

000 MATHMO: AUTOMATED MATHEMATICAL MODELING 001 002 THROUGH ADAPTIVE SEARCH 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Mathematical modeling is the process of understanding and predicting complex
011 real-world phenomena. Traditionally, it is a time-intensive effort reliant on deep hu-
012 man expertise and iterative refinement. Automating this intricate process, therefore,
013 offers the potential to significantly accelerate discovery and broaden the application
014 of mathematical modeling across diverse domains. Such automation, however,
015 must address inherent challenges, including fundamental modeling uncertainty, bal-
016 ancing multiple conflicting objectives, and incorporating subjective qualities into
017 assessing model utility. We approach this by conceptualizing mathematical model-
018 ing as a sequential decision-making problem under uncertainty. In response, we
019 introduce MATHMO, a novel adaptive search method designed to automatically navi-
020 gate the complex decisions in selecting mathematical frameworks, specifying model
021 formulations, and defining algorithmic procedures. Specifically, MATHMO employs
022 a principled bi-level search strategy—combining high-level exploration across
023 diverse frameworks and local intra-framework model refinements—leveraging
024 Large Language Models for exploration, surrogate evaluations, and incorporating
025 subjective preferences into the automated process. We demonstrate MATHMO’s
026 efficacy on diverse real-world tasks, where it successfully discovers Pareto-efficient
027 frontiers of models that balance varied objectives, including subjective criteria.

028 029 1 INTRODUCTION 030

031 Mathematical modeling is the art and science of translating complex real-world phenomena into
032 precise mathematical language, allowing us to represent, understand, and predict complex situations.
033 This capability is crucial in almost all aspects of life, from natural sciences and engineering to
034 economics and social systems (Turing, 1990; Sugihara and May, 1990; Banwarth-Kuhn and Sindi,
035 2020). Indeed, the ability to abstract complex realities into mathematical models is often considered
036 a key feature of intelligent civilization, enabling generalized problem-solving, knowledge transfer,
037 and accumulation of scientific understanding over time (Simon, 2019).

038 The time-consuming and expertise-driven nature of modeling, coupled with its inherent complexities,
039 makes its automation a highly appealing prospect. Automating this process could democratize access
040 to powerful analytical models and tools, efficiently exploring trade-offs between multiple models,
041 uncovering novel approaches or insights, and enhancing decision-making across diverse domains.

042 Several key characteristics distinguish the mathematical modeling process. Firstly, it is pursued in
043 the face of fundamental *uncertainty*: the optimal framework or model specification is often unknown
044 a priori, demanding iterative exploration involving building, testing, and refining models based
045 on feedback (Jakeman et al., 2006). Secondly, modeling frequently contends with *multiple, often*
046 *conflicting, objectives* (e.g., an optimization model’s solution quality versus runtime (Chandrasekaran
047 and Jordan, 2013)). Thus, the aim is typically not a single ‘best’ model, but a diverse frontier of models
048 representing different trade-offs for the human modeler to investigate. Lastly, *subjective qualities* like
049 interpretability and domain understanding further influence model utility beyond objective metrics
050 (Dirac, 1963), making it crucial for them to be captured in the automated modeling process.

051 We conceptualize automated mathematical modeling as a sequential decision-making problem. Here,
052 the modeler makes a series of choices: (1) the mathematical *framework* to employ, (2) the concrete
053 *model specification* including appropriate representations, parameterizations and assumptions, and
(3) the appropriate *computational algorithms* to obtain the desired outputs. The uncertainty lies in

054 not knowing which modeling decisions will yield useful models, thus requiring a principled approach
 055 of adaptive search that carefully balances exploration and exploitation given uncertainty.
 056

057 In response, we introduce MATHMO, a novel method designed to automate key aspects of the mathematical
 058 modeling pipeline. At a high level, our system takes a modeling problem description and
 059 a set of objective functions, and discovers a set of models that represent efficient trade-offs among
 060 these objectives. The adaptive search procedure in MATHMO operates under a *bi-level structure*. The
 061 upper-level performs adaptive resource allocation across different mathematical frameworks, while
 062 the low-level module initiates a local search mechanism to explore each framework’s model and
 063 algorithm space. Crucial to this process are *Large language models* (LLMs), which are employed to
 064 sample realizations from the search space and to perform surrogate evaluations to improve search efficiency.
 065 Additionally, they are used as models of subjective evaluations, thus incorporating subjective
 066 model preferences into the modeling pipeline. In each round, the generated model and its evaluations
 067 are observed, informing and adapting the next iteration of modeling decisions.

068 **Contributions.** The primary contributions of this work are threefold: **(1)** We formally define automated
 069 mathematical modeling, conceptualizing it as a sequential decision-making problem under
 070 uncertainty. **(2)** We present MATHMO, an adaptive search framework designed for this problem,
 071 capable of efficient exploration, balancing multiple objectives, and incorporating subjective modeling
 072 preferences. To the best of our knowledge, this is the first work to address this exciting problem
 073 area. **(3)** We demonstrate the efficacy of MATHMO on four diverse real-world modeling tasks (two
 074 prescriptive, two predictive), demonstrating its ability to discover Pareto-efficient frontiers of models.

075 2 PRELIMINARIES

076 2.1 FORMALISM

077 Mathematical modeling is primarily a *declarative* endeavor (Van Roy and Haridi, 2004). The core
 078 cognitive task involves translating a problem $p \in \mathcal{P}$ (potentially accompanied by a dataset \mathcal{D}_p) into a
 079 suitable formal representation, where the subsequent derivation of mathematical outputs is delegated
 080 to computational tools. We formalize this process of representation and derivation as a sequential
 081 decision-making problem under uncertainty (c.f. “Box’s Loop” (Box, 1979)):

- 082 1. Selecting a high-level approach or *framework* $f \in \mathcal{F}$, where the chosen framework encodes
 083 foundational assumptions and provides access to specialized mathematical tools and techniques.
- 084 2. Specifying an exact *model* $m \in \mathcal{M}(f)$ within this framework, which is a concrete mathematical
 085 formulation representing the system under study.
- 086 3. Developing a computational or *algorithmic procedure* $a \in \mathcal{A}(f, m)$ to this model to derive the
 087 desired *mathematical output* $o = a \circ m(\mathcal{D}_p) \in \mathcal{O}$.
- 088 4. Evaluating the mathematical output or model characteristics to inform subsequent refinements.

089 This sequence of modeling decisions, denoted by (f, m, a) , defines a structured search space $\Omega =$
 090 $\{(f, m, a) \mid f \in \mathcal{F}(d), m \in \mathcal{M}(f), a \in \mathcal{A}(f, m)\}$. We evaluate each modeling outcome using
 091 a vector of $k \in \mathbb{N}$ *objective functions*, $\mathcal{J}(m, a) = [\mathcal{J}_1(m, a), \dots, \mathcal{J}_k(m, a)]^T$. For generality,
 092 each objective $\mathcal{J}_i(\cdot, \cdot)$ can depend on both the mathematical outputs o (e.g., solution optimality)
 093 and characteristics of the model-output pair (e.g., runtime). The goal is then to identify modeling
 094 decisions $(f, m, a) \in \Omega$ that address the following multi-objective optimization problem:

$$095 \text{minimize}_{(\cdot, m, a) \in \Omega} \quad \mathcal{J}(m, a) = [\mathcal{J}_1(m, a), \dots, \mathcal{J}_k(m, a)]^T \quad (1)$$

096 As there typically does not exist a single model-algorithm pair that can minimize all objective
 097 functions simultaneously, we are more interested in finding the Pareto optimal models, i.e., models
 098 that cannot be improved in any of its objectives without degrading at least one. Mathematically, a
 099 Pareto optimal pair (m, a) is non-dominated, where a pair (m, a) is said to dominate another pair
 100 (m', a') if $\forall i \in [k], \mathcal{J}_i(m, a) \leq \mathcal{J}_i(m', a')$ and $\exists i \in [k], \mathcal{J}_i(m, a) < \mathcal{J}_i(m', a')$ (Miettinen, 1999).

101 **Key challenges.** This problem definition emits several noteworthy challenges:

- 102 1. **Efficient exploration under uncertainty.** A defining hallmark of the mathematical modeling
 103 process is the fundamental uncertainty. It involves making sequential choices (framework, model,
 104 algorithm) with uncertain outcomes in an interconnected and complex space. Efficient exploration,
 105 informed by prior beliefs and feedback, is thus crucial to navigating this search space.




Figure 1: **Overview of MATHMO.** Given a modeling problem description, MAMO employs a bi-level adaptive search strategy to identify a Pareto set of models presenting diverse trade-offs.

2. **Fundamental trade-offs.** Modeling inherently involves balancing conflicting objectives, such as accuracy and interpretability. Different modeling frameworks (e.g., deep learning versus linear models) often embody fundamentally distinct trade-off frontiers, necessitating exploration across frameworks, beyond locally within a framework, to identify a set of Pareto efficient models.
3. **Subjective qualities of models.** Beyond objective metrics, subjective qualities like Occam’s Razor, interpretability, or alignment with domain understanding are integral to a model’s utility.¹ Although mathematical proxies for these elusive qualities exist (e.g., sparsity, minimum description length), they are typically framework-specific and not directly comparable.

2.2 RELATED WORKS

Our work builds upon and extends several lines of work:

AutoML. The field of AutoML aims to automate applying machine learning to real-world problems (Thornton et al., 2013). This broad endeavor encompasses hyperparameter optimization (Snoek et al., 2012; Li et al., 2018), neural architecture search (Zoph and Le, 2016; Pham et al., 2018), automated feature engineering (Khurana et al., 2016), and discovering loss functions (Real et al., 2020) or optimization algorithms (Chen et al., 2023b). Typically, these approaches predefine a custom search space (sometimes referred to as a domain-specific language), and apply search techniques such as Bayesian optimization, evolutionary algorithms, or bandit-based search (He et al., 2021). Our work shares this automation goal but differs significantly in scope, focusing on constructing mathematical models and their algorithmic procedures, navigating a more complex, open-ended search space than the narrower, pre-defined search spaces targeted by conventional AutoML.

The advent of LLMs has presented new opportunities for automated search problems. These large-scale pretrained models function as highly flexible generators, enabling search over problem spaces expressible through natural language (Brown et al., 2020) and overcoming bottlenecks in representation and search. They have been employed as zeroth-order optimizers over numerical spaces (Yang et al., 2024) (e.g., for hyperparameter optimization (Liu et al., 2024)). More strikingly, LLMs show remarkable efficacy in symbolic and combinatorial domains, in search spaces of reward functions (Ma et al., 2024), neural architectures (Chen et al., 2023a), symbolic expressions (Liu et al., 2025; Shojaee et al., 2025), and algorithms (Romera-Paredes et al., 2024).

LLMs for mathematical formulations. Our work is most closely related to emerging research using LLMs to generate mathematical formulations. Here, “formulation” means specifying a model within a predefined framework. Instances include formulating statistical models (Li et al., 2024), game-theoretic models (Mensfelt et al., 2024), dynamic systems (Holt et al., 2024), and convex optimization models (Ahmaditeshni et al., 2024). These typically assume a known modeling approach. Our work generalizes this line of research by not assuming a predefined framework. Instead, it automatically searches across diverse mathematical frameworks, removing a priori assumptions and enabling efficient exploration of diverse trade-offs offered by fundamentally different frameworks (e.g., metaheuristic optimization vs. convex optimization). Furthermore, our work uniquely incorporates evaluation of subjective qualities, crucial to the modeling process.

3 PROPOSED FRAMEWORK

The space of potential mathematical models for any given problem is inherently vast and complex. Navigating this nested, heterogeneous space with a flat exploration strategy is prone to inefficiencies.

¹“It is more important to have beauty in one’s equations than to have them fit experiment”—Paul Dirac

162 We propose an adaptive search framework that exploits structure in the modeling process to decompose
 163 this complexity, with the aim of improving efficiency (Dempe, 2002).
 164

165 **3.1 BI-LEVEL ADAPTIVE SEARCH**
 166

167 Our approach is informed by key observations about mathematical modeling. Firstly, the set of viable
 168 high-level modeling frameworks for a given problem is typically much smaller than the vast space
 169 of concrete models and algorithmic instantiations. This allows for more readily applicable priors
 170 on framework-level performance characteristics and suitability, for instance, incorporating coarse
 171 priors on the trade-off between performance and interpretability for deep learning versus mechanistic
 172 models. Secondly, performance variations between frameworks generally dominate those within
 173 them. The fundamental trade-offs offered by an exact method (e.g., integer programming) versus an
 174 approximate one (e.g., a metaheuristic) are typically more significant than those between different
 175 formulations under the same integer programming paradigm.

176 Based on these insights, we introduce a bi-level separation in the search process. The **upper-level**
 177 **search** explores different modeling frameworks, while the **lower-level search** focuses on discovering
 178 effective model formulations and solver/algorithm designs within that chosen framework. This
 179 bi-level separation, which mirrors the cognitive workflow often employed by human modelers, is
 180 expected to confer several advantages. Explicitly separating framework-level decisions allows for
 181 more effective exploration of model trade-offs. Different frameworks often populate distinct regions
 182 of this frontier, and the bi-level formulation helps systematically identify such trade-offs. Furthermore,
 183 modeling choices within a given framework tend to be structurally similar. This relative homogeneity
 184 means that feedback from one model instance provides a stronger signal for guiding improvements to
 185 related formulations within the same framework. For instance, insights from incorporating a logistic
 186 growth term into one dynamical system model are more directly transferable to refining another
 187 dynamical system’s parameters than to designing the neural architecture of a deep model.

188 **3.2 FORMAL DESCRIPTION**
 189

190 Mirroring the sequential nature of modeling decisions, our framework is formalized as an adaptive
 191 search process. The search proceeds in iterations indexed by $t = 1, 2, \dots, T$. At each iteration t ,
 192 decisions are informed by the history of previously explored models and their evaluated performance.
 193 Let $\mathcal{S}_{t-1} = \{(m_{t'}, a_{t'}, r_{t'}) | t' < t\}$ denote this history, where, for notational simplicity, we represent
 194 the objective value as $r_{t'} = \mathcal{J}(m_{t'}, a_{t'}) \in \mathbb{R}^k$. For the history within a particular framework f , we
 195 define $\mathcal{S}_{t-1}^f \subseteq \mathcal{S}_{t-1}$. The iterative process is decomposed into two nested levels:
 196

197 **Upper-level problem.** At each iteration t , the upper-level decision involves selecting a modeling
 198 framework $f_t \in \mathcal{F}(p)$. This selection is guided by past performance across all explored frameworks:

$$f_t = \arg \max_{f \in \mathcal{F}(p)} \alpha(f; \mathcal{S}_{t-1}) \quad (2)$$

201 Here, α is a scalar *utility function* that estimates the potential value to explore framework f at time t ,
 202 given \mathcal{S}_{t-1} . This function, by quantifying preferences over frameworks, is crucial for managing the
 203 exploration-exploitation trade-off (Jones et al., 1998; Srinivas et al., 2010). For instance, α might
 204 prioritize frameworks that have recently yielded high-performing models (exploitation) or those less
 205 explored that could unveil novel regions of the Pareto frontier (exploration).

206 **Lower-level problem.** Once a framework f_t is selected by the upper level, the lower-level problem
 207 focuses on identifying a new model pair (m_t, a_t) within the space of $\mathcal{M}(f_t) \times \mathcal{A}(f_t, m)$. This local
 208 exploration leverages the historical performance of evaluated pairs within the framework:
 209

$$(m_t, a_t) = \arg \max_{(m, a) \in \mathcal{M} \times \mathcal{A}} \beta_{f_t}(m, a; \mathcal{S}_{t-1}^{f_t}) \quad (3)$$

212 Here, β_{f_t} is a framework-specific utility function that learns from evaluations of past pairs $\mathcal{S}_{t-1}^{f_t}$
 213 explored within framework f_t to estimate the potential value of a candidate pair (m, a) . Together,
 214 Equations (2) and (3) define an iterative loop that systematically explores the model space, leveraging
 215 the bi-level structure to balance broad exploration across frameworks with focused refinement within
 them. Once (f_t, m_t, a_t) are obtained, they are then evaluated to obtain r_t , and added to the history.

216 4 MATHMO: AUTOMATED MATHEMATICAL MODELING WITH LLMs
217218 In what follows, we describe the Automated Mathematical Modeler (MATHMO), our specific implementation
219 of the adaptive search framework (for an algorithmic overview, see Section D.1).
220221 4.1 LLM SEARCH OPERATORS
222223 Conventional automated search methods typically necessitate a clearly defined search space and a
224 formal solution representation, often through a domain-specific language (DSL) (Hutter et al., 2011;
225 2019). The domain of general mathematical modeling, however, presents a significant challenge:
226 the space of potential mathematical objects is extraordinarily vast and diverse, rendering the a priori
227 definition of a comprehensive DSL or a fully structured search space practically infeasible.
228229 To navigate this expansive and ill-defined landscape, MATHMO leverages Large Language Models
230 (LLMs) as core search operators. Modern LLMs, pre-trained on massive corpora of text and
231 code, encapsulate extensive knowledge across numerous domains, including mathematics, scientific
232 literature, and programming (Brown et al., 2020; Kaplan et al., 2020). This pre-training endows them
233 with strong implicit domain priors, which can be harnessed to guide the exploration of plausible
234 and potentially effective mathematical modeling choices, moving beyond the limitations of rigidly
235 defined search spaces. In MATHMO, LLMs fulfill two crucial roles in the bi-level search process:236 1. **Generative samplers.** LLMs are employed to sample from the space of frameworks, as well as
237 model and algorithmic specifications. In our implementation, specific models and algorithms are
238 represented as executable Python code, while high-level frameworks are represented as textual
239 descriptions. Conditioned on the problem description p , LLMs are prompted to sample suitable
240 modeling frameworks, denoted as $f \sim p_\theta(\cdot | p)$. For a selected framework f , the LLM generates
241 specific model and algorithmic instantiations, i.e., $(m, a) \sim p_\phi(\cdot \cdot | p, f, \mathcal{S}^f)$, conditioned on the
242 problem, framework, and past examples belonging to that framework.²
243 2. **Surrogate models.** LLMs also function as surrogate models to estimate the objective value of
244 proposed model-algorithm pairs and inform the utility functions to guide search. Specifically,
245 $\hat{r} \sim p_{\text{SM}}((m, a) | p, f, \mathcal{S}^f)$, where the subscript SM is employed to denote the surrogate model.
246 Additionally, we also use LLMs as Surrogate Models Of Subjective Evaluations (MOSE). This
247 surrogate, i.e., $\hat{r} = p_{\text{MOSE}}((m, a) | p)$, predicts subjective quality scores (e.g., human-perceived
248 interpretability) based on model representation and output, which are integrated into the overall
249 evaluation, allowing subjective inductive biases to be incorporated into search.250 The operation of LLMs in these roles relies on specific prompts. The details are provided in Section D,
251 but they follow a standard “skeleton” structure, incorporating the problem description p , current
252 context (e.g., selected framework f , and history \mathcal{S}), and the specific task for the LLM.
253254 4.2 UPPER-LEVEL PROBLEM: FRAMEWORK SELECTION
255256 The upper-level problem addresses the decision of which modeling framework f_t to commit to for
257 an additional step of exploration. In MATHMO, we employ the Pareto Upper Confidence Bound
258 (Pareto-UCB) strategy to realize the framework selection utility $\alpha(f, \mathcal{S}_{t-1}^f)$ (Equation (2)). This
259 method navigates the inherent multi-objective trade-offs by identifying a frontier of frameworks that
260 are optimistically non-dominated, thus balancing the need to explore new avenues with exploiting
261 proven ones (Drugan and Nowe, 2013; Xu and Klabjan, 2023).262 At the beginning of search ($t = 0$), an initial set of candidate frameworks is proposed by employing
263 the LLM-based sampler $f \sim p_\theta(\cdot | p)$. Each framework is initialized with an optimistic estimate
264 of its potential, by setting an infinite upper confidence bound (UCB) value. This ensures that
265 each framework is selected for at least one initial exploration cycle, providing data for subsequent,
266 more informed decisions. Specifically, for each framework f , the historical performance vectors
267 $\{r_{t'} | (m_{t'}, a_{t'}, r_{t'}) \in \mathcal{S}_{t-1}^f\}$ are used to estimate the empirical mean $\hat{\mu}_f \in \mathbb{R}^k$ and variance $\hat{\sigma}_f^2 \in \mathbb{R}^k$.
268 This is then used to calculate the UCB vector $\text{UCB}_f \in \mathbb{R}^k$, where each component is computed using
2692²Here ϕ, θ are used to denote the prompts that module that sampling distribution (Sumers et al., 2023).

270 a formula that considers both the estimated mean and its uncertainty, encouraging exploration:
 271

$$272 \quad \text{UCB}_{f,j} = \hat{u}_{f,j} + c \sqrt{\frac{\hat{\sigma}_{f,j}^2 \ln(N_{t-1})}{N_{f,t-1}}} + d \sqrt{\frac{\ln(N_{t-1})}{N_{f,t-1}}} \quad (4)$$

$$273$$

$$274$$

275 where $\hat{u}_{f,j}$ and $\hat{\sigma}_{f,j}^2$ are the estimated mean and variance of objective j of framework f , $N_{f,t-1}$ is
 276 the number of times framework f has been evaluated, and N_{t-1} is the total number of exploration
 277 steps (across all frameworks). c and d are hyperparameters that control the exploration bonus.

278 The set of UCB_f vectors, one for each framework, is then used to identify a subset of the promising
 279 frameworks. A framework f is considered part of the Pareto optimal set, if its UCB_f is non-dominated
 280 (i.e., $\text{UCB}_{f,j} \geq \text{UCB}_{f',j} \forall j \in [k]$ and $\exists j \in [k] : \text{UCB}_{f,j} > \text{UCB}_{f',j} \forall f'$, assuming maximization).
 281 From this Pareto-UCB set, one framework is randomly selected to be explored in the next iteration.
 282

283 4.3 LOWER-LEVEL PROBLEM: LOCAL EXPLORATION

$$284$$

285 Once the upper-level process selects a framework f_t , the lower-level is concerned with performing
 286 local exploration within the chosen framework to identify promising (m_t, a_t) for subsequent eval-
 287 uation. MATHMO achieves this through a three-stage process that involves first sampling candidate
 288 pairs, performing surrogate evaluations, and finally selecting the most promising one based on these
 289 predictions, in a process akin to Bayesian Optimization (Snoek et al., 2012; Liu et al., 2024).

290 First, a set of diverse candidate model-algorithm pairs are sampled, which we denote as $\tilde{\mathcal{S}}^f =$
 291 $\{(\tilde{m}^{(i)}, \tilde{a}^{(i)}) \mid i \in [l]\}$, where each $(\tilde{m}^{(i)}, \tilde{a}^{(i)}) \sim p_\phi(\cdot, \cdot \mid p, f_t, \mathcal{S}_{t-1}^{f_t})$. For each sampled candidate
 292 pair, we then estimate its k -dimensional objective vector using LLMs as a surrogate model: $\hat{r}^{(i)} =$
 293 $p_{\text{SM}}(\tilde{m}^{(i)}, \tilde{a}^{(i)} \mid p, f_t, \mathcal{S}_{t-1}^{f_t})$, which provides a low-cost prediction of how each candidate might
 294 perform if fully evaluated, with each of the k objectives estimated independently.
 295

296 Given the predicted objectives, MATHMO performs selection based on maximizing the estimated
 297 hypervolume (Guerreiro et al., 2020). Intuitively, a higher estimated hypervolume means the pair is
 298 more likely to dominate a larger portion of the k -dimensional objective space, relative to a reference
 299 point. Formally, $\text{HV}(\tilde{m}^{(i)}, \tilde{a}^{(i)}; r_{\text{ref}})$, where $r_{\text{ref}} \in \mathbb{R}^k$ is the reference point. To account for potentially
 300 different scales of the k objectives, the individual objective values are normalized to $[0, 1]$ before
 301 calculation, and r_{ref} is set as 1_k . The pair (m_t, a_t) is then selected as the candidate that yields the
 302 largest hypervolume: $(m_t, a_t) = \arg \max_{(\tilde{m}, \tilde{a}) \in \tilde{\mathcal{S}}} \text{HV}(\tilde{m}, \tilde{a}; r_{\text{ref}})$.
 303

304 4.4 MOSE: SURROGATE MODEL OF SUBJECTIVE EVALUATIONS

305 Mathematical modeling is not solely guided by objective performance metrics; subjective qualities,
 306 such as interpretability and alignment with domain knowledge, are crucial utility considerations for
 307 human modelers. Indeed, models generally reflect how we conceptualize and understand complex
 308 situations. Incorporating these aspects into an automated search is thus vital to ensure the generated
 309 models are useful, amenable to further analysis, and capable of communicating valuable insights.
 310 While framework-specific metrics like complexity or sparsity penalties can promote such qualities
 311 (e.g., in symbolic or linear regression), they are often not transferable across different paradigms,
 312 limiting their utility when the goal is to compare diverse models spanning multiple frameworks.
 313

314 To address this, we introduce MOSE as a generalized, cross-framework mechanism for integrating
 315 subjective criteria into model evaluation. We acknowledge that subjective qualities can be highly
 316 observer-dependent; MOSE therefore aims not to capture perfect objectivity but to provide consistent
 317 surrogate approximations. This approach is based on the observation that subjective qualities are
 318 often more reliably expressed through comparative evaluations than absolute scores—an insight also
 319 exploited in Reinforcement Learning from Human Feedback (Bradley and Terry, 1952; Christiano
 320 et al., 2017). Furthermore, we leverage the ability of advanced LLMs to simulate human judgments,
 321 enabling scalable preference elicitation without costly annotations, a technique proven effective in
 322 domains like alignment and diversity search (Bai et al., 2022; Bradley et al., 2024).

323 To ensure comparability of subjective evaluations across diverse discovered models, MOSE employs
 324 a predefined reference set of models, denoted as \mathcal{M}_{ref} . This set, generated at the start of the
 325 search, remains fixed as a consistent frame of reference. When evaluating a new model m_t on




Figure 2: **Pareto fronts and adaptive exploration.** (Top) Pareto fronts of models produced by MATHMO on TSP, JSS, Ecology, and Epidemiology tasks. (Bottom) Corresponding cumulative % of exploration effort allocated to each modeling framework throughout the search process.

a specific subjective quality (e.g., interpretability), MOSE performs pairwise comparisons against each reference model. Specifically, it predicts ‘1’ if the evaluated model is more preferable, and ‘0’ otherwise. The probabilities associated with these predictions are then averaged to obtain a score: $\hat{r} = \frac{1}{|\mathcal{M}_{\text{ref}}|} \sum_i p_{\text{MOSE}}(m_t \succ m_i | p)$. By averaging over a fixed reference set, this approach yields a consistent score in $[0, 1]$, and mitigates potential sensitivities to the choice of reference baseline.

5 EXPERIMENTS


Research questions. In this section, we present an empirical evaluation of MATHMO. Our experiments are designed to investigate the following research questions:

1. How does MATHMO perform on diverse, real-world tasks, particularly its ability to discover a range of models that effectively navigate different trade-offs (Section 5.1).
2. What are the contributions of specific algorithmic design decisions within MATHMO to search performance, as analyzed through controlled ablation studies (Section 5.2).
3. How effectively does MOSE capture and integrate subjective modeling preferences, such as interpretability, into the automated modeling process (Section 5.3).

Problems. To address these questions, we employ four distinct modeling problems, two prescriptive and two predictive, each presenting unique challenges and trade-offs: **Job Shop Scheduling (JSS):** This problem involves optimally scheduling jobs on machines, subject to precedence/resource constraints. We investigate the trade-off between **makespan** (total completion time) and **runtime**, using 10 instances of varying complexities (50-300 ops). **Traveling Salesman (TSP):** This NP-hard problem seeks the shortest route visiting each location exactly once. We examine the trade-off between **route cost** (total tour length) and **runtime** (time to find a tour), utilizing 10 instances of diverse complexities (30-50 locations). **Ecology:** This involves understanding and predicting population dynamics in ecological environments. We focus on the trade-off between **predictive performance** (measured by RMSE on unseen data) and **interpretability** (assessed by MOSE). We use a real-world two-species dataset. **Epidemiology:** The goal is to understand and simulate the spread of infectious diseases. We investigate the trade-off between **predictive performance** (RMSE) and **interpretability**, employing a real-world COVID-19 dataset from Italy.

In these problems, modeling trade-offs are crucial. For instance, operations managers might use slower, more optimal models for long-term planning (JSS, TSP), yet rapid, near-optimal solutions are invaluable for dynamic rescheduling or handling disruptions. Similarly, while ecologists and epidemiologists need high predictive accuracy, model interpretability is vital for gaining scientific insights into underlying mechanisms (e.g., population dynamics, disease transmission) and informing effective interventions like conservation strategies or public health policies.

For all experiments, we run MATHMO for 20 iterations. This process starts with proposing an initial set of 5 frameworks, and each model evaluation is subject to a time limit of 300 seconds. For MOSE,

Table 1: **Ablation study.** Hypervolume performance and relative improvements achieved by MATHMO.Figure 3: **Evaluation of MOSE.** Correlation analysis of MOSE interpretability scores against structural and functional complexity metrics on ecology task (**Left**) and epidemiology task (**Right**).

a reference set comprising 3 models is employed. We use gpt-4o-2024-05-13 as the LLM. Additional details on datasets/experimental setup are provided in Section E. **Additional results.** In the interest of space, we provide additional analysis of our method in Section C, including comparisons against baselines (Section C.1); sensitivity and robustness across runs (Section C.2); and insights on upper-level selection dynamics and lower-level exploration (Sections C.4 and C.5).

5.1 PERFORMANCE ON DIVERSE TASKS

The performance of MATHMO across the four modeling tasks is visualized in Figure 2. Panel (**Top**) of the figure displays the Pareto fronts of discovered models for each task, while panel (**Bottom**) illustrates the corresponding framework exploration dynamics (cumulative effort allocation).

Frameworks dominate trade-offs. A consistent finding is that different modeling frameworks tend to excel in different regions of the Pareto frontier, underscoring the importance of framework selection. For the JSS and TSP tasks, exact methods such as mathematical optimization and constraint programming yield solutions closer to optimality but incur significantly higher computational runtimes. Conversely, metaheuristics and custom heuristics provide solutions with much faster runtimes, with a trade-off in solution quality. This pattern extends to the Ecology and Epidemiology tasks. Here, time-series forecasting methods like vector autoregression achieve strong predictive performance but are assessed by MOSE as less interpretable. In contrast, frameworks such as dynamical systems (e.g., compartmental models) are considered more interpretable, though they exhibit higher RMSE.

Effective intra-framework refinement. MATHMO demonstrates effectively exploration within frameworks to refine solutions along the Pareto front, for instance, refining solver heuristics on JSS/TSP and redesigning simulated annealing-based metaheuristics. Similarly, for the Ecology and Epidemiology tasks, when exploring within the dynamical systems framework, MATHMO proposes variations in model specification (e.g., logistic growth terms, interaction terms) to improve predictive accuracy.

Adaptive exploration behavior. The framework exploration dynamics reveal adaptive search behavior, characteristic of UCB-family strategies. On JSS, TSP, and Epidemiology, we observe a clear pattern of initial broad exploration followed by focused exploitation on frameworks that contribute to the Pareto-UCB frontier. Interestingly, on the Ecology task, the exploration allocation remains more evenly distributed among several distinct frameworks (time-series, dynamical systems, rule-based models), which is consistent with the observation that these frameworks contribute unique, non-dominated solutions to different regions of the Pareto frontier.

5.2 CONTROLLED ANALYSIS OF ALGORITHMIC COMPONENTS

Next, we turn to understanding the contribution of the design decisions of MATHMO. For these purposes, we evaluate three ablations: (1) MATHMORAN, where the Pareto-UCB framework selection strategy is replaced with random selection. (2) MATHMOFLAT, which collapses the bi-level search into a flat search space, applying a globalized version of the local exploration mechanism. (3) MATHMONAIVE, which omits the surrogate-guided local exploration, relying solely on direct sampling.

432 The comparative performance of MATHMO and its ablated versions, as detailed by hypervolume in
 433 Table 1, reveals several key insights. The complete MATHMO generally achieves the best hypervolume,
 434 with one notable exception on the Ecology problem, where MATHMO_{FLAT} (without the bi-level struc-
 435 ture) found Pareto-dominant solutions by concentrating its search on dynamical systems. Employing
 436 random framework selection (MATHMO_{RAN}) resulted in an average hypervolume decrease of 2.5%,
 437 underscoring the value of adaptive exploration at the framework level. The removal of bi-level
 438 search structure had a more pronounced negative impact on model diversity; manual examination
 439 indicated this ablation tended to overconcentrate (allocating 95% of exploration effort to metaheuris-
 440 tics on JSP/TSP and 80% on ecology). MATHMO_{NAIVE}, which relies on repeated LLM sampling,
 441 exhibits the poorest performance overall. However, it interestingly achieved a better hypervolume
 442 than MATHMO_{FLAT} on the JSS task, suggesting that even naive, broad sampling by the LLM could
 443 sometimes provide better coverage of the Pareto frontier than an overly myopic flat search.
 444

445 5.3 SUBJECTIVE QUALITY EVALUATIONS

446 In this final experimental section, we assess how well MOSE captures aspects of model interpretability.
 447 Quantifying interpretability directly is challenging in the general sense, and the broader utility MOSE
 448 lies in its generalizability across different frameworks and problems. Fortunately, for the time-series
 449 problems in our benchmark, we can analyze MOSE’s scores against two commonly used proxy
 450 complexity metrics: (1) structural complexity measured by the number of free parameters (fewer
 451 parameters often correlate with more understandable models), and (2) functional complexity assessed
 452 using permutation entropy of the model’s predicted time-series. Permutation entropy quantifies the
 453 regularity and predictability of a time-series, with lower entropy suggesting simpler, more regular
 454 dynamics, while higher entropy indicates more chaotic patterns (Bandt and Pompe, 2002).

455 Our analysis, with detailed correlations presented in Figure 3, yields several insights into MOSE’s
 456 behavior. Firstly, MOSE scores tend to cluster by modeling framework (e.g., dynamical systems and
 457 rule-based models consistently receive higher interpretability compared to autoregressive forecasting).
 458 The scores exhibit a statistically significant negative correlation with structural complexity on both
 459 tasks: Spearman correlation $\rho = -0.678$ ($p = 1.31 \times 10^{-8}$) for Epidemiology and $\rho = -0.298$
 460 ($p = 0.0318$) for Ecology. The relationship with functional complexity appears more context-
 461 dependent. A significant negative correlation is observed on Ecology ($\rho = -0.261$, $p = 0.0313$), but
 462 almost no correlation is identified on Epidemiology. It is important to note that these proxy metrics
 463 (parameter count and permutation entropy) are themselves indirect measures of interpretability and
 464 are primarily employed here for analytical validation within these specific time-series contexts.
 465

466 6 DISCUSSIONS

467 Automated mathematical modeling represents an exciting frontier in applying artificial intelligence
 468 to complex social, scientific, and engineering problems. In this work, we advance this frontier by
 469 characterizing the process as a sequential decision-making problem under uncertainty. We proposed a
 470 novel adaptive search framework designed to navigate the complex modeling space, featuring mecha-
 471 nisms capable of efficient exploration, balancing multiple objectives, and incorporating subjective
 472 preferences. Our concrete instantiation, MATHMO, demonstrates the potential of leveraging LLMs as
 473 versatile search operators within a structured bi-level search architecture. Empirical results across
 474 real-world modeling tasks underscore MATHMO’s efficacy in discovering diverse frontiers of models,
 475 providing decision-makers with a rich set of alternatives offering different trade-offs.
 476

477 **Future directions.** This research opens numerous avenues for future exploration and enhancement.
 478 Currently, the set of modeling frameworks is sampled and fixed at the beginning of the search. Future
 479 work could explore dynamic framework generation, enabling the system to discover or construct
 480 new modeling paradigms based on accumulated insights, rather than being confined to an initial set.
 481 While designed with multi-objective scenarios in mind, the framework can naturally be applied to
 482 single-objective modeling tasks, and further investigation could optimize its performance for such
 483 cases. Beyond MOSE for subjective preferences, developing richer interactive mechanisms for human
 484 experts to influence the search could significantly enhance outcomes. Lastly, the framework’s general
 485 nature permits the development of more specialized utility functions and optimizations of LLMs as
 samplers or surrogate models, potentially yielding further performance gains. We hope that MATHMO
 and the proposed framework lay a useful foundation for future advancements in this domain.

486 REFERENCES
487

488 Ahmaditeshnizi, A., Gao, W., and Udell, M. (2024). Optimus: Scalable optimization modeling with
489 (mi) lp solvers and large language models. In *International Conference on Machine Learning*,
490 pages 577–596. PMLR.

491 Akinbami, L. J., Chen, T.-C., Davy, O., Ogden, C. L., Fink, S., Clark, J., Riddles, M. K., and
492 Mohadjer, L. K. (2022). National health and nutrition examination survey, 2017–march 2020
493 prepandemic file: sample design, estimation, and analytic guidelines.

494 Astorga, N., Liu, T., Xiao, Y., and van der Schaar, M. (2024). Autoformulation of mathematical
495 optimization models using llms. *arXiv preprint arXiv:2411.01679*.

496

497 Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini,
498 A., McKinnon, C., et al. (2022). Constitutional ai: Harmlessness from ai feedback. *arXiv preprint*
499 *arXiv:2212.08073*.

500 Bandt, C. and Pompe, B. (2002). Permutation entropy: a natural complexity measure for time series.
501 *Physical review letters*, 88(17):174102.

502

503 Banwarth-Kuhn, M. and Sindi, S. (2020). How and why to build a mathematical model: A case study
504 using prion aggregation. *Journal of Biological Chemistry*, 295(15):5022–5035.

505 Bellmore, M. and Nemhauser, G. L. (1968). The traveling salesman problem: a survey. *Operations
506 Research*, 16(3):538–558.

507

508 Blank, J. and Deb, K. (2020). pymoo: Multi-objective optimization in python. *IEEE Access*,
509 8:89497–89509.

510 Box, G. E. (1979). Robustness in the strategy of scientific model building. In *Robustness in statistics*,
511 pages 201–236. Elsevier.

512

513 Bradley, H., Dai, A., Teufel, H. B., Zhang, J., Oostermeijer, K., Bellagente, M., Clune, J., Stanley,
514 K., Schott, G., and Lehman, J. (2024). Quality-diversity through AI feedback. In *The Twelfth
515 International Conference on Learning Representations*.

516 Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs: I. the method of
517 paired comparisons. *Biometrika*, 39(3/4):324–345.

518

519 Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam,
520 P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural
521 information processing systems*, 33:1877–1901.

522 Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data
523 by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of
524 sciences*, 113(15):3932–3937.

525 Chandrasekaran, V. and Jordan, M. I. (2013). Computational and statistical tradeoffs via convex
526 relaxation. *Proceedings of the National Academy of Sciences*, 110(13):E1181–E1190.

527

528 Chen, A., Dohan, D., and So, D. (2023a). Evoprompting: Language models for code-level neural
529 architecture search. *Advances in neural information processing systems*, 36:7787–7817.

530

531 Chen, X., Liang, C., Huang, D., Real, E., Wang, K., Pham, H., Dong, X., Luong, T., Hsieh, C.-J., Lu,
532 Y., et al. (2023b). Symbolic discovery of optimization algorithms. *Advances in neural information
533 processing systems*, 36:49205–49233.

534 Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. (2017). Deep rein-
535 forcement learning from human preferences. *Advances in neural information processing systems*,
536 30.

537 Dempe, S. (2002). *Foundations of bilevel programming*. Springer Science & Business Media.

538

539 Dirac, P. A. M. (1963). The evolution of the physicist’s picture of nature. *Scientific American*,
208(5):45–53.

540 Dong, E., Du, H., and Gardner, L. (2020). An interactive web-based dashboard to track covid-19 in
 541 real time. *The Lancet infectious diseases*, 20(5):533–534.

542

543 Drugan, M. M. and Nowe, A. (2013). Designing multi-objective multi-armed bandits algorithms: A
 544 study. In *The 2013 international joint conference on neural networks (IJCNN)*, pages 1–8. IEEE.

545

546 Guerreiro, A. P., Fonseca, C. M., and Paquete, L. (2020). The hypervolume indicator: Problems and
 547 algorithms. *arXiv preprint arXiv:2005.00515*.

548

549 He, X., Zhao, K., and Chu, X. (2021). Automl: A survey of the state-of-the-art. *Knowledge-based
 549 systems*, 212:106622.

550

551 Holt, S., Liu, T., and van der Schaar, M. (2024). Automatically learning hybrid digital twins of
 552 dynamical systems. In *The Thirty-eighth Annual Conference on Neural Information Processing
 552 Systems*.

553

554 Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2011). Sequential model-based optimization for
 555 general algorithm configuration. In *Learning and intelligent optimization: 5th international
 556 conference, LION 5, rome, Italy, January 17-21, 2011. selected papers 5*, pages 507–523. Springer.

557

558 Hutter, F., Kotthoff, L., and Vanschoren, J. (2019). *Automated machine learning: methods, systems,
 558 challenges*. Springer Nature.

559

560 Hwang, Y., Kim, Y., Koo, J., Kang, T., Bae, H., and Jung, K. (2025). Llms can be easily confused by
 561 instructional distractions. *arXiv preprint arXiv:2502.04362*.

562

563 Jakeman, A. J., Letcher, R. A., and Norton, J. P. (2006). Ten iterative steps in development and
 564 evaluation of environmental models. *Environmental Modelling & Software*, 21(5):602–614.

565

566 Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive
 566 black-box functions. *Journal of Global optimization*, 13:455–492.

567

568 Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford,
 569 A., Wu, J., and Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint
 569 arXiv:2001.08361*.

570

571 Khurana, U., Turaga, D., Samulowitz, H., and Parthasarathy, S. (2016). Cognito: Automated feature
 572 engineering for supervised learning. In *2016 IEEE 16th international conference on data mining
 572 workshops (ICDMW)*, pages 1304–1307. IEEE.

573

574 Lawrance, S. (1984). Resource constrained project scheduling: an experimental investigation of
 575 heuristic scheduling techniques. *GSIA, Carnegie Mellon University*.

576

577 Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., and Talwalkar, A. (2018). Hyperband: A novel
 578 bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*,
 579 18(185):1–52.

580

581 Li, M. Y., Fox, E., and Goodman, N. (2024). Automated statistical model discovery with language
 581 models. In *International Conference on Machine Learning*, pages 27791–27807. PMLR.

582

583 Ling, C., Zhao, X., Zhang, X., Cheng, W., Liu, Y., Sun, Y., Oishi, M., Osaki, T., Matsuda, K., Ji, J.,
 584 et al. (2024). Uncertainty quantification for in-context learning of large language models. *arXiv
 585 preprint arXiv:2402.10189*.

586

587 Liu, T., Astorga, N., Seedat, N., and van der Schaar, M. (2024). Large language models to enhance
 587 bayesian optimization. In *The Twelfth International Conference on Learning Representations*.

588

589 Liu, T., Huynh, N., and van der Schaar, M. (2025). Decision tree induction through LLMs via
 590 semantically-aware evolution. In *The Thirteenth International Conference on Learning Represen-
 590 tations*.

591

592 Ma, Y. J., Liang, W., Wang, G., Huang, D.-A., Bastani, O., Jayaraman, D., Zhu, Y., Fan, L., and
 593 Anandkumar, A. (2024). Eureka: Human-level reward design via coding large language models.
 593 In *The Twelfth International Conference on Learning Representations*.

594 Mensfelt, A., Stathis, K., and Trencsenyi, V. (2024). Autoformalization of game descriptions using
 595 large language models. *arXiv preprint arXiv:2409.12300*.

596

597 Miettinen, K. (1999). *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business
 598 Media.

599

600 Pham, H., Guan, M., Zoph, B., Le, Q., and Dean, J. (2018). Efficient neural architecture search via
 601 parameters sharing. In *International conference on machine learning*, pages 4095–4104. PMLR.

602

603 Real, E., Liang, C., So, D., and Le, Q. (2020). Automl-zero: Evolving machine learning algorithms
 604 from scratch. In *International conference on machine learning*, pages 8007–8019. PMLR.

605

606 Requeima, J., Bronskill, J., Choi, D., Turner, R., and Duvenaud, D. K. (2024). Llm processes: Nu-
 607 mercial predictive distributions conditioned on natural language. *Advances in Neural Information
 608 Processing Systems*, 37:109609–109671.

609

610 Ries, L. A. G., Eisner, M., Kosary, C., Hankey, B., Miller, B., Clegg, L., Mariotto, A., Feuer, E., and
 611 Edwards, B. (1975). Seer cancer statistics review. *National Cancer Institute*, 2004.

612

613 Romera-Paredes, B., Barekatain, M., Novikov, A., Balog, M., Kumar, M. P., Dupont, E., Ruiz, F. J.,
 614 Ellenberg, J. S., Wang, P., Fawzi, O., et al. (2024). Mathematical discoveries from program search
 615 with large language models. *Nature*, 625(7995):468–475.

616

617 Shojaee, P., Meidani, K., Gupta, S., Farimani, A. B., and Reddy, C. K. (2025). LLM-SR: Scientific
 618 equation discovery via programming with large language models. In *The Thirteenth International
 619 Conference on Learning Representations*.

620

621 Simon, H. A. (2019). *The Sciences of the Artificial, reissue of the third edition with a new introduction
 622 by John Laird*. MIT press.

623

624 Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine
 625 learning algorithms. *Advances in neural information processing systems*, 25.

626

627 Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2010). Gaussian process optimization in
 628 the bandit setting: No regret and experimental design. In *Proceedings of the 27th International
 629 Conference on Machine Learning*, pages 1015–1022. Omnipress.

630

631 Stenseth, N. C., Falck, W., Bjørnstad, O. N., and Krebs, C. J. (1997). Population regulation in
 632 snowshoe hare and canadian lynx: asymmetric food web configurations between hare and lynx.
 633 *Proceedings of the National Academy of Sciences*, 94(10):5147–5152.

634

635 Sugihara, G. and May, R. M. (1990). Nonlinear forecasting as a way of distinguishing chaos from
 636 measurement error in time series. *Nature*, 344(6268):734–741.

637

638 Sumers, T., Yao, S., Narasimhan, K., and Griffiths, T. (2023). Cognitive architectures for language
 639 agents. *Transactions on Machine Learning Research*.

640

641 Thornton, C., Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2013). Auto-weka: Combined selection
 642 and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM
 643 SIGKDD international conference on Knowledge discovery and data mining*, pages 847–855.

644

645 Turing, A. M. (1990). The chemical basis of morphogenesis. *Bulletin of mathematical biology*,
 52:153–197.

646

647 van den Berg, N. I., Machado, D., Santos, S., Rocha, I., Chacón, J., Harcombe, W., Mitri, S., and Patil,
 648 K. R. (2022). Ecological modelling approaches for predicting emergent properties in microbial
 649 communities. *Nature ecology & evolution*, 6(7):855–865.

650

651 Van Roy, P. and Haridi, S. (2004). *Concepts, techniques, and models of computer programming*. MIT
 652 press.

653

654 Xiang, Y., Jia, Y., Chen, L., Guo, L., Shu, B., and Long, E. (2021). Covid-19 epidemic prediction
 655 and the impact of public health interventions: A review of covid-19 epidemic models. *Infectious
 656 Disease Modelling*, 6:324–342.

648 Xiang, Y., Yan, H., Gui, L., and He, Y. (2024). Addressing order sensitivity of in-context demon-
649 stration examples in causal language models. In *Findings of the Association for Computational*
650 *Linguistics ACL 2024*, pages 6467–6481.

651

652 Xiong, H., Shi, S., Ren, D., and Hu, J. (2022). A survey of job shop scheduling problem: The types
653 and models. *Computers & Operations Research*, 142:105731.

654

655 Xu, M. and Klabjan, D. (2023). Pareto regret analyses in multi-objective multi-armed bandit. In
656 *International Conference on Machine Learning*, pages 38499–38517. PMLR.

657

658 Yang, C., Wang, X., Lu, Y., Liu, H., Le, Q. V., Zhou, D., and Chen, X. (2024). Large language
659 models as optimizers. In *The Twelfth International Conference on Learning Representations*.

660

661 Yao, S., Liu, F., Lin, X., Lu, Z., Wang, Z., and Zhang, Q. (2025). Multi-objective evolution of heuristic
662 using large language model. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
663 volume 39, pages 27144–27152.

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

Appendix

Table of Contents

A Additional Rebuttal Results	15
A.1 Additional Domains: Large-Scale Medical Risk Prediction	15
A.2 Comparison to AutoML and Solver Baselines	15
A.3 Additional Evaluations of MOSE	16
B Additional Discussions	20
B.1 Real-World Modeling: Uncertainty, Trade-offs, and Subjective Criteria	20
B.2 Enhancing Components of the Adaptive Search Framework	21
C Additional Empirical Analyses	23
C.1 Comparison Against Baselines	23
C.2 Sensitivity Analysis	23
C.3 Comparison of Discovered Pareto Fronts	24
C.4 Insights: Upper-Level Framework Selection	25
C.5 Insights: Lower-Level Exploration Efficiency	26
D Additional Technical Details	27
D.1 Algorithm/Pseudocode	27
D.2 LLM Search Operators	27
D.3 Adaptive Search	28
D.4 Miscellaneous Implementation Details	30
E Additional Experimental Details	31
E.1 Datasets	31
E.2 Metrics	32

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745 The full implementation of MATHMO, along with the code necessary to reproduce all key results,
746 will be released on GitHub upon acceptance of the paper.

747

748

749

750

751

752

753

754

755




Figure 4: **NHANES/SEER Pareto fronts.** Pareto-efficient models discovered across AUROC–ECE trade-offs.

A ADDITIONAL REBUTTAL RESULTS

A.1 ADDITIONAL DOMAINS: LARGE-SCALE MEDICAL RISK PREDICTION

We evaluate MATHMO on two further domains (NHANES and SEER) to demonstrate generality beyond the tasks in Section 5:

- **NHANES** (Akinbami et al., 2022): A national health survey containing 86,000 records of demographic, behavioral, clinical, and environmental covariates, where the goal is to predict risk of myocardial infarction.
- **SEER** (Ries et al., 1975): A population-level cancer registry containing 100,000 patients with incidence, demographic, and survival information, where the goal is to predict risk for breast cancer.

These experiments are designed to evaluate: (1) they involve *substantively different* problem domains than our earlier tasks, (2) evaluation is significantly *more expensive* due to training large-scale models, and (3) they feature *different objective trade-offs*: here, discriminative performance (AUROC \uparrow) and probabilistic calibration (ECE \downarrow), which are central in clinical modeling. We also use a gpt5o LLM backbone to highlight that MATHMO is LLM-agnostic.

Analysis. Figure 4 illustrates that MATHMO consistently discovers diverse Pareto-efficient models spanning supervised learners, Bayesian models, survival analysis, and ensemble methods. These experiments show that MATHMO scales to high-cost real-world modeling settings and is able to autonomously identify high-quality trade-offs across heterogeneous modeling frameworks, reinforcing generality beyond the initial set of domains.

A.2 COMPARISON TO AUTOML AND SOLVER BASELINES

In this subsection, we compare MATHMO discovered models with mature solvers (on TSP/JSSP) and against AutoML baselines (on Ecology/Epidemiology).

Mature solvers. For TSP and JSSP, we benchmark against the highly optimized `Concorde` solver (branch-and-cut) and OR-Tools CP-SAT. These solvers are *framework-specific* and require strong manual modeling choices (variables, constraints, heuristics, tuning). In contrast, MATHMO automatically generates models across heterogeneous frameworks, producing a full Pareto frontier rather than a single optimized point. Including traditional solvers provides the strongest possible reference point, ensuring that the solutions generated by MATHMO are *sensible and well-grounded* relative to gold-standard human-engineered baselines.

Analysis. Tables 2 and 3 show that MATHMO attains competitive performance while simultaneously exploring solutions across optimization, heuristics, and metaheuristics, which populate distinct regions of the Pareto frontier. These comparisons validate that MATHMO produces reasonable solutions near




Figure 5: **SINDy vs. MATHMO.** Pareto fronts on Ecology (left) and Epidemiology (right): MATHMO dominates SINDy across accuracy and interpretability.

expert-tuned solvers, but with the added benefit of automatically discovering multiple qualitatively different trade-offs, which solvers cannot provide.

Table 2: **TSP baseline comparison.** Concorde vs. MATHMO and other baselines.

	Concorde	MATHMO	MEoH	FunSearch
(4.52, 0.0408)	(4.59, 0.6769)	(5.18, 0.1254)	(4.77, 0.1316)	
(3.57, 0.0140)	(3.57, 0.1838)	(3.70, 0.1307)	(3.90, 0.1289)	
(4.60, 0.0243)	(4.66, 0.6476)	(4.89, 0.1286)	(4.89, 0.1265)	
(5.11, 0.0305)	(5.19, 0.1566)	(5.66, 0.1307)	(5.40, 0.1418)	
(4.24, 0.0405)	(4.24, 0.1805)	(5.83, 0.1268)	(4.48, 0.1314)	

Table 3: **JSSP baseline comparison.** OR-Tools vs. MATHMO and other baselines.


	OR-Tools	MATHMO	MEoH	FunSearch
(1039.00, 0.0599)	(1039.00, 0.0230)	(1096.00, 0.0602)	(1039.00, 0.0536)	
(1218.00, 0.2954)	(1421.00, 0.0246)	(1503.00, 0.0640)	(1372.00, 0.1610)	
(1235.00, 13.4015)	(1235.00, 15.9644)	(1514.00, 0.0621)	(1644.00, 0.1747)	
(1721.00, 0.6016)	(1721.00, 1.1240)	(2175.00, 0.0615)	(1925.00, 0.1697)	
(1888.00, 0.2975)	(2127.00, 0.0233)	(2183.00, 0.0620)	(11487.00, 0.1949)	

Comparison to SINDy. For the time-series tasks (Ecology and epidemiology), we compare against SINDy (Brunton et al., 2016), a widely used symbolic regression technique to discover interpretable and predictive mechanistic models. As SINDy requires users to choose a basis-function library and sparsity threshold, it lends itself well to AutoML-based model selection. To keep budgets comparable, we allocate SINDy a consistent search budget of 20 configurations (4 libraries \times 5 sparsity thresholds), matching the per-framework iteration budget in MATHMO.

Analysis. Figure 5 shows that across both Ecology and Epidemiology, SINDy is entirely dominated: its models occupy a narrow interpretability band, while MATHMO’s models span a richer set of trade-offs through access to multiple modeling frameworks. **Takeaway:** MATHMO yields strictly better Pareto fronts due to its ability to explore beyond a single symbolic regression paradigm.

A.3 ADDITIONAL EVALUATIONS OF MOSE

Sensitivity analysis. MOSE is used during search as a surrogate for subjective human preferences (e.g., model interpretability). Surrogate stability is critical: if small changes to the reference set \mathcal{M}_{ref} produced inconsistent preferences, MATHMO’s search could be noisy or unreliable. Therefore, evaluating the sensitivity of MOSE directly addresses whether its interpretability judgments are robust.

Figure 6: **MOSE sensitivity.** Interpretability score distributions across four reference sets (\mathcal{M}_{ref}).

Analysis. We compute MOSE scores for 20 models across four independently sampled reference sets. The distributions in Figure 6 and correlations in Tables 4 and 5 show consistently high agreement ($r > 0.93$). MOSE produces stable and consistent judgments across reference sets, indicating that it is a reliable component for guiding subjective-objective trade-offs during search.

Human study. As MOSE is used to approximate human qualitative preferences, it is essential to verify that its judgments align with human experts. The goal is not to perform a full-scale user study, but to provide evidence that MOSE is directionally consistent with expert reasoning. We collect expert pairwise interpretability judgments on 25 randomly sampled model pairs. For each pair, we record the MOSE scores and the expert preference of interpretability.

Analysis. Agreement with MOSE is 79.2% for Ecology (Table 6) and 76.0% for Epidemiology (Table 7). The high agreement rates indicate that MOSE captures *meaningful subjective preferences* consistent with expert intuition, supporting its use as a scalable interpretability surrogate.

Table 4: Epidemiology - Correlations Between Reference Sets (\mathcal{M}_{ref})

Reference Set Pair	Pearson $r \uparrow$ (p -value)
(3, 2)	0.9695 (1.92e-12)
(3, 1)	0.9801 (4.30e-14)
(3, 0)	0.9673 (3.61e-12)
(2, 1)	0.9892 (1.85e-16)
(2, 0)	0.9507 (1.36e-10)
(1, 0)	0.9561 (4.87e-11)

Table 5: Ecology - Correlations Between Reference Sets (\mathcal{M}_{ref})

Reference Set Pair	Pearson $r \uparrow$ (p -value)
(3, 2)	0.9823 (1.53e-14)
(3, 1)	0.9525 (9.87e-11)
(3, 0)	0.9923 (8.84e-18)
(2, 1)	0.9351 (1.54e-09)
(2, 0)	0.9809 (3.06e-14)
(1, 0)	0.9567 (4.36e-11)

918
919
920
921
922
923
924
925
926
927
928
929
930

931 Table 6: **Human interpretability judgments (Ecology)**. Pairwise expert preferences compared
932 against MOSE predictions for 25 model pairs. Agreement indicates whether MOSE selects the same
933 model as the human expert.

934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

	MOSE(A)	MOSE(B)	Expert	Agree
	0.403	0.673	B	✓
	0.626	0.645	A	✗
	0.600	0.675	B	✓
	0.650	0.643	B	✗
	0.403	0.679	B	✓
	0.679	0.675	A	✓
	0.626	0.687	B	✓
	0.626	0.635	A	✗
	0.563	0.687	B	✓
	0.673	0.635	A	✓
	0.626	0.650	B	✓
	0.591	0.701	B	✓
	0.655	0.614	A	✓
	0.600	0.635	B	✓
	0.600	0.638	B	✓
	0.626	0.655	B	✓
	0.679	0.650	B	✗
	0.679	0.635	A	✓
	0.673	0.638	A	✓
	0.645	0.403	A	✓
	0.687	0.635	A	✓
	0.614	0.678	B	✓
	0.403	0.701	A	✗
	0.614	0.614	A	✓
	Agreement Rate		19/24 (79.2%)	

972
973
974
975
976
977
978
979
980
981
982
983

984 **Table 7: Human interpretability judgments (Epidemiology).** Expert pairwise evaluations compared
985 against MOSE decisions for 25 model pairs. Agreement indicates whether MOSE matches expert
986 preference.

987

	MOSE(A)	MOSE(B)	Expert	Agree
988	0.464	0.578	B	✓
989	0.524	0.525	B	✓
990	0.601	0.500	A	✓
991	0.663	0.507	A	✓
992	0.464	0.662	B	✓
993	0.662	0.500	A	✓
994	0.524	0.540	A	✗
995	0.524	0.502	B	✗
996	0.390	0.540	B	✓
997	0.578	0.502	B	✗
998	0.524	0.663	B	✓
999	0.550	0.556	B	✓
1000	0.658	0.548	A	✓
1001	0.601	0.502	A	✓
1002	0.601	0.551	A	✓
1003	0.524	0.658	B	✓
1004	0.662	0.663	A	✗
1005	0.662	0.502	B	✗
1006	0.578	0.551	A	✓
1007	0.525	0.464	A	✓
1008	0.556	0.540	A	✓
1009	0.540	0.502	B	✗
1010	0.548	0.488	A	✓
1011	0.464	0.556	B	✓
1012	0.555	0.548	A	✓
1013	Agreement Rate		19/25 (76.0%)	

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026 **B ADDITIONAL DISCUSSIONS**

1028 In the following section of the appendix, we offer additional discussion to motivate the need for
 1029 automated mathematical modeling (MATHMO), reflecting the multi-objective trade-offs, variety of
 1030 frameworks, and role of subjective criteria inherent in real-world modeling. We then outline potential
 1031 enhancements to key components of the adaptive search framework.

1033 **B.1 REAL-WORLD MODELING: UNCERTAINTY, TRADE-OFFS, AND SUBJECTIVE CRITERIA**
 1034 **Table 8: Overview of problems.** Illustrating multi-objective trade-offs, modeling diversity, uncer-
 1035 tainty, and subjective considerations.

1036 Problem	1037 Objectives	1038 Modeling Approaches	1039 Uncertainty in Modeling	1040 Subjective Criteria
1038 TSP	1039 \triangleright Journey cost, \triangleright Run- 1040 time	1041 \triangle ILP, \triangle Metaheuristics, \triangle Heuristics	1042 Instance size and structure affect per- 1043 formance; solver behavior is unpre- 1044 dictable	1045 -
1040 Job Shop Scheduling	1041 \triangleright Makespan, \triangleright Run- 1042 time	1043 \triangle Constraint program- 1044 ming, \triangle Metaheuristics, \triangle Heuristics	1045 Solution quality and runtime 1046 vary with problem characteristics; 1047 method choice is non-obvious	1048 -
1042 Ecology	1043 \triangleright Predictive accuracy, \triangleright Interpretability	1044 \triangle Differential equations, \triangle ARIMA, \triangle Graphical 1045 models	1046 Complex, noisy dynamics; uncertain 1047 model structure; missing or sparse 1048 data	1049 Alignment with 1050 ecological knowledge 1051 and interpretability
1044 Epidemiology	1045 \triangleright Forecast accuracy, \triangleright 1046 Interpretability	1047 \triangle Compartmental mod- 1048 els, \triangle Statistical models, \triangle Time-series models	1049 Highly sensitive to data quality and 1050 regime changes; difficult to validate 1051 assumptions	1052 Interpretability for 1053 public health communica- 1054 tion
1046 Medical Diagnosis	1047 \triangleright Accuracy, \triangleright Explain- 1048 ability, \triangleright Uncertainty 1049 calibration	1050 \triangle Deep learning, \triangle 1051 Rule-based systems, \triangle 1052 Probabilistic models	1053 Variation in populations, equipment, 1054 and labeling; generalization is uncer- 1055 tain	1056 Clinical trust and 1057 alignment with expert 1058 reasoning
1048 Portfolio Optimization	1049 \triangleright Expected return, \triangleright 1050 Risk, \triangleright Robustness	1051 \triangle Classical finance mod- 1052 els, \triangle ML models (e.g., 1053 RNNs)	1054 Financial time series are non- 1055 stationary; market conditions shift 1056 unpredictably	1057 Transparency, ex- 1058 plainability, and risk 1059 alignment
1050 Drug Response Modeling	1051 \triangleright Predictive accuracy, \triangleright Biological inter- 1052 pretability, \triangleright Safety	1053 \triangle Mechanistic models, \triangle Multi-omics ML, \triangle 1054 Hybrid causal models	1055 High patient variability; limited, ex- 1056 pensive data; strong prior assump- 1057 tions needed	1058 Regulatory approval 1059 and clinical inter- 1060 pretability

1052 This work addresses a relatively underexplored area: the automation of mathematical modeling under
 1053 real-world constraints. Our problem formulation and methodological framework are motivated by
 1054 three core challenges commonly encountered in practice: (1) uncertainty in modeling decisions—such
 1055 as which frameworks or assumptions are most appropriate; (2) multi-objective trade-offs between
 1056 competing criteria like accuracy, runtime, and interpretability; and (3) subjective, human-centric
 1057 considerations that are difficult to mathematically formalize but critical to real-world adoption.

1058 Our analysis focuses on domains such as job-shop scheduling, vehicle routing, ecology, and epi-
 1059 demiology, each of which presents unique modeling challenges and trade-offs, such as optimality
 1060 versus runtime or predictive accuracy versus interpretability. However, these issues are far from
 1061 domain-specific. In the following discussion, we illustrate how similar concerns arise across a wide
 1062 range of modeling scenarios, further underscoring the importance of a general-purpose, flexible
 1063 approach to automated modeling. An overview of the discussion is summarized in Table 8.

1064 **Medical diagnosis.** The task of predicting diseases or conditions based on clinical data such as
 1065 medical images, lab results, or patient history.

- 1067 **Trade-offs:** \triangleright Diagnostic accuracy (essential for minimizing missed or incorrect diagnoses) *vs.* \triangleright
 1068 explainability (clinicians need to understand model reasoning to trust and act on predictions) *vs.* \triangleright
 1069 uncertainty calibration (important for risk-aware decision-making, especially in borderline cases).
- 1070 **Possible modeling frameworks:** \triangle Deep learning (highly accurate, data-intensive, low inter-
 1071 pretability); \triangle rule-based expert systems (e.g., risk scores; interpretable but often underperform,
 1072 and lack uncertainty handling); \triangle probabilistic models (capture uncertainty but require strong
 1073 assumptions and are harder to scale).
- 1074 **Subjective criteria:** Clinicians value interpretability and alignment with domain knowledge—
 1075 understanding why a model makes a diagnosis is as important as the prediction itself.
- 1076 **Uncertainty in modeling:** Variation in patient populations, medical instrumentation (e.g., imaging
 1077 devices), and comorbidities makes it unclear which modeling assumptions will generalize. Ground
 1078 truth labels may also be noisy or inconsistent across annotators.

1079 **Financial portfolio optimization.** The process of allocating assets to maximize returns while
 1080 managing risk in dynamic market environments.

1080 1. **Trade-offs:** Expected return (central to investor objectives) *vs.* \triangleright risk (higher returns generally
 1081 entail higher volatility); \triangleright predictive performance (important for exploiting market inefficiencies)
 1082 *vs.* \triangleright robustness (models may overfit to past data and fail under new market regimes).
 1083 2. **Possible modeling frameworks:** \triangle Classical financial models (e.g., Markowitz, Black-Litterman;
 1084 principled but sensitive to input estimation errors); \triangle machine learning models (e.g., RNNs;
 1085 flexible, can learn patterns, but require large, clean datasets and may overfit or lack robustness).
 1086 3. **Uncertainty in modeling:** Market dynamics are non-stationary and hard to forecast. Expected re-
 1087 turns, volatilities, and correlations shift over time, making it unclear which models or assumptions
 1088 will remain valid. A model that performs well in one regime/time horizon may fail in another.

1089 **Drug response modeling.** Predicting how individual patients will respond to a given drug, often in
 1090 the context of personalized medicine or drug development.

1091 1. **Trade-offs:** \triangleright Predictive accuracy (critical for identifying effective treatments) *vs.* \triangleright biological
 1092 interpretability (important for understanding mechanisms and gaining trust); \triangleright short-term efficacy
 1093 (desired for immediate outcomes) *vs.* long-term safety (essential for regulatory approval and
 1094 patient well-being).
 1095 2. **Possible modeling frameworks:** \triangle Mechanistic models (e.g., PK/PD; grounded in biology,
 1096 interpretable, but slow and parameter-sensitive); \triangle multi-omics ML models (data-driven and
 1097 expressive, but opaque and difficult to validate); \triangle hybrid models (e.g., causal or semi-mechanistic;
 1098 combine strengths, but sensitive to misspecification).
 1099 3. **Subjective criteria:** Interpretability and biological plausibility are crucial for clinical trust and
 1100 regulatory acceptance.
 1101 4. **Uncertainty in modeling:** High variability across patients (e.g., in genetics or metabolism)
 1102 complicates model generalization. Data is often limited, expensive to obtain, and ethically
 1103 constrained.

1104 In all these examples, the goal is not to identify a single best model, but rather to present the human
 1105 user with a diverse set of viable modeling options—each representing different trade-offs across
 1106 relevant objectives such as accuracy, interpretability, robustness, and runtime. This is crucial because
 1107 the optimal modeling choice often depends on context-specific constraints, user preferences, and
 1108 shifting priorities. Framing this as a multi-objective optimization problem allows us to systematically
 1109 explore the space of trade-offs and approximate the Pareto frontier, enabling users to make informed
 1110 decisions based on their own criteria and operational needs.

1112 B.2 ENHANCING COMPONENTS OF THE ADAPTIVE SEARCH FRAMEWORK

1113 Our work presents a general framework for approaching automated mathematical modeling. Specifi-
 1114 cally, it decomposes the automated modeling process into a bi-level search: the upper level selects
 1115 among modeling frameworks while managing exploration–exploitation trade-offs, and the lower
 1116 level performs local search within each framework to identify high-performing model-algorithm
 1117 pairs. While our specific instantiation of this framework demonstrates effectiveness, we outline key
 1118 areas where individual components of this framework could be improved to enhance performance,
 1119 flexibility, and robustness.

1120 B.2.1 UPPER-LEVEL UTILITY FUNCTION

1121 We use Pareto-UCB to guide framework selection by approximating the Pareto frontier across multiple
 1122 objectives. This balances exploration and exploitation in a principled way. However:

1123 1. **Non-stationarity.** Pareto-UCB assumes a stationary reward distribution, which is perhaps un-
 1124 tenable in our setting—modeling performance is expected to improve over time due to ongoing
 1125 low-level exploration. This results in a shifting reward distribution and suggests that a dynamic or
 1126 non-stationary approach may be more appropriate.
 1127 2. **Independence assumption.** Frameworks are currently treated as independent in the selection
 1128 problem, but in reality, their performance is often correlated (e.g., if model-based control methods
 1129 perform well, similar model-based RL might too). Ideally, inter-framework correlations and
 1130 structure are captured somehow, although modeling or learning this correlation a-priori remains
 1131 challenging.

1134 3. **Alternative utility functions.** Optimism-based acquisition strategies like Pareto-UCB presents
 1135 a strong initial approach, but Bayesian acquisition functions (e.g., entropy search) may better
 1136 capture uncertainty and trade-offs in this multi-objective setting and are worth investigating in
 1137 future work.

1138 **B.2.2 LOWER-LEVEL UTILITY FUNCTION**

1139 Our lower-level utility function relies on an LLM-based surrogate to estimate the objective perfor-
 1140 mance of candidate models, using expected hypervolume improvement as the acquisition criterion.

1141 1. **Surrogate models.** Developing a surrogate model on this non-conventional search space (i.e.,
 1142 space of models expressed in natural language) is very challenging. While our work uses an
 1143 LLM-based surrogate model and demonstrated improved search efficiency as a result, this is a
 1144 fruitful area with open challenges in calibration and generalization.

1145 2. **Alternative search strategies.** Other local search methods, such as evolutionary algorithms or
 1146 learned reinforcement learning-based search policies, could offer more robust and generalizable
 1147 exploration under uncertainty.

1148 3. **Framework-specific search policies.** More interestingly, specialized framework-specific ex-
 1149 ploration strategies that exploit the structure unique to each framework could greatly improve
 1150 efficiency. For instance, leveraging the hierarchical structures in mathematical programming
 1151 ([Astorga et al., 2024](#)) or neural architecture search for deep learning models ([Zoph and Le, 2016](#)).
 1152 However, this requires engineering these framework-specific strategies beforehand.

1153 **B.2.3 LLM GENERATIVE SAMPLERS**

1154 LLMs are crucial for open-ended exploration in the space of frameworks, models, and algorithms.

1155 1. **Finetuned search operators.** In our work we used general purpose LLMs as samplers, although
 1156 there is no principled reason this could not be improved with specific finetuned search operators.

1157 2. **Dynamic sampling.** When it comes to framework sampling, our current approach samples and
 1158 fixes a set of frameworks at the beginning of search. Ideally, this could be improved with a
 1159 dynamic approach, which is less constraining, enables more efficient exploration budget allocation,
 1160 and encourages the emergence of hybrid or novel modeling paradigms.

1161 **B.2.4 LLM SURROGATE MODELS**

1162 LLMs, which operate directly on natural language inputs, offer a powerful alternative to traditional
 1163 surrogate models like Gaussian Processes that are limited to well-defined numerical feature spaces.
 1164 This flexibility enables surrogate modeling over open-ended model descriptions, expanding the
 1165 expressive range of the search process. However, LLM-based surrogates also come with notable
 1166 challenges. Their uncertainty estimates tend to be poorly calibrated ([Ling et al., 2024](#)), and their
 1167 outputs can be highly sensitive to prompt design and formatting ([Xiang et al., 2024; Hwang et al.,](#)
 1168 [2025](#)). Addressing these issues—through techniques such as prompt ensembling, temperature scaling,
 1169 and uncertainty-aware decoding—remains an important direction for improving both reliability and
 1170 performance in LLM-based surrogate evaluation.

1171 **B.2.5 MOSE: MODEL OF SUBJECTIVE EVALUATION**

1172 A general-purpose, cross-framework model of subjective evaluation introduces significant potential
 1173 for modeling human-like preferences in automated modeling pipelines. Our initial results showed
 1174 that MOSE scores are meaningfully correlated with structural and functional properties of candidate
 1175 models, suggesting alignment with domain-relevant heuristics. Nonetheless, the model remains in
 1176 an early stage and would benefit from further validation and tuning across diverse problem types.
 1177 Encouragingly, similar challenges have been tackled in reinforcement learning from human feedback
 1178 (RLHF), resulting in reliable and robust preference models ([Christiano et al., 2017](#)). In our setting,
 1179 MOSE could be further improved using paired preference labels gathered from domain experts,
 1180 enabling data-efficient finetuning and deeper alignment with subjective modeling criteria.

1188

C ADDITIONAL EMPIRICAL ANALYSES

1189
 1190 In this section, we supplement our main empirical results with additional analyses that provide deeper
 1191 insight into the performance of MATHMO. We start by focusing on two aspects: (i) comparison against
 1192 strong baselines in automated heuristic design, and (ii) sensitivity and stability analyses of MATHMO
 1193 with respect to framework selection and initialization. Subsequently, we focus on the dynamics of
 1194 upper-level framework selection and lower-level model exploration.
 1195

1196

C.1 COMPARISON AGAINST BASELINES

1197
 1198 We benchmark MATHMO against two recent and representative systems: **MEoH** (Yao et al., 2025) and
 1199 **FunSearch** (Romera-Paredes et al., 2024). Both are state-of-the-art approaches in automated heuristic
 1200 design using large language models (LLMs). Specifically, MEoH addresses multi-objective search
 1201 using evolutionary operators and a dominance–dissimilarity mechanism for diversity maintenance,
 1202 while FunSearch employs genetic programming with an island-model evolutionary strategy, tailored
 1203 primarily for single-objective problems. Although neither method is designed for general-purpose,
 1204 cross-framework mathematical modeling, they provide strong points of reference for evaluating the
 1205 relative effectiveness of MATHMO.
 1206

1207 **Experimental setup.** All methods are evaluated on four benchmark problems under a fixed budget
 1208 of 20 model–algorithm evaluations. For FunSearch, the multi-objective criteria are scalarized using
 1209 uniform weights to ensure a fair comparison. Table 9 summarizes the results.
 1210

Table 9: Performance comparison against baselines. Higher values are better.

Method	TSP	JSSP	Ecology	Epidemiology
MEoH	0.9480	0.8587	0.8360	0.9054
FunSearch	0.9772	0.8130	0.6240	0.6628
MATHMO	0.9877	0.9655	0.9576	0.9793

1211
 1212 MATHMO consistently outperforms both baselines across all domains. We attribute these gains to two
 1213 main factors:
 1214

- **Cross-framework modeling.** MATHMO explicitly searches over diverse modeling frameworks (e.g., dynamical systems, symbolic regression, constraint programming), whereas MEoH and FunSearch tend to restrict exploration to one or two frameworks. For example, in the JSSP domain, MATHMO leveraged constraint programming frameworks, yielding superior solutions not discovered by either baseline.
- **Surrogate-guided search.** MATHMO employs LLM-based surrogate models to guide candidate selection, improving sample efficiency and focusing evaluations on high-potential areas. In contrast, MEoH and FunSearch rely exclusively on direct evaluation of evolved models.

1215

C.2 SENSITIVITY ANALYSIS

1216 We also conducted a post-hoc sensitivity analysis by measuring the average drop in hypervolume
 1217 (HV) when each framework is removed from the search space. Results are summarized below:
 1218

- **TSP:** 7.38% (mathematical optimization), 7.20% (metaheuristics), 1.21% (heuristics)
- **JSSP:** 2.16% (constraint programming), 1.29% (metaheuristics)
- **Ecology:** 16.71% (symbolic regression), 9.62% (time-series forecasting), 0.35% (rule-based)
- **Epidemiology:** 15.00% (rule-based)

1219 These results indicate that certain frameworks contribute uniquely to the Pareto frontier, while others
 1220 are progressively deprioritized by the Pareto-UCB mechanism.
 1221


1242 Table 10: Inter-run consistency on VRP (5 runs). HV denotes hypervolume.
1243

Run	0	1	2	3	4
HV	0.989	0.974	0.980	0.989	0.987

1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
Average HV: 0.984 ± 0.005 .

Table 11: Framework contributions to the Pareto frontier (percentage HV contribution).

Framework	Run 0	Run 1	Run 2	Run 3	Run 4
Heuristics	9.02%	2.00%	9.54%	8.78%	10.21%
Mathematical optimization	7.18%	1.82%	5.44%	6.87%	7.34%
Metaheuristics	0.33%	—	0.14%	0.86%	0.23%

Figure 7: **Pareto front comparisons.** Visualizations of the discovered Pareto fronts for MATHMO, MATHMO_{FLAT}, and MATHMO_{NAIVE}. ▲ denotes non-dominated models; -- traces the estimated Pareto front; colors indicate distinct modeling frameworks. From top to bottom: **TSP**, **JSS**, **Ecology**, **Epidemiology**.

To further assess robustness, we conducted experiments on the vehicle routing problem (VRP) with multiple random initializations.

The results show that hypervolume remains stable across runs (0.984 ± 0.005), and framework contributions to the Pareto frontier are consistent across different initializations.

C.3 COMPARISON OF DISCOVERED PARETO FRONTS

We begin by analyzing the quality of the discovered Pareto frontiers, as well as the rate of hypervolume improvement. The Pareto front reflects the diversity and optimality of the trade-offs discovered during




Figure 8: **Hypervolume during search.** Temporal progression of normalized hypervolume throughout the search process for MATHMO and control variants: MATHMO_{FLAT} and MATHMO_{NAIVE}. From left to right: **TSP, JSS, Ecology, Epidemiology**.




Figure 9: **Upper-level framework selection.** Comparison of hypervolume progression between two upper-level selection strategies: MATHMO (adaptive) and MATHMORAN (random). From left to right: **TSP, JSS, Ecology, Epidemiology**.

search, while hypervolume progression captures the efficiency with which the method explores the multi-objective space.

We compare MATHMO against two ablations: MATHMO_{FLAT}, which removes the bi-level structure and performs sequential sampling of model-algorithm pairs without adaptive framework selection; and MATHMO_{NAIVE}, which further removes the sequential aspect entirely, instead sampling a set of model-algorithm pairs in parallel. For all three settings, we allow 20 iterations, with MATHMO_{NAIVE} generating 20 parallel samples.

Discovered pareto fronts. In Figure 7, we compare the Pareto frontiers discovered by each method. MATHMO consistently identifies the most complete and diverse set of Pareto-efficient solutions. For instance, on **TSP**, it uncovers models that offer a range of trade-offs between journey cost and runtime. In contrast, MATHMO_{FLAT} explores fewer frameworks and, while it identifies some competitive models in terms of journey cost, it fails to find low-runtime alternatives with modest sacrifices in optimality. More generally, we observe that purely sequential approaches like MATHMO_{FLAT} tend to focus locally on just 1–2 frameworks, limiting their coverage of the multi-objective space and resulting in fewer diverse trade-off solutions. This trend is also evident in **Ecology** and **Epidemiology**, where MATHMO discovers broader and more densely populated Pareto fronts that dominate across both objectives. An exception arises in **JSS**, where MATHMO_{NAIVE} discovers a Pareto front comparable to that of MATHMO. This is likely due to favorable random coverage and the nature of the problem landscape, where multiple high-performing models can be sampled without requiring adaptive search.

Search efficiency. Figure 8 shows the progression of normalized hypervolume across iterations. On **TSP**, **Ecology**, and **Epidemiology**, MATHMO achieves faster gains in hypervolume compared to the ablations, often requiring significantly fewer iterations to identify Pareto-efficient solutions. The case of **JSS** is again an outlier: although MATHMO_{NAIVE} eventually achieves comparable hypervolume, MATHMO reaches the same level in only 6 iterations, highlighting its superior sample efficiency.

C.4 INSIGHTS: UPPER-LEVEL FRAMEWORK SELECTION

We next examine the dynamics of upper-level framework selection, which in MATHMO is guided by the Pareto-UCB utility function. To isolate its impact, we compare against a control ablation, MATHMORAN, which is identical to MATHMO except that it replaces utility-guided selection with uniform random sampling over frameworks.




Figure 10: **Lower-level exploration efficiency.** (Top) Relative hypervolume improvement as a function of the number of model evaluations for MATHMO and RAN. (Bottom) NRMSE of surrogate predictions over time. From left to right: **TSP, JSS, Ecology, Epidemiology**

Observations. Figure 9 shows the progression of normalized hypervolume across iterations. On **TSP, Ecology**, and **Epidemiology**, we observe similar rates of improvement between MATHMO and MATHMO_{RAN} during the initial $t \leq 6$ iterations, reflecting the exploratory phase of the search. However, after this point, MATHMO exhibits a notably sustained increase in hypervolume, suggesting that it effectively focuses its search budget on frameworks with greater potential for Pareto improvement. This behavior illustrates the utility of guided upper-level selection: Pareto-UCB enables adaptive resource allocation toward promising regions of the search space.

In contrast, on **JSS**, both methods converge to comparable hypervolume levels at similar rates. This echoes the earlier observation in Pareto front comparisons, where MATHMO_{NAIVE} also performed well. These results suggest that the effectiveness of upper-level selection may be problem-dependent. For instance, in settings like **JSS** where the landscape is relatively flat or where good models are distributed across many frameworks, random selection may suffice to find competitive solutions.

C.5 INSIGHTS: LOWER-LEVEL EXPLORATION EFFICIENCY

We now turn to the efficiency of local exploration, specifically how well the surrogate model guides the selection of candidate model-algorithm pairs. To isolate this effect, we compare MATHMO against a control variant, RAN, which is identical in every respect except that it selects a candidate at random rather than using a surrogate to estimate and optimize objective performance.

To ensure a fair comparison, both methods are evaluated on the same historical set of models and the same set of candidate proposals at each step. This setup, averaged over 5 random seeds, controls for variation in the history and candidate pool, ensuring that any observed differences can be attributed solely to the decision-making strategy, i.e., surrogate-guided versus random selection.

Local search efficiency. Figure 10 (Top) shows relative hypervolume improvement as a function of the number of evaluated models. For each selected point, we compute the gain in hypervolume relative to the historical set alone; a larger value indicates that the newly acquired model improved the current Pareto front. Across all four benchmarks, **TSP, JSS, Ecology**, and **Epidemiology**, we observe that MATHMO consistently achieves higher relative hypervolume gains. This indicates that surrogate-guided selection is more effective at identifying models that advance the Pareto frontier, leading to more sample-efficient exploration compared to random selection.

Surrogate model performance. To better understand the surrogate’s behavior, we also measure its predictive accuracy in terms of normalized RMSE (NRMSE) on the candidate models, averaged across the two objectives. In **TSP, Ecology**, and **Epidemiology**, we find that NRMSE decreases with more historical data, suggesting that the LLM-based surrogate improves with more observations, consistent with the observation in Liu et al. (2024) that LLM-based surrogate estimation generalizes better with more context. In contrast, on **JSS**, surrogate accuracy worsens over time, with increasing NRMSE across iterations. Manual inspection revealed a possible explanation: that the dominant source of error lies in estimating runtime for metaheuristic algorithms—critical for identifying Pareto-

1404 improving trade-offs in this problem domain. These runtime behaviors are difficult to predict based
 1405 solely on surface-level descriptions, leading to poor surrogate performance. This likely explains why
 1406 both MATHMO and RAN achieve similar relative hypervolume improvements on **JSS**: as the surrogate
 1407 becomes less informative, its selection decisions approach random choice.

D ADDITIONAL TECHNICAL DETAILS

In this section of the Appendix, we provide additional details on the implementation of MATHMO.

D.1 ALGORITHM/PSEUDOCODE

Algorithm 1 Bi-level Adaptive Search Loop in MATHMO

- 1: **Input:** Problem description p , dataset \mathcal{D}_p , number of iterations T , number of frameworks F , number of candidates L
- 2: Sample initial frameworks: $f_i \sim p_\theta(\cdot | p)$, $\forall i \in [F]$
- 3: Initialize histories: $S_i^{(0)} = \emptyset$, $\forall i \in [F]$
- 4: **for** $t = 1$ to T **do**
- 5: **Upper-level:** Compute utility: $\alpha(f_i) = \text{Pareto-UCB}(f_i; S_i^{(t-1)})$, $\forall i \in [F]$
- 6: Select framework: $f_* = \arg \max_{f_i \in \mathcal{F}} \alpha(f_i)$
- 7: **Lower-level:** Sample candidate pairs: $(\tilde{m}_j, \tilde{a}_j) \sim p_\phi(\cdot, \cdot | p, f_*, S_*^{(t-1)})$, $\forall j \in [L]$
- 8: Estimate objectives: $\hat{r}_j = p_{SM}(\tilde{m}_j, \tilde{a}_j | p, f_*, S_*^{(t-1)})$ $\forall j \in [L]$
- 9: Select candidate pair: $(m^{(t)}, a^{(t)}) = \arg \max_{(m, a) \in \tilde{\mathcal{C}}} \text{HV}(\hat{r}_j; r_{\text{ref}})$
- 10: Solve and evaluate to obtain $r^{(t)}$
- 11: Update history: $S_*^{(t)} \leftarrow S_*^{(t-1)} \cup \{(m^{(t)}, a^{(t)}, r^{(t)})\}$
- 12: **end for**
- 13: **Output:** Pareto set: $\mathcal{P} = \text{Pareto} \left(\bigcup_{i=1}^F S_i^{(T)} \right)$

D.2 LLM SEARCH OPERATORS

To recap, MATHMO leverages LLMs as search operators, specifically for three distinct roles: sampling realizations, surrogate evaluations of candidate models, and as a model of subjective evaluation (MOSE).

1. **Generative sampler [frameworks].** Conditioned on the problem description p , LLMs are prompted to sample suitable modeling frameworks, which we denote as $f \sim p_\theta(\cdot | p)$. Specifically, the LLM is instructed to return proposed frameworks in a `JSON` structure containing two fields: “modeling_framework” (string) and “framework_description”. The prompt skeleton and output format are described in Figures 11 and 12 respectively. Note that the descriptions enclosed in {} represent placeholder values that are populated dynamically at runtime.
2. **Generative sampler [model and algorithm].** Conditioned on the problem description p , a selected framework f , and a set of previously evaluated models within that framework \mathcal{S}^f , the LLM generates a new model-algorithm pair, $(m, a) \sim p_\phi(\cdot, \cdot | p, f, \mathcal{S}^f)$. The output is returned in a `JSON` format with the fields: “model” (a `Python` code string), “dependencies” (a list of package names), and “explanation” (a rationale for the design). Prompt details and output structure are described in Figures 13 and 14.
3. **Surrogate evaluations [candidate model].** For each candidate model-algorithm pair (\tilde{m}, \tilde{a}) , LLMs are employed as surrogates to estimate multi-objective performance metrics, offering a low-cost approximation. This approach is motivated by the unstructured nature of the input space, which differs from traditional numerical or mixed-integer domains in Bayesian Optimization (Snoek et al., 2012), and is supported by recent successes of LLM-based surrogates in language-driven domains (Liu et al., 2024; Requeima et al., 2024). Each objective $\hat{r}_j \in \mathbb{R}$ for $j \in [k]$ is estimated independently using the LLM, based on input $(\tilde{m}, \tilde{a}, p, f, \mathcal{S}^f)$, where $\hat{r}_j = p_{SM}(\tilde{m}, \tilde{a} | p, f, \mathcal{S}^f)$. Multiple predictions are sampled in parallel to construct an empirical predictive distribution over the objectives. The prompt to achieve this is described in Figure 15.

1458
 1459 4. **MOSE Surrogate Model of Subjective Evaluations.** LLMs also serve as a surrogate for subjective
 1460 human judgment, enabling a generalized, cross-framework assessment mechanism for qualitative
 1461 criteria. Prior work has demonstrated that LLMs can effectively model human preferences
 1462 in alignment, safety, and prose diversity tasks (Bai et al., 2022; Bradley et al., 2024). In our
 1463 framework, MOSE uses an LLM to predict whether a proposed model m_t is subjectively preferred
 1464 over baseline models $m_i \in \mathcal{M}_{\text{ref}}$, given a problem description p . This is formalized as $p_{\text{MOSE}}(m_t \succ m_i | p)$. A prediction of ‘1’ indicates that m_t is subjectively superior. The associated token
 1465 probabilities are extracted. This process is repeated for each $m_i \in \mathcal{M}_{\text{ref}}$, where the preference
 1466 scores against each reference model are averaged to compute the final MOSE score. The prompt
 1467 structure depicted in Figure 16.

1468
 1469 You are an expert modeling assistant. Your task is to help the user
 1470 create a formal model to solve their problem.
 1471
 1472 ****Task:**** You will receive a description of the problem and the
 1473 desired objective(s) of the model. Your job is to propose a
 1474 modeling framework that can be used to solve the problem. You
 1475 should also provide a detailed explanation of your proposed
 1476 framework, including any assumptions or constraints that you are
 1477 making.
 1478
 1479 ****Problem description:****
 {GENERAL_PROBLEM_DESCRIPTION}
 1480
 1481 ****Problem instance descriptions:****
 {INSTANCE_DESCRIPTION}
 1482
 1483 ****Output format requirement:****
 - You must output your response as a single, valid JSON object.
 - No other text should precede or follow the JSON. The JSON object
 1484 must strictly follow this structure:
 {OUTPUT_FORMAT_REQUIREMENT}

Figure 11: Prompt structure for **framework proposal**.

1487
 1488
 1489
 1490
 1491 {
 1492 "modeling_framework" (string): "concise terminology to generally
 1493 describe the modeling framework (e.g., mathematical
 1494 optimization, dynamical systems)",
 1495 "framework_description" (string): "high-level description of the
 1496 proposed modeling framework"
 1497 }

Figure 12: Output format for **proposed frameworks**.

D.3 ADAPTIVE SEARCH

1503 Having detailed the implementation of LLM search operators in MATHMO, we now cover various
 1504 implementation details of the adaptive search process. Subsequently, we tabulate the key hyperpa-
 1505 rameters and describe the computational resources.

1506 **Upper-level: framework selection.** At the beginning of the search process ($t = 0$), an initial set
 1507 of w candidate frameworks is proposed independently using the LLM-based sampler $f \sim p_{\theta}(\cdot | p)$.
 1508 Then in each iteration:

1509 1. **Compute statistics.** For each framework, we compute summary statistics using the framework-
 1510 specific historical performance vectors $\{r_t | (m_t, a_t, r_t) \in \mathcal{S}^f\}$, specifically the empirical mean
 1511 $\mu_f \in \mathbb{R}^k$ and variance $\sigma_f^2 \in \mathbb{R}^k$.

```

1512 You are an expert modeling assistant. Your task is to help the user
1513 create a formal model to solve their problem.
1514
1515 **Task:** You will receive a description of the problem and the
1516 desired objective(s) of the model. Your job is to return the
1517 model that can generate the required output/solution in the
1518 output format specified. The model you generate should belong to
1519 the modeling framework specified.
1520
1521 **Problem description:**
1522 {GENERAL_PROBLEM_DESCRIPTION}
1523
1524 **Problem instance descriptions:**
1525 {INSTANCE_DESCRIPTION}
1526
1527 **Chosen modeling framework:**
1528 {MODELING_FRAMEWORK}
1529
1530 **Output format requirements:**
1531 {OUTPUT_FORMAT_REQUIREMENTS}
1532

```

Figure 13: Prompt structure for **model/algorithm proposal**.

```

1532 {
1533     "model" (Python code): "complete Python code of the model and
1534         algorithm generated to represent and solve the provided problem
1535         .",
1536     "dependencies" (list): "list of external Python package dependency"
1537         ,
1538     "model_explanation" (str): "detailed description of the generated
1539         model and algorithm"
1540 }

```

Figure 14: Output format for **proposed model/algorithm**.

2. **Compute UCB.** As our $\alpha(\cdot, \cdot)$ is implemented using the Pareto-UCB policy, we compute the UCB vector $UCB_f \in \mathbb{R}^k$ for each framework using Equation (2).
3. **Identify Pareto set.** Using the set of UCB vectors, the set of Pareto optimal (non-dominated) frameworks is identified.
4. **Selection.** If there exists more than one framework in the Pareto-UCB set, one framework is randomly selected to be explored next.

Lower-level: local exploration. In each iteration of lower-level exploration, the following steps occur:

1. **Proposal.** A set of candidate model-algorithm pairs are sampled, denoted as $\tilde{\mathcal{S}}^f = \{\tilde{m}^{(i)}, \tilde{a}^{(i)} | i \in [l]\}$.
2. **Surrogate estimation.** For each candidate pair, we obtain an estimated objective vector $\hat{r}^{(i)} = p_{SM}(\cdot | \tilde{m}^{(i)}, \tilde{a}^{(i)}; p, f, \mathcal{S}^f)$.
3. **Selection.** The pair (m_t, a_t) that yields the largest estimated hypervolume improvement $(m, a) = \arg \max_{(\tilde{m}, \tilde{a}) \in \tilde{\mathcal{S}}^f} HV(\tilde{m}, \tilde{a}; r_{ref})$ is selected to undergo evaluation.
4. **Evaluation.** We execute the model and algorithm as a subprocess. If any errors occurred during execution, the error trace is extracted and passed to the LLM to fix any mistakes and regenerate the model. If the model still does not execute after 3 MAX_ATTEMPTS, the process returns to the proposal step. Note that there is also a 300 second TIMEOUT imposed on model execution. If the model timeout, it is forcefully terminated.
5. **Observe objectives.** The objective vectors $r \in \mathbb{R}^k$ are obtained by evaluating the model and generated outputs, and the triplet are added to the set of models $\mathcal{S}^f \leftarrow \mathcal{S}^f \cup (m, a, r)$.

```

1566 You are an expert evaluation assistant. Your task is to help the user
1567 evaluate models that were generated for a particular problem.
1568
1569 **Task:** Your role is to help evaluate mathematical models designed
1570 for a specific problem. You will be given:
1571 - A description of the problem, the objective(s) the models are
1572 intended to achieve,
1573 - A history of previously generated models along with their
1574 performance metrics,
1575 - A candidate model for evaluation.
1576 Based on this information, predict the likely performance of the new
1577 candidate model.
1578
1579 **Problem description:**
1580 {GENERAL_PROBLEM_DESCRIPTION}
1581
1582 **Problem instance descriptions:**
1583 {INSTANCE_DESCRIPTION}
1584
1585 **History of proposed models:**
1586 {MODEL_HISTORY}
1587
1588 **Output format requirements:**
1589 {OUTPUT_FORMAT_REQUIREMENTS}

```

Figure 15: Prompt structure for **surrogate evaluations**.

```

1590 You are an expert at evaluating the subjective quality of models.
1591 Your task is to assess the subjective quality of a model based on
1592 its description against the baseline reference provided.
1593
1594 **Task:** You will receive a target model description (in code) and a
1595 baseline model description (also in code). Your job is to assess
1596 the {CRITERION} of the target model compared to the baseline
1597 reference. Provide a detailed assessment using the specified
1598 output format.
1599
1600 **Subjective criterion assessed:**
1601 {CRITERION_DESCRIPTION}
1602
1603 **Important instructions:**
1604 - Your assessment/explanation should be grounded in semantic meaning
1605 of the target model and the baseline reference.
1606 - The assessment should be based on your best intuition and semantic
1607 understanding of the models.
1608 - Then score the target model, returning 1 if the target model is
1609 more {CRITERION_VALUE} than the baseline reference model, or 0 if
1610 not.
1611
1612 **Target model:**
1613 {TARGET_MODEL}
1614
1615 **Baseline reference model:**
1616 {BASELINE_REF}

```

Figure 16: Prompt structure for **surrogate evaluations**.

D.4 MISCELLANEOUS IMPLEMENTATION DETAILS

Key hyperparameters. We detail the hyperparameters for implementing MATHMO in Table 12.

1620 **LLM.** We use `gpt-4o-2024-05-13` as the underlying LLM in all experiments.
 1621

1622 **Computer resources.** We run all experiments on an AMD EPYC 7V13 64-Core Processor.
 1623

1624 **Code and reproducibility.** The full implementation of **MATHMO**, along with the code necessary
 1625 to reproduce all key results, will be released on GitHub upon acceptance of the paper.
 1626

1627 **Table 12: Description of key hyperparameters.**

Hyperparameter	Description	Value
w	Number of frameworks	4
l	Number of candidate model-algorithm pairs proposed in each iteration of local exploration	3
q	Number of MC estimates for surrogate evaluations	3
(c, d)	Pareto-UCB exploration bonus hyperparameters	(1, 1) (default)
$ \mathcal{M}_{\text{ref}} $	Number of baseline models used in MOSE evaluations	3
T	Number of search iterations	20
τ	LLM hyperparameter (sampling temperature)	0.7 (default)
p	LLM hyperparameter (top-p sampling)	0.9 (default)
TIMEOUT	Max runtime (in seconds) allowed for each model-algorithm pair to execute	300

1636 E ADDITIONAL EXPERIMENTAL DETAILS

1637 In this section of the appendix, we will describe the datasets and metrics employed in our empirical
 1638 evaluations.
 1639

1640 E.1 DATASETS

1641 **Traveling Salesman (TSP).** The Traveling Salesman Problem (TSP) is a foundational combinatorial
 1642 optimization problem where the objective is to find the shortest possible route that visits a set of
 1643 cities exactly once and returns to the starting point. The two primary trade-offs in modeling TSP
 1644 are **journey cost** (total tour length) and **runtime** (solution time). Journey cost reflects the quality of
 1645 the solution and is crucial in applications like logistics and manufacturing, while runtime is vital in
 1646 scenarios requiring rapid decisions, such as dynamic routing. Common modeling techniques include:
 1647 (1) exact methods such as Integer Linear Programming (ILP), which guarantee globally optimal
 1648 solutions but scale poorly with problem size; (2) metaheuristic approaches like Genetic Algorithms,
 1649 which offer faster approximate solutions at the cost of optimality; and (3) domain-specific heuristics
 1650 such as nearest-neighbor or insertion algorithms, which are simple and computationally efficient
 1651 ([Bellmore and Nemhauser, 1968](#)). For our experiments, we generated 10 random Euclidean instances
 1652 with 30–50 cities each. Each instance was created by uniformly sampling city coordinates in a 2D
 1653 unit square, a standard method for generating synthetic TSP datasets.
 1654

1655 **Job Shop Scheduling (JSS).** Job Shop Scheduling is a canonical operations research problem that
 1656 involves assigning a sequence of jobs to a set of machines, where each job consists of a series
 1657 of operations with specific processing requirements. The primary objectives are to minimize the
 1658 **makespan** (i.e., the total time required to complete all jobs) and to reduce **runtime**, which becomes
 1659 critical in dynamic or large-scale industrial systems. A lower makespan increases throughput,
 1660 directly impacting productivity, while efficient computation ensures that schedules can be adapted
 1661 in real time. (1) Constraint Programming (CP), which provides precise encodings but may struggle
 1662 with scalability; (2) metaheuristic methods such as Tabu Search or Simulated Annealing, which
 1663 balance exploration and exploitation to find high-quality solutions efficiently; and (3) greedy or
 1664 rule-based heuristics, which offer speed and interpretability at the cost of optimality ([Xiong et al., 2022](#)).
 1665 We use 10 well-known benchmark instances from [Lawrance \(1984\)](#), which are widely
 1666 adopted in the scheduling literature. These instances span 10 to 30 jobs, 5 to 10 machines, and 50
 1667 to 300 operations. The dataset is available through the open-source `job_shop_lib` repository
 1668 (https://github.com/Pablooo22/job_shop_lib).
 1669

1670 **Ecology.** Ecological modeling seeks to understand and predict interactions among species and their
 1671 environments, often involving dynamic systems such as predator-prey relationships. A key modeling
 1672 trade-off in ecology lies between **predictive performance**—capturing future population dynamics
 1673 accurately—and **interpretability**, which is critical for gaining ecological insights and informing
 1674 conservation efforts. Common modeling approaches include: (1) differential equation systems
 1675 such as Lotka-Volterra models, which provide interpretable representations of species interactions;

(2) classical time-series models like ARIMA, which are effective for short-term forecasts; and (3) probabilistic graphical models, which capture structured uncertainty and latent ecological processes (van den Berg et al., 2022). We use a dataset containing Hare-Lynx populations (Stenseth et al., 1997), which records annual observations of the Snowshoe Hare and Canadian Lynx populations over multiple decades.

Epidemiology. Epidemiological modeling focuses on understanding and forecasting the spread of infectious diseases, often under constraints that demand both accurate **prediction** and clear **interpretability** for public health decision-making. Predictive performance ensures that interventions can be timed effectively, while interpretability allows stakeholders to understand transmission mechanisms and policy implications. Common approaches include: (1) compartmental models (e.g., SIR, SEIR), which capture the flow of individuals through disease states using differential equations; (2) statistical models such as Poisson and negative binomial regressions, which model count data under uncertainty; and (3) time-series forecasting techniques, including autoregressive and neural models, for flexible temporal prediction (Xiang et al., 2021). For our experiments, we use COVID-19 time series data from Italy, sourced from the COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (Dong et al., 2020). The dataset contains daily counts of confirmed cases, deaths, and recoveries.

We note that in all our experiments, the dataset is provided as an input only at runtime—after the model has been generated by the LLM. The LLMs themselves do not have access to the dataset contents during model generation; they are only given high-level metadata, such as the number of features, problem size (e.g., number of operations or cities), or time series length, to inform their proposals.

E.2 METRICS

Hypervolume. The hypervolume (HV) metric quantifies the volume of the objective space that is dominated by a set of solutions, relative to a fixed reference point. It serves as a standard measure in multi-objective optimization, capturing both convergence and diversity of the solution set. Formally, let $\mathcal{R} = r_1, \dots, r_n$ be a set of n k -dimensional objective vectors and let $r_{\text{ref}} \in \mathbb{R}^k$ be a reference point that is dominated by all vectors in \mathcal{R} . The hypervolume is defined as:

$$\text{HV}(\mathcal{R}; r_{\text{ref}}) = \lambda \left(\bigcup_{r \in \mathcal{R}} [r_1, r_{\text{ref},1}] \times \dots \times [r_k, r_{\text{ref},k}] \right) \quad (5)$$

where λ denotes the Lebesgue measure in \mathbb{R}^k . All objectives are first normalized to $[0, 1]$, and we set $r_{\text{ref}} = 1.1$ to ensure it lies outside the normalized Pareto front. We compute hypervolume using the `pymoo` library (Blank and Deb, 2020) (<https://pypi.org/project/pymoo/>).

Relative Hypervolume Improvement. To assess progress over time, we compute the relative hypervolume improvement, which quantifies the gain in hypervolume relative to the best value achieved at a previous timestep. Let HV_t and $\text{HV}_{t'}$ denote the HV at iteration t and t' , where $t' < t$.

$$\text{RHI}_t = \frac{\text{HV}_t - \text{HV}_{t'}}{\text{HV}_{t'}} \quad (6)$$

We employ this metric to compare the impact of search strategies over the course of search.

Normalized RMSE. Root Mean Squared Error (RMSE) is a standard regression metric that measures the average magnitude of prediction error. In our context, we use a normalized RMSE (NRMSE) to account for scale differences across objectives. Given a set of ground-truth values $\{y_i\}_{i=1}^n$ and predictions $\{\hat{y}_i\}_{i=1}^n$, RMSE is defined as:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (7)$$

We normalize this by the empirical range of the true values σ_y , yielding:

$$\text{NRMSE} = \frac{\text{RMSE}}{\sigma_y} \quad (8)$$

1728
 1729 **Permutation entropy.** Permutation entropy (PE) is a model-free measure of complexity for time
 1730 series or ordered sequences, capturing the unpredictability of local ordering patterns. Given a
 1731 time series $\{x_t\}_{t=1}^T$, the sequence is partitioned into overlapping windows of length d (embedding
 1732 dimension), and each window is mapped to a permutation pattern based on the relative ordering of its
 1733 elements. Let π_i denote the i -th unique pattern and $p(\pi_i)$ its empirical frequency. The PE is then
 1734 defined as:

$$H_d = - \sum_i p(\pi_i) \log p(\pi_i) \quad (9)$$

1735 which is often normalized by $\log(d!)$ to yield a value in $[0, 1]$. We compute PE using the `antropy`
 1736 library (<https://pypi.org/project/antropy/>). This metric provides a lens into the
 1737 structural complexity of sequences produced by time-series models

1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781