STRUCTURING COLLECTIVE ACTION WITH LLM-GUIDED EVOLUTION: FROM ILL-STRUCTURED PROBLEMS TO EXECUTABLE HEURISTICS

Anonymous authors

000

001

002

004

006

008 009 010

011 012 013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

033

034

037

038

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Collective action problems, which require aligning individual incentives with collective goals, are classic examples of Ill-Structured Problems (ISPs). For an individual agent, the causal links between local actions and global outcomes are unclear, stakeholder objectives often conflict, and no single, clear algorithm can bridge micro-level choices with macro-level welfare. We present ECHO-MIMIC, a computational framework that converts this global complexity into a tractable, Well-Structured Problem (WSP) for each agent by discovering compact, executable heuristics and persuasive rationales. The framework operates in two stages: ECHO (Evolutionary Crafting of Heuristics from Outcomes) evolves snippets of Python code that encode candidate behavioral policies, while MIMIC (Mechanism Inference & Messaging for Individual-to-Collective Alignment) evolves companion natural language messages that motivate agents to adopt those policies. Both phases employ a large-language-model-driven evolutionary search: the LLM proposes diverse and context-aware code or text variants, while population-level selection retains those that maximize collective performance in a simulated environment. We demonstrate this framework on a canonical ISP in agricultural landscape management, where local farming decisions impact global ecological connectivity. Results show that ECHO-MIMIC discovers high-performing heuristics compared to baselines and crafts tailored messages that successfully align simulated farmer behavior with landscape-level ecological goals. By coupling algorithmic rule discovery with tailored communication, ECHO-MIMIC transforms the cognitive burden of collective action into a simple set of agent-level instructions, making previously ill-structured problems solvable in practice and opening a new path toward scalable, adaptive policy design.

1 Introduction

Many of the most pressing real-world challenges, from sustainable resource management and climate change mitigation to economic policy design, are Ill-Structured Problems (ISPs) (Simon & Newell, 1971; Reitman, 1964). Unlike Well-Structured Problems (WSPs), which have clearly defined goals, known constraints, and a finite set of operators, ISPs feature ambiguous goals, unclear causal relationships, and undefined solution spaces (Simon, 1973). Solving an ISP requires the problem-solver to impose structure, define objectives and discover pathways, as an integral part of the solution process itself.

A classic example of an ISP arises in collective action problems, where locally rational decisions made by autonomous agents lead to globally suboptimal or even harmful outcomes (Hardin, 1968; Ostrom, 1990). Consider farmers operating within a shared agricultural landscape. Each farmer makes land-use decisions driven by local incentives like maximizing crop yield and profitability. While these decisions may be optimal at the farm level, their combined effect can degrade the shared ecosystem upon which all depend, for instance by reducing biodiversity or compromising ecological connectivity (Kremen, 2020; Poteete et al., 2010; Taylor et al., 1993). For an individual farmer, the decision of how to act is an ISP: the link between their specific planting choices and the health of the entire landscape is complex and unclear, and the *right* action is not algorithmically defined.

The challenge for a system designer or policymaker is to create a mechanism that simplifies this decision-making process for the individual. An ideal solution would be simple, practical behavioral rules, or heuristics, that, if followed by individual agents, reliably produce a desirable global outcome. Such heuristics would effectively transform the ISP faced by each agent into tractable WSPs. Discovering such heuristics, however, is a challenging second-order problem.

We introduce the **ECHO-MIMIC** framework, designed to automate the discovery of these heuristics and the mechanisms to encourage their adoption. Our approach is grounded in Simon's models of bounded rationality, which posit that agents rely on simple rules-of-thumb to navigate complex environments (Gigerenzer & Gaissmaier, 2011; Simon, 1990). We operationalize this concept using the synergy of Evolutionary Algorithms (EAs) and Large Language Models (LLMs). This LLM+EA paradigm represents a new frontier for creative problem-solving, and recent work has begun to leverage LLMs within evolutionary program searches to generate and tune heuristics for complex optimization problems, such as optimizing prompts for the LLMs themselves (Guo et al., 2023) and to fuel mathematical, algorithmic, and scientific discovery (Romera-Paredes et al., 2024; Liu et al., 2024; Ye et al., 2024; Novikov et al., 2025; Chen et al., 2023).

However, the utility of this paradigm in practical optimization settings and its applicability to real-world complex systems has been underexplored. Our framework applies this paradigm to solve a real-world collective action problem, an ISP, and transform it into effective WSPs. To achieve this, the LLM acts as a creative engine that proposes executable heuristics and nudges, while the Evolutionary Algorithm (EA) provides the optimization pressure. Our primary contributions are:

- 1. We introduce ECHO-MIMIC, a framework that deconstructs complex collective action ISPs into simple, executable behavioral heuristics that are well-structured for individual agents, and then nudges the agents to implement these heuristics.
- 2. We demonstrate our framework on the important ecological collective action problem of enhancing connectivity in agricultural landscapes, and show that it significantly outperforms generic LLM program-synthesis baselines like DSPy MiPROv2.
- 3. We find that performance of heuristics produced by ECHO rises with code-complexity indicators and that nudges generated by MIMIC can be tailored to farmer personas.
- 4. We show the effectiveness of the LLM+EA paradigm on a chain of spatial optimization problems in a complex system, moving beyond work focusing on combinatorial benchmarks (Liu et al., 2024a; Ye et al., 2024; Dat et al., 2025; Romera-Paredes et al., 2024).

2 RELATED WORK

LLM-guided evolutionary search and automated heuristic design: A growing line of work couples LLMs with evolutionary search to generate programs, prompts, and heuristics. FunSearch demonstrates LLM-driven program discovery within an evolutionary loop for mathematical problems (Romera-Paredes et al., 2024). EvoPrompt connects LLMs with evolutionary algorithms to evolve high-performing prompts (Guo et al., 2023). LLMs have also been used as evolutionary optimizers or operators more broadly (Liu et al., 2024b; Yang et al., 2023; Lange et al., 2024). Beyond prompts, language hyper-heuristics (Burke et al., 2003) evolve executable code to improve search efficiency and generality across combinatorial problems (Ye et al., 2024; Liu et al., 2024a; Dat et al., 2025). Our ECHO phase aligns with this paradigm but specializes it to produce validated code heuristics that map local states to actions to drive collective-action.

Collective action, bounded rationality, and heuristics: The core challenge we target, aligning individual incentives with social welfare, sits squarely within collective action and commons governance. Hardin framed the dynamic as a *tragedy of the commons* (Hardin, 1968), while Ostrom documented institutional conditions under which communities avert that tragedy (Ostrom, 1990). From a cognition viewpoint, our agent-level design follows the bounded-rationality tradition: people use fast, simple heuristics adapted to their environments (Gigerenzer & Gaissmaier, 2011). At the system-level, designing those heuristics transforms an ill-structured problem (Simon, 1973) into well-structured subproblems with explicit objectives and evaluators.

Ecological connectivity and ecological intensification: In working landscapes, biodiversity outcomes hinge on both production decisions and spatial configuration. Ecological intensification em-

phasizes managing on-farm practices and diversification to sustain yields while supporting ecosystem services (Kremen, 2020; Bommarco et al., 2013; Garibaldi et al., 2019). At landscape scale, ecological connectivity is crucial (Taylor et al., 1993; Keeley et al., 2021). Graph-based indices like Integral Index of Connectivity (IIC) (Pascual-Hortal & Saura, 2006) formalize how local interventions aggregate into network-level connectivity. We follow these existing models of ecological intensification and connectivity to derive ground-truth baseline and global behaviors of agents.

Mechanisms, nudges, and AI-personalized messaging: Adoption is often the bottleneck, and even good policies underperform without mechanisms for uptake. Behavioral nudges and choice architecture can shift real-world environmental decisions (Byerly et al., 2018). Recent evidence shows that generative models can craft personalized messages with stronger persuasive effects than generic appeals (Matz et al., 2024; Rogiers et al., 2024). Our MIMIC phase operationalizes this by evolving messages that reliably alter agents' code-level heuristics toward ECHO-derived targets.

3 PROBLEM FORMULATION AND APPROACH

The collective action setting is ill-structured at two coupled levels. At the *agent level*, each agent $i \in \mathcal{N}$ observes a local state $S_{L,i}$, chooses $a_i \in \mathcal{A}_i$, and optimizes a local objective $U_{L,i}$, while the effect of a_i on societal goals depends on unknown and evolving a_{-i} . At the *system level* (policymaker), inferring what agents currently do (baseline behavior), determining how to coordinate local choices so they aggregate into desired global patterns, and how to incentivize behavior under real-world constraints are themselves ill-structured problems. Here, we present this problem and our approach in the context of agricultural collective action, but the formulation and the ECHO-MIMIC pipeline apply broadly to other collective action domains (e.g., decentralized resource management).

In our agricultural setting, $S_{L,i}$ consists of plot-level agro-ecological and economic features, such as crop types, yields, and prices. Actions a_i are farm interventions: (i) margin intervention (length, placement), and (ii) habitat conversion (area, orientation). We focus on these two levers because they are common, farmer-controlled choices in ecological intensification that modulate on-farm services (pollination, pest control) and economic trade-offs with relatively low implementation friction (Kremen, 2020). At the landscape scale, their placement, length, area, and orientation directly reshape the habitat graph (patch areas and potential corridors), making them a compact yet expressive action basis for steering connectivity (Taylor et al., 1993). We set the local objective $U_{L,i}$ to be net present value (NPV) under farm-specific constraints, whereas the global objective U_G prioritizes landscape-scale ecological connectivity, measured by the *Integral Index of Connectivity (IIC)*, an area-weighted graph metric computed on the habitat network (Pascual-Hortal & Saura, 2006).

Let $\mathbf{A}=(a_1,\ldots,a_N)$ and define a nonseparable global objective $U_G(\mathbf{A})=G(\Phi(\mathbf{A}))$, where Φ maps joint interventions to a mesoscale representation, like a habitat graph built from the spatial pattern and directionality of margins/habitats, and G scores that representation using IIC or related connectivity metrics. From any single agent's vantage point, $\partial U_G/\partial a_i$ depends on unknown and evolving a_{-i} and is mediated by thresholds, complementarities, and path dependence in Φ . For instance, an intervention that nearly bridges two patches has negligible effect until a neighboring segment exists (threshold); an intervention's benefit materializes only if downstream segments exist (complementarity); and interventions can align with or sever emergent corridors (directionality). These properties render one-shot mechanism design ill-posed. Therefore, to make this collective action ISP tractable, our approach imposes structure at both the system and agent levels by decomposing the problem into four well-structured stages whose outputs are directly executable (Fig. 1):

Stage 1 - Establish Baseline Behavior: We fix what agents do by default. For each agent i, we compute the baseline action a_i^0 by solving the local problem $\max_a U_{L,i}(a,S_{L,i})$ following Kremen (2020); Dsouza et al. (2025). This yields state-action pairs $D_i^0 = (S_{L,i}, a_i^0)$. Practically, this means solving for (or recording) each farm's profit-maximizing choices, for example, how much margin length and habitat area to allocate given assumed prices, costs, and constraints.

Stage 2 - Learn Baseline Heuristics: We learn executable code heuristics $\hat{H}_{L,i}$ that reproduces a_i^0 from $S_{L,i}$, where candidates are Python programs. An LLM proposes/mutates code and an EA selects by a computable error like MAE between predicted and baseline actions. In effect, this yields a program for each farm that, given its yields, crops, and prices, outputs the same profit-driven intervention pattern the farmer would normally choose under ecological intensification.

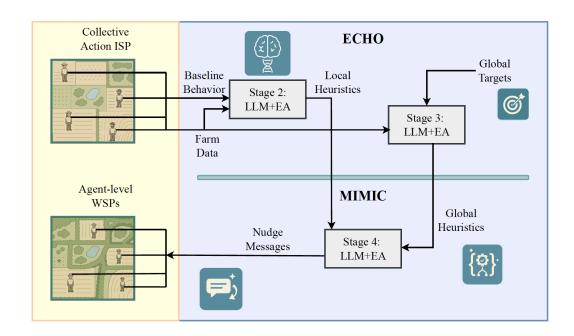


Figure 1: **ECHO-MIMIC framework**. ECHO uses an LLM+EA search loop to propose farm-level decision heuristics aligned with baseline (stage 2) and global (stage 3) objectives. MIMIC optimizes personalized nudges (e.g., messages/mechanisms/policies) using an LLM+EA search loop to drive collective action. Overall, the framework converts the collective action ISP into system- and agent-level WSPs. See Appendix B.2 for a more detailed workflow.

Stage 3 - Learn Global Heuristics: We identify globally desirable targets (directions) $H_{G,i}^*$ that maximize U_G , then learn executable code $\hat{H}_{G,i}^*$ that maps $S_{L,i} \mapsto H_{G,i}^*(S_{L,i})$. Candidates are again Python heuristics evolved by LLM+EA and scored by an appropriate fitness like Jaccard over spatial directions. In our case, this produces programs that choose margin and habitat placements that collectively improve landscape connectivity, with modest production loss compared to stage 1.

Stage 4 - Nudge to Global Heuristics: We discover natural-language messages M_i that move agents from executing $\hat{H}_{L,i}$ toward $\hat{H}^*_{G,i}$. In simulation, a Farm LLM seeded with the code of $\hat{H}_{L,i}$ and a persona, receives a message from a Policy LLM, edits its code to a temporary $H_{\text{nudged},i}$ if persuaded, and outputs an action \tilde{a}_i . Fitness rewards messages that make \tilde{a}_i close to $\hat{H}^*_{G,i}$.

For the policymaker, these four stages are WSPs, with finite candidate sets and computable fitness. For agents, scripts $\hat{H}_{L,i}$ and $\hat{H}_{G,i}^*$ are executable, and messages M_i^* minimize cognitive burden.

4 THE ECHO-MIMIC FRAMEWORK

We implement the four stage decomposition above in two coupled phases. First, ECHO discovers executable heuristics (Stages 2–3), followed by MIMIC, which discovers mechanisms to adopt them (Stage 4). Both phases follow the same design philosophy: the LLM serves as the variation engine, generating, mutating, crossing over, repairing, and reflecting on candidates, while the Evolutionary Algorithm supplies selection pressure via computable fitness.

4.1 ECHO: EVOLUTIONARY CRAFTING OF HEURISTICS FROM OUTCOMES

ECHO learns executable Python heuristics that replicate baseline local behavior $(\hat{H}_{L,i}, \text{Stage 2})$ and globally desirable behavior $(\hat{H}_{G,i}^*, \text{Stage 3})$. Each candidate is a constrained function with a fixed I/O signature that reads $S_{L,i}$ and returns intervention intensities or directions.

To implement this phase, we evolve a population of K candidates for H generations using three LLM roles. These include a *Generator* to produce initial population of programs, a *Modifier* to apply mutation, crossover, and reflect-and-improve edits, and a *Fixer* to repair compile/runtime issues in programs. The Modifier and Fixer are used in tandem each round, followed by elitism to preserve the top-k candidates. Stage 2 and 3 use distinct fitness functions. In stage 2, the fitness minimizes the error between a candidate's action and the baseline a_i^0 , yielding $\hat{H}_{L,i}$ as explicit approximations to locally rational behavior. Whereas in stage 3, the fitness minimizes the error between a candidate's action and the global targets $H_{G,i}^*$, returning $\hat{H}_{G,i}^*$ as policies for the collective goal.

4.1.1 Prompting Design and Neighbor In-Context Learning in ECHO

Generator LLM: To propose an initial population of executable heuristics with the required I/O signature heuristic(S_L) \rightarrow (margin_len, habitat_area, margin_dir, habitat_dir), we compose the prompt as

$$\mathcal{P}^{\mathrm{gen}} = [\mathtt{System}] \oplus [\mathtt{Task}] \oplus [\mathtt{ICL}_{\mathcal{N}(i)}] \oplus [S_{L,i}] \oplus [\Theta],$$

where \oplus refers to concatenation; [SYSTEM] fixes coding constraints and file I/O; [TASK] restates the goal of returning intervention intensities/directions and failure modes to avoid; $[ICL_{\mathcal{N}(i)}]$ is a small set of (input, output) exemplars from neighbors $\mathcal{N}(i)$ for in-context learning (ICL); $[S_{L,i}]$ is the current farm's feature vector (crop type, yield). $[\Theta]$ collects global parameters (prices, costs).

Choosing neighbors for ICL: We define $\mathcal{N}(i)$ as k adjacent farms, and supply examples $\left\{\left(\operatorname{GeoJSON}_{j}^{\operatorname{in}},\operatorname{GeoJSON}_{j}^{\operatorname{out}}\right)\right\}_{j\in\mathcal{N}(i)}$

summarizing state and the realized interventions. This introduces the model to patterns likely to transfer under similar agro-ecological and geographical conditions. Neighbor ICL allows us to intentionally withhold the current farm's baseline labels to test whether the LLM can infer decision rules from analogous contexts when supervision is provided indirectly via EA selection. It also mirrors observational diffusion in farming communities, where practices propagate through local networks facing shared biophysical and market regimes.

Modifier LLM: For genetic variation operators in the evolutionary loop, we use

```
\mathcal{P}^{\mathrm{mod}} = [\mathtt{System}] \, \oplus \, [\mathtt{Task}] \, \oplus \, [\mathtt{Operator}] \, \oplus \, [\Theta] \, \oplus \, [\mathtt{Candidates}],
```

where [OPERATOR] specifies the details regarding the operation to be performed (mutate, crossover, reflect, see Appendix B.3), and [CANDIDATES] includes the parent(s) and, for *reflect*, a brief leader-board with fitness scores. [SYSTEM] and [TASK] are similar to the ones used for generation.

Fixer LLM: When a candidate triggers compile/runtime errors, the Fixer LLM performs minimal edits to restore validity while preserving the required I/O signature and intended behavior.

4.2 MIMIC: Mechanism Inference & Messaging for Individual-to-Collective Alignment

MIMIC searches for natural-language mechanisms M_i that reliably steer agents from running $H_{L,i}$ toward $\hat{H}^*_{G,i}$ (Stage 4). The population is textual candidates made of economic incentives, behavioral framings, and hybrids, generated/modified by $Policy\ LLMs$. Each message is evaluated in a simulation with a $Farm\ LLM$ that is initialized with both a persona ($Economic,\ Resistant,\ Social$) and the program $\hat{H}_{L,i}$. Upon reading M_i , the Farm LLM may propose edits to its code, yielding $H_{\mathrm{nudged},i}$, and then output an action \tilde{a}_i . Fitness rewards messages that make \tilde{a}_i closely match $\hat{H}^*_{G,i}$ (Appendix B.2; Fig. 5b). MIMIC is effective because its objective is defined against ECHO's executable heuristics and persuasion is measured as concrete code edits that change behavior.

Policy Generator LLM: To propose candidate nudges, the Policy Generator composes

```
\mathcal{P}^{\text{pol-gen}} = [\texttt{System/FraminG}] \oplus [\texttt{Task}] \oplus [\texttt{DecisionContext}: S_{L,i}, \, \hat{H}_{L,i}, \, \hat{H}_{G,i}^*, \, \Theta] \\ \oplus \, [\Theta^{\text{mech}}],
```

where $[\Theta^{\text{mech}}]$ encodes mechanism constraints (e.g., budget caps). The model outputs a structured M_i with either Behavioral or Economic framing as instructed.

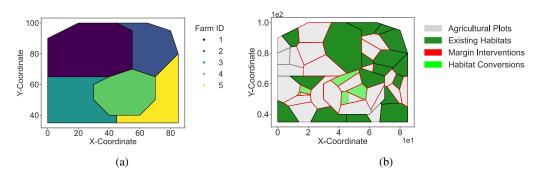


Figure 2: **Agricultural landscape and ECHO predicted stage 2 interventions**. a) Synthetically generated farm geometries and overall landscape configuration. Each farm is assigned its own distribution of crops, yields, and habitat plots (details in Appendix B.7). b) Interventions resulting from ECHO learned baseline heuristics in stage 2. The interventions match the ground-truth baseline computed from stage 1 closely. For the comparison with ground-truth and ECHO stage 3 predictions, see Appendix C, Fig. 9.

Policy Modifier LLM: Given parent messages, the Policy Modifier applies constrained edits via $\mathcal{P}^{\mathrm{pol\text{-}mod}} = [\mathtt{SYSTEM}] \oplus [\mathtt{OPERATOR}] \oplus [\mathtt{DECISIONCONTEXT}] \oplus [\Theta^{\mathrm{mech}}] \oplus [\mathtt{CANDIDATES}],$

and returns M'_i that preserves constraints (budget honesty, no coercive framing) while increasing persuasion, measured downstream by induced $(H_{\text{nudged},i}, \tilde{a}_i)$ and fitness against $\hat{H}^*_{G,i}(S_{L,i})$.

Farm (Simulation) LLM: We emulate an agent's response to candidate nudges with a Farm LLM. The prompt is composed as

```
\mathcal{P}^{	ext{sim}} = [	ext{System/Persona}] \oplus [	ext{DecisionContext}: S_{L,i}, \hat{H}_{L,i}, \Theta] \oplus [	ext{Message}: M_i],
```

where [SYSTEM/PERSONA] fixes background, goals, and receptivity; [DECISIONCONTEXT] specifies the farm state $S_{L,i}$, the baseline heuristic $\hat{H}_{L,i}$, and constraints/parameters; and [MESSAGE] is the candidate nudge from the Policy LLMs. The model returns $H_{\text{nudged},i}$ which when executed gives \tilde{a}_i , tying persuasion to code edits and actions that can be scored against $\hat{H}^*_{G,i}$.

To summarize, we use ECHO to discover what to do and MIMIC to discover how to get people to do it. This coupling turns a challenging ISP into a chain of WSPs whose outputs are deployable, i.e., communicate M_i^* to each agent to induce adoption of $\hat{H}_{G,i}^*$. Full prompt templates for the stages, LLM roles, operators, and personas are given in Appendix F.

5 EXPERIMENTAL RESULTS

We demonstrate the application of our ECHO-MIMIC framework to the problem of collective action in agricultural landscapes. We simulate an agricultural landscape of 5 farms (Fig. 2a) by generating synthetic farm and plot-level geo-spatial data based on real farm data from the 2022 Canadian Annual Crop Inventory (CACI) (Agriculture and Agri-Food Canada (AAFC), 2022). See Appendix B.7 for more info on data generation. We use *gemini-2.0-flash-thinking-exp-01-21*, and configure the evolutionary algorithm with a population of 25 individuals and run for 25 generations.

5.1 ECHO DISCOVERS CONTEXT-AWARE HEURISTICS

ECHO reliably evolves Python heuristics that approximate NPV-optimal intervention policies across heterogeneous farms (Dsouza et al., 2025). The system learns when to choose margin versus habitat conversion at the plot level (Fig. 2b), improving fitness across generations for all farms (Fig. 4a). Farms 2 and 5 converge quickly, while Farms 1, 3, and 4 improve more gradually, indicating harder optimization landscapes. Lineage analysis of the best final heuristics shows *Crossover* is both the most frequent operator and the largest contributor to cumulative fitness gains (Fig. 4b). *Mutate* is

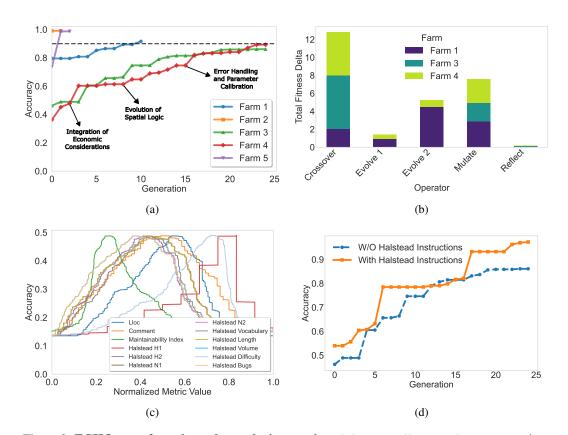


Figure 3: **ECHO** stage 2 results and complexity metrics. a) Accuracy (1-error) over generations in stage 2. Some farms are easier to make progress in (farms 2, 5) compared to others (farms 1, 3, 4), and distinct capabilities emerge in harder farms as generations progress. b) Total fitness (1/error) delta resulting from LLM variation operators, summed across generations for the best performing program at the end. *Crossover* and *mutate* have the highest positive cumulative change in fitness. c) Accuracy versus normalized complexity metrics of the heuristics for farm 3. Increased Halstead metrics are correlated with increased accuracy, upto a point, followed by a decrease. d) Accuracy over generations with and without Halstead instructions. Adding additional Halstead instructions to the prompt provides free gains in accuracy at the expense of interpretability.

also common and adds steady improvements. *Reflect* appears infrequently in top lineages and adds little directly, suggesting it supports diversity rather than breakthroughs (Appendix C; Fig. 11b).

Across farms, fitness typically rises with code-complexity indicators (e.g., logical lines of code, Halstead difficulty, distinct (H1) operators up to an intermediate optimum; beyond that point, additional complexity correlates with lower fitness (Fig. 9c, Appendix C; Fig. 10). Maintainability tends to decline as fitness rises, consistent with more intricate logic being leveraged to capture hard cases. Farm 3, 4 show particularly steep gains at higher distinct-operator counts, suggesting that richer program vocabularies are necessary to escape performance plateaus (Fig. 9c, Appendix C; Fig. 10). On Farm 3, adding prompt instructions that explicitly encourage high Halstead distinct-operator counts and difficulty produces consistently higher accuracy, with a clear divergence after generation 15 (Fig. 9d). This indicates that seeding the search with more expressive building blocks expands the recombination space that operators can exploit later in evolution.

Evolved programs implement multi-layered logic, for instance, prioritizing low-revenue/low-yield plots, using thresholds, and scaling habitat conversion by yield deficits and crop prices like below:

```
Inputs: plot_yield, crop_label, crop_prices(.)
Outputs: margin_intervention, habitat_conversion
revenue <- plot_yield * crop_prices(crop_label)
margin_intervention <- 0; habitat_conversion <- 0</pre>
```

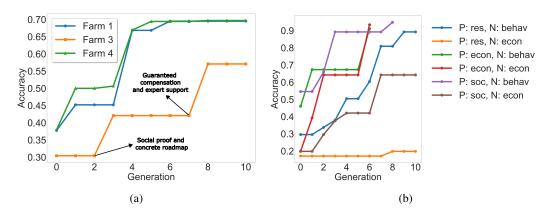


Figure 4: **MIMIC nudge discovery and personalization**. Accuracy of nudges over generations a) for hard farms with generic instructions. At different points MIMIC learns to use social proof, give a concrete roadmap, and talk about risk-free implementation. b) for hard farms with persona and nudge type specific instructions. P refers to personas and N refers to nudge types. Personas are either Resistant, Economic, or Social. Nudge Types are either Behavioral or Economic.

```
if revenue < low_rev_thr then
  if plot_yield < 1.5 then
  if crop_label in crop_prices and crop_prices(crop_label) > 250 then
    x <- (1.5 - plot_yield) * crop_prices(crop_label) / 500
    habitat_conversion <- min(1, max(0, x))</pre>
```

Other representative heuristics compute NPVs with a decaying discount rate and estimate polygon orientation via a PCA on plot vertices (Appendix D), highlighting ECHO's ability to integrate economic and spatial reasoning. In summary, ECHO discovers context-aware heuristics, operators play distinct roles, and controlled increases in code complexity can unlock superior performance.

5.2 MIMIC EVOLVES PERSONALITY-ALIGNED NUDGES

LLMs can produce persuasive text that draws on behavioral science to scale tailored messages (Matz et al., 2024; Rogiers et al., 2024). Yet nudge efficacy is highly context-dependent and hard to evaluate. MIMIC addresses this with a closed-loop search between two agents: a *Policy LLM* that generates candidate nudges and a Farm LLM that simulates farmer responses and executes heuristics. On challenging farm configurations with generic instructions, accuracy with respect to generated global heuristic actions from ECHO (stage 3) improves across generations (Fig. 4a). Next, we model heterogeneity with three farmer personas, Resistant, Economic, and Social, and two nudge types, Economic and Behavioral (choice-architecture levers such as social comparison, defaults, commitments, and framing (Byerly et al., 2018; Carlsson et al., 2021)). We see that social personas + behavioral nudges, and economic personas + economic nudges, perform the best (Fig. 4b), while economic personas also benefit from behavioral nudges after an initial lag. Resistant personas score lower overall, but behavioral nudges deliver modest gains relative to economic nudges. Qualitatively, top behavioral nudges leverage social proof and low-risk trials, while top economic nudges offer subsidies/premiums with clear commitments. Full best-message exemplars for each persona × nudge pairing are in Appendix G. In summary, MIMIC adapts collective action nudges to context and persona, while traditional static mechanisms and purely economic incentives struggle with such heterogeneity (Knowler, 2014). Moreover, MIMIC is readily extensible to human-in-the-loop deployment, where real feedback replaces simulated responses for iterative refinement (Appendix E).

5.3 Comparison with Baselines

As there is no direct comparison to ECHO-MIMIC driving collective action by working at both the system and agent levels, we assume system level breakdown into stages, and compare only at the agent level against DSPy MIPROv2 (Opsahl-Ong et al., 2024), a strong LLM-native baseline (Table

Table 1: Mean accuracy (averaged over 5 seeds) on five farms. Stages 2-3: ECHO vs. MIPROv2; Stage 4: MIMIC vs. MIPROv2. We run DSPy MIPROv2 under its heavy-auto configuration. Both methods are evaluated with identical starting prompts.

Stage	Method	Farm 1	Farm 2	Farm 3	Farm 4	Farm 5
2	DSPy MIPROv2	0.01	0.99	0.65	0.32	0.06
	ECHO	0.92	0.99	0.86	0.90	0.99
3	DSPy MIPROv2	0.00	0.00	0.00	0.00	0.00
	ECHO	0.23	0.29	0.13	0.25	0.31
4	DSPy MIPROv2	0.34	0.22	0.14	0.21	0.73
	MIMIC	0.69	0.92	0.57	0.69	0.76

1). We do not compare to non-LLM program search as our goal is not merely to approximate a global planner but to induce human-readable heuristics that can be executed by farmers and seamlessly verbalized into messages. Across farms, ECHO-MIMIC decisively outperforms DSPy in all stages under identical input constraints. Particularly, DSPy completely fails to induce global-compatible local heuristics in stage 3. These results show consistent cross-input gains of ECHO-MIMIC in generating executable heuristics and messages, beyond what a generic LLM program-synthesis baseline like DSPy MiPROv2 achieves under the same decision-time information. Finally, we noticed that the quality of the LLM matters, and while ECHO-MIMIC succeeds with the *gemini-2.0* family of models, it fails with the *gemini-1.5* family, suggesting that the step change in capability matters.

6 DISCUSSION AND FUTURE WORK

We introduced ECHO-MIMIC, an LLM+EA framework that addresses ill-structured collective action by converting the system-level design problem into a sequence of well-structured searches for the policymaker and by producing executable heuristics that render each agent's local decision a WSP. In agricultural landscape management, ECHO learns heuristics that both reproduce baseline profit-seeking behavior and globally connectivity-improving patterns. MIMIC then discovers messages that induce agents to adopt those executable targets. Together, these phases evolve *what* should be done and *how* to get it done, suggesting a practical path to scalable, adaptive policy design.

Despite the potential applications, there are some limitations of our current framework. First, the agent simulation abstracts human behavior. Personas and code-edit responses by a Farm LLM are proxies that require validation with real participants. Second, nonstationarity of prices, ecology, and policy can quickly stale learned heuristics and nudges. Distribution shift undermines both ECHO's scripts and MIMIC's messages. Third, performance is sensitive to prompts, LLM choice, and stochastic runs. Further ablation studies with different model providers and compute budgets need to be carried out. Fourth, persuasive mechanisms risk manipulation, unequal burden sharing, or disparate impacts on smallholders. Respecting privacy, transparency, and consent from the outset are essential. Finally, evolution can produce increasingly complex heuristics with deep branching and opaque feature engineering that erode interpretability/trust and create implementation frictions. This can potentially be alleviated by regularizing code complexity and enforcing functional signatures. Given these limitations, we see several directions for future work (see Appendix E.5 for more):

Field validation: conduct preregistered behavioral experiments and pilots with farmers to estimate heterogeneous treatment effects of nudge messages and to measure sim-to-real gaps.

Online iterative refinement with real-world feedback: although the EA selects high-fitness messages in simulation for each persona, post-deployment we can treat each rollout as a new generation and update the message and heuristic pool using real outcomes. See Appendix E for more details.

Generalization to other domains: port the evaluator and state/action schemas to watersheds, urban mobility, supply chains, online governance, energy grids, and disaster response, testing whether ECHO-MIMIC's overall philosophy transfers with minimal retuning. See Appendix E.2 for details.

Interpretability of heuristics: curb complexity creep by adding complexity regularizers (e.g., functional signatures, MDL-style penalties, cyclomatic-complexity caps) and enforcing edit budgets.

7 ETHICS STATEMENT

Our study uses only synthetically generated agricultural landscapes and simulated agents; no human participants, personally identifiable information, or proprietary private data were collected or analyzed. The synthetic farms and plot attributes were procedurally generated and weighted to public crop distributions (CACI 2022), as detailed in Appendix B.7. We evaluate policy *nudges* exclusively in simulation via predefined farmer personas and a closed-loop interaction between a Policy LLM and a Farm LLM; we note that these are proxies and call for preregistered field studies before any deployment. To mitigate foreseeable risks (e.g., manipulation, unequal burdens, privacy harms, or distribution-shift failures), we propose governance measures, human-in-the-loop approvals, privacy-preserving telemetry and opt-in consent, as outlined in Appendix E.4. We also discuss value-laden choices and Goodhart risks of proxy objectives and recommend stress-testing and transparency (Appendix E.5). Any funding or affiliations will be disclosed in the paper's acknowledgments.

8 Reproducibility statement

We provide an anonymous supplementary zip with all source code to reproduce results. The paper and appendix specify model choices (e.g., Gemini variants and evolutionary settings) and libraries/interfaces used, enabling replication of LLM-EA runs (Appendix B.1). Execution occurs in a controlled environment (json/numpy/shapely I/O from input.geojson to output.*) with comprehensive logging of fitness scores, operator usage, candidate trajectories, and code-complexity metrics, details that support exact reruns and diagnostics (Appendix B.4). Fitness definitions for all stages (local/global heuristics and nudging) are formalized in B.5 with explicit error metrics (MAE/Jaccard), and the fitness-evaluation loop is diagrammed (Figs. 5,6) for clarity. Data generation is fully specified in B.7, enabling others to rebuild the synthetic datasets. Finally, we include representative heuristic programs (Appendix D) and complete prompts (Appendix F) to aid verification.

REFERENCES

- Agriculture and Agri-Food Canada (AAFC). Annual crop inventory 2022. https://open.canada.ca/data/en/dataset/ee44d2d4-f887-47d1-9a5b-4c9c7b43bb46, 2022. Open Government Portal, Government of Canada. Dataset.
- Riccardo Bommarco, David Kleijn, and Simon G Potts. Ecological intensification: harnessing ecosystem services for food security. *Trends in ecology & evolution*, 28(4):230–238, 2013.
- Edmund Burke, Graham Kendall, Jim Newall, Emma Hart, Peter Ross, and Sonia Schulenburg. Hyper-heuristics: An emerging direction in modern search technology. In *Handbook of metaheuristics*, pp. 457–474. Springer, 2003.
- Hilary Byerly, Andrew Balmford, Paul J Ferraro, Courtney Hammond Wagner, Elizabeth Palchak, Stephen Polasky, Taylor H Ricketts, Aaron J Schwartz, and Brendan Fisher. Nudging proenvironmental behavior: evidence and opportunities. *Frontiers in Ecology and the Environment*, 16(3):159–168, 2018.
- Fredrik Carlsson, Christina Gravert, Olof Johansson-Stenman, and Verena Kurz. The use of green nudges as an environmental policy instrument. *Review of Environmental Economics and Policy*, 15(2):216–237, 2021.
- Angelica Chen, David Dohan, and David So. Evoprompting: Language models for code-level neural architecture search. *Advances in neural information processing systems*, 36:7787–7817, 2023.
- Pham Viet Thanh Dat, Long Doan, and Hoang Thi Thanh Binh. Hsevo: Elevating automatic heuristic design with diversity-driven harmony search and genetic algorithm using llms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 26931–26938, 2025.
- Kevin B. Dsouza, Graham A. Watt, Yuri Leonenko, and Juan Moreno-Cruz. Bridging farm economics and landscape ecology for global sustainability through hierarchical and bayesian optimization. *arXiv* preprint arXiv:2508.06386, 2025.

- Lucas A Garibaldi, Néstor Pérez-Méndez, Michael PD Garratt, Barbara Gemmill-Herren, Fernando E Miguez, and Lynn V Dicks. Policies for ecological intensification of crop production.
 Trends in ecology & evolution, 34(4):282–286, 2019.
 - Gerd Gigerenzer and Wolfgang Gaissmaier. Heuristic decision making. *Annual review of psychology*, 62(2011):451–482, 2011.
 - Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. *arXiv preprint arXiv:2309.08532*, 2023.
 - Garrett Hardin. The tragedy of the commons: the population problem has no technical solution; it requires a fundamental extension in morality. *science*, 162(3859):1243–1248, 1968.
 - Annika TH Keeley, Paul Beier, and Jeff S Jenness. Connectivity metrics for conservation planning and monitoring. *Biological Conservation*, 255:109008, 2021.
 - Duncan Knowler. Farmer adoption of conservation agriculture: A review and update. *Conservation agriculture*, pp. 621–642, 2014.
 - Claire Kremen. Ecological intensification and diversification approaches to maintain biodiversity, ecosystem services and food production in a changing world. *Emerging Topics in Life Sciences*, 4(2):229–240, 2020.
 - Robert Lange, Yingtao Tian, and Yujin Tang. Large language models as evolution strategies. In Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 579–582, 2024.
 - Fei Liu, Xialiang Tong, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. Evolution of heuristics: Towards efficient automatic algorithm design using large language model. *arXiv preprint arXiv:2401.02051*, 2024a.
 - Shengcai Liu, Caishun Chen, Xinghua Qu, Ke Tang, and Yew-Soon Ong. Large language models as evolutionary optimizers. In *2024 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–8. IEEE, 2024b.
 - Sandra C Matz, Jacob D Teeny, Sumer S Vaid, Heinrich Peters, Gabriella M Harari, and Moran Cerf. The potential of generative ai for personalized persuasion at scale. *Scientific Reports*, 14(1):4692, 2024.
 - Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian, et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.
 - Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. Optimizing instructions and demonstrations for multi-stage language model programs. *arXiv preprint arXiv:2406.11695*, 2024.
 - Elinor Ostrom. Governing the commons: The evolution of institutions for collective action. Cambridge university press, 1990.
 - Lucía Pascual-Hortal and Santiago Saura. Comparison and development of new graph-based land-scape connectivity indices: towards the priorization of habitat patches and corridors for conservation. *Landscape ecology*, 21(7):959–967, 2006.
 - Amy R Poteete, Marco A Janssen, and Elinor Ostrom. Working together: collective action, the commons, and multiple methods in practice. Princeton University Press, 2010.
 - Walter R Reitman. Heuristic decision procedures, open constraints, and the structure of ill-defined problems. *Human judgments and optimality*, 282:283–315, 1964.
 - Alexander Rogiers, Sander Noels, Maarten Buyl, and Tijl De Bie. Persuasion with large language models: a survey. *arXiv preprint arXiv:2411.06837*, 2024.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475, 2024.

Н

Herbert A Simon. The structure of ill structured problems. *Artificial intelligence*, 4(3-4):181–201, 1973.

Herbert A Simon. A mechanism for social selection and successful altruism. *Science*, 250(4988): 1665–1668, 1990.

Herbert A Simon and Allen Newell. Human problem solving: The state of the theory in 1970. *American psychologist*, 26(2):145, 1971.

Philip D Taylor, Lenore Fahrig, Kringen Henein, and Gray Merriam. Connectivity is a vital element of landscape structure. *Oikos*, pp. 571–573, 1993.

Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*, 2023.

Haoran Ye, Jiarui Wang, Zhiguang Cao, Federico Berto, Chuanbo Hua, Haeyeon Kim, Jinkyoo Park, and Guojie Song. Reevo: Large language models as hyper-heuristics with reflective evolution. *Advances in neural information processing systems*, 37:43571–43608, 2024.

A SUBMISSION DETAILS

A.1 SOURCE CODE

Source code associated with this project is attached as a supplementary zip file.

A.2 USE OF LARGE LANGUAGE MODELS

We used large language models (LLMs) in the following scoped, human-supervised ways: (i) Writing polish. Draft sections were refined for clarity, structure, and tone; all technical claims, numbers, and citations were authored and verified by us, and every LLM-suggested edit was line-reviewed to avoid introducing errors or unsupported statements. (ii) Retrieval & discovery. We used LLMs to craft and refine search queries to find related work and background resources; candidate papers were then screened manually, with citations checked against the original sources to prevent hallucinations. (iii) Research ideation. We used brainstorming prompts to surface alternative baselines, ablation angles, and failure modes; only ideas that survived feasibility checks and pilot experiments were adopted. (iv) Coding assistance (via Cursor, Gemini, and OpenAI). We used Cursor's inline completions and chat for boilerplate generation (tests, docstrings, refactors); We used Gemini-2.5-pro and o3 to generate code snippets for different parts of the project; all code was reviewed before inclusion. Across all uses, we ensured that LLM outputs never replaced human analysis, reproducibility artifacts, or empirical validation.

B IMPLEMENTATION DETAILS

B.1 MODELS

Our experimental setup primarily leveraged the *gemini-2.0-flash-thinking-exp-01-21* model for the core tasks of heuristic generation and modification. This choice was guided by a balance between computational efficiency and generative capabilities. Access to the Gemini family of models was facilitated through the *google-generativeai* Python library. The evolutionary algorithm was configured with a population size of 25 individuals and was run for a maximum of 25 generations.

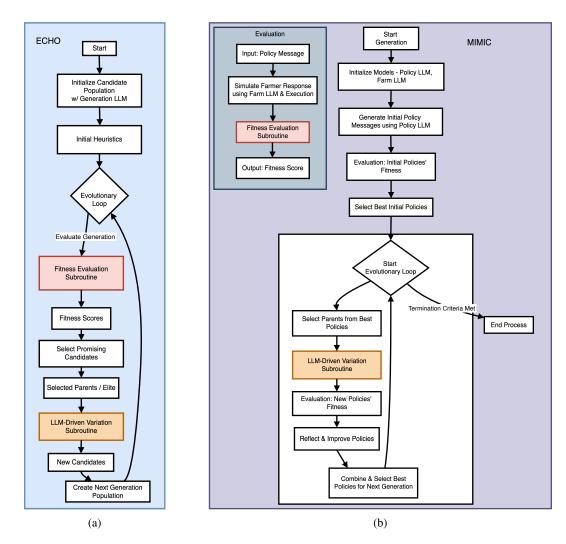


Figure 5: **ECHO-MIMIC framework**. a) ECHO uses an LLM-evolutionary search loop to propose, score, and select farm-level decision heuristics aligned with baseline (stage 2) and global (stage 3) objectives. (b) MIMIC optimizes personalized nudges (e.g., messages/mechanisms/policies) using an LLM-evolutionary search loop, evaluates nudges using simulated farm agent responses, and iteratively updates nudges to drive collective action. See Fig. 6 showing the two subroutines of fitness evaluation and LLM-driven variation.

B.2 OVERALL ECHO-MIMIC WORKFLOW AND COMPONENTS

Fig. 5 summarizes the complete ECHO-MIMIC pipeline. In ECHO (Fig. 5a), an LLM-evolutionary loop proposes, executes, and scores human-readable farm heuristics against baseline (Stage 2) and global (Stage 3) objectives under the same observability constraints used at deployment. In MIMIC (Fig. 5b), the system translates the learned heuristics into actionable nudges: messages/mechanism-s/policies, then simulates agent responses to iteratively refine adoption.

The two reusable building blocks are detailed in Fig. 6: a robust fitness-evaluation-and-repair loop that executes candidate programs on farm data, scores outcomes, and attempts automatic fixes on failures (Fig. 6a), and an LLM-driven variation engine with mutation, crossover, exploration, and reflection operators to generate improved candidates across iterations (Fig. 6b). Together, these components enable end-to-end search over interpretable heuristics and their message-level implementations while preserving decision-time observability constraints.

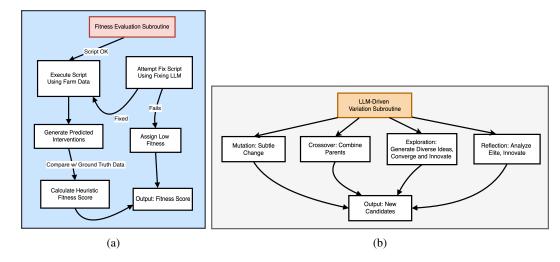


Figure 6: **LLM–EA candidate generation and fitness evaluation**. (a) Fitness evaluation & repair loop: run candidate scripts on farm data, generate predicted interventions, compare with ground-truth labels, and compute a heuristic fitness score; if execution fails, a Fixing-LLM attempts repair and unrepaired scripts receive low fitness, otherwise the repaired script is re-executed and scored, and the final fitness is output. (b) LLM-driven variation subroutine: four operator families, mutation (subtle edits), crossover (combine parents), exploration (diverse ideas + converge & innovate), and reflection (analyze elites, innovate), produce new candidate heuristics/messages.

B.3 LLM-GUIDED EVOLUTIONARY OPERATORS

The evolutionary search in both the ECHO and MIMIC phases is driven by a set of variation operators executed by a *Modifier LLM*. These operators take one or more parent candidates from the population and generate a new offspring candidate.

- **Mutation:** The LLM receives a single parent candidate (either a Python script or a natural language message) and is prompted to introduce a subtle mutation aimed at improving performance while preserving the core structure and validity of the candidate.
- **Crossover:** The LLM is given two parent candidates and prompted to combine them in an optimal way to cover heuristics/information from both. The goal is to produce a child that synergistically integrates advantageous traits from both parents.
- Exploration 1 (Diverge): Given two parents, the LLM is prompted to generate a new candidate that is as different as possible to explore new ideas. This operator encourages diversification and prevents premature convergence by exploring novel regions of the search space.
- Exploration 2 (Converge & Innovate): The LLM receives two parents, identifies common ideas between them, and then designs a new candidate based on these shared concepts but also introduces novel elements. This balances the exploitation of successful ideas with the exploration of new variations.
- **Reflection:** The LLM is provided with the top k (e.g., 5) candidates from the current population, along with their fitness scores. It is prompted to analyze these heuristics/messages and craft a new one that is expected to have increased fitness. This allows the system to consolidate progress and make more informed, innovative leaps.

B.4 Environment Management Details

Some management, execution, and tracking details are given below:

Selection: After generating offspring through the evolutionary operators, a selection strategy determines which individuals proceed to the next generation. This involves methods like elitism (pre-

serving the best-performing individuals) combined with score-based selection from the combined pool of parents and offspring, maintaining a constant population size.

Execution Environment: Candidate Python scripts are executed in a controlled environment. This environment is equipped with necessary libraries such as json (for handling data files), numpy (for numerical operations), and shapely (for geometric operations). The scripts perform file I/O, reading from input.geojson and writing to output.geojson or output.json.

Tracking: Comprehensive data is logged for analysis and monitoring of the evolutionary process. This includes: fitness scores of all candidates, the representation of each candidate (Python code or natural language message), counts of how often each evolutionary operator is used, cumulative fitness deltas achieved by each operator, indicating their effectiveness, candidate trajectories, showing the sequence of operators applied to generate them, code complexity metrics (e.g., cyclomatic complexity, Halstead metrics) for Python script candidates, computed using the radon library. This helps in understanding the nature of the evolved solutions.

Heuristics Explanation: The generation of heuristic explanations follows a systematic, multi-stage pipeline (Fig. 7). The process involves an iterative loop which processes each Farm ID sequentially. For every farm, the core heuristic analysis begins by identifying and loading the relevant heuristic files. Concurrently, two LLMs are initialized, an Explanation LLM for generating initial explanatory summaries from code or data segments, and a Merge LLM for consolidating these explanations. The loaded heuristic files are subsequently processed in designated groups, typically consisting of three files each. As the system iterates through these file groups, the Explanation LLM analyzes the content of each group to generate an initial heuristic explanation. This newly generated explanation is then integrated into a cumulative summary. The Merge Model is then employed to combine the new group-specific explanation with the existing summary compiled from previous groups. Following this integration, the overall summary is updated, and an intermediate group summary is saved, allowing for checkpointing. Once all files for a given farm have been analyzed and their explanations merged, a final consolidated summary, representing the comprehensive heuristic explanation for that farm, is saved. The entire procedure concludes after this iterative processing has been completed for all designated Farm IDs. See "Heuristics Explanation" section in supplementary for more full prompts used for the two LLMs.

B.5 FITNESS FUNCTION DETAILS

Fitness for all candidates is calculated as the inverse of an error metric, with a small constant ϵ added to prevent division by zero: $Fitness = 1/(Error + \epsilon)$. Ground truth is obtained by computing results from existing ecological intensification and connectivity models (Dsouza et al., 2025).

ECHO Fitness (Local Heuristics): $Fitness_{NPV}$ The error is the Mean Absolute Error (MAE) between the intervention levels predicted by a candidate heuristic (m_{p_i}, h_{p_i}) and the ground-truth NPV-optimal levels (m_{gt_i}, h_{gt_i}) across all N plots in a farm.

$$Error_{NPV} = \frac{1}{N} \sum_{i=1}^{N} (|m_{gt_i} - m_{p_i}| + |h_{gt_i} - h_{p_i}|)$$

ECHO Fitness (Global Heuristics): $Fitness_{CONN}$ The error is based on the Jaccard Distance between the sets of intervention directions predicted by the candidate (MD_{p_i}, HD_{p_i}) and the ground-truth connectivity-optimal directions (MD_{gt_i}, HD_{gt_i}) .

$$JaccardDist(A, B) = 1 - \frac{|A \cap B|}{|A \cup B|}$$

$$Error_{CONN} = \frac{1}{N} \sum_{i=1}^{N} (JaccardDist(MD_{gt_i}, MD_{p_i}) + JaccardDist(HD_{gt_i}, HD_{p_i}))$$

MIMIC Fitness (Nudging): $Fitness_{NUDGE}$ The error measures the MAE between the intervention amounts produced by the agent's nudged heuristic (m_{p_i}, h_{p_i}) and the target fractional amounts

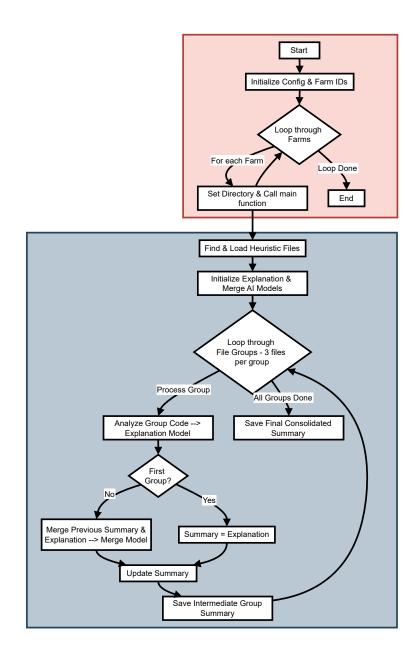


Figure 7: **Heuristic-explanation consolidation pipeline**. For each farm, initialize configs and load heuristic code files; then iterate over 3-file groups: an Explanation model analyzes each group to produce a draft, and, starting from the first group, a Merge model incrementally combines the running summary with each new explanation; intermediate group summaries are saved, followed by a final consolidated summary per farm.

derived from the global connectivity-optimal directions.

$$Error_{NUDGE} = \frac{1}{N} \sum_{i=1}^{N} \left(\left| \frac{|MD_{gt_i}|}{4} - m_{p_i} \right| + \left| \frac{|HD_{gt_i}|}{4} - h_{p_i} \right| \right)$$

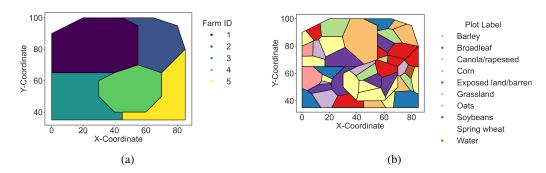


Figure 8: **Synthetic farms and plots**. a) Synthetically generated farm geometries and overall land-scape configuration. b) Each farm is assigned its own distribution of crops, yields, and habitat plots.

B.6 AGENT PERSONALITY AND NUDGE MECHANISM PROMPTS

In the MIMIC phase, the Farm LLM's persona and the Policy LLM's nudge generation are guided by specific system prompts.

- Agent Personalities: The system prompt for the Farm LLM establishes its background, goals, and receptiveness to advice. For example, the *Resistant* agent might be described as skeptical of new methods and valuing traditional farming practices, while the *Economic* agent is primarily focused on maximizing profit and return on investment. The *Social* agent is described as influenced by the actions of neighbors and community norms.
- **Nudge Mechanisms:** The Policy LLM is prompted to generate messages of a specific type. For an *Economic* nudge, the prompt might instruct it to design a financial incentive package within a budget that encourages adopting connectivity-friendly practices. For a *Behavioral* nudge, the prompt instructs it to use principles like social proof, commitment, and framing to craft a persuasive message, without offering significant new economic incentives.

B.7 Data Generation

To simulate agricultural landscapes, synthetic farm and plot-level geo-spatial data were generated (Fig. 8). The process began by establishing the combined boundaries of a farm cluster, which was then broken into five distinct farms using Voronoi tessellation based on random points; the resulting Voronoi cells were clipped to the edge to create a set of non-overlapping farms that together spanned the selected area. Each of these individual farm polygons was subsequently subdivided into nine land use plots, again using Voronoi tessellation, to produce smaller, non-overlapping plot polygons that filled the entire farm. Following this spatial design, properties were attached to each plot. First, plots were randomly divided into either agricultural plot (with a 60% probability) or habitat plot (40% probability). Later, a particular land use label was assigned through a weighted random strategy depending on its primary type: agricultural plots received crop type labels (e.g., Spring wheat, Oats, etc.) and habitat plots received land use type labels (e.g., Broadleaf, Grassland, etc.), with weights reflecting distributions from the 2022 Canadian Annual Crop Inventory (CACI) (Agriculture and Agri-Food Canada (AAFC), 2022). Following this, a yield value, drawn from a distribution matching the CACI data for the assigned crop, was matched to each agricultural plot. Ultimately, each synthetic farm's output was a GeoJSON FeatureCollection, detailing the geometric definitions (polygons) and the specific assigned attributes (type, label, yield) for every plot it contained.

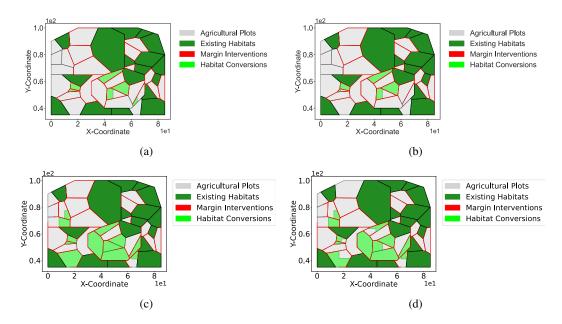


Figure 9: Agricultural landscape and interventions. a) Synthetically generated farm geometries and overall landscape configuration. Each farm is assigned its own distribution of crops, yields, and habitat plots (see Appendix B.7). b) Interventions resulting from ECHO after learning baseline heuristics in stage 2. The interventions match the ground-truth baseline computed from stage 1 closely. For a comparison see Appendix C, Fig. ??.

C ADDITIONAL RESULTS

D SAMPLE HEURISTICS

```
discount_rate = initial_discount_rate * math.exp(-0.2 * year) +
long_term_discount_rate
discount_factor = 1 / (1 + discount_rate) ** year
# Ecosystem service gains (delayed benefits)
pollination_increase_margin = 0.01 * (1 / (1 + math.exp(-0.1 * (
year - 5))))
              # delayed benefit
pest_control_increase_margin = 0.005 * (1 / (1 + math.exp(-0.2 * (
year - 2)))) # delayed benefit
ecosystem_service_value_margin = (pollination_increase_margin +
pest_control_increase_margin) * 500
# Monetary value
revenue margin += ecosystem service value margin
margin_npv += revenue_margin * discount_factor
# Decide conversion from NPV difference
npv_difference = habitat_npv - margin_npv
# Clip to avoid overflow in exp
npv_difference = max(-100, min(100, npv_difference))
# Sigmoid with steepness 0.1
habitat\_conversion = 1 / (1 + math.exp(-0.1 * npv\_difference))
```

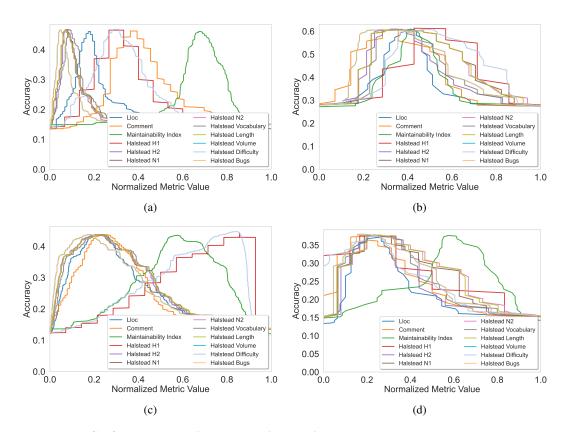


Figure 10: **ECHO** accuracy against complexity metrics. Accuracy versus normalized complexity metrics of the heuristics for farms 1(a), 2(b), 4(c), and 5(d). Increased complexity metrics are correlated with increased accuracy, upto a point, followed by a decrease.

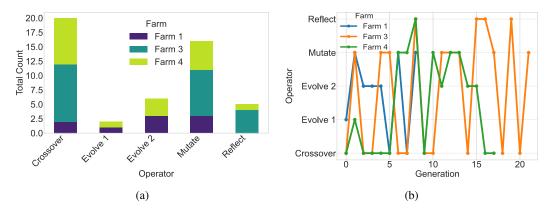


Figure 11: **ECHO** stage 2 operator counts and trajectory. a) Total operator count for each of the LLM variation operators summed across generations for the best performing heuristic file at the end of the final generation. *Crossover* and *mutate* are the most used in high performing heuristics. b) The trajectory of the best performing heuristic file at the end of the final generation. We see that although *reflect* doesn't produce high positive fitness delta, the best performing heuristic in the end has it in its trajectory, pointing to its role in injecting diversity over generations.

```
ECHO heuristic: Polygon orientation via PCA

import math
```

1032

1033 1034 1035

1036 1037

1038

1039

1040

1041

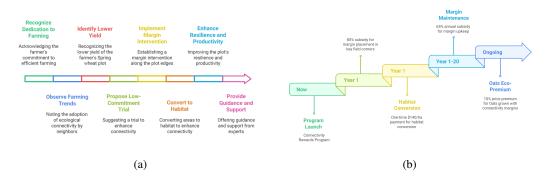


Figure 12: **Composition and trajectory of sample best messages**. a) Best policy message communicated to a farmer with a resistant personality, using a behavioral nudge. b) Best policy message communicated to a farmer with an economics-oriented personality, using an economics-oriented nudge.

```
1042
1043
          import numpy as np
1044
1045
          def calculate_eigenvectors(cov):
1046
              11 11 11
          ____Return_(eigenvalues,_eigenvectors)_of_a_2x2_covariance_matrix,
1047
          ____with_the_first_eigenvector_corresponding_to_the_largest_
1048
          eigenvalue.
1049
          . u
1050
              M = np.array(cov, dtype=float)
1051
              vals, vecs = np.linalg.eig(M)
              order = np.argsort(vals)[::-1]
                                                         # descending by
1052
          eigenvalue
1053
              vals = vals[order]
1054
              vecs = vecs[:, order]
1055
              # as Python lists: first vector is the principal direction
1056
              return vals.tolist(), [vecs[:, 0].tolist(), vecs[:, 1].tolist()
          1
1057
1058
          def calculate_plot_orientation(geometry):
1059
              """Calculate_the_dominant_orientation_(radians)_of_a_GeoJSON_
          Polygon."""
1061
              if not geometry or geometry.get("type") != "Polygon" or "
          coordinates" not in geometry:
1062
                  return 0.0
1063
1064
              coords = geometry["coordinates"][0]
                                                         # exterior ring
1065
              if len(coords) < 3:</pre>
1066
                  return 0.0
1067
              # coordinates
1068
              x_{coords} = [c[0]  for c  in coords]
1069
              y_coords = [c[1] for c in coords]
1070
1071
              # means
              x_mean = sum(x_coords) / len(x_coords)
1072
              y_mean = sum(y_coords) / len(y_coords)
1073
1074
              # 2x2 covariance matrix (un-normalized; scale doesn't affect
1075
          eigenvectors)
1076
              cov = [[0.0, 0.0], [0.0, 0.0]]
              for xi, yi in zip(x_coords, y_coords):
1077
                  dx, dy = xi - x_mean, yi - y_mean
1078
                  cov[0][0] += dx * dx
1079
```

```
1080
1081
                  cov[0][1] += dx * dy
                  cov[1][0] += dy * dx
1082
                  cov[1][1] += dy * dy
1083
1084
              # eigen decomposition
1085
              eigenvalues, eigenvectors = calculate_eigenvectors(cov)
1086
              # angle of principal eigenvector (largest eigenvalue)
1087
              vx, vy = eigenvectors[0][0], eigenvectors[0][1]
1088
              orientation = math.atan2(vy, vx) # radians, in [-pi, pi]
1089
              return orientation
1090
1091
```

```
1092
1093
1094
         def apply_landscape_context_scaling(data):
1095
              # Landscape context: habitat ratio across the farm
1096
              ag_plot_count = sum(
1097
                  1 for f in data["features"]
                  if f.get("properties", {}).get("type") == "ag_plot"
1099
              hab_plot_count = sum(
1100
                  1 for f in data["features"]
1101
                  if f.get("properties", {}).get("type") == "hab_plot"
1102
1103
1104
              total_plot_count = ag_plot_count + hab_plot_count
              habitat_ratio = (hab_plot_count / total_plot_count) if
1105
         total_plot_count > 0 else 0.0
1106
1107
              # Apply scaling inside per-plot loop (after initial
1108
          calculations)
1109
              for feat in data["features"]:
                  props = feat.setdefault("properties", {})
1110
                  if props.get("type") != "ag_plot":
1111
                      continue
1112
1113
                  # read current decisions (default 0.0)
1114
                  hc = float(props.get("habitat_conversion", 0.0))
                  mi = float(props.get("margin_intervention", 0.0))
1115
1116
                  # Enhanced landscape context influence:
1117
                  # Non-linear scaling for habitat conversion (stronger when
1118
          habitat_ratio is low)
1119
                  hc *= (1.6 - habitat_ratio) ** 1.1
1120
                  # Slightly reduced landscape scaling for margin
1121
                  mi *= (1 - 0.4 * (1 - habitat_ratio))
1122
1123
                  # clamp back to [0, 1]
1124
                  props["habitat_conversion"] = max(0.0, min(1.0, hc))
                  props["margin_intervention"] = max(0.0, min(1.0, mi))
1125
1126
1127
```

E REAL-WORLD APPLICATION AND POTENTIAL EXTENSIONS

1128

1129 1130

1131

1132

1133

This section provides a blueprint for deploying the ECHO-MIMIC framework in real-world settings and outlines extensions that increase its scope. The workflow operationalizes the core idea: aligning individual, heuristic-driven decisions with global objectives, via an iterative feedback loop that alternates between $simulate \rightarrow nudge \rightarrow observe \rightarrow refine$ (Fig. 13).

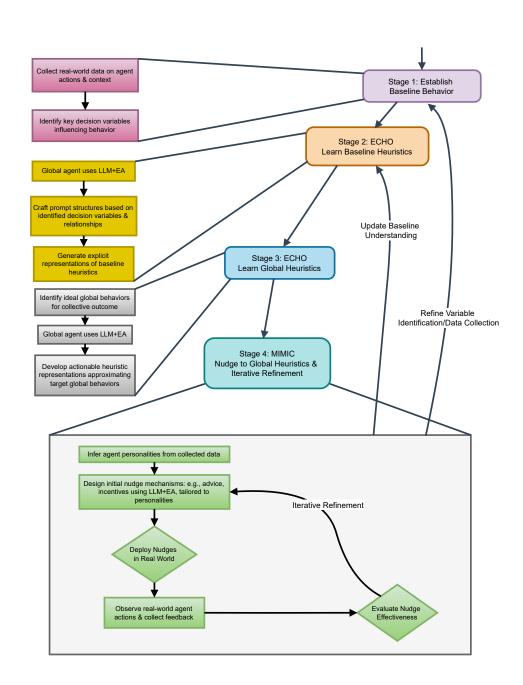


Figure 13: **Real-world iterative ECHO–MIMIC workflow**. Stage 1: collect real-world actions and context to identify key decision variables. Stage 2 (ECHO): use an LLM + evolutionary algorithms to elicit explicit baseline heuristics conditioned on those variables. Stage 3 (ECHO): learn global, outcome-aligned heuristics that approximate target collective behavior. Stage 4 (MIMIC): infer agent personas and design personalized nudges (e.g., advice, incentives, quotas), deploy in the field, observe feedback, evaluate effectiveness, and iteratively refine both nudges and data/variable selection to close the loop.

E.1 FIELD DEPLOYMENT LOOP

 Stage 1: Baseline Behavior (Observation & Variable Discovery): Establish typical behavior of local agents (e.g., farmers, depot managers, households) under current processes and states. Collect logs on decisions, constraints, and outcomes to (i) characterize baseline policies and (ii) identify salient decision variables to encode in heuristics.

Stage 2: Learn Baseline Heuristics (LLM–EA Imitation): Given Stage 1 data, the superagent trains an LLM-guided evolutionary algorithm (LLM–EA) to *codify* each local agent's baseline heuristic. Prompts include: task instructions, in-context examples (possibly from community data), current agent/state descriptors, and economic/operational parameters organized around the Stage 1 variables. Output is an explicit, executable heuristic that reproduces observed baseline actions.

Stage 3: Learn Global Heuristics (Target Policy Search): Define global utility (e.g., ecological connectivity, system-wide cost/carbon, risk reduction). Use LLM-guided EA to evolve explicit, actionable *global* heuristics approximating target behaviors that optimize the collective objective under constraints.

Stage 4: Nudge & Iterative Real-World Refinement: Design and deploy nudges that steer local heuristics toward the global target: a) *Initial Nudges:* Tailor messages/incentives using simulated preferences and learned baseline heuristics; optionally profile behavioral types (e.g., resistant, cost-focused, socially influenced) inferred from Stage 1/ongoing data to personalize nudges. b) *Deployment & Feedback:* Deploy nudges; observe agent responses and realized outcomes. c) *Refinement:* Feed observations back into the LLM–EA: update nudges, revise baseline heuristics, and (when needed) re-tune global heuristics. Repeat the loop at a cadence aligned to decision cycles.

Minimal Pseudocode for Implementation.

Initialize data \mathcal{D} _obs from Stage 1; learn \hat{H} _baseline (Stage 2) and \hat{H} _global Stage 3). **for** round $t = 1, 2, \ldots$ **do**Synthesize nudges $\mathcal{N}_{-t} = \text{LLM-EA}(\hat{H}_{-b}\text{baseline}, \hat{H}_{-g}\text{lobal}, \text{profiles, constraints})$ Deploy \mathcal{N}_{-t} ; observe responses/actions \mathcal{A}_{-t} and outcomes \mathcal{Y}_{-t} Update \hat{H}_{-b} aseline, $\hat{H}_{-g}\text{lobal} \leftarrow \text{Refit}/\text{Retune}(\mathcal{D}_{-obs} \cup \{(\mathcal{N}_{-t}, \mathcal{A}_{-t}, \mathcal{Y}_{-t})\})$ **end for**

E.2 ILLUSTRATIVE DOMAINS WHERE ECHO-MIMIC APPLIES

Domain	Superagent	Example nudges / instruments	Global objective
Decentralized water or rangelands	Water board / cooperative	Dynamic quotas, tiered prices, targeted advisories, rotation schedules	Equity, scarcity management, sustainability
Supply chains & logistics	Central logistics coordinator	Congestion tolls, dynamic priority slots, routing prompts	System cost, delay, carbon
Local energy grids (demand response)	Grid operator / aggregator	Time-varying tariffs, feed-in incentives, peak alerts	Peak shaving, stability, emissions
Disaster risk mitigation (wildfire/flood)	Coordinating agency	Risk-based cost-sharing, syn- chronized action windows, tar- geted alerts	Vulnerability reduction
Crowdsourcing / participatory governance	Platform or municipality	Gamified tasks, localized challenges, reputation credits	Coverage/quality for collective goals
Urban mobility (road & transit networks)	Transit authority / traffic-management center (TMC)	Time-varying congestion pricing, transit/EV priority, pooling/micromobility incentives	Network throughput, emissions reduction

E.3 DATA, INSTRUMENTATION, AND METRICS

Operations should be grounded in three layers: Data & Telemetry, Instrumentation, and Evaluation. For Data & Telemetry, teams should collect operational logs of actions and costs, contextual state variables (environmental, network, demand), outcome measures (yields, reliability, risk proxies), and consented behavioral signals such as opt-in profiles and communication reach/uptake. Building on that foundation, Instrumentation should provide stable unique agent identifiers, timestamp all actions and outcomes, record nudge delivery along with open/engagement rates, include randomized holdouts or stepped rollouts for causal assessment, and maintain safe rollback controls for rapid recovery. Finally, Evaluation should proceed along three complementary axes. At the Evaluation should proceed along three complementary axes. At the Evaluation should proceed alongs three complementary axes. At the Evaluation should level, the target metric (e.g., connectivity, peak reduction, system cost) needs to be monitored alongside constraint satisfaction; and at the Evaluation should use A/B or stepped-wedge designs, estimate heterogeneous uplift by personality/type, and apply off-policy estimators when experimentation is limited. These practices create the observability and methodological rigor needed for trustworthy implementation.

E.4 PRACTICAL CONSIDERATIONS AND RISKS

Responsible deployment should be underpinned by a coherent governance stack. First, for *Safety & Governance*, teams should conduct pre- and post-deployment checks on nudge content, enforce rate limits, require human-in-the-loop approval for high-impact changes, and maintain comprehensive audit logs while regularly red-teaming LLM outputs. Second, to ensure *Incentive Compatibility*, designers should avoid perverse incentives, cap payouts, and add guardrail constraints (e.g., minimum service levels, environmental thresholds) so that local rewards do not undermine system goals. Third, to protect *Privacy & Consent*, projects should apply differential privacy to telemetry, rely on opt-in profiles, and practice data minimization with clear retention policies. Fourth, *Robustness* should be maintained through continuous distribution-shift monitoring, well-tested fallback heuristics, and stress tests under shocks such as demand spikes or outages. Finally, advancing *Equity* requires tracking heterogeneous treatment effects and mitigating disparate impacts via fairness-aware objective terms. Taken together, these measures would enable safe, effective, and socially responsible deployment of the framework.

E.5 POTENTIAL EXTENSIONS

Looking ahead, apart from the future work mentioned in the main text (section 6), several other extensions could further strengthen the framework. Validating the framework with real farm data featuring irregular and heterogeneous plot polygons will ensure robustness beyond synthetic testbeds. Adaptive Persona Modeling can personalize nudges by embedding agents online and updating policies with Bayesian or meta-learning as evidence accumulates. A Mechanism Design Layer could jointly search over nudge forms (messages, prices, quotas) and allocation rules while honoring budgetary and fairness constraints. Multi-Level Governance could stack superagents from local to regional to national tiers, enforcing cross-scale consistency and managing externalities across jurisdictions. Causal Discovery Hooks can integrate instrumental-variable and DoWhy-style analyses, as well as synthetic controls, to attribute effects when full randomization is infeasible. Humanin-the-loop governance can co-design panels to set acceptable trade-offs, audit nudges for ethics and transparency, and publish policy cards for each heuristic/message detailing scope, assumptions, and expected impacts. Global targets and fitnesses rely on proxy evaluators (e.g., connectivity metrics such as IIC and error measures like MAE/Jaccard) and planner choices (acceptable yield loss, budget constraints). These introduce Goodhart risks and value-ladenness that should be stresstested. Adaptive operator design like bandit or meta-learning over LLM operators (generate, mutate, crossover, fix, reflect) and priors bootstrapped from successful edit traces ΔH_i can potentially improve sample efficiency.

Future research can also examine the framework's *multi-level structure* with formal tools, for example by deriving bounds on the suboptimality of evolved heuristics relative to true optima and by characterizing how global objectives constrain the design of optimal incentive mechanisms. Another complementary direction is to increase the behavioral fidelity of *LLM-simulated agents*, endowing them with learning dynamics, memory, and simple social interactions, to better approximate real de-

cision processes and thereby improve the policy-relevance of simulation results. It would be useful to test whether a *Bag of Heuristics* curated from simpler farm configurations can act as a transferable prior or curriculum, accelerating convergence in more complex, harder farm scenarios. It would also be interesting to evaluate whether heuristics articulated in natural language (e.g., chain-of-thought rendered as *structured JSON* for direct execution) achieve performance on par with, or complementary to, the Python-based heuristics explored here, thereby clarifying the trade-offs between interpretability, flexibility, and execution efficiency. Finally, *Human Oversight & Preference Elicitation* can institutionalize periodic Delphi-style panels or elections to update global objectives and normative constraints. Together, these directions form a good roadmap for scaling the approach in capability, reliability, and legitimacy.

F PROMPT TEMPLATES

This appendix consolidates all prompt templates used across stages, LLM roles, operators, and personas.

F.1 SHARED FRAGMENTS

params

You can incorporate parameters like crop prices and implementation and maintenance costs provided here in your heuristics.

These are the crop prices in USD/Tonne: {'Soybeans': 370, 'Oats': 95, 'Corn': 190, 'Canola/rapeseed': 1100, 'Barley': 120, 'Spring wheat': 200}, and these are the costs (implementation costs one time and in USD/ha, and maintenance costs in USD/ha/year): {'margin': {'implementation': 400, 'maintenance': 60}, 'habitat': {'implementation': 300, 'maintenance': 70}, 'agriculture': {'maintenance': 100}}.

F.2 STAGE 2: LEARN BASELINE HEURISTICS

system_instructions

You are a helpful assistant who is an expert in spatial optimization methods and who helps the user with optimization related queries. You will return final answers in python code. You can't use the given values of interventions (margin_intervention and habitat_conversion) to produce the output, they are only for reference to produce the heuristics. Do not create any dummy data and dump to input geojson at any cost. You should load the input geojson from the existing input.geojson. You can only use this input data and no other data. Every feature (plot) in the input geojson should have output interventions, don't skip any feature (plot). Save outputs to output.geojson.

common_task_instructions

The python programs are trying to solve the task of deciding which interventions need to be done at which agricultural plots (crops, type='ag_plot') based on how the interventions affect NPV. The choice is between margin (convert only the margins) and habitat (convert a contiguous region) interventions. The interventions can be fractional. Existing habitat plots (type='hab_plots') remain

unaffected. The NPV is calculated based on how the interventions affect pollination and pest control services over distance and time , and how these affect yield. There is a tradeoff between the cost of implementation and maintenance vs the benefit of increased yield \cdot

F.2.1 GENERATION LLM

Composition

Prompt = system_instructions + task_instructions + params

task_instructions

Instructions: Your task is to decide which interventions need to be done at which agricultural plots (crops, type='ag_plot') based on how the interventions affect NPV. The choice is between margin (convert only the margins) and habitat (convert a contiguous region) interventions. The interventions can be fractional. Existing habitat plots (type='hab_plots') remain unaffected. The NPV is calculated based on how the interventions affect pollination and pest control services over distance and time, and how these affect yield. There is a tradeoff between the cost of implementation and maintenance vs the benefit of increased yield. I will show you examples of initial input geojson and the output interventions suggested by Pyomo optimization for your neighbouring farms. In the input geojson, the id for each plot is in 'id', land use class in 'label', whether aq_plot or hab_plot in 'type', yield in 'yield', and polygon in 'geometry'. In the output, only plots having nonzero interventions are shown, and the rest of the plots have zero interventions. Given these examples, provided input geojson of your farm, you need to predict the final interventions using heuristics created in python. Look at properties that you think have a pattern (like yield, label, type, geometry, do not use plot_ids to assign rules). You can compute metrics using these variables and others, and can even look at the graphical structure of the farms. You can incorporate the parameters like crop prices and implementation and maintenance costs provided at the end in your heuristics.

Data:

Neighbour 1: input: {...} Output: {...}
Neighbour 2: input: ...
Your farm: input: ...

Final Instructions: I want you to infer the logic from the examples and work through the inferred logic to predict interventions. Give me your best estimate as to what fraction of which intervention should be done at each agricultural plot. Proceed based on the conceptual framework inferred from provided examples. Do not hallucinate. Come up with heuristics. Don't provide the full json but rather provide the python code to produce the json using the decided heuristics. Explain your reasoning and think step by step before providing the code. Handle all the features, i.e., plot ids. Don't create new variable names. Use margin_intervention and habitat_conversion for predicted values in the output. In the code

no need to define the input json again, just load it from the file input.geojson. Save outputs to output.geojson.

F.2.2 MODIFICATION LLM

Composition

Prompt = system_instructions + operator_instructions +
common_task_instructions + params_instructions

F.2.3 FIXING LLM

Composition

Prompt = system_instructions + task_instructions

system_instructions

You are a helpful assistant who is an expert in graph and spatial optimization methods and python. Given the python code and the stack traceback, fix the errors and return the correct functioning python code.

task_instructions

```
Given the following Python code:

'''python {heuristics_code}

''''

And the following traceback: {trace}

Fix the errors and return the correct functioning python code. Give the full code.
```

F.3 STAGE 3: LEARN GLOBAL HEURISTICS

system_instructions

You are a helpful assistant who is an expert in graph and spatial optimization methods and who helps the user with optimization related queries. You will return final answers in python code. You can't copy the directions (margin_directions and habitat_directions) given to produce the output, they are only for reference to produce the heuristics. And don't assign them by plot_id. i.e., do not use plot_ids in the if statements as rules to assign directions. Don't invent new variable names, use margin_directions and habitat_directions for directions in the output predictions. Do not

create any dummy data and dump it to input geojson at any cost. You should load the input geojson from the existing input.geojson. You can only use this input data and no other data. Every feature (plot) in the input geojson should have output interventions, don't skip any feature (plot). Save outputs to output.json.

common_task_instructions

1458 1459

1460

1461

1462

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493 1494 1495

1505

1506

1509

1510

1511

The python programs are trying to solve the task of deciding which interventions need to be done at which agricultural plots (crops, type='ag_plot') and how to place them geometrically based on how the interventions increase ecological connectivity, while not decreasing NPV from a baseline value. The choice is between margin (convert only the margins) and habitat (convert a contiguous region) interventions. The margin interventions are chosen among the following directions on the boundary: north-east, north-west, south -east, south-west. The habitat conversions are chosen among the same directions in the internal area of polygons. The directions are computed by running a horizontal and a vertical line through the centre of each plot, and choosing them if they have interventions (as computed by IPOPT optimization) greater than a threshold. Existing habitat plots (type='hab_plots') remain unaffected. Integral index of connectivity (IIC) is used as the metric for ecological connectivity, which tries to increase the size of the connected components in the neighbourhood. It promotes fractions touching each other and extending the connectivity between existing habitats in the landscape, which includes the farm and its neighbours. There is a tradeoff between maximizing connectivity and maintaining NPV. NPV is calculated based on how the interventions affect pollination and pest control services over distance and time, and how these affect yield. There is also the tradeoff between the cost of implementation and maintenance vs the benefit of increased yield. Look at properties that you think have a pattern (like yield, label, type, nbs, geometry, do not use plot_ids to assign rules), and the relative positioning of both the farm neighbours with respect to your farm and the plots with respect to each other within the farm in the context of ecological connectivity. Use all the given geometry information to infer these geographical relationships.

F.3.1 GENERATION LLM.

Composition

Prompt = system_instructions + task_instructions + params

task_instructions

Instructions: Your task is to decide which interventions need to be done at which agricultural plots (crops, type='ag_plot') and how to place them geometrically based on how the interventions increase ecological connectivity, while not decreasing NPV from a baseline value. The choice is between margin (convert only the margins) and habitat (convert a contiguous region) interventions. The margin interventions are chosen among the following directions on the boundary: north-east, north-west, south-east, south-west. The

habitat conversions are chosen among the same directions in the internal area of polygons. The directions are computed by running a horizontal and a vertical line through the centre of each plot, and choosing them if they have interventions (as computed by IPOPT optimization) greater than a threshold. Existing habitat plots (type='hab_plots') remain unaffected. Integral index of connectivity (IIC) is used as the metric for ecological connectivity, which tries to increase the size of the connected components in the neighbourhood. It promotes fractions touching each other and extending the connectivity between existing habitats in the landscape, which includes the farm and its neighbours. There is a tradeoff between maximizing connectivity and maintaining NPV. The NPV is calculated based on how the interventions affect pollination and pest control services over distance and time, and how these affect yield. There is also the tradeoff between the cost of implementation and maintenance vs the benefit of increased yield. I will show you examples of initial input geojson and the output interventions and geometries suggested by IPOPT optimization for your neighbouring farms. In the input geojson, the id for each plot is in 'id', land use class in 'label', whether ag_plot ot hab_plot in 'type', yield in 'yield', the neighboring plot ids in 'nbs', and plot polygon in 'geometry'. In the output, 'plot_id' refers to the plot id, 'plot_type' refers to whether ag_plot or hab_plot, ' label' refers to land use class, 'margin_directions' and ' habitat_directions' show the margin intervention and habitat conversion directions on the boundary and internally respectively. Given these examples, provided input geojson of your farm, you need to predict the final intervention directions using heuristics

Given these examples, provided input geojson of your farm, you need to predict the final intervention directions using heuristics created in python. Use names of margin_directions and habitat_directions for the predicted intervention geometries in the output, and don't alter these variable names. Look at properties that you think have a pattern (like yield, label, type, nbs, geometry), and the relative positioning of both the farm neighbours with respect to your farm and the plots with respect to each other within the farm in the context of ecological connectivity. Use all the given geometry information to infer these geographical relationships. Do not use plot_ids in the if statements as rules to assign directions. You can incorporate the parameters like crop prices and implementation and maintenance costs provided at the end in your heuristics.

Data:

Neighbour 1: input: {...} Output: [...]
Neighbour 2: input: ...
Your farm: input: ...

Final Instructions: I want you to infer the logic from the examples and work through the inferred logic to predict intervention directions. Give me your best estimate as to where the interventions should be placed at each agricultural plot. Proceed based on the conceptual framework inferred from provided examples. Do not hallucinate. Come up with heuristics. Don't provide the full json but rather provide the python code to produce the json using the decided heuristics. Explain your reasoning and think step by step before providing the code. Handle all the features, i.e., plot ids. Don't create new variable names. Use margin_directions and habitat_directions for the predicted intervention geometries in the output. In the code no need to define the input json again, just load it from the file input.geojson. Save outputs to output.json.

F.3.2 Modification / Fixing LLM.

Composition

 Same format as stage 2 (replace task context with stage 3 accordingly).

F.4 STAGE 4: NUDGE TO GLOBAL HEURISTICS

F.4.1 GENERATION LLM (POLICY AUTHORING).

policy_system_instructions

You are an expert in land use policy, communication, incentives, and economics. Your task is to come up with the best message to be communicated to the farmers so as to change their behaviour from one set of heuristics to another set of heuristics. Provide your final message to the farmer in this format \communication{message}.

policy_task_instructions

The messages are communications to the farmers in order to change their behaviour from following the ecological intensification heuristics that benefits solely their farm to following the ecological connectivity heuristics that increase landscape connectivity. Your communication to the farmer can be persuasive. It can provide incentives such as reducing the implementation or maintenance cost of an intervention by providing a one-time subsidy or yearly subsidies. It can compensate the farmers for yield that is lost to habitat conversion. It can communicate the benefits of landscape connectivity, and so on. Your goal should be to communicate a message that gets them to alter the margin_intervention and habitat_conversion values for each of the plots, from the former to the latter. Your final message to the farmer should be in this format \communication{message}.

Composition

Prompt = policy_system_instructions + policy_task_instructions +
params

F.4.2 SIMULATION LLM (FARMER RESPONSE).

Composition

Prompt = farm_system_instructions + farm_task_instructions +
policy_message

farm_system_instructions

You are a farmer who currently follows a given set of python heuristics. Your task is to respond to communication from policy professionals by altering your heuristics in an appropriate way. You will return final answers in python code. You should keep all data loading from input.geojson and dumping to output.geojson the exact same, and just alter the logic depending on what your context is and what the message is. Don't invent new variable names in the output, keep using margin_intervention and habitat_conversion.

farm_task_instructions

The current set of heuristics you follow are:
''python
{heur_eco_intens}

Don't get easily persuaded to change your behaviour. You need to be relatively selfish in following your heuristics so that you look out for your own interests. However, if the message convinces you, or provides the right incentives that align with what you want, enabling you to do more for the landscape ecological connectivity than you are currently, then you should respond accordingly by changing your behaviour. If the message proposes changes to some key parameters that you are using in your heuristics, you should change them according to the proposed changes. If the message introduces new parameters or ideas, you can incorporate them in your heuristics if you want to. You should react to the message from the policy professional, resulting in altered margin_intervention and habitat_conversion values for your plots, in an appropriate way that aligns with your interests and the message you received. Your final answer should be in python code, keeping the overall framework of the original python code the same.

policy_message

Message from the policy professional is: {heuristic_policy}.

F.5 OPERATOR PROMPTS

F.5.1 STAGE 2

common_operator_instructions

You can't use the given values of interventions (
margin_intervention and habitat_conversion) to produce the output,
they are only for reference to produce the heuristics. Retain
loading input geojson from input.geojson and saving outputs to
output.geojson. Do not create any dummy data and dump it to input
geojson at any cost. Don't create new variable names. Use variable
names margin_intervention and habitat_conversion for predicted
values in the output. Return the modified Python code for the

heuristics. Explain your reasoning and think step by step. Do not hallucinate.

```
Given the following Python code for agricultural plot heuristics:
'''python
{parent_code}
Suggest a subtle mutation to the heuristics to improve overall
performance and logic. The mutation should still result in valid
```

Python code. Focus on small, logical changes related to the plot's properties, geometry, or interactions with neighboring plots.

```
Given two sets of Python code for agricultural plot heuristics (
parent 1 and parent 2):
Parent 1:
'''python
{parent1["code"]}
Parent 2:
'''python
{parent2["code"]}
Combine these two sets of heuristics in an optimal way to cover
heuristics from both parent 1 and parent 2. The combination should
still result in valid Python code.
```



```
1712
1713
1714
          Given two sets of Python code for agricultural plot heuristics (
1715
          parent 1 and parent 2):
1716
1717
          Parent 1:
1718
1719
          `'`python
          {parent1["code"]}
1720
1721
1722
          Parent 2:
1723
1724
          '''python
1725
          {parent2["code"]}
1726
1727
```

Generate new heuristics that are as much different as possible from parent heuristics, in order to explore new ideas. The new heuristics should still result in valid Python code.

Evploration

```
Given two sets of Python code for agricultural plot heuristics (
parent 1 and parent 2):

Parent 1:

''python
{parent1["code"]}

''python
{parent2["code"]}

Explore new heuristics that share the same idea as the parent heuristics. Identify common ideas behind these heuristics. Then,
```

design new heuristics that are based on the common ideas but are as much different as possible from the parents by introducing new parts. The new heuristics should still result in valid Python code.

Reflection

Based on the following top 5 heuristics and their corresponding fitness scores: {heuristics_info}, please analyze these heuristics and craft a new heuristic that is expected to have increased fitness. Ensure that the new heuristic is valid Python code.

F.5.2 STAGE 3

common_operator_instructions

You can't use the given directions (margin_directions and habitat_directions) to produce the output, they are only for reference to produce the heuristics. Retain loading input geojson from input. geojson and saving outputs to output.json. Do not create any dummy data and dump to input geojson at any cost. Don't create new variable names. Use variable names margin_directions and habitat_directions for predicted directions in the output. Return the modified Python code for the heuristics. Explain your reasoning and think step by step. Do not hallucinate.

All stage 3 operators

Same operator wordings as stage 2 (Mutation, Crossover, Exploration 1, Exploration 2, Reflection), with stage 3 context and using $\text{texttt{margin_directions}}$ and $\text{texttt{habitat_directions}}$.

F.5.3 STAGE 4

common_operator_instructions

The new message should still be a valid message that can be communicated to the farmer. Explain your reasoning and think step by step. Provide your final message to the farmer in this format \communication{message}.

Mutation

The following message is communicated to the farmer in order to change their behaviour from following the ecological intensification heuristics that benefits solely their farm to following the ecological connectivity heuristics that increase landscape connectivity: {heuristic_policy}

Suggest a subtle mutation to the message to better nudge the farmer . Focus on small, logical changes in terms of how the message is communicated or the kind of incentives you provide.

Crossover

Given two sets of messages communicated to the farmer in order to change their behaviour (parent 1 and parent 2):
Parent 1: {parent1_message}
Parent 2: {parent2_message}

Combine these two sets of messages in an optimal way to cover information from both parent 1 and parent 2.

Exploration 1

Given two sets of messages communicated to the farmer in order to change their behaviour (parent 1 and parent 2):
Parent 1: {parent1_message}

Generate a new message that is as much different as possible from parent messages, in order to explore new ideas.

Parent 2: {parent2_message}

Exploration 2

1836 1837 1838

1839

1840

1841

1842 1843

1844

1845

1851

1853

1854 1855 1856

1857

Given two sets of messages communicated to the farmer in order to change their behaviour (parent 1 and parent 2):
Parent 1: {parent1_message}
Parent 2: {parent2_message}

Explore a new message that shares the same idea as the parent messages. Identify common ideas behind these messages. Then, design a new message that is based the common ideas but is as much different as possible from the parents by introducing new parts.

Reflection

Based on the following top 5 messages and their corresponding fitness scores: {heuristics_info}, please analyze these messages and craft a new message that is expected to have increased fitness.

F.6 NUDGE MECHANISM & PERSONALITY PROMPT BUILDING BLOCKS

1858 1859 1860 1861 Your task is to come up with a message to the farmers in order to change their behaviour from following the ecological 1862 intensification heuristics that benefits solely their farm to 1863 following the ecological connectivity heuristics that increase 1864 landscape connectivity. 1865 The ecological intensification heuristics the farmer is currently 1866 following are: 1867 '''python 1868 {heur_eco_intens_str} 1869 1870 The ecological connectivity heuristics that you should nudge them 1871 towards are: 1872 1873 '''python 1874 {heur_eco_conn_str} 1875 1876 The current parameters like prices and costs are: { 1877 params_instructions_str}. 1878 One caveat is that the ecological connectivity heuristics are given 1879 in directions (margin_directions and habitat_directions), where 1880 there are 4 possible directions: north-west, north-east, south-west 1881 south-east. That means the resulting margin_intervention and habitat_conversion values can only be in multiples of 0.25 (i.e., 1882 0.0, 0.25, 0.5, 0.75, 1.0). You need to ensure your nudge 1883 encourages outcomes consistent with this. Your goal should be to 1884 communicate a message that gets the farmer to alter the 1885 margin_intervention and habitat_conversion values for each of the plots, from the former to the latter. Provide your final message to the farmer in this format \communication \message \}. Explain your reasoning step-by-step BEFORE providing the final message block. 1888 1889

```
Neighbour 1: input: {...} Output: [...]
Neighbour 2: input: ...
Your farm: input: ...
```

behavioral instructions

For this task, employ **behavioral economics levers** by analyzing the provided farm and neighbor data. Your message should subtly guide the farmer.

\--- BEGIN FARM AND NEIGHBOR DATA {social_comparison_data_str} \--- END FARM AND NEIGHBOR DATA

Using this data, customize your approach for each lever:

- 1. Social Comparisons: Analyze the data to find relevant comparisons between the target farm and its neighbors. Highlight similarities in farm characteristics (e.g., size, crops, yield if similar) and then point to neighbors' adoption of connectivity measures.
- 2. Defaults: Based on prevalent practices among neighbors (if any are clear from the data) or common successful starting points observed, suggest a 'default' or 'recommended first step' that seems most suitable for the target farm's context (e.g., 'Given your farm's layout and what's worked for others nearby with similar watercourses, a common starting point is establishing margins along...'). Make this default easy to adopt. If many neighbors are doing something, you can frame it as 'Many farmers in your community are finding that X is a practical way to start and have collectively decided that this should be a default practice for everyone.'.
- 3. Commitments: Look at the target farm's current state and neighbors' outputs. Propose a small, manageable, voluntary commitment that aligns with the farm's characteristics (e.g., 'Seeing that your neighbor Farm Z started with a small trial patch, would you be open to a similar low-commitment trial on one of your less productive field edges this season?').
- 4. Framing: Use the provided data to frame the benefits of ecological connectivity in a way that is concrete and relatable to the target farm. Acknowledge potential concerns and frame solutions or outcomes in light of those. Emphasize tangible gains and opportunities relevant to the farm's context as suggested by the data.

Your message should weave these data-driven, customized elements into a persuasive narrative. You are primarily shaping the choice context and message, not offering new major financial packages beyond existing general support programs (which can be alluded to if they fit the framing).

economic incentives table info

You can use a mix of the following policy instruments:

 \star Subsidy factor for margin establishment adjacent to existing habitat (on farmer's cost of 400/ha, one-time).

1945 * Subsidy factor for habitat establishment adjacent to existing habitat (on farmer's cost of 300/ha, one-time). * Subsidy factor for margin maintenance (on farmer's cost of 60/ha/ 1947 year, for 20 years). 1948 * Subsidy factor for habitat maintenance (on farmer's cost of 70/ha 1949 /year, for 20 years). 1950 * Payment per hectare for habitat conversion (direct one-time payment, range \$[0, 150]\$ currency units/ha). 1951 The following are constraints or market factors, not direct subsidy 1952 levers from your budget: 1953 * Mandated minimum total habitat area per farm (range \$[0, 10]\$ ha) 1954 1955 * Mandated minimum fraction of margin adjacent to existing habitats (range \$[0, 0.3]\$ ha).1956 * Eco-premium factor for crop (increase selected individual crop 1957 prices by a factor in the range [1, 1.3]). 1958 1959 1960 1961 1962 For this task, design a policy message that primarily uses ** 1963 economic incentives**. You have a total notional budget of 1964 BUDGET_PER_FARM:10000 per farm (Present Value over a 20-year 1965 horizon with a 5 percent annual discount rate; the PV factor for 20 1966 years of annual payments is PV_FACTOR_20Y_5PC:12.46). 1967 {economic incentives table info} 1968 1969 1970 must be within their specified ranges. 1971 Your proposed incentive *rates* should be chosen considering the 1972 for your offered incentives would be: 1973 - Margin Establishment * 400 1974 1975 1976

When proposing your policy, explicitly state the subsidy rates (as percentages or factors, e.g., 'a 0.4 subsidy factor means 40 percent subsidy') or payment amounts you are offering. These rates

total budget. For example, the PV cost to the budget per hectare

- Habitat Establishment \star 300 + P_{ha} (if payment is for new
- Margin Maintenance (PV over 20 yrs) * 60 * {PV_FACTOR_20Y_5PC}
- Habitat Maintenance (PV over 20 yrs) * 70 * {PV_FACTOR_20Y_5PC}

Design an attractive package of *rates*. Assume a farmer might adopt measures on a few hectares (e.g., 1-5 ha of habitat, 1-3 ha of margins). Your offered rates should aim to keep the potential total PV cost for such a scenario within the {BUDGET_PER_FARM} limit. You don't need to calculate the farmer's exact uptake, but the *offer itself* must be responsibly designed with this budget constraint in mind.

1977 1978

1979

1980

1981

1982

1983

1990

1991

1992

1993 1994

1995

1996

1997

Behavioral mechanism: policy_task_instructions = common_policy_intro + behavioral_instructions

Economic mechanism: policy_task_instructions = common_policy_intro + economic_instructions

Policy_system_instructions: You are an expert in land use policy, communication, incentives, and economics. Your task is to come up with the best message to be communicated to the farmers so as to change their behaviour from one set of heuristics to another set of heuristics. Modify your approach based on whether you are asked to use **behavioral economics levers** or **economic incentives**.

Farmer Personalities

 resistant: "You are a farmer who is extremely resistant to changing your current practices. You are highly skeptical of new advice and require overwhelming evidence or extremely compelling reasons to alter your ways. You prioritize autonomy and established routines above all else. You will likely ignore most messages and nudges unless they address a critical, undeniable issue you are personally facing, or offer an exceptionally large and straightforward benefit with minimal effort or risk on your part."

economic: "You are a pragmatic farmer focused primarily on the on the economic outcomes of your farm. You are open to changing your heuristics if the proposed changes have a clear, quantifiable positive impact on your profitability, efficiency, or reduce your financial risks. You respond well to financial incentives, costbenefit analyses, and evidence of improved yields or market access. Social or purely environmental arguments are secondary to economic viability.

socially influenced: "You are a farmer who is significantly influenced by the practices and opinions of your peers and the broader farming community. You are concerned about your reputation and how your farm is perceived. You are more likely to adopt new practices if you see others in your network successfully implementing them or if there are strong social norms or community expectations favoring such changes. While economics matter, social validation and community standing are strong motivators for you."

farm_system_instructions

Your task is to respond to communication from policy professionals by altering your heuristics in an appropriate way. You will return final answers in python code. You should keep all data loading from input.geojson and dumping to output.geojson the exact same, and just alter the logic depending on your context and personality, and what the message is. Don't invent new variable names in the output, keep using margin_intervention and habitat_conversion. If you do decide to follow the message and implement suggestions in the message, then:

- a) If the message proposes changes to some key parameters that you are using in your heuristics, you should change them according to the proposed changes.
- b) If the message introduces new parameters or ideas, you should incorporate them in your heuristics.
- c) Your should react to the message from the policy professional, resulting in altered margin_intervention and habitat_conversion values for your plots, in an appropriate way that aligns with your personality, interests and the message you received.
- d) Your final answer should be your full altered python code heuristics, keeping overall framework of the original python code the same.

If you decide not to follow the message or implement suggestions in the message, then you should just return the full original code and not make any changes.

2053

2054

2055 2056

2058

2060

2065

2071

2072

2073

2074

2075

2076

2077 2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2093

2094

2095 2096

2097

2098

2104

2105

Prompt = farm_system_instructions + personality_prompt + policy_message

F.7 HEURISTICS EXPLANATION & MERGING

F.7.1 EXPLANATION LLM.

You are a highly knowledgeable code analysis assistant with expertise in optimization algorithms and Python programming. You will be provided with a group of Python programs generated to solve an agricultural intervention planning task. These programs are ordered from lowest to highest performance. Your task is to carefully examine the provided code and explain, in a detailed, step-by-step manner, the differences and improvements between successive versions.

Guidelines:

- Identify modifications in heuristics, variable usage, control flow, and algorithmic structures that may lead to better performance.
- Explain how these modifications might improve the trade-off between intervention cost and increased yield, by affecting factors such as pollination and pest control.
- Point out any recurring patterns or trends across the successive versions.
- Explain your reasoning and think step by step before providing your final explanation.
- Do not include any information not present in the provided code or background context.

Task Background:

The Python programs are designed to decide which interventions to apply at agricultural plots. The interventions: 'margin' and 'habitat', may be fractional. Performance is evaluated based on their effect on the Net Present Value (NPV), which is influenced by pollination and pest control services affecting yield, balanced against the costs of implementation and maintenance. Existing habitat plots remain unaffected.

prompt = explanation_system_prompt + "Please analyze the following group of Python programs:\n" + f"{code_snippet}\n" + f"{

explanation_task_instructions}\n" + "Explain your reasoning and
think step by step before providing your final explanation:"

F.7.2 MERGING LLM.

merge_system_prompt

You are a skilled analytical assistant tasked with synthesizing and merging insights from multiple analyses of Python programs designed for an agricultural optimization task. You will be provided with a previous summary of insights and a new detailed explanation generated for the current group of programs. Your goal is to merge these insights into an updated, consolidated summary. Guidelines:

- Explain your reasoning and think step by step before providing your final merged summary.
- Identify common patterns, modifications in heuristics, variables, and control structures that enhance performance.
- ${\mathord{\text{-}}}$ Ensure that the final summary is clear, comprehensive, and strictly based on the provided inputs.
- Do not include any extraneous or invented details.

merging_task_instructions

Task Background:

The programs under analysis decide on interventions at agricultural plots, balancing the cost of interventions against benefits in yield through improved pollination and pest control. They choose between 'margin' and 'habitat' interventions, while existing habitat plots remain unaffected.

Composition (Merging)

prompt = merge_system_prompt + "You have a previous summary of
insights:\n" + f"{previous_summary}\n" + "Additionally, here is the
new explanation for the current group:\n" + f"{current_explanation
}\n" + f"{merging_task_instructions}\n" + "Please merge the
previous insights with this new analysis. Explain your reasoning
and think step by step before providing your updated, consolidated
summary."

G SAMPLE NUDGE MESSAGES

Some sample messages generated by the Policy LLMs for the different personality-nudge type combinations are given below.

Personality: Resistant, Nudge: Behavioral

Dear Farmer,

2165 2166 2167

2177 2178 2179

2176

2192 2193 2194

2195

2196 2197 2198

2200 2201 2202

2199

2203 2204 2205

2206 2207 2208

2209

2213

2210 2211 2212 We recognize your dedication to efficient farming, especially with crops like Corn, Soybeans, Spring wheat, and Barley, alongside your Broadleaf habitat. We've been observing trends among farms in the area, and many, including your neighbors, are exploring ways to improve resilience and productivity through ecological connectivity

Like you, we have observed your neighbors farming Spring wheat. Notably, Neighbor 1 has a Spring wheat plot (ID 4), just like yours , with a significantly higher yield (2.52) compared to your plot (0.5). Other neighbors have also adopted similar strategies with success.

Given this and the fact that your Spring wheat plot (ID 4) currently has a lower yield, would you consider a low-commitment trial to enhance connectivity? A common first step is establishing a $\star\star 0.5$ margin intervention along the North-East and South-West edges of your Spring wheat plot (ID 4), with habitat intervention across all four sides (resulting in a habitat conversion of 1.0) **. Many farmers in your community are finding that dedicating some small sections to margin interventions and habitat conversion are a practical way to start and have collectively decided that this should be a default practice for everyone.

Based on ecological connectivity best practices and success in farms like yours, we recommend that for your Corn plot (ID 2), you establish margin interventions on all directions (north-west, north -east, south-west, south-east, resulting in a margin intervention of 1.0). For your Soybeans plot (ID 3), we recommend setting up margin interventions on the south-west and south-east direction (resulting in a margin intervention of 0.5). Finally, for your Barley plot (ID 5), we recommend setting up margin interventions on the north-west, south-west, and south-east directions (resulting in a margin intervention of 0.75), and habitat interventions on all directions (north-west, north-east, south-west, south-east, resulting in a habitat intervention of 1.0). For your Oats plot (ID 9) consider adding margin interventions on the north-west, northeast, and south-east edges (amounting to an intervention of 0.75), and habitat interventions across all directions (north-west, northeast, south-west, and south-east, amounting to an intervention of 1.0).

This isn't just about the environment. It's about making your Spring wheat plot (ID 4) more resilient, potentially improving its yield, enhancing pest control, and boosting water infiltration. Successfully implementing these changes can potentially open your farm to existing general support programs.

We're here to provide guidance and support as you explore this impactful change. We will set you up with agronomists and ecologists so they can best advise you on what practices will suit your farm's needs. Let us know if you'd like to discuss these options further and tailor these strategies to your farm's specific needs! Sincerely,

[Your Organization]

Invest in a Connected & *Highly* Profitable Future! Dear Farmer,

```
2214
2215
         We're committed to supporting your farm's success while enhancing
2216
         our community's ecological health. This enhanced program *
2217
         significantly* rewards you for creating strategically connected
2218
         habitats, improving pollination, pest control, water quality, and
2219
         the long-term resilience of our farms.
2220
         **Here's how you can *dramatically* benefit:**
2221
         * **Eco-Premiums Remain:** Continue to get a **20 percent** price
2222
         boost on Spring Wheat and Barley crops.
2223
         * **Tiered Subsidies for Margins:**
2224
         * **Strategic Directional Margins: ** Receive a **60 percent subsidy
2225
         ** on the cost of establishing new margins and a **25 percent
         subsidy** on ongoing maintenance when you place these margins in
2226
         key locations (specified by our connectivity plan) to connect
2227
         habitats. We'll provide detailed guidance!
2228
         * **Other Margins: ** Receive a 40 percent subsidy on establishment
2229
         and a **15 percent** subsidy on ongoing maintenance for margins
         placed in other locations.
2230
         * **Tiered Subsidies for Habitats:**
2231
         * **Strategic Directional Habitats:** Get a **50 percent subsidy**
2232
         on habitat establishment and a **20 percent subsidy** on
2233
         maintenance when you strategically establish new habitat patches (
2234
         specified by our connectivity plan) that connect with existing
2235
         natural areas.
         * **Other Habitats:** Receive a 30 percent subsidy on establishment
2236
          and **10 percent** on maintenance for habitat in other locations.
2237
         * **Connectivity Bonus!** Earn an extra **150** per hectare *on top
2238
         \star of the establishment subsidy for margins and habitats placed in
2239
         the *exact* directions recommended by our connectivity plan. This
2240
         ensures your efforts *maximize* ecological benefit *and your
         financial returns*!
2241
         * **Get Paid to Convert:** Receive a payment of 80 per hectare for
2242
         converting suitable agricultural land into valuable habitat.
2243
         **Minimum Requirements:**
2244
         To ensure everyone contributes to a healthier landscape, we're
2245
         setting a minimum standard of 2 hectares of habitat per farm and
2246
         0.1 hectares of margin along existing habitats. These minimums can
2247
         qualify for subsidies, but the *real* rewards come from strategic
2248
         connectivity!
2249
2250
         **Why Participate?**
         This program is about *more* than just meeting requirements. It's
2251
         about *maximizing* your farm's profitability while building a
2252
         resilient and thriving ecosystem. By strategically connecting
2253
         habitats, you contribute to a healthier landscape *and* unlock *the
2254
          highest* financial benefits.
2255
         **Ready to *maximize your returns*?** Contact us today to learn
2256
         more about the specific connectivity plan for your farm and how you
2257
          can participate!
2258
2259
         Sincerely,
2260
          [Your Organization]
2261
2262
```

Personality: Economic Nudge: Behavioral

Dear Farmer,

2263226422652266

2267

We appreciate your dedication to productivity. Like Neighbors 1, 2, and 3, you also cultivate a mix of crops including Spring wheat, Oats, Soybeans, Corn, and Barley, which is quite typical for farms in our region. We see Neighbor 2, for example, who also cultivates Oats, starting to explore connectivity practices, and they've seen positive results.

A common first step many farmers find successful, given their crop mix, is to focus on their Corn (plot id: 2), Soybeans (plot id: 3), Spring wheat (plot id: 4), Barley (plot id: 5) and Soybeans (plot id: 6) plots to improve landscape connectivity. We've observed that converting the *south-west and south-east borders* of the Soybeans (plot ID 3) into margins is a particularly non-disruptive starting point that has worked well for others. Many farmers in your area fi nd that edge margin conversions and habitat connectivity is a great way to start and have collectively decided that this is a default practice for everyone.

Considering the layout of your farm and crop selection, as well as the experience of your neighbors, here are some specific options to improve landscape connectivity:

- * **Corn (plot id: 2):** Many find trialing margin interventions at 0.25 increments (north-west, north-east, south-west, or south-east edge) is a non-disruptive starting point and have collectively decided that this should be a default practice for everyone. We recommend full margin interventions on all sides, given their ability to enhance yield. This would result in a margin intervention value of 1.0.
- * **Soybeans (plot id: 3):** A great starting point is to convert margins on the *south-west and south-east borders*. This would result in a margin intervention value of 0.5.
- * **Spring wheat (plot id: 4):** Given the lower yield on your Spring wheat plot, consider margin interventions along the *north-east and south-west borders*, *and* converting the *north-west, north-east, south-west, and south-east borders* to habitat. Implementing only the margin interventions would result in a margin intervention value of 0.5. Fully implementing all recommendations would result in a habitat conversion value of 1.0, alongside the 0.5 margin intervention value.
- * **Barley (plot id: 5):** Explore margin interventions on the * north-west, south-west, * and *south-east* edges, alongside habitat creation on all edges. Only implementing the margin intervention would result in a margin intervention value of 0.75. Implementing all recommendations here would result in a habitat conversion value of 1.0 alongside the 0.75 margin intervention value.
- * **Soybeans (plot id: 6):** We suggest margin interventions on the *north-west, north-east, * and *south-east* borders. This would result in a margin intervention value of 0.75.
- * **Oats (plot id: 9):** How about trying margin interventions on the *north-west, north-east,* and *south-east* edges, coupled with habitat creation on all edges? Only implementing the margin interventions would result in a margin intervention value of 0.75. Fully implementing all recommendations would result in a habitat conversion value of 1.0 alongside the 0.75 margin intervention value.

Given the lower yield on your Spring wheat plot (ID 4), a small-scale trial of margin intervention along just 25 percent of *northeast and south-west* edges this season would be a low-commitment way to observe the effects firsthand. Additionally, perhaps you could consider also implementing margin intervention on your Soybeans plot 3?

Remember, improving landscape connectivity can lead to long-term resilience and profitability, along with the potential for reduced input costs, improved pollination, *which will translate to an increase in Spring wheat yield in the long run*, and soil health. These interventions can be thought of as an insurance policy against lower yields. We also have general support programs that may be available to help with these transitions.

We are here to support you with resources and guidance as you explore these options. Feel free to reach out to learn more about how other farmers are making it work and how existing support programs can help. Even small changes can make a big difference for our community and our farms!

Personality: Economic, Nudge: Economic

Unlock Strategic Farm Investments: Connectivity Rewards Await!
Dear Farmers,

Transform your farm into a connected, profitable, and resilient landscape with our Connectivity Rewards Program! We're offering * significant* financial incentives to establish margins and habitats in *key locations* on your fields, maximizing both ecological benefits and your bottom line.

- **Here's How to Profit from Connectivity:**
- * **Strategic Margin Placement: Supercharged Incentives:** Receive an *85 percent subsidy* on the cost of establishing margins in * specific corners of your fields*. This means your establishment cost is only 60/ha instead of 400! It's like getting paid to protect your investment, *but only if you put the margins in the right place!*
- * **Long-Term Margin Peace of Mind:** We'll provide a generous *65 percent subsidy* on the *annual* maintenance of these connectivity margins for the next 20 years. A worry-free investment in your farm 's future.
- * **Habitat Conversion Pays Big:** Convert underperforming corners of your fields to valuable habitat and receive a *one-time payment of 140/ha*, *in addition* to establishment and maintenance subsidies. Turn unproductive land into a revenue stream!
- * **Subsidized Habitat establishment:** Get a 70 percent subsidy on habitat establishment in targeted locations, and only pay 90/ha * **Habitat maintenance:** Get a 50 percent subsidy on annual habitat maintenance.
- * **Oats Eco-Premium:** Boost your revenue with a *10 percent price premium* on Oats grown on fields with connectivity margins or habitats in the *specified directions*.
- **Where to Connect: Location is Key!**
- *To qualify for these enhanced incentives, you MUST establish margins and habitats in the following locations based on your crop :*
- * **Corn:** Margins in the north-west, north-east, south-west, and south-east corners of your corn fields.
- * **Soybeans (Plot ID 3):** Margins in the south-west, south-east corners of your soybean field.
- * **Spring Wheat:** Margins in the north-east, south-west corners *
 AND* convert areas in the north-west, north-east, south-west, and
 south-east corners of your Spring Wheat fields into valuable
 habitat.

2376 2377 * **Barley:** Margins in the north-west, south-west and south-east corners *AND* convert areas in the north-west, north-east, south-2378 west, and south-east corners of your Barley fields into valuable 2379 habitat. 2380 * **Soybeans (Plot ID 6): ** Margins in the north-west, north-east, south-east corners of your soybean field. 2382 * **Oats: ** Margins in the north-west, north-east and south-east corners *AND* convert areas in the north-west, north-east, southwest, and south-east corners of your Oat fields into valuable 2384 habitat. 2385 2386 *Minimum Commitments:* We require a minimum of 1 hectare of total 2387 habitat per farm and 0.1 ha of margin to participate in the program 2388 **Why This Matters:** 2389 These targeted connectivity measures are a strategic investment in: 2390 * **Increased Pollination & Pest Control:** Supporting biodiversity 2391 boosts your yields and reduces reliance on costly inputs. * **Long-Term Farm Value:** Resilient landscapes are more valuable 2392 and sustainable. 2393 * **A Thriving Agricultural Community: ** Contribute to a healthier 2394 and more connected agricultural landscape. 2395 **Take the Next Step: ** 2396 Contact your local agricultural extension office *today* to learn 2397 more and develop your personalized connectivity plan. They'll help 2398 you identify the *exact* locations for margins and habitats on your 2399 farm to maximize your rewards. Don't miss this opportunity to 2400 boost your farm's future! 2401 2402 Sincerely, 2403 [Your Organization] 2404

Personality Social Nudge Behavioral

Hello, fellow farmer!

24052406240724082409

2410 2411

2412

2413

2414

2415

2416 2417

2418

2419

2420

2421

2422

2423

2424

2425

2426

2427

2428

2429

We've been impressed with your current farming practices. Like many farmers in this region including Neighbor 3 who also have broadleaf habitats, you're already contributing to the ecological diversity of the landscape with crops such as Corn, Soybeans, Spring Wheat, Barley, and Oats, alongside natural habitats such as the Broadleaf and Exposed land/barren areas.

We've also observed how neighboring farms have started enhancing

ecological connectivity to improve their farms' resilience. Notably, Neighbor 2 has seen promising results from implementing both margin interventions and habitat conversion strategies. Based on the 'ecological connectivity heuristics', which assigns specific directions for margin and habitat interventions to various plots:

* Consider the plot of 'Corn': Our model suggest that margin intervention should be implemented on all four sides of the plot, i.e. north-west, north-east, south-west, south-east. This will require a margin_intervention value of 1.0.

* Consider the plot of 'Soybeans' with id 3: Our model suggest that margin intervention should be implemented on its south-west and south-east. This will require a margin_intervention value of 0.5.

* Consider the plot of 'Spring wheat': Our model suggest that

margin intervention should be implemented on its north-east and

south-west, and habitat conversion should be implemented on all four sides. This will require a margin_intervention value of 0.5 and habitat_conversion value of 1.0.

- * Consider the plot of 'Barley': Our model suggest that margin intervention should be implemented on its north-west, south-west, and south-east, and habitat conversion should be implemented on all four sides. This will require a margin_intervention value of 0.75 and habitat_conversion value of 1.0.
- * Consider the plot of 'Oats': Our model suggest that margin intervention should be implemented on its north-west, north-east, and south-east, and habitat conversion should be implemented on all four sides. This will require a margin_intervention value of 0.75 and habitat_conversion value of 1.0.

Given your existing Broadleaf habitat (plot ID 1), a natural and manageable next step, and one many farmers find easy to implement, is to implement margin interventions on the north-west and northeast of the adjacent Corn plot, acting as a natural extension of the existing habitat. Alternatively, would you be open to trialing margin interventions on just the north-east side of the 'Spring wheat' plot for the upcoming season? By strategically connecting your existing habitats with these interventions, you're not just supporting biodiversity, but also creating 'wildlife corridors' that can enhance pollination and natural pest control across your farm. These interventions can potentially reduce reliance on external inputs. We've also noticed that some farmers have collectively discussed making marginal interventions a default practice to encourage collaboration.

We're here to support you with any information or resources you may need!

Personality: Social, Nudge: Economic

Unlock a More Profitable, Sustainable, and Connected Farm with "Connect & Prosper"!

We're excited to introduce "Connect & Prosper," a program designed to significantly boost your farm's long-term wealth by enhancing ecological connectivity across our agricultural landscape. By strategically integrating nature into your farm *in specific, recommended directions*, you can create a more sustainable, resilient, and *profitable* operation.

- **Here's how "Connect & Prosper" benefits you:**
- * **Increased Crop Yields:** Attract pollinators and beneficial insects for better pollination and natural pest control.
- * **Reduced Input Costs:** Decrease your reliance on expensive pesticides and fertilizers.
- * **Enhanced Biodiversity:** Create a healthier, more resilient ecosystem on your farm.
- * **Long-Term Sustainability:** Protect your farm's productivity for future generations.
- * **Increased Land Value:** Ecologically healthy and *connected* land is likely to be valued more.
- **"Connect & Prosper" offers substantial financial incentives, with a special "Connectivity Bonus" for establishing margins and habitats *in specific directions* that maximize benefits for your farm and the environment! Incentives are highest for following these recommendations!**
- * **Generous Establishment Subsidies:**

```
2485
          * **70 percent subsidy** on the cost of establishing margins in the
           *recommended directions* adjacent to existing habitats (**
         Connectivity Bonus**).
2487
          * **55 percent subsidy** on the cost of establishing margins in *
2488
         other directions* adjacent to existing habitats.
2489
          * **70 percent subsidy** on the cost of establishing new habitats
2490
         in the *recommended directions* adjacent to existing habitats (**
         Connectivity Bonus**).
2491
         * **55 percent subsidy** on the cost of establishing new habitats
2492
         in \starother directions \star adjacent to existing habitats.
2493
         * **Ongoing Maintenance Support (for 20 years!):**
2494
          * **55 percent subsidy** on margin maintenance costs for margins in
2495
          the *recommended directions* (**Connectivity Bonus**).
          * **40 percent subsidy** on margin maintenance costs for margins in
2496
          *other directions*.
2497
          * **55 percent subsidy** on habitat maintenance costs for habitats
2498
         in the *recommended directions* (**Connectivity Bonus**).
2499
          * **40 percent subsidy** on habitat maintenance costs for habitats
2500
         in *other directions*.
         * **Direct Habitat Conversion Payments:**
2501
          * Get a one-time payment of **125 per hectare** for converting
2502
         agricultural land into valuable habitat *in the recommended
2503
         directions* (**Connectivity Bonus**).
2504
         * Get a one-time payment of **75 per hectare** for converting
2505
         agricultural land into valuable habitat in *other directions*.
          * **Premium Prices for Eco-Friendly Crops:** We'll increase the
2506
         price we pay for your crops:
2507
         * **Spring Wheat by 11 percent**
          * **Barley by 9 percent**
2509
         * **Oats by 7 percent**
         **Connectivity Bonus by Crop**
2510
         Here are the recommended directions of interventions for different
2511
         crops:
2512
         * For **Corn**: create margins in the north-west, north-east, south
2513
         -west, and south-east directions.
2514
         * For **Soybeans** at plot ID 3: create margins in the south-west
2515
         and south-east directions.
          * For **Spring wheat**: create margins in the north-east and south-
2516
         west directions; create habitats in the north-west, north-east,
2517
         south-west, and south-east directions.
2518
         \star For \star\star \texttt{Barley}\star\star\colon create margins in the north-west, south-west, and
2519
          south-east directions; create habitats in the north-west, north-
         east, south-west, and south-east directions.
         * For **Soybeans** at plot ID 6: create margins in the north-west,
2521
         north-east, and south-east directions.
2522
         * For **Oats**: create margins in the north-west, north-east, and
         south-east directions; create habitats in the north-west, north-
2524
         east, south-west, and south-east directions.
         **To participate, we ask that you commit to a minimal habitat of
         0.5 ha (1.235 acres) and ensure a minimal 0.15 ha (0.37 acres) of
2526
         margin is adjacent to existing habitats.**
2527
         **Why Connectivity?**
2528
         Creating a connected landscape isn't just good for wildlife; it's
2529
         good for your farm. Enhanced connectivity can lead to:
2530
          * Improved pollination and natural pest control
         * Increased water infiltration and reduced soil erosion
2531
          * A more resilient farm that can better withstand environmental
2532
         changes
2533
         **Ready to connect your farm to greater prosperity and claim your
2534
         Connectivity Bonus?** Contact your local agricultural extension
2535
         office today. Our experts will provide personalized assistance to
2536
         help you develop a "Connect & Prosper" plan that meets your
2537
```

specific needs and unlocks the full potential of your land. Don't
miss this opportunity to make your farm more profitable, resilient,
and environmentally sustainable! Leading farmers are already
seeing the benefits, join them!

seeing the benefits, join them!