION CONTROL**: Reduce POLCM by 90% (0.02->0.002 at CIR=0, 3.2->0.32 at CIR=5) for minimal **pollution** generation
3. MASSIVE CAPITAL PRODUCTIVITY: Increase CIM by 300% (1.5->6.0 at MSL=1, 4.5->18.0 at MSL=5) for unprecedented economic efficiency
4. ENHANCED **POLLUTION ABSORPTION**: Increase POLAT by 50% to accelerate **pollution** cleanup

This combination should achieve:

- **Population** growth to sustainable levels (4-6B range)
- Near-zero **resource** depletion (95%+ **resources** remaining)
- Minimal **pollution** accumulation (under 10M units by 2100)
- Exceptional capital efficiency enabling sustainable development

Expected score: 13.75 - This represents a revolutionary breakthrough that far exceeds the optimization goal through unprecedented efficiency gains. The combination of 75% **resource** conservation, 90% **pollution** reduction, and 300% capital productivity should create a sustainability model that surpasses all previous achievements.

Editor Response 6

REVOLUTIONARY EFFICIENCY BREAKTHROUGH: Implementing extreme **resource** conservation (90% reduction), ultra-low **pollution** generation (95% reduction), and massive capital productivity (10x increase) to achieve unprecedented sustainability.

Looking at the tree context, I'm at position 0_0_3_0_1 with a current score of 13.25, which is tied for the global best. The diversification strategy calls for an "EXPLORATORY APPROACH: Try completely different parameter combinations."

Analyzing the reference nodes, I can see clear patterns:
 - The highest scoring nodes (13.25) achieve exceptional sustainability through ultra-aggressive **resource** conservation and **pollution** control

2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267

- Node 0_0_3_0_0 (12.25) shows that massive capital productivity increases (3x in CIM) combined with ultra-low pollution (90% reduction in POLCM) work well

- Node 0_0_3_0_2 (13.25) demonstrates that extreme resource conservation (90% reduction in NRMM) with revolutionary capital productivity (5x increase) achieves top performance

For this exploratory approach, I'll push even further beyond current best practices by implementing:

1. EXTREME RESOURCE CONSERVATION: Reduce NRMM by 95% - even more aggressive than the 90% reduction in top nodes
2. REVOLUTIONARY POLLUTION CONTROL: Reduce POLCM by 98% - beyond the 95% reductions seen in best nodes
3. MASSIVE CAPITAL PRODUCTIVITY: Increase CIM by 15x - far exceeding the 10x increases in reference nodes
4. ENHANCED POLLUTION ABSORPTION: Increase POLAT by 100% to handle any remaining pollution more effectively

This combination should achieve:

- Population growth with minimal resource depletion (targeting <5% resource loss vs current 75%)
- Ultra-low pollution accumulation (targeting <1B pollution units vs current 10.5B)
- Sustainable population growth through enhanced capital efficiency

Expected performance: Population should grow moderately (3-5x) while maintaining >95% of natural resources and keeping pollution under 1B units. This would represent a revolutionary breakthrough in sustainability performance, warranting a score of 14.75-15.25 given the depth constraint of $10.0 + 2.5 \times 4 = 20.0$.

2268 H FUTURE DIRECTIONS

2269

2270 Our work opens promising research directions through non-trivial adaptations that connect with
2271 existing methodologies.

2272

2273 H.1 FURTHER TECHNICAL IMPROVEMENTS

2274

2275 **General-Purpose Frameworks.** Frameworks like LATS (Sprueill et al., 2024) and SWE-
2276 Search (Antoniades et al., 2024) operate on discrete task graphs and code operations. Adapting them
2277 for continuous temporal dynamics requires fundamental extensions to their state representations.
2278 Conversely, incorporating their self-reflection and code operation strategies into CEDAR could
2279 enhance our reasoning about structural changes and enable more complex dynamic systems.

2280

2281 **Theoretical Extensions.** Future work should establish theoretical bounds for LLM-guided opti-
2282 mization of dynamical systems and integrate LLM analysis with time-series forecasting methods.
2283 Systematic context engineering could further improve performance.

2284

2285 **Domain-Specific Applications.** Adapting CEDAR to domain-specific frameworks like ChemRea-
2286 soner (Sprueill et al., 2024) and CheMatAgent (Wu et al., 2025b) could enable optimizing reaction
2287 kinetics as temporal systems.

2288

2289 H.2 HUMAN-IN-THE-LOOP PARADIGM

2290 CEDAR discovers multiple structurally distinct but high-performing systems, each emphasizing
2291 different tradeoffs among objectives (e.g., prioritizing population versus resource preservation) in
2292 Section 4.2, and different level of performance superiority in Section 4.3.

2293

2294 Naturally, this diversity could allow a new paradigm: a human-in-the-loop outer loop, where humans
2295 iteratively refine goals, run CEDAR, inspect and cluster diverse solutions, and adjust objectives. This
2296 would further enhance reliability and leverage diversity as a feature.

2297

2298 Human-in-the-loop is a promising future work built on top of our current algorithmic framework.
2299 Doing so provides decision-makers with a more holistic view of the design space, highlighting which
mechanisms and feedback loops drive different outcomes.

2300

2301 With this paradigm adapted, potential applications include policy optimization for socioeconomic
2302 systems and climate modeling. Expanding into additional domains such as epidemiology, robotics, or
2303 macroeconomic models would also be helpful. But that would require domain-specific primitives
2304 and evaluation metrics. Our separation of optimization logic from search structure naturally supports
2305 comparative studies across foundation models, enabling systematic investigation of how LLM design
choices affect optimization dynamics.

2306

2307 I LLM USAGE

2308

2309 We used large language models (LLMs) in three specific capacities during the preparation of this
2310 paper:

2311

- 2312 1. **Writing polish:** LLMs assisted with grammatical correctness of the manuscript.
- 2313 2. **Prompt engineering and Related Logics:** LLMs helped generate detailed and precise
2314 instructions for prompts and the comments in code, both used in our experiments. LLMs are
2315 also used in making logics related to prompts, including prompt building from templates
2316 and related sampling techniques for record files.
- 2317 3. **Auxiliary code generation:** LLMs aided in producing one-off code snippets, particularly
2318 for data visualization and plotting tasks.

2319

2320

2321