

AUTOMATED ALGORITHMIC DISCOVERY FOR GRAVITATIONAL-WAVE DETECTION GUIDED BY LLM-INFORMED EVOLUTIONARY MONTE CARLO TREE SEARCH*

Anonymous authors

Paper under double-blind review

ABSTRACT

Gravitational-wave signal detection with unknown source parameters buried in dynamic detector noise remains a formidable computational challenge. Existing approaches face core limitations from restrictive assumptions: traditional methods rely on predefined theoretical priors, while neural networks introduce hidden biases and lack interpretability. We propose Evolutionary Monte Carlo Tree Search (Evo-MCTS), the first integration of large language model (LLM) guidance with domain-aware physical constraints for automated gravitational wave detection. This framework systematically explores algorithmic solution spaces through tree-structured search enhanced by evolutionary optimization, combining MCTS for strategic exploration with evolutionary algorithms for solution refinement. The LLM component provides domain-aware heuristics while maintaining interpretability through explicit algorithmic pathway generation. Experimental validation demonstrates substantial performance improvements, achieving a 20.2% improvement over state-of-the-art gravitational wave detection algorithms on the MLGWSC-1 benchmark dataset and a remarkable 59.1% improvement over other LLM-based algorithm optimization frameworks. Beyond performance improvements, our framework establishes a transferable methodology for automated algorithmic discovery across computational science domains.

1 INTRODUCTION

The pursuit of scientific discovery increasingly demands computational approaches that can navigate complex, high-dimensional data spaces while maintaining physical interpretability [Zheng et al. \(2025a\)](#); [Wang et al. \(2023\)](#); [Karniadakis et al. \(2021\)](#); [Baker et al. \(2019\)](#). In gravitational wave astronomy, this computational challenge manifests as a fundamental algorithmic problem: detection systems must identify faint astrophysical signals buried in detector noise while leveraging theoretical predictions from general relativity [Abbott et al. \(2016; 2019\)](#), yet remain adaptable to unexpected signal morphologies that could reveal new physics beyond current theoretical models.

Gravitational wave detection serves as an exemplary case study for automated algorithm discovery, embodying the fundamental tension between model-driven precision and data-driven flexibility [Abbott et al. \(2020\)](#). Current methodologies represent three complementary yet insufficient paradigms: Matched filtering techniques [Owen \(1996\)](#); [Cutler & Flanagan \(1994\)](#) achieve optimal sensitivity under Gaussian stationary noise assumptions but critically depend on accurate prior knowledge of signal morphologies. Non-template methods [Klimenko et al. \(2016b\)](#) offer model-independent detection but sacrifice sensitivity for generality. Deep neural networks [George & Huerta \(2018\)](#); [Gabbard et al. \(2018\)](#); [Huerta et al. \(2021\)](#) provide computational efficiency but operate as black-box architectures that obscure decision-making logic and introduce hidden biases [Nagarajan &](#)

*See also: <https://arxiv.org/abs/2508.03661> and <https://iphysresearch.github.io/evo-mcts/>

054 [Messenger \(2025\)](#). Each paradigm represents different trade-offs between sensi-
055 tivity, generality, and interpretability, yet none adequately addresses the challenge
056 of discovering novel algorithmic approaches that could transcend these traditional
057 limitations.

058 The fundamental bottleneck in modern scientific computing is algorithmic sophis-
059 tication rather than computational power or data availability [Baker et al. \(2019\)](#);
060 [Karniadakis et al. \(2021\)](#). Current gravitational wave detection pipelines represent
061 decades of manual engineering yet still miss potentially discoverable signals due
062 to three critical limitations: (i) *combinatorial explosion* of possible signal process-
063 ing combinations that exceeds human exploration capacity [Elsken et al. \(2019\)](#);
064 [Eiben & Smith \(2015\)](#), (ii) *cognitive biases* that constrain designers to familiar
065 paradigms [Kahneman \(2011\)](#), and (iii) *local optimization traps* where manual re-
066 finement leads to incremental improvements while missing global optima [Wolpert
& Macready \(2002\)](#). These limitations necessitate automated algorithm discovery
067 approaches that can navigate vast design spaces while maintaining scientific rigor.
068

069 Automated algorithm discovery in scientific domains presents unique challenges
070 requiring simultaneous navigation of *semantic coherence* (algorithms must respect
071 physical laws), *exploration efficiency* (vast search spaces with expensive evalua-
072 tions), and *interpretability preservation* (discoverable and validatable by experts).
073 The interdependent nature of these challenges suggests effective solutions must
074 address all three simultaneously. Large Language Models (LLMs) emerge as
075 natural candidates for semantic coherence, having been trained on scientific liter-
076 ature and capable of incorporating domain knowledge directly into code gener-
077 ation [Chen et al. \(2021\)](#); [Lewkowycz et al. \(2022\)](#). Monte Carlo Tree Search
078 (MCTS) provides structured exploration framework, creating hierarchical mem-
079 ory while balancing exploration with exploitation [Browne et al. \(2012\)](#); [Wang
et al. \(2020b\)](#); [Li et al. \(2025\)](#). Their synergy creates a self-reinforcing cycle:
080 LLMs provide semantic understanding and domain knowledge integration, MCTS
081 provides structured exploration and cumulative learning, enabling evolutionary
082 operations that preserve algorithmic coherence while systematically exploring sci-
083 entific algorithm spaces.

084 Existing approaches to algorithm discovery span multiple paradigms. Tradition-
085 al approaches including genetic programming [Koza \(1994\)](#), neural architec-
086 ture search [Elsken et al. \(2019\)](#), and evolutionary computation [Eiben & Smith
\(2015\)](#) suffer from critical limitations: generating syntactically invalid code, lack-
087 ing domain knowledge integration, or focusing on network topology rather than
088 algorithmic logic. Recent advances have integrated LLMs with structured search
089 and evolutionary methods, including FunSearch [Romera-Paredes et al. \(2024\)](#),
090 EoH [Liu et al. \(2024\)](#), AEL [Liu et al. \(2023\)](#), ReEvo [Ye et al. \(2024\)](#), and
091 MCTS-AHD [Zheng et al. \(2025b\)](#) for algorithm discovery, which combine LLMs’
092 code generation capabilities with systematic optimization frameworks. However,
093 the existing LLM-based frameworks focus on combinatorial optimization tasks
094 with discrete decision sequences and well-defined mathematical formulations.
095 Scientific signal processing problems like gravitational wave detection present
096 fundamentally different challenges, requiring continuous parameter optimization,
097 domain-specific physical constraints, and interpretable algorithmic pathways that
098 can be validated against theoretical predictions.

099 We propose Evo-MCTS (Evolutionary Monte Carlo Tree Search), a framework
100 that realizes synergistic integration through three key innovations: (i) *Reflective
Code Synthesis* that leverages LLM domain knowledge for performance-driven
101 algorithm generation, adapting to optimization landscapes while maintaining sci-
102 entific validity, (ii) *Multi-Scale Evolutionary Operations* (Parent/Sibling/Path-
103 wise Crossover, Point Mutation) that operate on structured code representations
104 through MCTS tree traversal, enabling semantic-aware algorithmic transformations,
105 and (iii) *Interpretable Algorithm Pathways* that emerge naturally from the
106 tree structure, enabling post-hoc analysis of algorithmic evolution while provid-
107 ing cumulative learning for future discoveries. The framework addresses interde-

pendent challenges through architectural design—each element enhances others’ capabilities while compensating for individual limitations.

We demonstrate Evo-MCTS effectiveness through comprehensive evaluation on gravitational wave detection benchmarks, achieving a 20.2% improvement over state-of-the-art algorithms on the MLGWSC-1 benchmark dataset and a remarkable 59.1% improvement over other LLM-based algorithm optimization frameworks. The framework demonstrates systematic breakthrough discoveries across independent executions, with high-performing variants consistently incorporating sophisticated signal processing techniques including Multi-resolution Thresholding, Continuous Wavelet Transform with Ricker wavelets, Tikhonov Regularization, and Curvature Boosting, among others. These discoveries generate human-interpretable algorithmic pathways that reveal distinct performance patterns organized by functional categories, identifying novel algorithmic combinations that human designers might overlook while providing empirical guidance for component selection in complex signal processing pipelines.

Beyond gravitational wave astronomy, this work establishes a general paradigm for AI-guided algorithmic discovery in scientific computing. The Evo-MCTS framework’s ability to generate interpretable algorithmic pathways while maintaining high performance makes it particularly valuable for scientific applications where understanding algorithmic reasoning is as important as achieving optimal performance. Our approach opens new avenues for automated scientific discovery across physics, chemistry, biology, and engineering disciplines, providing a transferable methodology that respects domain-specific constraints while systematically exploring algorithmic possibilities beyond human intuition.

2 RESULTS

2.1 FRAMEWORK ARCHITECTURE FOR AUTOMATED ALGORITHM DISCOVERY

Algorithmic Discovery Pipeline. The Evo-MCTS framework systematically transforms raw time-series data into a comprehensive catalog of optimized detection algorithms through an automated discovery pipeline that integrates domain knowledge encoded in large language models (Figure 1a, see Methods Section 4.1 for formal problem definition). To understand this automated transformation process, the framework can be conceptualized through two complementary perspectives (Figure 1b): as an MCTS-guided tree search where nodes represent complete algorithmic implementations and edges encode LLM-driven transformations (detailed implementation in Methods Section 4.2), or as an evolutionary algorithm where populations of algorithms undergo sophisticated selection, crossover, and mutation operations guided by domain knowledge (detailed implementation in Methods Section 4.3).

MCTS-Guided Evolutionary Exploration. The core innovation of our Evo-MCTS framework lies in reformulating algorithm design as a tree search problem, where each node represents an executable algorithm and edges correspond to code transformations (Figure 1c). Starting from a seed algorithm, the framework employs four specialized evolutionary operations to expand the search tree:

- *Parent Crossover (PC)*: Combines algorithmic features from parent nodes to generate offspring that inherit successful detection strategies while exploring new combinations.
- *Sibling Crossover (SC)*: Enables horizontal knowledge transfer between algorithms at the same tree depth, promoting diversity while maintaining comparable complexity levels.
- *Path-wise Crossover (PWC)*: Synthesizes information across complete root-to-leaf trajectories, capturing long-range dependencies and enabling global optimization strategies.

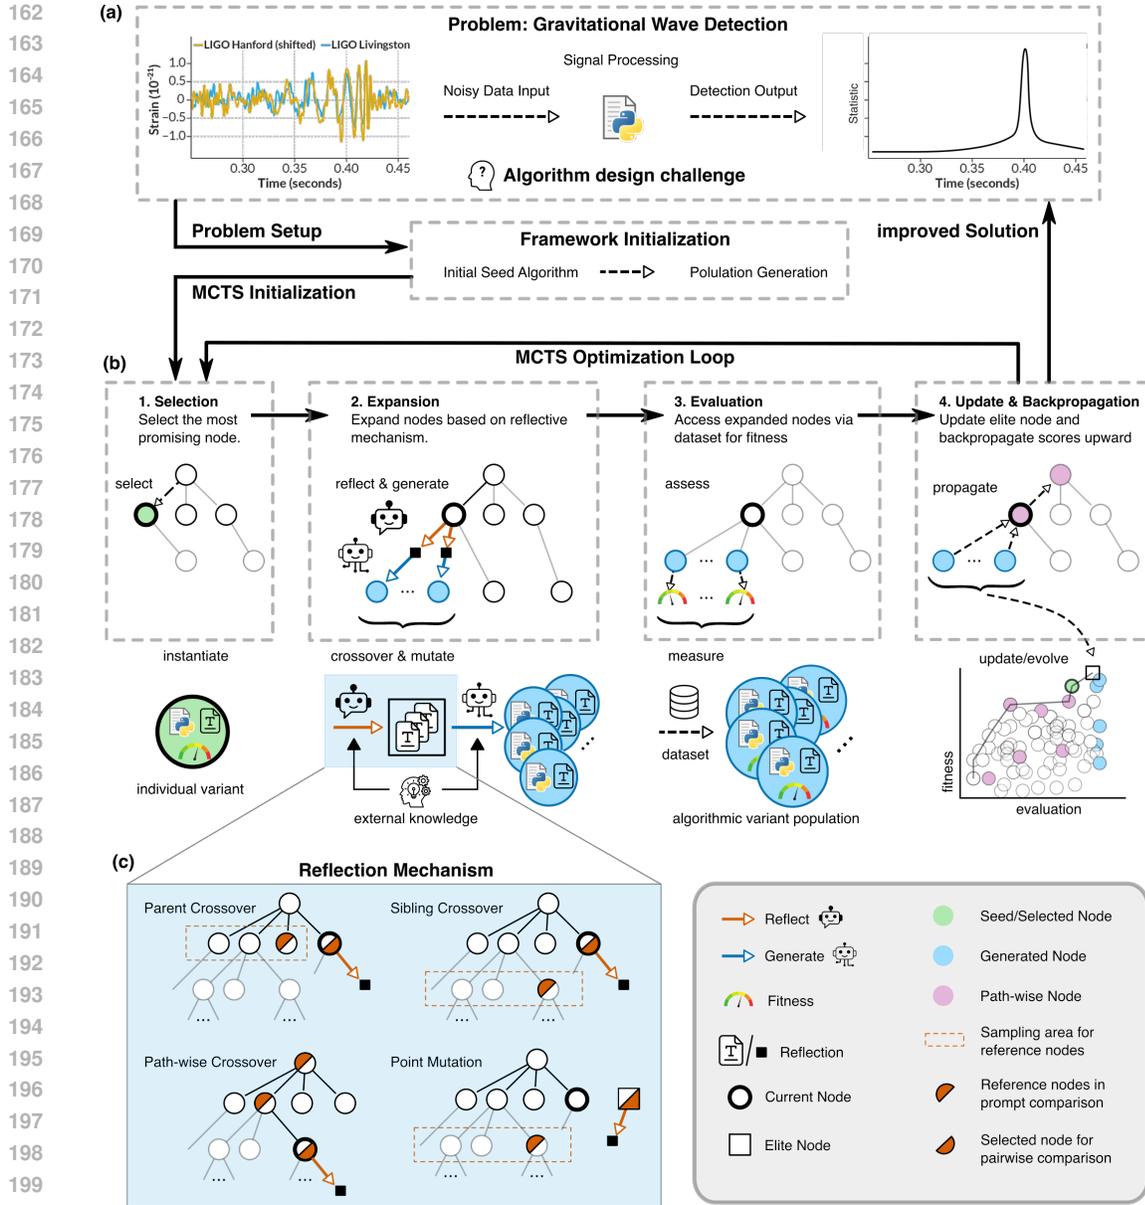


Figure 1: LLM-Informed Evolutionary Monte Carlo Tree Search Framework for Automated Algorithm Discovery. (a) Overview of the algorithm discovery pipeline. Starting from raw gravitational wave strain data (left), the framework applies automated algorithmic transformations through LLM-generated code synthesis (center) to produce optimized detection statistics (right). (b) Core architectural components showing the integration of MCTS exploration with evolutionary optimization through dual perspectives of tree search and population evolution. (b.1) UCT-based node selection from initial algorithmic variants including seed algorithms and individual variants, each represented as nodes containing baseline signal processing code. (b.2) MCTS expansion phase where new algorithmic variants are generated through evolutionary operations. Each node contains executable Python code implementing specific detection strategies. (b.3) Algorithm evaluation phase where generated variants are tested against benchmark data to compute fitness scores, determining performance-based selection for subsequent iterations. (b.4) MCTS backpropagation and elite node updates after multiple evolutionary cycles, propagating performance feedback through the tree structure and maintaining diverse high-performing detection strategies. (c) Detailed view of the reflection mechanism during MCTS expansion, showing four evolutionary operations: Parent Crossover, Sibling Crossover, Path-wise Crossover, and Point Mutation.

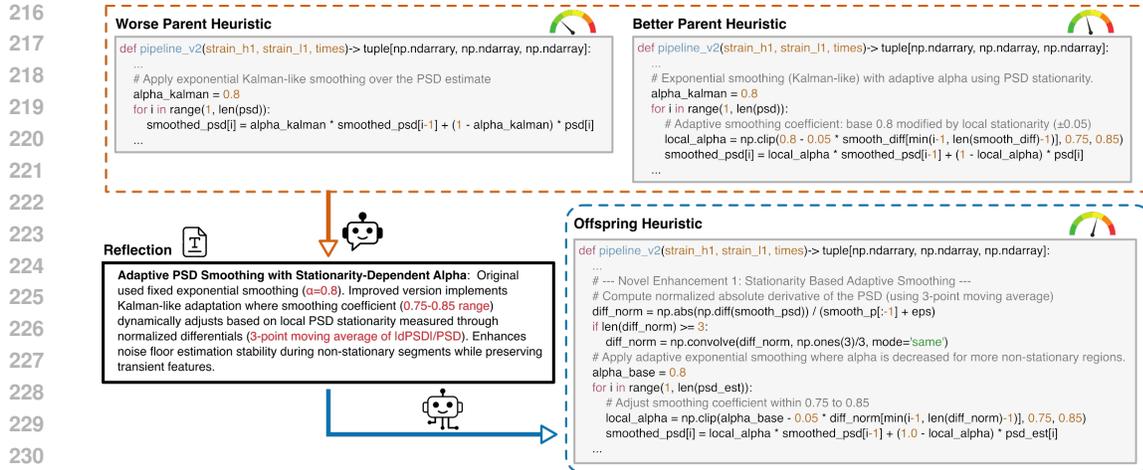


Figure 2: **LLM-Driven Algorithmic Evolution Through Reflective Code Synthesis.** Demonstration of a single Parent Crossover evolutionary step showing the transformation from two parent algorithms to an enhanced offspring algorithm. **(Top row)** Code segments from two parent nodes highlighting complementary algorithmic components that will be combined through the crossover operation. **(Bottom left, black box)** Reflective analysis process showing how the LLM identifies strengths and limitations in the parent algorithms, synthesizing insights about their respective detection strategies and potential synergies. **(Bottom right)** Generated offspring algorithm code incorporating successful elements from both parents while addressing identified limitations through domain-aware synthesis. This example illustrates the framework’s capability to generate physically-motivated algorithmic improvements through automated reasoning, demonstrating how LLM-guided reflection enables discovery of sophisticated signal processing techniques by combining and enhancing existing algorithmic components without manual intervention. The complete reflection prompts and additional evolution examples are provided in Supplementary Material A.1.

- *Point Mutation (PM)*: Introduces targeted modifications to individual algorithms based on performance analysis, leveraging insights from elite algorithms to enable fine-grained optimization of specific components.

These operations are fundamentally different from traditional genetic algorithms because they operate on structured code representations rather than abstract encodings, and modifications are guided by LLM-based reasoning rather than random perturbations (see Methods Section 4.3 and Supplementary Material A.1 for operation details).

Reflection-Driven Code Synthesis. Central to the Evo-MCTS framework’s effectiveness is the reflection mechanism that analyzes algorithm performance patterns and guides subsequent explorations (Figure 1b-(2) and Figure 1c). This dual-component system comprises: (i) a performance reflection module that identifies strengths and weaknesses in current algorithms through systematic evaluation across diverse signal conditions, and (ii) a code synthesis module that leverages these insights to generate improved implementations (An example is shown in Figure 2). The reflection process operates at multiple scales—from individual algorithmic components to complete detection pipelines—ensuring both local optimization and global coherence (detailed prompts and examples in Supplementary Information Section A.1).

2.2 PERFORMANCE BENCHMARKING ON MLGWSC-1 DATASET

We evaluated our Evo-MCTS framework’s algorithmic discovery capabilities using the first official Machine Learning Gravitational-Wave Search Mock Data Challenge (MLGWSC-1) benchmark dataset Schäfer et al. (2023)—an internationally recognized and rigorous evaluation standard established by the gravi-

270 tational wave detection community. The MLGWSC-1 benchmark represents a
271 comprehensive and challenging assessment framework that includes mock detec-
272 tor data with simulated gravitational wave signals embedded in realistic detector
273 noise, enabling systematic comparison of detection algorithms across diverse sig-
274 nal morphologies and noise conditions.

275 **Framework Adaptation for Domain-Specific Evaluation.** Figure 3a illustrates
276 the Evo-MCTS framework’s adaptation to the MLGWSC-1 evaluation protocol,
277 showcasing the critical integration between automated algorithm generation and
278 standardized performance assessment. The system transforms raw dual-channel
279 strain data through our evolved algorithms, producing detection catalogs that are
280 evaluated against the ground truth signal catalog. The evaluation pipeline incor-
281 porates the area under the curve (AUC) metric as the primary performance indi-
282 cator, computed from sensitivity distance versus false alarm rate curves spanning
283 4-1000 events per month. This comprehensive metric balances detection sensi-
284 tivity against false alarm rates—a critical trade-off in gravitational wave detection
285 applications—ensuring that algorithmic optimization targets practical deployment
286 requirements rather than narrow performance metrics (detailed experimental con-
287 figuration provided in Methods Section 4.5).

288 **Progressive Complexity and Performance Improvements.** Figure 3b presents
289 the comprehensive optimization trajectory across 877 total evaluations from five
290 independent Evo-MCTS runs, revealing systematic algorithmic discovery with
291 progressive increases in both algorithmic sophistication and detection perfor-
292 mance. The optimization process exhibits four distinct phase transitions (PT 1-4),
293 each marking algorithmic breakthroughs that represent the discovery of increas-
294 ingly sophisticated algorithmic components building upon previous innovations.
295 The combined fitness trajectory demonstrates the framework’s capability to navi-
296 gate the complex algorithmic design space while maintaining consistent improve-
297 ment patterns across multiple independent runs, systematically exploring and in-
298 tegrating complex detection strategies (detailed analysis of algorithmic evolution
299 patterns provided in Supplementary Material A.3).

300 **Maintaining Solution Diversity while Converging toward Optimal Perfor-**
301 **mance.** The diversity analysis reveals sophisticated exploration patterns that bal-
302 ance algorithmic variety with performance optimization throughout the search
303 process. Peak diversity occurs during the intermediate optimization phase be-
304 tween PT 2 and PT 3, where both Shannon diversity index and Complexity In-
305 dex of Diversity reach maximum values, indicating the framework maintains di-
306 versity in both component selection and structural complexity while systemati-
307 cally exploring combinations of successful components before converging on opti-
308 mal configurations. Fitness-stratified analysis demonstrates systematic conver-
309 gence: diversity decreases progressively from broad initial exploration in lower-
310 performing variants to focused refinement in highest-performing configurations,
311 confirming effective balance between exploration and exploitation (diversity met-
312 ric definitions and computational details provided in Methods Section 4.5).

312 **Superior Performance Against State-of-the-Art Methods.** Figure 3c demon-
313 strates the Evo-MCTS framework’s superior sensitivity performance against seven
314 benchmark gravitational wave detection algorithms on MLGWSC-1, Set 4 dataset.
315 The PT-4 configuration achieves a 20.2% improvement over state-of-the-art meth-
316 ods including Sage Nagarajan & Messenger (2025), Virgo-AUTH Nousi et al.
317 (2023); Schäfer et al. (2023), PyCBC Nitz et al. (2023); Schäfer et al. (2023),
318 TPI FSU Jena Zelenka et al. (2024); Schäfer et al. (2023), cWB Drago et al.
319 (2021); Schäfer et al. (2023), MFCNN Wang et al. (2020a); Schäfer et al. (2023),
320 and CNN-Coinc Schäfer & Nitz (2022); Schäfer et al. (2022; 2023). Most
321 significantly, at the stringent low false alarm rate of approximately 4 events
322 per month—the regime most critical for confident astrophysical detection—PT-4
323 achieves a 23.4% performance improvement over the highest-performing bench-
mark (Sage), with progressive algorithmic sophistication demonstrated through
four milestone configurations (PT-1 through PT-4) across the false alarm rate

range of 4-1000 events per month (detailed performance metrics provided in Supplementary Material A.3).

Interpretable Nonlinear Algorithm Discovery. The systematic improvement from PT-1 to PT-4 demonstrates our framework’s capability to discover sophisticated nonlinear filtering algorithms that surpass traditional approaches across different algorithmic paradigms. While traditional matched filtering pipelines like PyCBC Nitz et al. (2023) represent the theoretical optimum for Gaussian stationary noise conditions Finn (1992); Cutler & Flanagan (1994), real detector noise exhibits non-Gaussian and non-stationary characteristics that limit the effectiveness of linear correlation operations despite incorporating nonlinear post-processing stages Usman et al. (2016); Dal Canton et al. (2021). Our evolved algorithms achieve superior performance through intrinsically nonlinear transformations that adapt dynamically to these realistic noise conditions, effectively addressing the fundamental limitations of matched filtering in practical detector environments Abbott et al. (2016; 2019).

Compared to coherent WaveBurst (cWB) Klimentko et al. (2016b;a), which shares our template-free philosophy and employs nonlinear wavelet-based detection methods, our framework demonstrates the effectiveness of systematic algorithmic exploration over heuristic design approaches. Both methods recognize that optimal detection strategies must transcend linear processing assumptions, but our automated discovery process identifies algorithmic configurations that cWB’s manually-designed heuristics cannot achieve. The performance gains reflect the fundamental advantage of exhaustive exploration over expert intuition in complex optimization landscapes Liu et al. (2024); Zhang et al. (2024).

Furthermore, our framework achieves a 20.2% performance improvement over state-of-the-art AI-based methods despite their inherent nonlinear processing capabilities, demonstrating that interpretable algorithmic discovery can achieve superior detection performance while maintaining complete transparency in decision logic. Unlike deep learning approaches that operate as black boxes with millions of parameters LeCun et al. (2015), our evolved algorithms provide explicit mathematical formulations that enable physical interpretation of detection mechanisms Rudin (2019); Molnar (2020). This interpretability advantage, combined with the progressive complexity enhancement observed across phase transitions, establishes the framework’s capability to discover sophisticated yet transparent algorithmic solutions that bridge the gap between model-driven precision and data-driven flexibility (detailed quantitative analysis and algorithmic specifications provided in Supplementary Material A.3).

2.3 INTERPRETABILITY ANALYSIS

2.3.1 ALGORITHM PERFORMANCE AND GENERALIZATION ANALYSIS

Generalization capability and robustness of optimized algorithms. We evaluated 877 algorithmic configurations on an independent 1-day test dataset, distinct from the 7-day training corpus, under a 0.2-second trigger arrival time uncertainty constraint to assess robustness and interpretability (data specifications in Supplementary Material A.2).

Figure 4a demonstrates strong training-test performance correlation ($r = 0.84$) across all algorithms, with each point representing AUC-based fitness scores for individual configurations. This correlation validates robust generalization despite significant domain shift between datasets, with performance variation reflecting the non-stationary, non-Gaussian noise characteristics inherent in realistic gravitational wave detection environments. The 0.2-second timing constraint ensures temporal precision essential for astrophysical parameter estimation while preserving detection sensitivity, empirically determined through systematic constraint-correlation analysis (Supplementary Material A.4). These results confirm that optimized algorithms achieve genuine performance improvements transferable to independent datasets, validating our evolutionary framework’s practical utility for real-world gravitational wave detection applications.

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

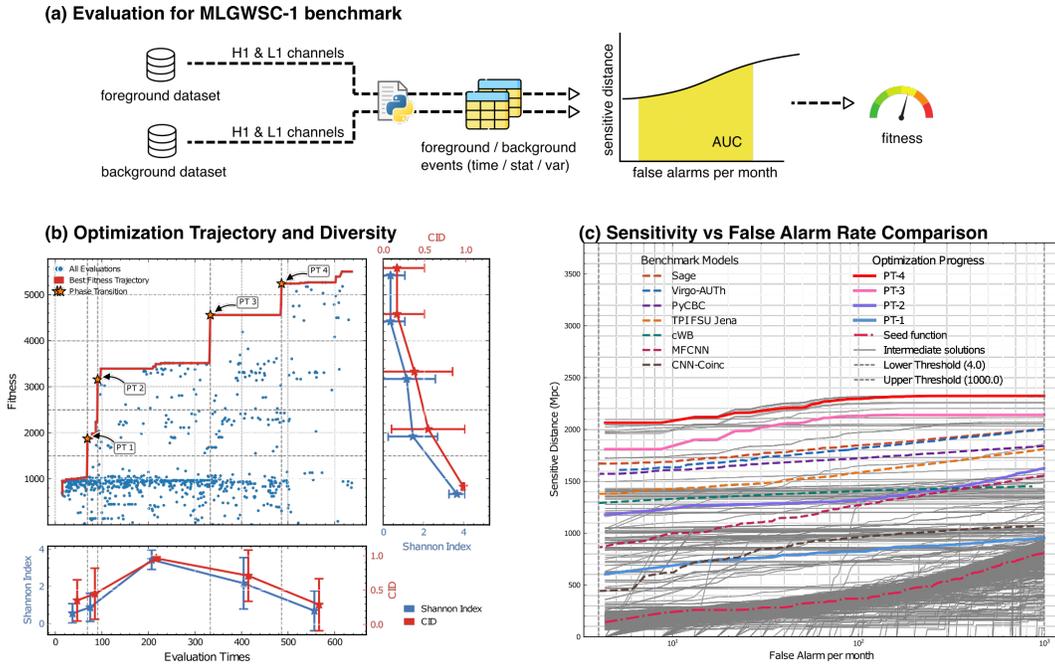


Figure 3: Comprehensive Performance Analysis on MLGWSC-1 Benchmark Dataset. (a) EvoMCTS framework adaptation pipeline showing domain-specific fitness evaluation using area under the curve (AUC) metric for sensitivity distance versus false alarm rate curves, demonstrating integration with standardized gravitational wave detection evaluation protocols. The pipeline processes dual-channel strain data from Hanford(H1) and Livingston(L1) detectors through evolved algorithms to produce detection statistics evaluated against ground truth catalogs. (b) Objective optimization trajectory and diversity analysis across 877 evaluations from 5 independent EvoMCTS runs. Combined fitness trajectory (blue dots) with best objective envelope (red line) showing four phase transitions (PT 1-4, orange stars) marking algorithmic breakthroughs with fitness gains ≥ 400 units. Maximum fitness of 5,241 units achieved, representing 6-fold improvement from baseline. Diversity metrics include Shannon diversity index (blue, left axis) and Complexity Index of Diversity (CID, red, right axis) with error bars showing standard deviation across runs. Right panel shows fitness-stratified diversity analysis revealing systematic exploration patterns across performance levels. (c) Sensitivity versus false alarm rate comparison on MLGWSC-1, Set 4 dataset. Optimization milestones PT-1 through PT-4 show progressive improvement, with PT-4 achieving AUC of 5,241 units outperforming seven benchmark algorithms (Sage, Virgo-AUTH, PyCBC, TPI FSU Jena, cWB, MFCNN, CNN-Coinc). Grey curves represent intermediate solutions explored during optimization, while red dotted line shows seed function baseline. Vertical dashed lines indicate evaluation range boundaries (4-1000 events per month). Results demonstrate systematic algorithmic discovery with superior sensitivity performance, controllable thresholds, and clear interpretability through progressive complexity enhancement. The optimization process systematically explored diverse parameter spaces while converging toward optimal configurations, integrating nonlinear transformations, adaptive parameter selection, and sophisticated statistical analysis methods that remain fully interpretable throughout the discovery process.

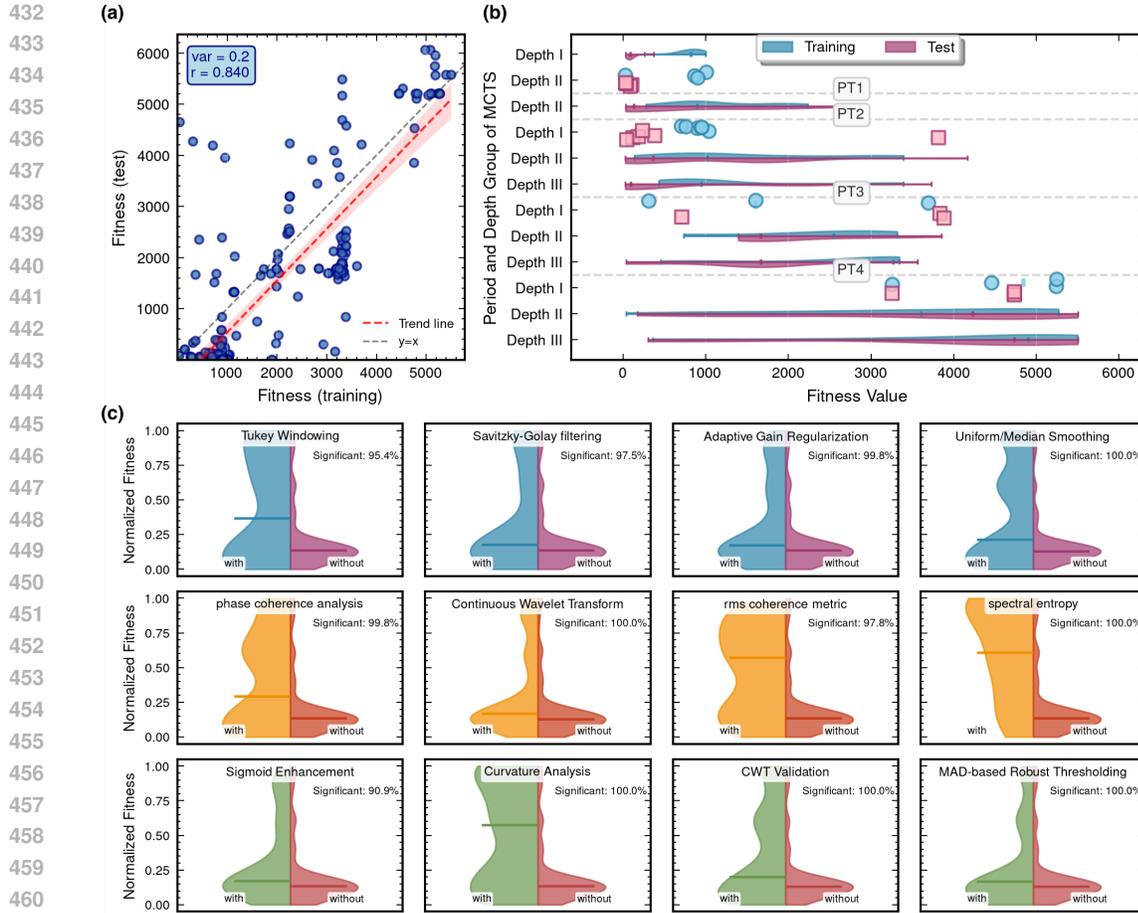


Figure 4: **Comprehensive Algorithm Performance Analysis.** (a) Training versus test performance correlation for 877 algorithmic configurations evaluated under 0.2-second trigger arrival time uncertainty constraint. Each point represents an individual algorithm's AUC-based fitness scores on training (7-day dataset) and test (1-day independent dataset) data. Linear correlation coefficient $r = 0.840$ indicates strong generalization capability, while variance reflects expected performance variation due to non-stationary, non-Gaussian noise characteristics. Red dashed line shows the empirical trend relationship, while grey dashed line represents perfect correlation ($y=x$). High-performing algorithms (fitness > 4000) demonstrate particularly robust generalization across different noise realizations and signal parameters. (b) MCTS depth-stratified performance analysis across optimization phases. Fitness distribution of algorithms organized by MCTS tree depth groups (Depth I: depths 1-4, Depth II: depths 5-7, Depth III: depths 8-10) and phase transitions (PT1-PT4). Training performance (teal) and test performance (pink) are shown with violin plots for sample sizes $n \geq 10$ and scatter plots (circles/rectangles) for $n < 10$. The analysis reveals systematic migration of high-fitness algorithms toward deeper tree layers as optimization progresses, with elite algorithms (fitness $> 5,000$) emerging exclusively in deeper layers during PT4. Enhanced generalization capability is observed in deeper layers during later optimization phases, as evidenced by improved training-test performance alignment in Depth III compared to shallower depth groups. (c) Algorithmic component impact analysis. Violin plots comparing normalized fitness distributions between algorithms with specific techniques (left) versus without (right). Techniques categorized as conditioning methods (teal), time-frequency analysis (orange), and trigger detection (green). Technique effectiveness is determined by distributional separation: wider gaps between left and right distributions indicate stronger performance impact. Conditioning techniques (Savitzky-Golay filtering, Adaptive Gain Regularization) and trigger detection methods (Curvature Analysis, Continuous Wavelet Transform Validation) demonstrate the most substantial improvements through clear distributional shifts toward higher fitness values. Statistical validation across 1,000 resampling iterations confirms significance ($p < 0.001$) and practical importance.

MCTS Depth-Stratified Performance Analysis. We analyzed the relationship between MCTS tree depth and algorithm fitness across optimization phases. Figure 4b reveals systematic evolution patterns in performance and generalization capability.

The analysis demonstrates clear progression in algorithmic quality through successive phase transitions. During PT1, algorithms are predominantly in shallow depth groups with modest fitness values ($< 2,000$). As optimization progresses to PT2 and PT3, high-performing algorithms (fitness $> 3,000$) increasingly emerge in deeper tree layers, indicating successful identification and refinement of promising algorithmic directions.

Elite algorithms (fitness $> 5,000$) emerge exclusively in deeper tree layers during PT4, suggesting sophisticated solutions require extensive refinement through multiple decision levels. Training (teal) and test (pink) performance distributions show robust generalization across all depth groups, with algorithms maintaining relative performance rankings regardless of tree depth.

High-performing algorithms become increasingly rare but more consistent in deeper layers, reflecting natural convergence toward superior solutions. Critically, deeper layers exhibit superior generalization capability during later optimization phases, with training-test performance gaps narrowing significantly in Depth III compared to shallower groups. This improved generalization suggests extensive algorithmic refinement enhances both performance and robustness across different observational conditions, confirming the MCTS framework effectively balances exploration breadth with exploitation depth.

Algorithmic Component Impact Analysis. We conducted comprehensive technique impact analysis using controlled comparative methodology, systematically evaluating algorithms with specific signal processing techniques against matched controls. (details in Supplementary Material A.5).

Figure 4c reveals distinct performance impacts across algorithmic components. Conditioning techniques demonstrate the most pronounced positive effects, with Savitzky-Golay filtering showing clear distributional separation and asymmetric violin plots shifted toward higher fitness values. This technique has been applied in gravitational wave counterpart studies for smoothing two-dimensional dispersed images from the Hubble Space Telescope, effectively removing high-frequency structure while preserving underlying emission patterns Troja et al. (2017). These findings establish quantitative benchmarks for algorithmic component selection in automated gravitational wave detection systems.

2.3.2 ALGORITHMIC EVOLUTION PATHWAY AND DISCOVERY MECHANISM ANALYSIS

MCTS Tree Structure and Knowledge Propagation. To understand the mechanistic basis of algorithmic discovery, we conducted comprehensive analysis of the complete MCTS exploration pathway leading to the PT4 algorithm (node 486, fitness = 5241.4). Figure 5a presents the full tree structure encompassing all nodes associated with the selected algorithm, revealing systematic patterns in knowledge accumulation and technique integration across multiple optimization phases (full MCTS tree data and visualization available in Supplementary Information Section S6).

Critical to understanding the framework's discovery mechanism is the identification of five key algorithmic breakthroughs that emerge at specific nodes and propagate through subsequent generations. These innovations demonstrate systematic knowledge accumulation, with breakthrough techniques subsequently incorporated into descendant algorithms through evolutionary operations. The propagation patterns reveal progressive sophistication through iterative refinement and technique combination, with superior algorithmic components showing robust inheritance and integration capabilities that directly influence fitness improvements across generations. The complete tree structure demonstrates effective LLM-guided exploration that balances exploitation of promising directions with ex-

540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

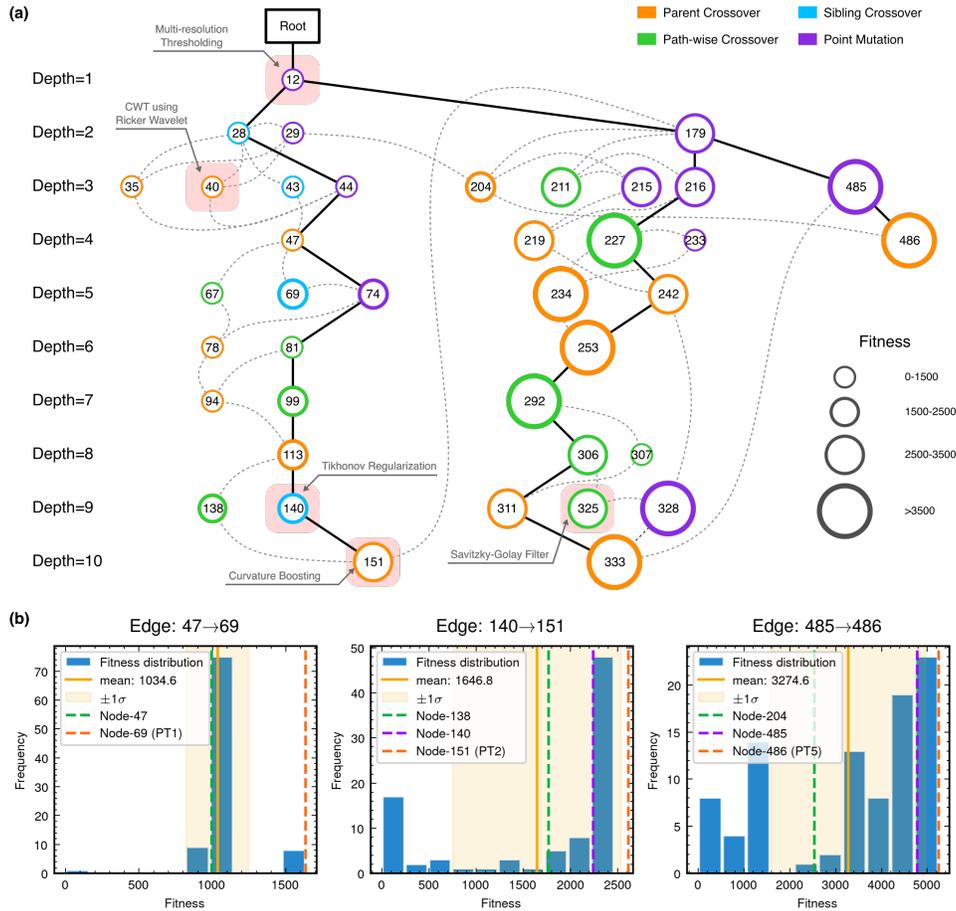


Figure 5: MCTS Algorithmic Evolution Pathway and Edge Robustness Analysis. (a) Complete MCTS tree structure showing all nodes associated with the PT4 algorithm (node 486, fitness=5241.4) discovered in an optimization run. Node sizes encode fitness values (larger circles = higher performance), with evaluation times displayed inside circles. Node colors indicate expansion operation types: Parent Crossover (orange), Sibling Crossover (cyan), Path-wise Crossover (green), and Point Mutation (purple). Solid black lines represent the selected MCTS exploration path, while dashed gray lines indicate nodes referenced in expansion prompts for knowledge synthesis. Five key algorithmic breakthroughs are annotated: Multi-resolution Thresholding (first appearing at node 12), CWT using Ricker Wavelet (node 28), Tikhonov Regularization (node 140), Curvature Boosting (node 151), and Savitzky-Golay Filter (node 333). These techniques propagate through subsequent generations, demonstrating systematic knowledge accumulation and refinement. The tree visualization reveals how sophisticated detection algorithms emerge through progressive technique integration across multiple MCTS depth levels. (b) Edge robustness analysis for three critical evolutionary transitions. Each subplot shows fitness distributions from 100 independent re-executions of specific edges: Edge 47→69 (early breakthrough, mean fitness 1034.6, 89.25% variants exceeding preceding node performance), Edge 140→151 (intermediate advancement, mean fitness 1646.8, 52.81% achieving superior fitness with 100% regularization technique inheritance), and Edge 485→486 (final optimization stage, mean fitness 3274.6, 70.65% variants outperforming node 204, 25.00% surpassing node 485). Vertical reference lines indicate the original node fitness values and key ancestral nodes. The distributions demonstrate the stochastic nature of LLM-driven code generation while confirming consistent discovery of high-performance algorithmic variants with robust knowledge transfer across independent executions.

594 ploration of novel algorithmic territories, leading to high-performance detection
595 strategies significantly exceeding conventional approaches.

596 **Edge Robustness and Stochastic Validation.** Through comprehensive re-
597 execution analysis of three pivotal evolutionary transitions using 100 independent
598 runs each, we validated the consistency of breakthrough innovations in Figure 5b.
599 These analyses confirm that breakthrough algorithmic innovations emerge through
600 systematic discovery processes rather than fortuitous random variations, demon-
601 strating stable technique inheritance and high-probability performance improve-
602 ments across all re-executions despite increased algorithmic complexity, validat-
603 ing the framework’s reliability for automated algorithm discovery and providing
604 confidence in the generalizability of discovered techniques to broader gravitational
605 wave detection challenges.

606 2.4 FRAMEWORK MECHANISM ANALYSIS

608 To validate Evo-MCTS’s effectiveness, we conducted systematic mechanism anal-
609 ysis across three critical dimensions: integrated optimization architecture, LLM
610 model selection, and domain knowledge incorporation. These analyses reveal
611 that superior performance emerges from synergistic component interactions rather
612 than simple additive effects - addressing the vast search space problem through
613 MCTS-guided reflection structuring, ensuring code generation quality via optimal
614 LLM selection, and maintaining scientific relevance through domain knowledge
615 integration. This multi-faceted approach demonstrates that successful automated
616 discovery requires intelligent integration of search strategy, reasoning capability,
617 and domain expertise beyond mere computational power.

618 **Integrated Architecture Validation.** Figure 6a demonstrates the substantial ben-
619 efits of combining evolutionary optimization with Monte Carlo Tree Search in
620 LLM-guided algorithm discovery. Evo-MCTS achieves superior performance
621 compared to its constituent components operating in isolation - pure MCTS-
622 AHD [Zheng et al. \(2025b\)](#) and pure evolutionary optimization ([ReEvo Ye et al. \(2024\)](#)),
623 all of which leverage LLMs for code generation and algorithmic reason-
624 ing.

625 The performance hierarchy reveals critical insights about LLM-based optimiza-
626 tion strategy effectiveness. Pure evolutionary approaches struggle with the vast
627 search space and become trapped in local optima despite LLM guidance, exhibit-
628 ing high variance with frequent suboptimal exploration in the LLM-generated al-
629 gorithm space. MCTS-AHD provides improvement through systematic search
630 space organization but with limited diversity in LLM-generated solutions. The
631 full Evo-MCTS combines the best of both worlds: maintaining population diver-
632 sity through evolutionary mechanisms while focusing computational resources on
633 promising algorithmic directions through tree search guidance, all while maxi-
634 mizing the utilization of LLM reasoning capabilities. This integration achieves
635 a remarkable 59.1% improvement over MCTS-AHD alone, demonstrating that
636 combining population-based diversity maintenance with tree-structured exploita-
637 tion creates emergent optimization capabilities that exceed the sum of individual
638 LLM-powered components.

639 **LLM Model Selection and Robustness Analysis.** Figure 6b investigates the
640 impact of different foundation models on algorithmic discovery performance, re-
641 vealing significant variations in code generation capability across state-of-the-art
642 language models. The analysis establishes `o3-mini-medium` as our fiducial
643 model configuration.

644 The performance hierarchy reveals insights about the relationship between model
645 architecture and scientific code generation capability. Particularly intriguing is
646 the superior performance of `claude-3-7-sonnet-20250219-thinking`
647 over `o1-2024-12-17`, despite both being reasoning-enhanced models. This
648 suggests that Claude’s specific training methodologies and architectural choices
649 may be better suited for sustained algorithmic reasoning tasks in gravitational
650 wave detection algorithm development. The substantial performance gap between

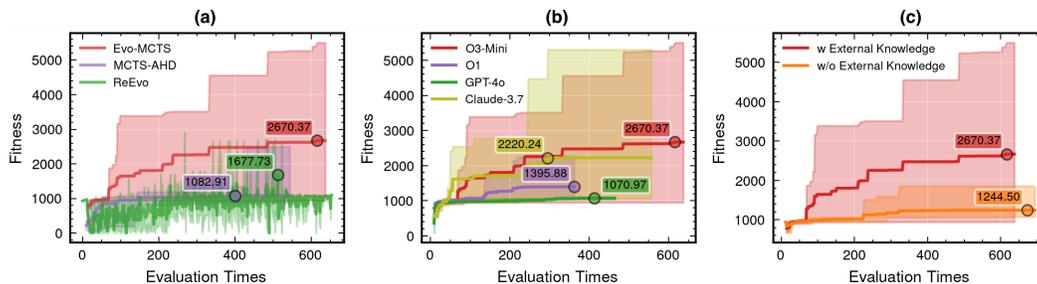


Figure 6: **Framework Mechanism Analysis and Component Contributions.** (a) Integrated architecture validation comparing Evo-MCTS (red) against constituent components: MCTS-AHD (purple, 1,677.73 fitness) and ReEvo (green, 1,082.61 fitness). Evo-MCTS achieves superior performance (2,670.37 fitness) through synergistic combination of evolutionary population dynamics and tree-structured search. (b) LLM model selection analysis showing performance variation across foundation models: o3-mini-medium (red, 2,670.37), claude-3-7-sonnet-20250219-thinking (yellow, 2,220.24), o1-2024-12-17 (purple, 1,395.88), and gpt-4o-2024-11-20 (green, 1,070.97). Results demonstrate the superior performance of reasoning-enhanced models, with o3-mini-medium achieving 150% improvement over general-purpose models. (c) Domain knowledge integration impact comparing frameworks with external knowledge (red, 2,670.37) versus without domain-specific guidance (orange, 1,244.50). The 115% performance improvement demonstrates the essential role of scientific domain expertise in automated algorithm discovery. All curves represent averages over at least five independent runs with shaded regions indicating standard deviation. Results validate the framework’s three core design principles: integrated optimization architecture, optimal model selection, and domain knowledge incorporation.

reasoning-enhanced models and general-purpose models underscores the critical importance of model architecture selection for scientific code generation tasks.

This demonstrates that framework effectiveness depends not merely on having access to large language models, but on selecting models with appropriate reasoning architectures for complex scientific applications. The robustness analysis shows consistent performance rankings across multiple runs, validating our model selection strategy while revealing that different LLM architectures exhibit distinct strengths in algorithmic reasoning and code synthesis.

Domain Knowledge Integration Impact. Figure 6c presents one of the most striking results: the dramatic impact of domain knowledge integration on algorithmic discovery performance. The comparison between frameworks with and without external knowledge reveals a performance difference of 115%, demonstrating that domain-specific guidance is essential for effective automated algorithm discovery in specialized scientific domains.

The framework without external knowledge exhibits relatively flat optimization trajectories, while domain knowledge integration demonstrates sustained improvement by serving as a constraint mechanism that guides the vast search space toward physically meaningful solutions. The emphasis on non-linear processing significantly enhances gravitational wave signal detection in real-world non-Gaussian, non-stationary noise environments, successfully leveraging scientific expertise to accelerate discovery beyond pure computational search (domain knowledge templates and integration methodology detailed in Supplementary Material A.I).

3 DISCUSSION

The Evo-MCTS framework demonstrates that automated algorithm discovery can achieve substantial performance improvements in gravitational wave detection through systematic exploration of algorithmic spaces. Our results establish key

insights extending beyond gravitational wave detection to broader scientific algorithm discovery applications.

Physical Insights and Algorithmic Discovery. The evolved algorithms reveal nonlinear processing strategies that effectively address fundamental limitations of both template-based and non-template approaches in realistic detector environments. The discovered multi-scale wavelet decompositions, adaptive thresholding mechanisms, and sophisticated statistical analysis methods dynamically adapt to non-Gaussian and non-stationary noise characteristics [Abbott et al. \(2016\)](#); [Dal Canton et al. \(2021\)](#). These algorithmic discoveries provide new perspectives on optimal gravitational wave detection, demonstrating that theoretical optimality of template-based methods under idealized conditions [Finn \(1992\)](#); [Cutler & Flanagan \(1994\)](#) can be substantially enhanced through interpretable nonlinear transformations that bridge template and non-template paradigms in practical scenarios.

Current Limitations. The framework focuses on static optimization scenarios and requires extensions for dynamic environments involving real-time parameter estimation. MLGWSC-1’s simplified testing environment, limited evaluation metrics, and potential overfitting risks do not fully capture operational detector complexity, as analyzed in the Supplementary Material [A.7](#). The framework’s reliance on large language models introduces dependencies on model architecture and prompt engineering, with 150% performance variation across implementations requiring expert validation modules for physical reasonableness verification. Our approach provides a novel framework for algorithmic optimization rather than a complete production-ready pipeline, with discovered algorithms serving as proof-of-concept demonstrations requiring further validation before operational deployment.

Future Directions. The evolutionary MCTS approach demonstrates broad potential for scientific algorithm discovery beyond gravitational waves. As Browne et al. note, MCTS “can be used with little or no domain knowledge, and has succeeded on difficult problems where other techniques have failed” [Browne et al. \(2012\)](#). The framework’s domain-agnostic architecture suggests applications in molecular optimization for drug discovery, materials design algorithms, and signal detection across astrophysical domains. The demonstrated interpretability advantages enable hybrid human-AI systems where algorithmic discoveries inform theoretical understanding while domain insights guide optimization, potentially accelerating scientific algorithm development across multiple disciplines.

4 METHODS

4.1 PROBLEM FORMULATION

This section formalizes gravitational wave algorithm discovery as a constrained optimization problem and introduces the LLM-guided evolutionary framework architecture.

Gravitational Wave Detection Problem. Given dual-detector strain data $\mathbf{d}(t) \in \mathbb{R}^{2 \times N}$, where $\mathbf{d}(t) = [d_H(t), d_L(t)]^T$ represents the strain data from Hanford and Livingston interferometers sampled at $f_s = 2048$ Hz over finite observation windows of length N samples, we seek to discover an optimal detection algorithm that maximizes performance while satisfying operational constraints.

Optimization Objective. The algorithm discovery problem is formulated as:

$$a^* = \arg \max_{a \in \mathcal{A}} \mathcal{F}(a, \mathbf{d}) \quad (1)$$

$$\text{subject to } \|\Delta t_{\text{arrival}}\| \leq 0.2 \text{ seconds} \quad (2)$$

$$T_{\text{comp}}(a) \leq T_{\text{max}} \quad (3)$$

$$E(a) \leq E_{\text{max}} \quad (4)$$

$$a : \mathbb{R}^{2 \times N} \rightarrow \mathbb{R}^3 \quad (5)$$

where:

- $\mathcal{A} = \{a \mid a \text{ is executable and satisfies constraints}\}$ represents the space of executable detection algorithms
- $\mathcal{F}(a, \mathbf{d}) = \int_{FAR_{\min}}^{FAR_{\max}} d_L(FAR; a, \mathbf{d}) d(FAR)$ measures detection performance as the area under the sensitive distance versus false alarm rate curve following the MLGWS-1 protocol Schäfer et al. (2023)
- $\Delta t_{\text{arrival}}$ denotes trigger arrival time uncertainty for astrophysical parameter estimation
- $T_{\text{comp}}(a)$ and $E(a)$ represent computational time and error handling trial count constraints with thresholds T_{max} and E_{max}
- Each algorithm a maps dual-detector strain data to a three-column detection catalog table containing peak times, signal ranking statistics, and timing uncertainties

The fitness function \mathcal{F} evaluates algorithms against ground truth labels \mathbf{y}_{true} using the area under the curve (AUC) of sensitive distance versus false alarms per month, where $d_L(FAR; a, \mathbf{d})$ represents the sensitive distance at a given false alarm rate FAR , providing more physically meaningful performance assessment than traditional Receiver Operating Characteristic curves for gravitational wave detection applications.

Algorithm Discovery Framework. We model the discovery process as an iterative search procedure that combines Monte Carlo Tree Search exploration with evolutionary population dynamics:

$$\mathbf{P}_{t+1} = \text{Evolve}(\mathbf{P}_t, \mathcal{L}, \mathcal{K}_{\text{GW}})$$

where $\mathbf{P}_t = \{a_1^{(t)}, a_2^{(t)}, \dots, a_k^{(t)}\}$ represents the algorithm population at iteration t .

Key Components:

- **LLM Code Generation:** $\mathcal{L} : (\text{prompt}, \text{context}) \rightarrow \text{code}$ represents the language model for code generation, with operation-specific prompting strategies $\sigma : \mathcal{O} \rightarrow \mathcal{P}$ mapping evolutionary operations $\mathcal{O} = \{\text{PC}, \text{SC}, \text{PWC}, \text{PM}\}$ to specialized prompt templates \mathcal{P}
- **Domain Knowledge:** $\mathcal{K}_{\text{GW}} = (\mathcal{K}_{\text{physics}}, \mathcal{K}_{\text{signal}}, \mathcal{K}_{\text{constraint}})$ encapsulates gravitational wave detection expertise including physical principles, signal processing techniques, and computational constraints
- **MCTS Selection:** Node selection follows Upper Confidence bounds applied to Trees (UCT) with adaptive exploration:

$$\pi_{\text{UCT}}(n) = \arg \max_{c \in \text{children}(n)} [\text{normalized_fitness}(c) + \text{adaptive_exploration}(c)]$$

where the complete implementation with fitness normalization, adaptive exploration constants, and epsilon regularization is detailed in Section 4.4.

Evolutionary Operations. The framework employs four evolutionary operations for algorithmic transformation, each utilizing operation-specific prompting strategies with the same underlying language model:

$$\text{PC}(\mathbf{P}_t) : a_{\text{new}} = \mathcal{L}(\text{prompt}_{\text{PC}}(a_p, a_r), \mathcal{K}_{\text{GW}}) \quad (6)$$

$$\text{SC}(\mathbf{P}_t) : a_{\text{new}} = \mathcal{L}(\text{prompt}_{\text{SC}}(a_c, \{a_{s_i}\}), \mathcal{K}_{\text{GW}}) \quad (7)$$

$$\text{PWC}(\mathbf{P}_t) : a_{\text{new}} = \mathcal{L}(\text{prompt}_{\text{PWC}}(\{a_{d_i}\}), \mathcal{K}_{\text{GW}}) \quad (8)$$

$$\text{PM}(\mathbf{P}_t) : a_{\text{new}} = \mathcal{L}(\text{prompt}_{\text{PM}}(a_c, a_e), \mathcal{K}_{\text{GW}}) \quad (9)$$

where PC = Parent Crossover, SC = Sibling Crossover, PWC = Path-wise Crossover, PM = Point Mutation. Here, a_p and a_r denote parent and reference algorithms, a_c represents the current algorithm, $\{a_{s_i}\}$ are sibling algorithms, a_e is an elite algorithm, and $\{a_{d_i}\}$ are distant algorithm sets.

Each operation employs tailored prompt templates $\text{prompt}_{\text{op}}(\cdot)$ that guide the language model to generate appropriate algorithmic variants: Parent Crossover

prompts emphasize combining successful features from parent algorithms, Sibling Crossover focuses on lateral exploration within similar algorithmic families, Path-wise Crossover promotes paradigm shifts by integrating distant algorithmic approaches, and Point Mutation targets localized refinements of existing implementations. Domain knowledge \mathcal{K}_{GW} ensures adherence to gravitational wave detection principles and physical constraints across all operations.

Reflection and Adaptation. The system incorporates analytical reasoning through specialized reflection using the DeepSeek-R1 model:

$$\mathcal{K}_{\text{GW}}^{(t+1)} = \mathcal{K}_{\text{GW}}^{(t)} \cup \{\text{insights}(R(\text{history}_t, \text{performance}_t, \mathcal{K}_{\text{physics}}))\}$$

where $R : \mathcal{H} \times \mathcal{F}^* \times \mathcal{K}_{\text{physics}} \rightarrow \mathcal{I}$ represents the reflection function mapping MCTS history \mathcal{H} , fitness evaluations \mathcal{F}^* , and physical knowledge to actionable insights \mathcal{I} .

This reflection mechanism analyzes performance patterns across the MCTS tree to identify successful algorithmic principles and guide subsequent evolutionary operations toward promising regions of the solution space.

Population Management. At each iteration, the algorithm population is updated through elite preservation with selection pressure β :

$$\mathbf{P}_{t+1} = \text{Elite}(\mathbf{P}_t \cup \{a_{\text{new}}\}, k, \beta) \quad (10)$$

where $\text{Elite}(\cdot, k, \beta)$ selects the top- k algorithms based on fitness scores $\mathcal{F}(a_i, \mathbf{d})$ with selection probability: $p_{\text{select}}(a_i) = \frac{\exp(\beta \cdot \mathcal{F}(a_i, \mathbf{d}))}{\sum_{j=1}^{|P_t|} \exp(\beta \cdot \mathcal{F}(a_j, \mathbf{d}))}$. This formulation establishes gravitational wave algorithm discovery as a constrained optimization problem in the space of executable detection algorithms, solved through LLM-guided evolutionary search with MCTS exploration and domain knowledge integration.

4.2 LLM INTEGRATION FOR CODE GENERATION

The framework leverages state-of-the-art language models to transform algorithmic concepts into executable code, implementing a multi-model strategy that capitalizes on the complementary strengths of different architectures. This subsection details the model selection, prompting strategies, and error handling mechanisms that enable robust algorithmic discovery.

Model Architecture and Task Allocation. The framework employs a heterogeneous ensemble of four language models: o3-mini-medium, o1-2024-12-17, gpt-4o-2024-11-20, and claude-3-7-sonnet-20250219-thinking for code generation tasks. For reflection operations, we utilize deepseek-r1-250120 exclusively due to its analytical reasoning capabilities.

Prompting Strategy and Temperature Control. All models operate with temperature 1.0 to optimize the trade-off between algorithmic diversity and code validity.

The prompting framework employs depth-aware adaptation mechanisms. Task descriptions clarify optimization objectives, while depth information guides exploration scope: shallow nodes emphasize paradigm shifts, deeper nodes focus on parameter optimization. External domain knowledge integration provides optimization directives referencing established signal processing principles and computational constraints. This adaptive architecture enables systematic solution space exploration while maintaining gravitational wave detection coherence (complete templates in Supplementary Material A.1).

Error Handling and Iterative Refinement. The framework implements a robust error recovery mechanism to handle code generation failures. When syntax errors or runtime exceptions occur during algorithm evaluation, the system captures detailed error information including stack traces and execution context. This diagnostic information is then incorporated into subsequent conversation rounds with the LLM to generate corrected implementations. The system attempts up

864 to three correction iterations per failed algorithm. If all correction attempts fail,
 865 the corresponding node expansion is skipped to maintain computational efficiency.
 866 This approach ensures that the majority of generated algorithms remain executable
 867 while preventing infinite correction loops that could stall the evolutionary process
 868 (complete templates in Supplementary Material A.1).

869 **Post-Generation Analysis** Following successful code generation, each algorithm
 870 undergoes a post-thought analysis phase that extracts key design principles and
 871 compresses algorithmic representations. This reflection process creates human-
 872 readable summaries while reducing token consumption to prevent context window
 873 overflow in subsequent LLM interactions.

874 The analysis captures algorithmic innovations, signal processing techniques, per-
 875 formance expectations, and computational characteristics in compressed form.
 876 This enables efficient reference to previous discoveries without overwhelming
 877 the LLM context, facilitating continued exploration while maintaining algorithmic
 878 memory across generations (complete templates in Supplementary Material A.1).

879 **Domain Knowledge Integration.** The framework incorporates gravitational
 880 wave domain expertise through three structured prompt categories: initialization
 881 prompts defining matched filtering principles and noise characteristics, evolution
 882 prompts encouraging nonlinear transformations and adaptive thresholding, and
 883 reflection prompts evaluating sensitivity and computational efficiency (complete
 884 templates in Supplementary Material A.1).

885 The knowledge base prioritizes nonlinear processing techniques including adap-
 886 tive thresholds, phase space reconstruction, and robust estimators. Implementa-
 887 tion emphasizes adaptive parameters over fixed values while maintaining compu-
 888 tational efficiency.

889 This approach ensures physical constraint compliance while enabling exploration
 890 beyond conventional linear methods.

891 4.3 EVOLUTIONARY OPERATIONS DESIGN AND SEED ALGORITHM

892 **Baseline Algorithm Architecture.** The optimization process begins with a del-
 893iberately simple, linear seed function that establishes a baseline for algorithmic
 894 improvement (detailed implementation in Supplementary Material A.1). This
 895 seed algorithm implements a conventional signal processing pipeline consisting
 896 of three sequential operations that represent standard approaches in gravitational
 897 wave data analysis.

898 The first operation performs frequency-domain whitening to normalize the detec-
 899 tor noise characteristics:

$$900 X_{\text{white}}(f) = \frac{X(f)}{\sqrt{S(f)}} \quad (11)$$

902 where $X(f)$ represents the Fourier transform of the input strain data, and $S(f)$
 903 is the power spectral density estimate obtained via Welch’s method using a 4096-
 904 sample window with 50% overlap and Hann windowing. This whitening operation
 905 ensures that all frequency components contribute equally to subsequent analysis,
 906 compensating for the detector’s frequency-dependent noise characteristics.

907 The second operation applies time-frequency decomposition using the short-time
 908 Fourier transform (STFT) to capture transient signal characteristics:

$$909 S_{xx}(f, \tau) = \left\| \sum_{n=0}^{N-1} x(n + \tau) w(n) e^{-j2\pi f n / N} \right\|^2 \quad (12)$$

913 where $w(n)$ is a window function (256 samples with 128-sample overlap), τ is the
 914 time shift, and N is the window length. The algorithm independently processes
 915 both Hanford (H1) and Livingston (L1) detector data, then combines the resulting
 916 spectrograms using simple averaging:

$$917 \text{TF}_{\text{metric}} = \frac{1}{2} \langle S_{xx}^{\text{H1}} + S_{xx}^{\text{L1}} \rangle_f$$

918 where $\langle \cdot \rangle_f$ denotes averaging over frequency bins.

919 The final operation identifies candidate events through basic peak detection with
 920 fixed thresholds. The algorithm estimates background levels using the median and
 921 applies simple peak finding with predetermined height and prominence criteria.
 922 This approach represents a minimalist detection strategy that lacks the sophisti-
 923 cation necessary for robust gravitational wave identification, particularly in the
 924 presence of non-Gaussian noise transients and weak signals.

925 **Initial Population Generation.** The evolutionary framework initializes with
 926 a single seed algorithm, then generates 8 diverse variants through systematic
 927 prompting variations (Figure 1a). Each variant maintains identical input-output
 928 interfaces while implementing distinct signal processing approaches: alternative
 929 whitening schemes, varied time-frequency decomposition methods, and different
 930 peak detection strategies. Following initial generation, two Point Mutation opera-
 931 tions are applied to create additional variants, resulting in a total population of 10
 932 algorithms that form the depth-1 initial population for MCTS exploration.

933 **Elite Preservation Strategy.** The framework maintains an elite individual rep-
 934 resenting the best-performing algorithm discovered throughout evolution. This
 935 elite serves as a performance benchmark for new variants and provides genetic
 936 material during Point Mutation operations, which specifically leverage elite char-
 937 acteristics to guide targeted algorithmic improvements (Figure 1c). The elite is
 938 updated whenever a new algorithm demonstrates superior performance, ensuring
 939 monotonic progress while maintaining access to the current best solution (Fig-
 940 ure 1b.4).

941 **Evolutionary Operation Framework.** Each MCTS expansion level follows a
 942 structured sequence of evolutionary operations: Parent Crossover (PC) executes
 943 5 times, Path-wise Crossover (PWC) executes 2 times, Sibling Crossover (SC)
 944 executes once, and Point Mutation (PM) executes twice. As observed in Figure 5a,
 945 this sequence balances four distinct algorithmic improvement mechanisms: (i)
 946 vertical knowledge transfer through PC operations that combine features from
 947 algorithms at different tree depths, (ii) long-range dependency capture via PWC
 948 operations that synthesize information along complete evolutionary trajectories,
 949 (iii) horizontal exploration using SC operations that facilitate exploration between
 950 algorithms of similar complexity based on nodes already generated at the same
 951 tree level, and (iv) fine-grained optimization through PM operations that introduce
 952 targeted modifications based on performance analysis. The systematic application
 953 of these operations ensures comprehensive exploration of the algorithmic solution
 954 space while maintaining computational efficiency through controlled expansion
 955 rates.

956 4.4 MONTE CARLO TREE SEARCH IMPLEMENTATION

957 **MCTS Framework and Tree Policy.** The framework implements a standard
 958 Monte Carlo Tree Search algorithm adapted for algorithmic discovery, where each
 959 node represents an executable algorithm and tree expansion corresponds to algo-
 960 rithmic evolution (Figure 1b). The MCTS operates through four canonical phases:
 961 selection, expansion, simulation, and backpropagation. However, our implemen-
 962 tation modifies the traditional simulation phase by replacing random rollouts with
 963 direct algorithm evaluation on the gravitational wave detection task.

964 The selection phase traverses the tree from root to leaf using the Upper Confidence
 965 Bound applied to Trees (UCT) policy, balancing exploitation of high-performing
 966 algorithms with exploration of under-visited branches (Figure 1c). Unlike tradi-
 967 tional MCTS applications where leaf nodes represent terminal game states, our
 968 leaf nodes represent algorithms that can be further evolved through the four evo-
 969 lutionary operations (PC, SC, PWC, PM).

970 **UCT Score Calculation and Adaptive Exploration.** The UCT score for each
 971 node combines exploitation and exploration terms with an adaptive exploration
 strategy that accounts for the finite evaluation budget:

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

$$\text{UCT}(n) = \frac{Q(n) - Q_{\min}}{Q_{\max} - Q_{\min} + \epsilon} + c \cdot \sqrt{\frac{\ln(N(p) + 1)}{N(n) + \epsilon}} \quad (13)$$

where $Q(n)$ represents the normalized fitness value of node n , Q_{\min} and Q_{\max} are the minimum and maximum fitness values observed across all nodes, $N(p)$ and $N(n)$ are the visit counts of the parent and current node respectively, ϵ is a small constant preventing division by zero, and c is the exploration constant.

The exploration constant c adapts dynamically based on the remaining evaluation budget:

$$c = c_0 \cdot \max\left(1 - \frac{t}{T}, 0\right) \quad (14)$$

where c_0 is the initial exploration constant, t is the current evaluation count, and T is the maximum evaluation budget. This adaptive mechanism ensures that the algorithm emphasizes exploration early in the search when the budget is abundant, then gradually shifts toward exploitation as evaluations are consumed.

Fitness Normalization and Q-Value Management. The fitness values (corresponding to the objective function in MCTS) are normalized to the range $[0, 1]$ to ensure consistent UCT calculations regardless of the absolute scale of performance metrics. The normalization uses running minimum and maximum values:

$$Q_{\text{normalized}} = \frac{Q_{\text{raw}} - Q_{\min}}{Q_{\max} - Q_{\min} + \epsilon} \quad (15)$$

This normalization is crucial for maintaining meaningful exploration-exploitation balance, as it prevents algorithms with vastly different performance scales from skewing the selection process.

Backpropagation and Value Updates. The backpropagation phase updates node values along the path from the newly evaluated leaf to the root. Our implementation uses a discount factor approach that balances immediate performance with long-term potential:

$$Q(p) = Q(p) \cdot (1 - \gamma) + \max_{c \in \text{children}(p)} Q(c) \cdot \gamma \quad (16)$$

where $Q(p)$ is the parent node's Q-value, γ is the discount factor, and the maximum is taken over all child nodes. This update rule ensures that parent nodes reflect the performance of their best children while maintaining some memory of their previous estimates.

The backpropagation also maintains global statistics by updating the minimum and maximum Q-values across the entire tree, enabling consistent normalization for future UCT calculations. Additionally, the algorithm maintains a ranked list of all observed fitness values to support advanced selection strategies and performance analysis.

Tree Expansion Strategy. Node expansion occurs when the UCT selection phase reaches a leaf node that has been visited multiple times, indicating sufficient confidence in its potential. The expansion strategy creates one new child node per expansion operation, chosen through the evolutionary operations framework. This conservative expansion approach prevents explosive tree growth while ensuring thorough exploration of promising regions.

Each newly created node inherits structural information from its parent, including the algorithmic context and performance history. The initial Q-value for new nodes is set to the parent's Q-value, providing a reasonable starting estimate that will be refined through subsequent evaluations.

Memory Management and Tree Pruning. To maintain computational efficiency with limited memory resources, the implementation includes mechanisms for selective tree pruning. Nodes that demonstrate consistently poor performance relative to their siblings are marked for potential removal, while maintaining sufficient diversity to avoid premature convergence.

1026 The tree structure also maintains subtree references that enable efficient traversal
1027 and analysis of evolutionary pathways. These references support the PWC oper-
1028 ation by providing access to complete root-to-leaf trajectories without requiring
1029 expensive tree traversals.

1030 **Convergence Detection and Termination.** The MCTS continues until either the
1031 evaluation budget is exhausted or convergence is detected through analysis of the
1032 Q-value distribution. Convergence is identified when the top-performing algo-
1033 rithms show minimal improvement over a specified number of iterations, indicat-
1034 ing that the search has reached a local optimum within the current exploration
1035 strategy.

1036 This MCTS implementation creates a systematic framework for exploring the
1037 algorithmic space while maintaining computational efficiency and ensuring re-
1038 producible results. The combination of adaptive exploration, normalized fitness
1039 values, and efficient tree management enables the discovery of high-performing
1040 algorithms within reasonable computational budgets.

1041 4.5 EXPERIMENTAL SETUP

1042 **MLGWSC-1 Dataset and Evaluation Framework.** We evaluated our Evo-
1043 MCTS framework using the first Machine Learning Gravitational-Wave Search
1044 Mock Data Challenge (MLGWSC-1) benchmark, which provides a standardized
1045 evaluation environment for gravitational wave detection algorithms Schäfer et al.
1046 (2023). The benchmark includes four datasets with increasing complexity, of
1047 which we utilized Dataset 4 for our primary evaluation as it represents the most
1048 realistic scenario using actual detector noise from the O3a observing run.

1049 Dataset 4 incorporates real noise from both Hanford (H1) and Livingston (L1) de-
1050 tectors, filtered to include only segments with the DATA data-quality flag active
1051 while excluding problematic categories (CBC_CAT1, CBC_CAT2, CBC_HW_INJ,
1052 and BURST_HW_INJ). Overlapping segments span a minimum duration of 2
1053 hours, with signal injection parameters specifically designed to ensure that at least
1054 33% of injected signals have an optimal network SNR < 4 and thus cannot be
1055 detected, creating a challenging evaluation scenario that includes both detectable
1056 and sub-threshold events.

1057 **Algorithm Input/Output Interface.** Following the MLGWSC-1 specification,
1058 each algorithm processes HDF5 files containing raw detector data from both H1
1059 and L1 detectors. The input files contain grouped datasets organized by inte-
1060 ger start times, with each dataset including strain data and metadata attributes
1061 (`start_time`, `delta_t`). Our evolved algorithms output HDF5 files containing
1062 exactly three datasets of equal length: (i) `time` - GPS times of suspected grav-
1063 itational wave events, (ii) `stat` - ranking statistics where larger values indicate
1064 higher detection confidence, and (iii) `var` - timing accuracy tolerance repre-
1065 senting the maximum allowed separation between predicted and true event times.

1066 **Evaluation Metrics and Performance Assessment.** The framework employs
1067 two primary evaluation metrics consistent with MLGWSC-1 standards. The false-
1068 alarm rate (FAR) measures the expected frequency of false-positive events exceed-
1069 ing a given ranking statistic threshold, calculated by applying algorithms to pure
1070 noise data and dividing event counts by total analyzed time. The sensitive dis-
1071 tance quantifies detection capability at specified false-alarm rates, representing
1072 the maximum distance at which sources can be reliably detected. For uniformly
1073 distributed signals in volume, this reduces to the fraction of detected signals mul-
1074 tiplied by the volume of a sphere with radius equal to the maximum injected dis-
1075 tance. The area under the curve (AUC) metric integrates sensitive distance across
1076 the false-alarm rate range, providing a comprehensive performance indicator that
1077 balances detection sensitivity against false-alarm tolerance.

1078 **Training and Test Set Configuration.** To ensure efficient optimization while
1079 maintaining statistical rigor, we partitioned the MLGWSC-1 Dataset 4 into train-
ing and test subsets. The training set comprises the first 7 days of data from early
injection indices, enabling rapid algorithmic evaluation during the optimization
process. This temporal partitioning ensures that the minimum false-alarm rate

1080 boundary is set at 4 events per month, corresponding to the statistical require-
1081 ments for meaningful AUC calculation. The test set consists of 1 day of data
1082 selected from later temporal segments, providing independent validation of op-
1083 timized algorithms (detailed data partitioning specifications provided in Supple-
1084 mentary Material A.2).

1085 **Diversity Metrics in Evolutionary Computation.** We implemented sophisti-
1086 cated diversity measurement following established practices in evolutionary com-
1087 putation [Dat et al. \(2024\)](#). Population encoding involves three preprocessing steps:
1088 (i) removing comments and docstrings using abstract syntax tree parsing to focus
1089 on functional algorithmic content, (ii) standardizing code snippets into common
1090 coding style following PEP 8 conventions to eliminate stylistic variations, and (iii)
1091 converting normalized code snippets to vector representations using the CodeT5+
1092 embedding model to enable quantitative similarity analysis.

1093 We employ two complementary diversity metrics: the Shannon Diversity Index,
1094 calculated as $H = -\sum_{i=1}^n p_i \log p_i$ where p_i represents the frequency of the i -th
1095 unique algorithmic variant in the population, and the Complexity Index of Diver-
1096 sity (CID), computed as $CID = \frac{1}{n} \sum_{i=1}^n \frac{\|x_i - \bar{x}\|}{\|\bar{x}\| + \epsilon}$ where x_i represents the em-
1097 bedding vector of the i -th algorithm, \bar{x} is the population centroid, and ϵ prevents
1098 division by zero. These metrics provide complementary perspectives on popula-
1099 tion diversity: Shannon index captures algorithmic variety while CID measures
1100 structural complexity differences.

1101 **LLM API Configuration and Model Selection.** Our framework inte-
1102 grates multiple state-of-the-art language models through official APIs
1103 from OpenAI, Anthropic, and DeepSeek [OpenAI et al. \(2024\)](#); [Open-
1104 nAI \(2024\)](#); [Anthropic \(2025\)](#); [DeepSeek-AI et al. \(2025\)](#). For code
1105 generation tasks, we employ thinking-enhanced models including
1106 o3-mini-medium, o1-2024-12-17, gpt-4o-2024-11-20, and
1107 claude-3-7-sonnet-20250219-thinking, selected for their demon-
1108 strated capabilities in multi-step reasoning and code synthesis [Chen et al.
1109 \(2021\)](#); [Bubeck et al. \(2023\)](#). For reflection mechanisms, we exclusively utilize
1110 deepseek-r1-250120 due to its superior performance in analytical reason-
1111 ing tasks [DeepSeek-AI et al. \(2025\)](#), particularly its ability to identify subtle
1112 performance patterns and propose targeted algorithmic improvements based on
1113 empirical observations. We designate o3-mini-medium as our fiducial model
1114 configuration, providing the baseline reference for performance comparisons
1115 and ensuring consistency across experimental runs. All models operate with
1116 temperature parameter set to 1.0 to balance creative exploration with syntactic
1117 reliability [Ouyang et al. \(2022\)](#).

1118 **Hyperparameter Configuration and Experimental Design.** The Evo-MCTS
1119 framework employs carefully tuned hyperparameters optimized for gravitational
1120 wave detection algorithm discovery. Initial population size is set to 10 algo-
1121 rithms, balancing computational efficiency with adequate diversity for effective
1122 exploration. MCTS depth is limited to 10 levels, providing sufficient hierarchical
1123 structure for complex algorithmic development while maintaining computational
1124 tractability. Each experimental configuration undergoes 5 independent runs with
1125 different random seeds to ensure statistical robustness and enable confidence in-
1126 terval estimation.

1127 The 5-run experimental design serves multiple purposes: (i) quantifying algo-
1128 rithmic discovery reliability across different initialization conditions, (ii) enabling
1129 statistical analysis of optimization trajectories and convergence patterns, and (iii)
1130 providing robust diversity measurements that account for stochastic variations in
1131 the evolutionary process. Comprehensive results from all individual runs are doc-
1132 umented in Supplementary Material A.3, enabling detailed analysis of inter-run
1133 variability and optimization consistency.

1134 **Ablation Studies and Comparative Analysis.** To validate the effectiveness of
1135 our integrated Evo-MCTS approach, we conduct comprehensive ablation studies
1136 comparing against two established frameworks: ReEvo [Ye et al. \(2024\)](#) (pure evo-
1137 lutionary mechanism) and MCTS-AHD [Zheng et al. \(2025b\)](#) (pure MCTS mech-

anism). These comparisons isolate the contributions of different algorithmic components while maintaining consistent evaluation protocols.

ReEvo represents the evolutionary optimization baseline, employing genetic algorithms with traditional crossover and mutation operations powered by large language models but without tree-structured exploration. The framework implements population-based optimization through iterative code generation and selection, focusing on evolutionary diversity maintenance and fitness-driven selection pressure. MCTS-AHD implements pure Monte Carlo Tree Search for automated heuristic design without evolutionary population dynamics or reflection mechanisms. This approach emphasizes tree-based exploration with UCT-guided node selection but lacks the multi-generational insight synthesis and population diversity maintenance that characterizes evolutionary approaches.

By comparing Evo-MCTS against these component frameworks using identical LLM integration protocols and evaluation budgets, we quantify the synergistic benefits of combining evolutionary population dynamics with structured tree search exploration. The experimental design maintains consistent evaluation metrics, hardware configurations, and statistical analysis procedures across all three frameworks, with computational fairness ensured through equivalent LLM evaluation counts as the unified iteration metric. This approach enables direct performance comparison while isolating the specific contributions of different optimization strategies under identical resource constraints. The comprehensive multi-run analysis results across all frameworks are presented in Figure 8, demonstrating the statistical robustness of our comparative evaluation.

Edge Robustness Analysis Protocol. To validate the consistency of algorithmic breakthroughs, we developed a systematic edge re-execution protocol. For each selected evolutionary transition, we perform 100 independent re-executions using identical prompt templates while maintaining all experimental conditions constant except for the LLM sampling parameters. Each re-execution employs the same parent algorithms, evolutionary operation type, and external knowledge integration templates used in the original optimization run.

The re-execution protocol preserves all deterministic components (fitness evaluation, data preprocessing, statistical analysis) while allowing natural variation in LLM-generated code through different random seeds. This approach enables quantification of discovery mechanism robustness while accounting for the inherent stochasticity in large language model outputs. Statistical analysis employs standard descriptive metrics (mean, standard deviation) combined with confidence interval estimation to characterize the reliability of breakthrough discovery patterns, as illustrated in Figure 5(b).

Computational Resources and Parallelization. The numerical calculations in this study were carried out on the ORISE Supercomputer, equipped with 32-core x86 processors operating and 4 GPGPU accelerators per node, enabling efficient parallel execution of LLM API calls and algorithm evaluations. The Evo-MCTS framework employs asynchronous parallelization to maximize resource utilization, allowing multiple LLM requests to be processed concurrently while maintaining synchronization for tree updates and performance analysis. This parallelization strategy significantly reduces overall computational time while ensuring that all evaluations are performed under consistent conditions.

5 DATA AVAILABILITY

The MLGWSC-1 Dataset 4 used in this study is publicly available at <https://github.com/gwastro/ml-mock-data-challenge-1/>. The dataset includes real detector noise from LIGO Hanford (H1) and Livingston (L1) observatories during the O3a observing run.

The ReEvo framework implementation is available at <https://github.com/ai4co/reevo> with the original codebase and documentation. The MCTS-AHD framework can be accessed through its official repository at <https://github.com/zz1358m/MCTS-AHD-master>. Both frameworks were

1188 used for comparative analysis following their original specifications and
1189 hyperparameter configurations.

1190 Training and test data partitions, along with detailed preprocessing specifications,
1191 are provided in the Supplementary Material A.2. All experimental results and
1192 algorithm performance metrics are available upon reasonable request to the corre-
1193 sponding author.

1195 6 CODE AVAILABILITY

1196 The complete source code for the Evo-MCTS framework is publicly available at
1197 <https://github.com/iphysresearch/evo-mcts>. The repository includes all implementa-
1198 tion details, experimental configurations, and reproducibility instructions. Ad-
1199 ditionally, a permanent archived version is available through the Zenodo reposi-
1200 tory at <https://zenodo.org/record/100000000000000000>. The code is written
1201 in Python and is fully reproducible with the provided environment and depen-
1202 dencies. The repository contains comprehensive documentation, example usage
1203 scripts, and all necessary configuration files to replicate the experimental results
1204 presented in this study.

1206 ACKNOWLEDGMENTS

1207 We thank the LIGO Scientific Collaboration for providing the MLGWSC-1
1208 dataset and establishing the evaluation framework that enabled this research. We
1209 acknowledge the computational resources provided by the ORISE Supercomputer
1210 facility, which were essential for the large-scale LLM-based optimization experi-
1211 ments.

1212 We are grateful to OpenAI, Anthropic, and DeepSeek for providing API access to
1213 their language models, which formed the core of our algorithmic discovery frame-
1214 work. We also thank the developers of the ReEvo and MCTS-AHD frameworks
1215 for making their implementations publicly available for comparative analysis.

1216 H.W. is supported by the National Key Research and Development Program of
1217 China (Grant No. 2021YFC2203004), the National Natural Science Foundation
1218 of China (NSFC) (Grant Nos. 12405076, 12247187, 12147103), the National
1219 Astronomical Data Center (Grant No. NADC2023YDS-01), and the Fundamental
1220 Research Funds for the Central Universities.

1222 7 DATA AVAILABILITY

1223 The MLGWSC-1 Dataset 4 used in this study is publicly available at
1224 <https://github.com/gwastro/ml-mock-data-challenge-1/>. The dataset includes real
1225 detector noise from LIGO Hanford (H1) and Livingston (L1) observatories during
1226 the O3a observing run.

1227 The ReEvo framework implementation is available at
1228 <https://github.com/ai4co/reevo> with the original codebase and documenta-
1229 tion. The MCTS-AHD framework can be accessed through its official repository
1230 at <https://github.com/zz1358m/MCTS-AHD-master>. Both frameworks were
1231 used for comparative analysis following their original specifications and
1232 hyperparameter configurations.

1233 Training and test data partitions, along with detailed preprocessing specifications,
1234 are provided in the Supplementary Information Section S2. All experimental re-
1235 sults and algorithm performance metrics are available upon reasonable request to
1236 the corresponding author.

1238 8 CODE AVAILABILITY

1239 The complete source code for the Evo-MCTS framework is publicly available at
1240 <https://github.com/iphysresearch/evo-mcts>. The repository includes all implementa-
1241 tion details, experimental configurations, and reproducibility instructions. The

code is written in Python and is fully reproducible with the provided environment and dependencies. The repository contains comprehensive documentation, example usage scripts, and all necessary configuration files to replicate the experimental results presented in this study.

REFERENCES

- B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016. doi: 10.1103/PhysRevLett.116.061102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- B. P. Abbott et al. Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Phys. Rev. X*, 9:031040, Sep 2019. doi: 10.1103/PhysRevX.9.031040. URL <https://link.aps.org/doi/10.1103/PhysRevX.9.031040>.
- B. P. Abbott et al. A guide to ligo-virgo detector noise and extraction of transient gravitational-wave signals. *Classical and Quantum Gravity*, 37(5):055002, feb 2020. doi: 10.1088/1361-6382/ab685e. URL <https://dx.doi.org/10.1088/1361-6382/ab685e>.
- Anthropic. Claude 3.7 sonnet and claude code. <https://www.anthropic.com/news/claude-3-7-sonnet>, 2025. Accessed: 2025.
- Nathan Baker, Frank Alexander, Timo Bremer, Aric Hagberg, Yannis Kevrekidis, Habib Najm, Manish Parashar, Abani Patra, James Sethian, Stefan Wild, et al. Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence. Technical report, USDOE Office of Science (SC), Washington, D.C. (United States), 02 2019. URL <https://www.osti.gov/biblio/1478744>.
- Cameron B. Browne, Edward Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1): 1–43, March 2012. ISSN 1943-068X, 1943-0698. doi: 10.1109/TCIAIG.2012.2186810.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4, 2023. URL <https://arxiv.org/abs/2303.12712>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebguss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL <https://arxiv.org/abs/2107.03374>.

- 1296 Curt Cutler and Éanna E. Flanagan. Gravitational waves from merging com-
1297 pact binaries: How accurately can one extract the binary’s parameters from
1298 the inspiral waveform? *Phys. Rev. D*, 49:2658–2697, Mar 1994. doi:
1299 10.1103/PhysRevD.49.2658. URL [https://link.aps.org/doi/10.](https://link.aps.org/doi/10.1103/PhysRevD.49.2658)
1300 [1103/PhysRevD.49.2658](https://link.aps.org/doi/10.1103/PhysRevD.49.2658).
- 1301
1302 Tito Dal Canton, Alexander H. Nitz, Bhooshan Gadre, Gareth S. Cabourn Davies,
1303 Verónica Villa-Ortega, Thomas Dent, Ian Harry, and Liting Xiao. Real-time
1304 search for compact binary mergers in advanced ligo and virgo’s third ob-
1305 serving run using pycbc live. *The Astrophysical Journal*, 923(2):254, dec
1306 2021. doi: 10.3847/1538-4357/ac2f9a. URL [https://dx.doi.org/10.](https://dx.doi.org/10.3847/1538-4357/ac2f9a)
1307 [3847/1538-4357/ac2f9a](https://dx.doi.org/10.3847/1538-4357/ac2f9a).
- 1308 Pham Vu Tuan Dat, Long Doan, and Huynh Thi Thanh Binh. Hsevo: Elevating au-
1309 tomatic heuristic design with diversity-driven harmony search and genetic algo-
1310 rithm using llms, 2024. URL <https://arxiv.org/abs/2412.14995>.
- 1311
1312 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu
1313 Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang
1314 Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu
1315 Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng,
1316 Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
1317 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli
1318 Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Han-
1319 wei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu,
1320 Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang
1321 Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai
1322 Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang,
1323 Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia,
1324 Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang,
1325 Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu
1326 Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J.
1327 Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen,
1328 Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S.
1329 Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun,
1330 T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun
1331 Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan
1332 Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao
1333 Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue
1334 Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan
1335 Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X.
1336 Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang,
1337 Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong
1338 Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian
1339 Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxi-
1340 ang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang,
1341 Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yut-
1342 ing Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie,
1343 Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui
1344 Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen
1345 Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incenti-
1346 zing reasoning capability in llms via reinforcement learning, 2025. URL
1347 <https://arxiv.org/abs/2501.12948>.
- 1348
1349 Marco Drago, Sergey Klimenko, Claudia Lazzaro, Edoardo Milotti, Guenakh
1350 Mitselmakher, Valentin Necula, Brendan O’Brian, Giovanni Andrea Prodi,
1351 Francesco Salemi, Marek Szczepanczyk, Shubhanshu Tiwari, Vaibhav Ti-
1352 wari, Gayathri V, Gabriele Vedovato, and Igor Yakushin. coherent wave-
1353 burst, a pipeline for unmodeled gravitational-wave data analysis. *SoftwareX*,
1354 14:100678, 2021. ISSN 2352-7110. doi: <https://doi.org/10.1016/j.softx>.

- 1350 2021.100678. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S2352711021000236)
1351 [article/pii/S2352711021000236](https://www.sciencedirect.com/science/article/pii/S2352711021000236).
1352
- 1353 Agoston E. Eiben and Jim Smith. From evolutionary computation to the evolution
1354 of things. *Nature*, 521(7553):476–482, May 2015. ISSN 1476-4687. doi:
1355 10.1038/nature14544.
- 1356 Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture
1357 search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.
1358 URL <http://jmlr.org/papers/v20/18-598.html>.
1359
- 1360 Lee S. Finn. Detection, measurement, and gravitational radiation. *Phys. Rev. D*,
1361 46:5236–5249, Dec 1992. doi: 10.1103/PhysRevD.46.5236. URL [https://](https://link.aps.org/doi/10.1103/PhysRevD.46.5236)
1362 link.aps.org/doi/10.1103/PhysRevD.46.5236.
1363
- 1364 Hunter Gabbard, Michael Williams, Fergus Hayes, and Chris Messenger.
1365 Matching matched filtering with deep networks for gravitational-wave as-
1366 tronomy. *Phys. Rev. Lett.*, 120:141103, Apr 2018. doi: 10.1103/
1367 PhysRevLett.120.141103. URL [https://link.aps.org/doi/10.](https://link.aps.org/doi/10.1103/PhysRevLett.120.141103)
1368 [1103/PhysRevLett.120.141103](https://link.aps.org/doi/10.1103/PhysRevLett.120.141103).
1369
- 1370 Daniel George and E. A. Huerta. Deep neural networks to enable real-time mul-
1371 timessenger astrophysics. *Phys. Rev. D*, 97:044039, Feb 2018. doi: 10.1103/
1372 PhysRevD.97.044039. URL [https://link.aps.org/doi/10.1103/](https://link.aps.org/doi/10.1103/PhysRevD.97.044039)
1373 [PhysRevD.97.044039](https://link.aps.org/doi/10.1103/PhysRevD.97.044039).
1374
- 1375 E. A. Huerta, Asad Khan, Xiaobo Huang, Minyang Tian, Maksim Levental, Ryan
1376 Chard, Wei Wei, Maeve Heflin, Daniel S. Katz, Volodymyr Kindratenko, Dawei
1377 Mu, Ben Blaiszik, and Ian Foster. Accelerated, scalable and reproducible AI-
1378 driven gravitational wave detection. *Nature Astronomy*, 5(10):1062–1068, Oc-
1379 tober 2021. ISSN 2397-3366. doi: 10.1038/s41550-021-01405-0.
- 1380 Daniel Kahneman. *Thinking, Fast and Slow*. Macmillan, New York, 2011. ISBN
1381 978-0-374-27563-1.
- 1382 George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan
1383 Wang, and Liu Yang. Physics-informed machine learning. *Nature Re-*
1384 *views Physics*, 3(6):422–440, June 2021. ISSN 2522-5820. doi: 10.1038/
1385 s42254-021-00314-5.
- 1386 S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Laz-
1387 zaro, K. Ackley, S. Tiwari, C. F. Da Silva, and G. Mitselmakher. Method
1388 for detection and reconstruction of gravitational wave transients with networks
1389 of advanced detectors. *Phys. Rev. D*, 93:042004, Feb 2016a. doi: 10.1103/
1390 PhysRevD.93.042004. URL [https://link.aps.org/doi/10.1103/](https://link.aps.org/doi/10.1103/PhysRevD.93.042004)
1391 [PhysRevD.93.042004](https://link.aps.org/doi/10.1103/PhysRevD.93.042004).
1392
- 1393 S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Laz-
1394 zaro, K. Ackley, S. Tiwari, C. F. Da Silva, and G. Mitselmakher. Method
1395 for detection and reconstruction of gravitational wave transients with networks
1396 of advanced detectors. *Phys. Rev. D*, 93:042004, Feb 2016b. doi: 10.1103/
1397 PhysRevD.93.042004. URL [https://link.aps.org/doi/10.1103/](https://link.aps.org/doi/10.1103/PhysRevD.93.042004)
1398 [PhysRevD.93.042004](https://link.aps.org/doi/10.1103/PhysRevD.93.042004).
1399
- 1400 John R. Koza. Genetic programming as a means for programming computers by
1401 natural selection. *Statistics and Computing*, 4(2):87–112, June 1994. ISSN
1402 1573-1375. doi: 10.1007/BF00175355.
- 1403 Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521
(7553):436–444, May 2015. ISSN 1476-4687. doi: 10.1038/nature14539.

- 1404 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk
1405 Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag,
1406 Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant
1407 Misra. Solving quantitative reasoning problems with language models. 35:
1408 3843–3857, 2022. doi: 10.48550/arXiv.2206.14858.
- 1409 Yang Li, Dong Du, Linfeng Song, Chen Li, Weikang Wang, Tao Yang, and Haitao
1410 Mi. Hunyuanprover: A scalable data synthesis framework and guided tree
1411 search for automated theorem proving, 2025. URL [https://arxiv.org/
1412 abs/2412.20735](https://arxiv.org/abs/2412.20735).
- 1413 Fei Liu, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang. Algorithm evolu-
1414 tion using large language model, 2023. URL [https://arxiv.org/abs/
1415 2311.15249](https://arxiv.org/abs/2311.15249).
- 1416 Fei Liu, Xialiang Tong, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao
1417 Lu, and Qingfu Zhang. Evolution of heuristics: Towards efficient automatic
1418 algorithm design using large language model, 2024. URL [https://arxiv.
1419 org/abs/2401.02051](https://arxiv.org/abs/2401.02051).
- 1420 Cody Messick, Kent Blackburn, Patrick Brady, Patrick Brockill, Kipp Cannon,
1421 Romain Cariou, Sarah Caudill, Sydney J. Chamberlin, Jolien D. E. Creighton,
1422 Ryan Everett, Chad Hanna, Drew Keppel, Ryan N. Lang, Tjonnie G. F. Li, Dun-
1423 can Meacher, Alex Nielsen, Chris Pankow, Stephen Privitera, Hong Qi, Surabhi
1424 Sachdev, Laleh Sadeghian, Leo Singer, E. Gareth Thomas, Leslie Wade, Made-
1425 line Wade, Alan Weinstein, and Karsten Wiesner. Analysis framework for
1426 the prompt discovery of compact binary mergers in gravitational-wave data.
1427 *Phys. Rev. D*, 95:042001, Feb 2017. doi: 10.1103/PhysRevD.95.042001. URL
1428 <https://link.aps.org/doi/10.1103/PhysRevD.95.042001>.
- 1429 Christoph Molnar. *Interpretable Machine Learning*. Leanpub, Munich, Germany,
1430 2nd edition, 2020. ISBN 9780244768522. URL [https://christophm.
1431 github.io/interpretable-ml-book/](https://christophm.github.io/interpretable-ml-book/).
- 1432 Narenraju Nagarajan and Christopher Messenger. Identifying and mitigating ma-
1433 chine learning biases for the gravitational-wave detection problem, 2025. URL
1434 <https://arxiv.org/abs/2501.13846>.
- 1435 Alexander H. Nitz, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Shichao Wu,
1436 Marlin Schäfer, Rahul Dhurkunde, and Collin D. Capano. 4-ogc: Catalog of
1437 gravitational waves from compact binary mergers. *The Astrophysical Jour-
1438 nal*, 946(2):59, mar 2023. doi: 10.3847/1538-4357/aca591. URL [https:
1439 //dx.doi.org/10.3847/1538-4357/aca591](https://dx.doi.org/10.3847/1538-4357/aca591).
- 1440 Paraskevi Nousi, Alexandra E. Koloniari, Nikolaos Passalis, Panagiotis Iosif,
1441 Nikolaos Stergioulas, and Anastasios Tefas. Deep residual networks for grav-
1442 itational wave detection. *Phys. Rev. D*, 108:024022, Jul 2023. doi: 10.
1443 1103/PhysRevD.108.024022. URL [https://link.aps.org/doi/10.
1444 1103/PhysRevD.108.024022](https://link.aps.org/doi/10.1103/PhysRevD.108.024022).
- 1445 OpenAI. Learning to reason with llms. [https://openai.com/index/learning-to-
1446 reason-with-llms/](https://openai.com/index/learning-to-reason-with-llms/), 2024. Accessed: 2024.
- 1447 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge
1448 Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam
1449 Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Val-
1450 erie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum,
1451 Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner,
1452 Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg
1453 Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie
1454 Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael,
1455 Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen,
1456

- 1458 Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won
 1459 Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux,
 1460 Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan,
 1461 Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou,
 1462 David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte,
 1463 Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun
 1464 Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Graf-
 1465 stein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo,
 1466 Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Hei-
 1467 decke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon
 1468 Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain,
 1469 Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin,
 1470 Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Ka-
 1471 mali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kil-
 1472 patrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner,
 1473 Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kon-
 1474 drich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael
 1475 Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming
 1476 Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez,
 1477 Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor
 1478 Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob
 1479 McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake
 1480 McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey
 1481 Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Moss-
 1482 ing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro
 1483 Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh,
 1484 Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ash-
 1485 ley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex
 1486 Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Bel-
 1487 bute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Poko-
 1488 rny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris
 1489 Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh,
 1490 Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted,
 1491 Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar,
 1492 Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam,
 1493 Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam,
 1494 Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian
 1495 Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski
 1496 Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B.
 1497 Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle,
 1498 Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vi-
 1499 jayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang,
 1500 Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Pe-
 1501 ter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens
 1502 Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff
 1503 Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wo-
 1504 jciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao,
 1505 Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 tech-
 1506 nical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- 1504 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela
 1505 Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John
 1506 Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda
 1507 Askeel, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. Training
 1508 language models to follow instructions with human feedback. 35:27730–27744,
 1509 2022. doi: 10.48550/arXiv.2203.02155.
- 1510 Benjamin J. Owen. Search templates for gravitational waves from inspiraling
 1511 binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, Jun 1996.

- 1512 doi: 10.1103/PhysRevD.53.6749. URL [https://link.aps.org/doi/](https://link.aps.org/doi/10.1103/PhysRevD.53.6749)
1513 [10.1103/PhysRevD.53.6749](https://link.aps.org/doi/10.1103/PhysRevD.53.6749).
1514
- 1515 Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov,
1516 Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan
1517 S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein
1518 Fawzi. Mathematical discoveries from program search with large language
1519 models. *Nature*, 625(7995):468–475, January 2024. ISSN 1476-4687. doi:
1520 10.1038/s41586-023-06924-6.
- 1521 Cynthia Rudin. Stop explaining black box machine learning models for high
1522 stakes decisions and use interpretable models instead. *Nature Machine In-*
1523 *telligence*, 1(5):206–215, May 2019. ISSN 2522-5839. doi: 10.1038/
1524 s42256-019-0048-x.
- 1525
- 1526 Marlin B. Schäfer and Alexander H. Nitz. From one to many: A deep learning
1527 coincident gravitational-wave search. *Phys. Rev. D*, 105:043003, Feb 2022. doi:
1528 10.1103/PhysRevD.105.043003. URL [https://link.aps.org/doi/](https://link.aps.org/doi/10.1103/PhysRevD.105.043003)
1529 [10.1103/PhysRevD.105.043003](https://link.aps.org/doi/10.1103/PhysRevD.105.043003).
- 1530 Marlin B. Schäfer, Ondřej Zelenka, Alexander H. Nitz, Frank Ohme, and
1531 Bernd Brügmann. Training strategies for deep learning gravitational-wave
1532 searches. *Phys. Rev. D*, 105:043002, Feb 2022. doi: 10.1103/PhysRevD.105.
1533 043002. URL [https://link.aps.org/doi/10.1103/PhysRevD.](https://link.aps.org/doi/10.1103/PhysRevD.105.043002)
1534 [105.043002](https://link.aps.org/doi/10.1103/PhysRevD.105.043002).
- 1535
- 1536 Marlin B. Schäfer, Ondřej Zelenka, Alexander H. Nitz, He Wang, Shichao Wu,
1537 Zong-Kuan Guo, Zhoujian Cao, Zhixiang Ren, Paraskevi Nousi, Nikolaos Stergioulas,
1538 Panagiotis Iosif, Alexandra E. Koloniari, Anastasios Tefas, Nikolaos Passalis,
1539 Francesco Salemi, Gabriele Vedovato, Sergey Klimenko, Tanmaya Mishra,
1540 Bernd Brügmann, Elena Cuoco, E. A. Huerta, Chris Messenger, and Frank Ohme.
1541 First machine learning gravitational-wave search mock data challenge. *Phys. Rev. D*,
1542 107:023021, Jan 2023. doi: 10.1103/PhysRevD.107.023021. URL [https://link.aps.org/doi/10.1103/PhysRevD.](https://link.aps.org/doi/10.1103/PhysRevD.107.023021)
1543 [107.023021](https://link.aps.org/doi/10.1103/PhysRevD.107.023021).
- 1544
- 1545 E. Troja, L. Piro, H. van Eerten, R. T. Wollaeger, M. Im, O. D. Fox, N. R. Butler,
1546 S. B. Cenko, T. Sakamoto, C. L. Fryer, R. Ricci, A. Lien, R. E. Ryan, O. Korobkin,
1547 S.-K. Lee, J. M. Burgess, W. H. Lee, A. M. Watson, C. Choi, S. Covino, P. D’Avanzo,
1548 C. J. Fontes, J. Becerra González, H. G. Khandrika, J. Kim, S.-L. Kim, C.-U. Lee,
1549 H. M. Lee, A. Kuttyrev, G. Lim, R. Sánchez-Ramírez, S. Veilleux, M. H. Wieringa,
1550 and Y. Yoon. The X-ray counterpart to the gravitational-wave event GW170817. *Nature*,
1551 551(7678):71–74, November 2017. ISSN 1476-4687. doi: 10.1038/nature24290.
- 1552
- 1553 Samantha A Usman, Alexander H Nitz, Ian W Harry, Christopher M Biwer, Duncan
1554 A Brown, Miriam Cabero, Collin D Capano, Tito Dal Canton, Thomas Dent,
1555 Stephen Fairhurst, Marcel S Kehl, Drew Keppel, Badri Krishnan, Amber Lenon,
1556 Andrew Lundgren, Alex B Nielsen, Larne P Pekowsky, Harald P Pfeiffer, Peter R
1557 Saulson, Matthew West, and Joshua L Willis. The pycbc search for gravitational
1558 waves from compact binary coalescence. *Classical and Quantum Gravity*, 33(21):215004,
1559 oct 2016. doi: 10.1088/0264-9381/33/21/215004. URL [https://dx.doi.org/10.1088/0264-9381/33/](https://dx.doi.org/10.1088/0264-9381/33/21/215004)
1560 [21/215004](https://dx.doi.org/10.1088/0264-9381/33/21/215004).
- 1561
- 1562 Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming
1563 Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima
1564 Anandkumar, Karianne Bergen, Carla P. Gomes, Shirley Ho, Pushmeet Kohli,
1565 Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun Manrai, Debora Marks, Bharath
Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Veličković, Max

- 1566 Welling, Linfeng Zhang, Connor W. Coley, Yoshua Bengio, and Marinka Zit-
1567 nik. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):
1568 47–60, August 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06221-2.
- 1569
- 1570 He Wang, Shichao Wu, Zhoujian Cao, Xiaolin Liu, and Jian-Yang Zhu.
1571 Gravitational-wave signal recognition of ligo data by deep learning. *Phys.*
1572 *Rev. D*, 101:104003, May 2020a. doi: 10.1103/PhysRevD.101.104003. URL
1573 <https://link.aps.org/doi/10.1103/PhysRevD.101.104003>.
- 1574 Linnan Wang, Yiyang Zhao, Yuu Jinnai, Yuandong Tian, and Rodrigo Fonseca.
1575 Neural architecture search using deep neural networks and monte carlo tree
1576 search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(06):
1577 9983–9991, Apr. 2020b. doi: 10.1609/aaai.v34i06.6554. URL [https://](https://ojs.aaai.org/index.php/AAAI/article/view/6554)
1578 ojs.aaai.org/index.php/AAAI/article/view/6554.
- 1579
- 1580 David H Wolpert and William G Macready. No free lunch theorems for optimiza-
1581 tion. *IEEE transactions on evolutionary computation*, 1(1):67–82, 2002.
- 1582 Haoran Ye, Jiarui Wang, Zhiguang Cao, Federico Berto, Chuanbo Hua, Haeyeon
1583 Kim, Jinkyoo Park, and Guojie Song. Reevo: Large language models as hyper-
1584 heuristics with reflective evolution, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2402.01145)
1585 [abs/2402.01145](https://arxiv.org/abs/2402.01145).
- 1586 Ondřej Zelenka, Bernd Brüggemann, and Frank Ohme. Convolutional neural net-
1587 works for signal detection in real ligo data. *Phys. Rev. D*, 110:024024, Jul
1588 2024. doi: 10.1103/PhysRevD.110.024024. URL [https://link.aps.](https://link.aps.org/doi/10.1103/PhysRevD.110.024024)
1589 [org/doi/10.1103/PhysRevD.110.024024](https://link.aps.org/doi/10.1103/PhysRevD.110.024024).
- 1590
- 1591 Rui Zhang, Fei Liu, Xi Lin, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. Un-
1592 derstanding the importance of evolutionary search in automated heuristic de-
1593 sign with large language models, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2407.10873)
1594 [2407.10873](https://arxiv.org/abs/2407.10873).
- 1595 Yizhen Zheng, Huan Yee Koh, Jiaxin Ju, Anh T. N. Nguyen, Lauren T. May, Geof-
1596 frey I. Webb, and Shirui Pan. Large language models for scientific discovery in
1597 molecular property prediction. *Nature Machine Intelligence*, February 2025a.
1598 ISSN 2522-5839. doi: 10.1038/s42256-025-00994-z.
- 1599
- 1600 Zhi Zheng, Zhuoliang Xie, Zhenkun Wang, and Bryan Hooi. Monte carlo tree
1601 search for comprehensive exploration in llm-based automatic heuristic design,
1602 2025b. URL <https://arxiv.org/abs/2501.08603>.

1603 A SUPPLEMENTARY MATERIAL

1604 A.1 LLM PROMPTING TEMPLATES

1605 This section provides comprehensive details of the prompting strategies employed
1606 across different phases of the algorithmic discovery process. The templates are de-
1607 signed to guide language models through systematic reasoning while incorporat-
1608 ing domain-specific knowledge and maintaining consistency across evolutionary
1609 operations.

1610 **System Context and Task Definition.** All interactions with the LLM ensemble
1611 begin with a standardized system prompt that establishes the expert role and
1612 problem context:

1613 You are an expert in gravitational wave signal detection
1614 algorithms. Your task is to design heuristics that can
1615 effectively solve optimization problems. The task involves
1616 constructing a pipeline for gravitational wave signal
1617 detection. This pipeline will encompass data conditioning
1618 and time-frequency transformations as part of the signal
1619 processing workflow. The input will consist of raw, finite-

1620 length dual-channel gravitational wave data from the H1 and
 1621 L1 detectors. The pipeline will be tested on segmented data
 1622 spanning several weeks, with each segment having variable
 1623 length (7000s–30000s). Each segment's dual-channel data will
 1624 be directly used as input. The ultimate goal is to produce
 1625 a catalog of potential gravitational wave signals, where
 1626 each trigger includes information such as GPS time, ranking
 1627 statistic, and the timing accuracy of the prediction. This
 1628 systematic approach is essential for effectively identifying
 and cataloging candidate gravitational wave signals.

1629 This system prompt serves multiple purposes: (i) establishing domain expertise
 1630 expectations, (ii) defining the specific optimization context, (iii) specifying input
 1631 data characteristics, and (iv) clarifying the expected output format and evaluation
 1632 criteria.

1634 A.1.1 INITIAL ALGORITHM GENERATION PROMPTS

1635 **Seed Function Template and Analysis Framework.** The initial algorithm gen-
 1636 eration process begins with a structured analysis of the seed function to estab-
 1637 lish baseline understanding. The seed function analysis template guides the LLM
 1638 through systematic examination of the foundational algorithm:

```

1639 ## Seed Function Analysis Task
1640 Analyze the foundational algorithm's design strategy to
1641 establish baseline understanding for Monte Carlo Tree Search
1642 (MCTS) exploration. This first-level analysis will guide
1643 subsequent optimization directions.

1644 ## Seed Function Implementation
1645 ```python
1646 {prompt_seed_func}
1647 ```

1648 - **Technical implementation details**:: {prompt_other_inf}
1649 - **Performance impact rationale**:: {prompt_inout_inf}

1650 ## Context for Analysis
1651 This initial analysis at MCTS depth first-level should:
1652 - Identify core algorithmic mechanisms
1653 - Extract fundamental processing stages
1654 - Surface high-level optimization opportunities
1655 - Establish baseline for diversity generation
1656 {external_knowledge}

1657 ## Analysis Requirements
1658 1. Characterize the seed's core approach in one sentence
1659 containing:
1660 - Primary computational strategy
1661 - Key transformation stages
1662 - Fundamental signal processing techniques
1663 - Overall optimization philosophy

1664 2. Focus on architectural-level characteristics rather than
1665 implementation details

1666 3. Description must fit within single braces and avoid:
1667 - Code references
1668 - Parameter-level details
1669 - Performance assessments
1670 - Comparative statements

1671 ## Output Format Rules
1672 - Return optimization strategies within SINGLE BRACE
1673 - Ensure entire response can be parseable by regex:
  \\\{\\{(.*)\\}\\} with DOTALL flag

```

Seed Algorithm Specification. The seed function implements a three-stage linear signal processing pipeline that serves as the evolutionary starting point:

- Stage 1: Data Conditioning and Whitening

```

1  def data_conditioning(strain_h1: np.ndarray, strain_l1: np.ndarray, times: np.ndarray) ->
  ↳ tuple[np.ndarray, np.ndarray, np.ndarray]:
2      window_length = 4096
3      dt = times[1] - times[0]
4      fs = 1.0 / dt
5
6      def whiten_strain(strain):
7          strain_zeromean = strain - np.mean(strain)
8          freqs, psd = signal.welch(strain_zeromean, fs=fs, nperseg=window_length,
9                                  window='hann', noverlap=window_length//2)
10         smoothed_psd = np.convolve(psd, np.ones(32) / 32, mode='same')
11         smoothed_psd = np.maximum(smoothed_psd, np.finfo(float).tiny)
12         white_fft = np.fft.rfft(strain_zeromean) /
  ↳ np.sqrt(np.interp(np.fft.rfftfreq(len(strain_zeromean), d=dt), freqs,
  ↳ smoothed_psd))
13         return np.fft.irfft(white_fft)
14
15     whitened_h1 = whiten_strain(strain_h1)
16     whitened_l1 = whiten_strain(strain_l1)
17
18     return whitened_h1, whitened_l1, times

```

- Stage 2: Time-Frequency Decomposition

```

1  def compute_metric_series(h1_data: np.ndarray, l1_data: np.ndarray, time_series: np.ndarray)
  ↳ -> tuple[np.ndarray, np.ndarray]:
2      fs = 1 / (time_series[1] - time_series[0])
3      f_h1, t_h1, Sxx_h1 = signal.spectrogram(h1_data, fs=fs, nperseg=256, noverlap=128,
  ↳ mode='magnitude', detrend=False)
4      f_l1, t_l1, Sxx_l1 = signal.spectrogram(l1_data, fs=fs, nperseg=256, noverlap=128,
  ↳ mode='magnitude', detrend=False)
5      tf_metric = np.mean((Sxx_h1**2 + Sxx_l1**2) / 2, axis=0)
6      gps_mid_time = time_series[0] + (time_series[-1] - time_series[0]) / 2
7      metric_times = gps_mid_time + (t_h1 - t_h1[-1]) / 2
8
9      return tf_metric, metric_times

```

- Stage 3: Peak Detection and Trigger Generation

```

1  def calculate_statistics(tf_metric, t_h1):
2      background_level = np.median(tf_metric)
3      peaks, _ = signal.find_peaks(tf_metric, height=background_level * 1.0, distance=2,
  ↳ prominence=background_level * 0.3)
4      peak_times = t_h1[peaks]
5      peak_heights = tf_metric[peaks]
6      peak_deltat = np.full(len(peak_times), 10.0) # Fixed uncertainty value
7      return peak_times, peak_heights, peak_deltat

```

Template Variables and Customization. The prompting template incorporates several customizable variables that enable systematic variation generation:

- `prompt_seed_func`: Complete seed function implementation
- `prompt_other_inf`: Technical implementation details including sampling rates, window parameters, and algorithmic constraints
- `prompt_inout_inf`: Performance impact rationale explaining the relationship between input characteristics and expected output quality
- `external_knowledge`: Domain-specific knowledge injection including gravitational wave physics, detector characteristics, and signal morphology constraints

Output Format Requirements. All generated algorithms must conform to the standardized interface:

```

1  def pipeline_v(N)(strain_h1: np.ndarray, strain_l1: np.ndarray, times: np.ndarray) ->
  ↳ tuple[np.ndarray, np.ndarray, np.ndarray]:
2      # Algorithm implementation for seed function

```

```

1728     3     # ...
1729     4     # Stage 1: Data Conditioning and Whitening
1730     5     whitened_h1, whitened_l1, data_times = data_conditioning(strain_h1, strain_l1, times)
1731     6     # Stage 2: Time-Frequency Decomposition
1732     7     tf_metric, metric_times = compute_metric_series(whitened_h1, whitened_l1, data_times)
1733     8     # Stage 3: Peak Detection and Trigger Generation
1734     9     peak_times, peak_heights, peak_deltat = calculate_statistics(tf_metric, metric_times)
1735    10     return peak_times, peak_heights, peak_deltat

```

This interface consistency ensures that all generated algorithms can be evaluated within the same framework while enabling diverse internal implementations.

A.1.2 PARENT CROSSOVER IMPLEMENTATION

The Parent Crossover (PC) operation represents a sophisticated genetic operation that combines algorithmic components from two reference implementations at different levels of the MCTS hierarchy. This operation is designed to preserve successful characteristics from both parent algorithms while introducing novel enhancements that exceed simple interpolation.

Template Structure and Crossover Strategy. The PC operation employs a structured template that guides the LLM through systematic analysis and synthesis of two parent algorithms:

```

1747     ## Task Overview
1748     Develop a novel algorithm that strategically combines
1749     components from two reference implementations while
1750     introducing innovative enhancements. The solution must
1751     demonstrate measurable improvements beyond simple
1752     interpolation of existing approaches.
1753     Current Depth Level: [Level {depth}]
1754
1754     ## Implementation Analysis
1755     ### Code Comparison
1756     1. VERSION A (Baseline Implementation):
1757     ```python
1758     {worse_code}
1759     ```
1760     2. VERSION B (Enhanced Implementation):
1761     ```python
1762     {better_code}
1763     ```
1764     ### Strengths to Combine
1765     ```text
1766     {reflection}
1767     ```
1768     Key Synthesis Requirements:
1769     - Preserve 2 distinct advantages from Version A
1770     - Incorporate 3 critical enhancements from Version B
1771     - Identify 1 synergistic improvement opportunity
1772
1772     ## Architecture Strategy
1773     {external_knowledge}
1774
1774     ### Depth-Specific Synthesis Guidelines (Depth={depth})
1775     1. Structural Synthesis (Depth 1-2):
1776     - Create hybrid control flow combining best elements from
1777     both versions
1778     - Example: "Combine Version A's iteration structure with
1779     Version B's termination conditions"
1780     - Forbid direct replication of either version's
1781     architecture
1782
1782     2. Implementation Fusion (Depth 3-4):

```

- 1782 - Develop novel parameter hybridization techniques
 - 1783 - Example: "Blend Version A's exploration mechanism with
 - 1784 Version B's exploitation strategy"
 - 1785 - Require at least one innovative combination per
 - 1786 functional module
- 1787
- 1788 3. Mathematical Innovation (Depth 5+):
- 1789 - Derive new computational operators through version
 - 1790 synthesis
 - 1791 - Example: "Fuse Version A's approximation method with
 - 1792 Version B's error correction"
 - 1793 - Mandate 10-20% computational complexity reduction

1793 This template structure ensures that the crossover operation is not merely concate-
1794 native but involves intelligent analysis and strategic combination of algorithmic
1795 strengths.

1796 **Depth-Adaptive Synthesis Guidelines.** The PC operation implements depth-
1797 specific strategies that adapt the crossover complexity based on the current posi-
1798 tion in the MCTS tree:

- 1799 • Structural Synthesis (Depth 1-2):
 - 1800 - Focuses on combining high-level architectural elements from both parent
 - 1801 algorithms
 - 1802 - Creates hybrid control flow structures that merge the best organizational
 - 1803 patterns
 - 1804 - Example directive: "Combine Version A's iteration structure with Version
 - 1805 B's termination conditions"
 - 1806 - Explicitly forbids direct replication of either parent's complete architec-
 - 1807 ture
- 1808 • Implementation Fusion (Depth 3-4):
 - 1809 - Emphasizes parameter hybridization and functional module integration
 - 1810 - Develops novel approaches to blend algorithmic strategies
 - 1811 - Example directive: "Blend Version A's exploration mechanism with Ver-
 - 1812 sion B's exploitation strategy"
 - 1813 - Requires at least one innovative combination per functional module
 - 1814
- 1815 • Mathematical Innovation (Depth 5+):
 - 1816 - Derives new computational operators through sophisticated version syn-
 - 1817 thesis
 - 1818 - Focuses on mathematical justification for algorithmic improvements
 - 1819 - Example directive: "Fuse Version A's approximation method with Ver-
 - 1820 sion B's error correction"
 - 1821 - Mandates 10-20% computational complexity reduction alongside perfor-
 - 1822 mance gains

1823 **Innovation Requirements and Quality Assurance.** The PC operation enforces
1824 strict innovation standards to ensure that generated algorithms represent genuine
1825 improvements:

- 1826 • Core Innovation Targets:
 - 1827 - Synthesize 3+ novel elements not present in either parent version
 - 1828 - Resolve 2 fundamental limitations identified through comparative analy-
 - 1829 sis
 - 1830 - Introduce 1 breakthrough enhancement with rigorous mathematical jus-
 - 1831 tification
 - 1832 - Demonstrate non-trivial performance gains over both parent algorithms
 - 1833 - Prohibit direct replication of complete code blocks from either parent
 - 1834

1835 **Reflection Generation Process.** Before conducting the crossover synthesis, the
system generates analytical insights through a depth-adaptive reflection template:

```

1836     ## Task Objective
1837     Analyze optimization patterns across algorithm versions and
1838     generate depth-specific improvement strategies. Current MCTS
1839     Depth: depth/max_depth={depth}/{max_depth}
1840
1841     ## Depth-Specific Focus
1842     - Shallow (Depth 1-2): Structural patterns & control flow
1843     - Medium (Depth 3-4): Implementation techniques &
1844     parameterization
1845     - Deep (Depth 5+): Mathematical formulations & computational
1846     primitives
1847
1848     ## Algorithm Comparison
1849     - Original (Suboptimal)
1850     ```python
1851     {code_worse}
1852     ```
1853     - Improved (Optimized)
1854     ```python
1855     {code_better}
1856     ```
1857
1858     ## Depth-Adaptive Analysis
1859     ### 1. Core Pattern Extraction
1860     For {depth}-level analysis:
1861     - Shallow: Compare control structures/algorithmic paradigms
1862     - Medium: Analyze parameter configurations/function
1863     compositions
1864     - Deep: Examine mathematical operators/numerical methods
1865
1866     ### 2. Optimization Principle Generation
1867     Generate 3-5 transferable rules that:
1868     - Directly address {depth}-specific limitations
1869     - Contain concrete parameter values from improved version
1870     - Maintain functional equivalence
1871
1872     ## Output Format Rules
1873     - Return optimization strategies within SINGLE BRACE
1874     - Ensure entire response can be parseable by regex:
1875     \{\{(.*)\}\} with DOTALL flag
1876
1877     This reflection generation produces the reflection variable used in the main
1878     crossover template.
1879
1880     Reflection-Guided Analysis. The crossover process incorporates a reflection
1881     component that analyzes the strengths and weaknesses of both parent algorithms:
1882
1883     ## Requirements
1884     1. Core Innovation Targets:
1885     - Synthesize 3+ novel elements not present in either
1886     version
1887     - Resolve 2 fundamental limitations identified in
1888     analysis
1889     - Introduce 1 breakthrough enhancement with mathematical
1890     justification
1891     - Demonstrate non-trivial performance gain over both
1892     versions
1893     - Prohibit direct replication of complete code blocks
1894
1895     This reflection analysis is generated through the deepseek-r1-250120 model
1896     and provides crucial insights that guide the synthesis process. The reflection identifies:
1897
1898     • Computational advantages in each parent algorithm
1899     • Structural design patterns that contribute to performance

```

- Potential synergistic combinations that could yield emergent benefits
- Limitation patterns that should be addressed in the offspring

Output Format and Validation. The PC operation enforces a standardized output format that ensures both human readability and automated processing:

2. Output Format:

- Place the core design idea in a sentence within a brace BEFORE the function definition
- For the core design idea format: `\{\{A hybrid gravitational wave detection pipeline...\}}`
- Implement as Python function: `{func_name}`
- Inputs: `{input_count}` parameter(s) (`{joined_inputs}`)
- Outputs: `{output_count}` return value(s) (`{joined_outputs}`)
- Follow: `{inout_inf}`
- Constraints: `{other_inf}`
- IMPORTANT: All output code MUST be valid Python syntax. Do not place description text inside curly braces within the function body.
- Example of correct format:


```

\{\{Core design description here\}\}
```python
def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.
ndarray, times: np.ndarray) -> tuple[np.ndarray, np.
ndarray, np.ndarray]:
 """Core design description can alternatively be placed
 here as a docstring"""
 # Function implementation...
...

```

### A.1.3 SIBLING CROSSOVER IMPLEMENTATION

The Sibling Crossover (SC) operation implements a sophisticated two-phase approach that leverages peer algorithm insights to generate improved offspring. Unlike Parent Crossover, which combines algorithms from different hierarchical levels, SC focuses on horizontal knowledge transfer between algorithms at the same MCTS depth, promoting diversity while maintaining comparable complexity levels.

**Two-Phase Architecture.** The SC operation employs a unique two-stage process: first generating optimization hints through multi-level reflection analysis (Phase 1), then implementing concrete algorithmic improvements based on these insights (Phase 2). This separation enables more targeted optimization by allowing the system to first identify improvement opportunities before implementing solutions.

**Phase 1: Multi-Level Reflection Analysis.** The first phase generates depth-specific optimization hints by synthesizing insights from sibling algorithms and parent-level analysis:

```

Task Overview
Generate depth-specific optimization hints for gravitational
wave detection algorithms by synthesizing multi-level
reflections.
Current Optimization Depth: {parent_depth}/{max_depth} (
shallow: structural patterns, medium: implementation
techniques, deep: mathematical details)

Contextual Insights
1. Peer Algorithm Reflections (Depth {parent_depth}):
- Formatted as performance-annotated entries: [No.N
Brother Reflection | Score: X]<reflection>
- Time-ordered weighting (newest=highest priority) with
objective score-based ranking
- Includes full technical post-mortems from immediate
ancestors
{parent_reflections}

```

```

1944
1945 2. Father Algorithm Analysis (Depth {father_depth}):
1946 {father_reflection}
1947
1948 ## Hint Generation Requirements
1949 1. Produce 3-5 executable optimization directives that:
1950 - Integrate cross-depth insights from peer
1951 implementations
1952 - Target {parent_depth}-level (shallow: structural
1953 patterns, medium: implementation techniques, deep:
1954 mathematical details) components for improvement
1955 - Formulate mathematically sound enhancements
1956 - Align with gravitational wave data processing
1957 objectives
1958
1959 2. Output Format Rules
1960 - Return optimization strategies within SINGLE BRACE
1961 - Ensure entire response can be parseable by regex:
1962 \{\{(.*)\}\} with DOTALL flag
1963 - Focus on {parent_depth}-appropriate modifications
1964 - Emphasize time-domain processing optimizations
1965
1966 ## Critical Constraints
1967 - Each directive must correspond to concrete code changes
1968 - Explicitly connect to reflection insights where applicable
1969 - Maintain strict {parent_depth}-level focus in all
1970 suggestions
1971 - Exclude explanatory text within the hint brace
1972 - Prioritize modifications matching current depth's
1973 optimization type
1974
1975 Sibling Selection and Weighting Strategy. The SC operation employs a sophis-
1976 ticated parent selection mechanism that prioritizes high-performing sibling algo-
1977 rithms:
1978
1979

1980 1 # Select parents based on objective value weights
1981 2 other = [ind for ind in pop if ind['code'] != father['code']]
1982 3 weights = [1.0 / (-ind['fitness'] + 1e-10) for ind in other] # Lower objective = higher weight
1983 4 normalized_weights = [w / sum(weights) for w in weights]
1984 5 parents = random.choices(other, weights=normalized_weights, k=min(self.m, len(other)))
1985

1986
1987 This weighting strategy ensures that successful algorithmic patterns from high-
1988 performing siblings are more likely to influence the offspring generation process.
1989
1990 Depth-Adaptive Optimization Focus. The first phase implements depth-specific
1991 optimization strategies that adapt to the current position in the MCTS tree:
1992
1993 • Shallow Depth (1-2): Focuses on structural patterns and control flow restruc-
1994 turing
1995 • Medium Depth (3-4): Emphasizes implementation techniques and numerical
1996 optimizations
1997 • Deep Depth (5+): Concentrates on mathematical details and advanced com-
1998 putational methods
1999
2000 Phase 2: Concrete Algorithm Implementation. The second phase transforms
2001 the optimization hints into executable algorithms:
2002
2003 ## Algorithm Optimization Task
2004 Develop an enhanced gravitational signal processing
2005 algorithm for interferometer data analysis by implementing
2006 concrete improvements from multi-level code analysis.
2007
2008 ## Technical Context
2009 1. Optimization Depth Specifications:
2010 - Current Focus Level: {depth} (max_depth={max_depth})

```

```

1998 (1-2: Control flow restructuring, 3-4: Numerical
1999 computation optimizations, 5+: Advanced linear algebra
2000 methods)
2001 - Code Analysis Insights from Prior Level:
2002 ```text
2003 {reflection}
2004 ```
2005
2006 2. Base Implementation Details:
2007 [Functional Purpose] {algorithm_description}
2008 [Core Implementation]
2009 ```python
2010 {algorithm_code}
2011 ```
2012
2013 ## Implementation Directives (Depth {depth}):
2014 - Shallow (1-2): Restructure control flow using reflection
2015 suggestion (e.g., split data conditioning/analysis phases)
2016 - Medium (3-4): Apply numerical optimizations from
2017 reflection (e.g., FFT window size optimization)
2018 - Deep (5+): Implement matrix computation improvements from
2019 reflection (e.g., regularized inverse covariance)
2020
2021 {external_knowledge}
2022
2023 ## Output Format
2024 - Place the core design idea in a sentence within a brace
2025 BEFORE the function definition
2026 - For the core design idea format: \\{{A hybrid
2027 gravitational wave detection pipeline...}}
2028 - Implement as Python function: {func_name}
2029 - Inputs: {input_count} parameter(s) ({joined_inputs})
2030 - Outputs: {output_count} return value(s) ({joined_outputs})
2031 - Follow: {inout_inf}
2032 - Constraints: {other_inf}
2033 - IMPORTANT: All output code MUST be valid Python syntax. Do
2034 not place description text inside curly braces within the
2035 function body.
2036 - Example of correct format:
2037 \\{{Core design description here}}
2038 ```python
2039 def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.
2040 ndarray, times: np.ndarray) -> tuple[np.ndarray, np.
2041 ndarray, np.ndarray]:
2042 """Core design description can alternatively be placed
2043 here as a docstring"""
2044 # Function implementation...
2045 ```
2046
2047 ## Important Notes
2048 - Focus on algorithmic improvements rather than code style
2049 changes
2050 - Ensure the new implementation directly addresses the
2051 reflection insights

```

**Reflection Processing and Template Variables.** The SC operation processes multiple sources of algorithmic insight through structured template variables:

- `parent_reflections`: Performance-annotated reflections from peer algorithms, formatted as ranked entries with objective scores
- `father_reflection`: Analysis from the immediate parent algorithm at depth-1
- `reflection`: Synthesized optimization hints generated in Phase 1
- `algorithm_description`: Functional description of the base algorithm

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

- `algorithm_code`: Complete implementation of the parent algorithm

**Quality Assurance and Validation.** The two-phase approach enables comprehensive quality control:

- Phase 1 Validation:
  - Ensures reflection insights are depth-appropriate
  - Validates mathematical soundness of optimization suggestions
  - Confirms alignment with gravitational wave processing objectives
- Phase 2 Validation:
  - Verifies syntactic correctness of generated code
  - Confirms interface compliance with standardized function signatures
  - Tests algorithmic improvements against reflection insights
  - Validates computational efficiency claims

**Temporal Weighting and Performance Ranking.** The SC operation implements sophisticated temporal weighting that prioritizes recent algorithmic discoveries while maintaining objective score-based ranking:

- Time-ordered weighting: Newer algorithms receive higher priority in the reflection synthesis
- Performance-based ranking: Algorithms with better objective scores contribute more heavily to the optimization hints
- Cross-depth integration: Insights from both peer algorithms and parent-level analysis are systematically combined

This comprehensive approach ensures that SC operations generate algorithms that not only improve upon their immediate ancestors but also incorporate the collective intelligence of high-performing siblings, leading to more robust and efficient gravitational wave detection strategies.

#### A.1.4 POINT MUTATION IMPLEMENTATION

Point Mutation (PM) operations introduce targeted modifications to individual algorithms based on performance analysis, implementing two distinct approaches that offer different levels of sophistication and computational investment. The framework provides both single-stage direct improvement and two-stage reflection-driven enhancement strategies.

**Single-Stage Point Mutation: Direct Algorithm Improvement.** The operation implements a straightforward approach that directly compares an original algorithm with a high-performing elite algorithm to generate improvements. This method prioritizes computational efficiency while maintaining effective algorithmic enhancement.

Template Structure:

```
Task Overview
You will analyze an original algorithm, an improved version
of it, and create a new enhanced algorithm. Below are the
key components:

Algorithm Details
1. ORIGINAL ALGORITHM:
 - Description: {original_algorithm_description}
 - Code:
  ```python
  {original_algorithm_code}
  ```
 - Objective Value: {original_objective_value}

2. BETTER ALGORITHM (Reference Implementation):
 - Description: {better_algorithm_description}
 - Code:
  ```python
  {better_algorithm_code}
```

```

2106     ...
2107         - Objective Value: {better_objective_value}
2108         - Improvement Insights:
2109     ```text
2110     {better_algorithm_reflection}
2111     ...
2112     ## Implementation Requirements
2113     1. Analyze the differences between the original and better
2114     algorithms
2115     2. Create a new algorithm that:
2116         - Incorporates successful elements from the better
2117         algorithm
2118         - Addresses limitations revealed in the improvement
2119         insights
2120         - Produces better results than the original algorithm
2121     3. Output format requirements:
2122         - Place the core design idea in a sentence within a brace
2123         BEFORE the function definition
2124         - For the core design idea format: \\\{A hybrid
2125         gravitational wave detection pipeline...}
2126         - Implement as Python function: {func_name}
2127         - Inputs: {input_count} parameter(s) ({joined_inputs})
2128         - Outputs: {output_count} return value(s) ({
2129         joined_outputs})
2130         - Follow: {inout_inf}
2131         - Constraints: {other_inf}
2132         - IMPORTANT: All output code MUST be valid Python syntax.
2133         Do not place description text inside curly braces within
2134         the function body.
2135         - Example of correct format:
2136             \\\{Core design description here}
2137             ```python
2138             def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.
2139             ndarray, times: np.ndarray) -> tuple[np.ndarray, np.
2140             ndarray, np.ndarray]:
2141                 """Core design description can alternatively be
2142                 placed here as a docstring"""
2143                 # Function implementation...
2144             ...
2145     {external_knowledge}
2146     ## Important Notes
2147     - Focus on algorithmic improvements rather than code style
2148     changes
2149     - Ensure the new implementation directly addresses the
2150     reflection insights
2151     Two-Stage Point Mutation: Reflection-Driven Enhancement. The operation
2152     implements a sophisticated two-phase approach that mirrors the sibling crossover
2153     methodology but focuses on individual algorithm improvement rather than hori-
2154     zontal knowledge transfer.
2155     Phase 1: Strategic Reflection Generation. The first phase synthesizes insights
2156     from multiple sources to generate comprehensive optimization guidelines:
2157     ## Task Overview
2158     Generate optimized technical guidelines for gravitational
2159     wave detection algorithms through systematic analysis of
2160     multi-generational reflection insights. Focus on enhancing
2161     data conditioning pipelines, time-frequency analysis methods
2162     , noise suppression techniques, and H1-L1 detector coherence
2163     optimization. Produce executable directives addressing:
2164     waveform recognition precision, computational complexity
2165     management, and non-stationary noise differentiation while
2166     maintaining strict API compliance.

```

```

2160
2161 ## Input Context
2162 1. NEW INSIGHTS FROM RECENT ITERATIONS:
2163   - Formatted as performance-annotated entries: [Parent N
2164     Reflection | Score: X]<reflection>
2165   - Time-ordered weighting (newest=highest priority) with
2166     objective score-based ranking
2167   - Includes full technical post-mortems from immediate
2168     ancestors
2169   {parent_reflections}
2170
2171 2. LONG-TERM REFLECTION REPOSITORY:
2172   - Contains battle-tested insights from top 1% performers
2173   - 3x weighting factor for architectural-level insights
2174   - Curated through 3-stage filtration:
2175     1. Statistical significance validation
2176     2. Cross-generational effectiveness verification
2177     3. Compatibility check with current detector
2178     configurations
2179   {elite_reflection}
2180
2181 ## Implementation Requirements
2182 1. Perform weighted synthesis of reflections
2183 2. Generate 3-5 technically-grounded optimization directives
2184 3. Prioritize:
2185   - Mitigation of historical implementation flaws
2186   - Amplification of proven effective patterns
2187   - Weighted integration of multi-generational insights
2188
2189 ## Output Format
2190 - Return all guidelines within SINGLE BRACE
2191 - Ensure entire response can be parseable by regex:
2192   \{\{(.*)\}\} with DOTALL flag
2193 - Concrete technical directives only
2194 - No explanatory text or formatting
2195
2196 Phase 2: Concrete Algorithm Implementation. The second phase transforms
2197 the strategic insights into executable algorithmic improvements:
2198
2199 ## Task Overview
2200 Leverage insights from prior strategic reflection to
2201 architecturally enhance the gravitational wave detection
2202 algorithm. Develop improvements that directly address
2203 identified limitations in CRITICAL REFLECTION INSIGHTS while
2204 preserving core functionality through:
2205
2206 1. Stage-level architectural modifications informed by
2207 reflection analysis
2208 2. Reflection-driven noise reduction and coherence
2209 enhancement strategies
2210 3. Time-frequency analysis variations targeting specific
2211 weaknesses identified
2212 4. H1-L1 synthesis improvements based on cross-detector
2213 insights
2214
2215 Generate architecturally distinct variants that implement
2216 reflection-derived concepts through fundamental structural
2217 changes.
2218
2219 ## Input Context
2220 1. CRITICAL REFLECTION INSIGHTS (Improvement Basis):
2221 ```text
2222 {reflection}
2223 ```

```

```

2214     2. REFERENCE IMPLEMENTATION:
2215     [Description] {elite_algorithm_description}
2216     [Baseline Code]
2217     ```python
2218     {elite_algorithm_code}
2219     ```
2220     ## Implementation Requirements
2221     1. Execute reflection-guided analysis:
2222         - Map reflection insights to specific code components
2223         - Identify 2-3 architectural limitations in current
2224           implementation
2225     2. Propose improvements that directly convert reflection
2226       insights into:
2227         - Enhanced signal path architecture
2228         - Novel noise handling structures
2229         - Optimized computational patterns
2230         - Advanced detector synergy mechanisms
2231     3. Maintain strict interface compatibility with existing
2232       system integration
2233
2234     {external_knowledge}
2235
2236     ## Output Format
2237     - Place the core design idea in a sentence within a brace
2238     BEFORE the function definition
2239     - For the core design idea format: \\{{A hybrid
2240       gravitational wave detection pipeline...}}
2241     - Implement as Python function: {func_name}
2242     - Inputs: {input_count} parameter(s) ({{joined_inputs}})
2243     - Outputs: {output_count} return value(s) ({{joined_outputs}})
2244     - Follow: {inout_inf}
2245     - Constraints: {other_inf}
2246     - IMPORTANT: All output code MUST be valid Python syntax. Do
2247       not place description text inside curly braces within the
2248       function body.
2249     - Example of correct format:
2250       \\{{Core design description here}}
2251       ```python
2252       def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.
2253         ndarray, times: np.ndarray) -> tuple[np.ndarray, np.
2254         ndarray, np.ndarray]:
2255           """Core design description can alternatively be placed
2256             here as a docstring"""
2257           # Function implementation...
2258           ...
2259
2260     ## Important Notes
2261     - Focus on algorithmic improvements rather than code style
2262     changes
2263     - Ensure the new implementation directly addresses the
2264     reflection insights
2265
2266     Selection Strategies and Elite Integration. Both PM operations leverage the
2267     elite offspring as a performance benchmark and source of successful algorithmic
     patterns. The key distinction lies in their selection strategies:
     

- Single-Stage Selection Strategy:
  - Selects a single parent algorithm from the population (excluding the elite)
  - Directly compares parent performance with elite offspring
  - Implements immediate improvement through direct analysis
- Two-Stage Selection Strategy:
  - Selects multiple parent algorithms for comprehensive reflection analysis
  - Incorporates both recent algorithmic insights and long-term elite patterns

```

2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321

- Implements sophisticated multi-generational knowledge synthesis

Computational Efficiency Considerations. The two PM approaches offer different computational trade-offs:

- Single-Stage Advantages:
 - Single-stage processing reduces computational overhead
 - Direct comparison enables rapid algorithm improvement
 - Simplified prompting reduces LLM token consumption
- Two-Stage Advantages:
 - Two-stage processing enables more sophisticated optimization
 - Multi-generational insight integration leads to more robust improvements
 - Reflection-driven approach produces more interpretable algorithmic modifications

Template Variable Integration. Both PM operations incorporate comprehensive template variables that enable systematic algorithmic improvement:

- Common Variables:
 - `func_name`, `input_count`, `output_count`: Interface specification
 - `joined_inputs`, `joined_outputs`: Parameter documentation
 - `better_algorithm_description`,
`better_algorithm_code`: Elite algorithm details
 - `original_objective_value`, `better_objective_value`: Performance metrics
 - `better_algorithm_reflection`: Elite algorithm insights
- Two-Stage Variables:
 - `parent_reflections`: Multi-parent reflection synthesis
 - `elite_reflection`: Long-term elite insights
 - `reflection`: Generated optimization guidelines

Quality Assurance and Validation. Both PM operations implement rigorous validation procedures:

- Single-Stage Validation:
 - Direct performance comparison with both parent and elite algorithms
 - Verification of improvement insight integration
 - Confirmation of interface compliance
- Two-Stage Validation:
 - Two-stage validation covering both reflection generation and implementation
 - Cross-generational consistency checking
 - Architectural improvement verification

The dual PM approach provides flexibility in algorithmic improvement strategies, enabling the framework to adapt to different optimization scenarios while maintaining consistent quality standards and interface compliance.

A.1.5 PATH-WISE CROSSOVER IMPLEMENTATION

Path-wise Crossover (PWC) operations synthesize information along complete root-to-leaf trajectories in the MCTS tree, capturing long-range dependencies and enabling global optimization strategies. The framework implements two distinct PWC approaches that differ in their analytical methodologies: reflection-based synthesis and comprehensive algorithm analysis.

Reflection-Based Path-wise Crossover: Multi-Algorithm Insight Synthesis. The operation implements a two-stage process that analyzes reflection patterns across multiple algorithms in a complete MCTS path to identify generalizable optimization principles.

Phase 1: Cross-Algorithm Pattern Analysis. The first phase extracts recurring technical strategies from multiple algorithm reflections:

```

2322     ## Task Overview
2323     Analyze and synthesize technical reflections from multiple
2324     algorithm iterations to identify cross-algorithm
2325     optimization patterns and guide next-generation algorithm
2326     design. Prioritize extraction of generalizable technical
2327     principles over implementation-specific details.
2328     Current Optimization Depth: depth/max_depth={depth}/{
2329     max_depth} (shallow: structural patterns, medium:
2330     implementation techniques, deep: mathematical details)
2331
2331     ## Input Context
2332     Analyzing {num_algorithms} algorithm reflections from MCTS
2333     exploration trajectories. Technical reflections follow depth
2334     -specific analysis requirements. Structural format: [No.N
2335     algorithm's reflection (depth: X)]<reflection>
2336     {algorithm_reflections}
2337
2337     ## Reflection Requirements
2338     1. Pattern Identification (Key Observed Patterns):
2339     - Extract 2-3 recurring technical strategies (e.g. "Multi
2340     -scale wavelet decomposition" not "used Morlet wavelet")
2341     - Categorize by analysis level:
2342     * Structural: Component architecture (e.g. "Parallel
2343     filter banks")
2344     * Implementation: Algorithmic choices (e.g. "Adaptive
2345     thresholding")
2346     * Mathematical: Core transforms (e.g. "Orthogonal
2347     matching pursuit")
2348
2348     2. Technical Pathway Analysis (Promising Technical
2349     Pathways):
2350     - Identify under-utilized but theoretically sound
2351     approaches (e.g. "Sparse representation in frequency
2352     domain")
2353     - Specify required technical components without code
2354     details (e.g. "Requires: Overcomplete basis construction
2355     ")
2356
2356     3. Optimization Principles (Strategic Optimization
2357     Principles):
2358     - Formulate depth-specific guidelines (e.g. "At
2359     mathematical level: Maximize time-frequency product  $\leq 0.5$ ")
2360     - Relate physical constraints to algorithmic parameters (
2361     e.g. "Wavelet duration should match typical glitch
2362     durations")
2363
2363     4. Specificity Balance:
2364     - Technical specificity: Name mathematical concepts (e.g.
2365     "Gabor uncertainty") and signal processing domains
2366     - Implementation avoidance: Omit code structures (e.g. "
2367     Avoid: 'Use 3 nested loops'")
2368
2368     ## Output Format Rules
2369     - Return optimization strategies within SINGLE BRACE
2370     - Ensure entire response can be parseable by regex:
2371     \{\{(.*)\}\} with DOTALL flag
2372     - Do not include markdown formatting or additional
2373     explanations
2374
2374     Phase 2: Algorithm Implementation. The second phase transforms the synthe-
2375     sized insights into concrete algorithmic improvements:
2376
2376     ## Task Overview

```

```

2376
2377     Develop an enhanced gravitational wave detection algorithm
2378     through targeted modifications addressing specific technical
2379     shortcomings identified in the reflection analysis.
2380
2381     ## Input Context
2382     [Critical Reflection Insights]
2383     ```text
2384     {reflection}
2385     ...
2386
2387     [Baseline Implementation]
2388     [Functional Description] {algorithm_description}
2389     [Current Codebase]
2390     ```python
2391     {algorithm_code}
2392     ...
2393
2394     {external_knowledge}
2395
2396     ## Output Format
2397     - Place the core design idea in a sentence within a brace
2398     BEFORE the function definition
2399     - For the core design idea format: \\{{A hybrid
2400     gravitational wave detection pipeline...}}
2401     - Implement as Python function: {func_name}
2402     - Inputs: {input_count} parameter(s) ({{joined_inputs}})
2403     - Outputs: {output_count} return value(s) ({{joined_outputs}})
2404     - Follow: {inout_inf}
2405     - Constraints: {other_inf}
2406     - IMPORTANT: All output code MUST be valid Python syntax. Do
2407     not place description text inside curly braces within the
2408     function body.
2409     - Example of correct format:
2410     \\{{Core design description here}}
2411     ```python
2412     def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.
2413     ndarray, times: np.ndarray) -> tuple[np.ndarray, np.
2414     ndarray, np.ndarray]:
2415         """Core design description can alternatively be placed
2416         here as a docstring"""
2417         # Function implementation...
2418     ...
2419
2420     ## Important Notes
2421     - Focus on algorithmic improvements rather than code style
2422     changes
2423     - Ensure the new implementation directly addresses the
2424     reflection insights
2425
2426     Comprehensive Algorithm Analysis Path-wise Crossover: Multi-Level Technical Synthesis. The operation implements a more sophisticated analytical approach that examines complete algorithm implementations across different depth levels.
2427
2428     Phase 1: Multi-Level Technical Analysis. The first phase conducts comprehensive analysis of algorithms along the complete MCTS path:
2429
2430     ## Task Objective
2431     Synthesize technical insights from algorithm evolution MCTS
2432     path to guide targeted improvements. Current Analysis Level:
2433     depth/max_depth={depth}/{max_depth} (1-2: structural, 3-4:
2434     implementation, 5+: mathematical)
2435
2436     ## Depth-Specific Focus
2437     - Shallow (Depth 1-2): Structural patterns & control flow

```

```

2430     - Medium (Depth 3-4): Implementation techniques &
2431     parameterization
2432     - Deep (Depth 5+): Mathematical formulations & computational
2433     primitives
2434
2435     ## Input Context
2436     Analyzing {num_algorithms} algorithm reflections from MCTS
2437     exploration trajectories. Technical reflections follow depth
2438     -specific analysis requirements. Structural format: [No.N
2439     algorithm's reflection (depth: X)]<description><objective><
2440     code>
2441     {parent_info}
2442
2443     ## Synthesis Process
2444     1. Cross-Level Insight Integration:
2445         - Identify key recurring technical strategies across
2446         abstraction levels
2447         - Note level-specific constraints affecting current
2448         implementations
2449
2450     2. Domain Compliance Verification:
2451         - Validate approaches against gravitational wave signal
2452         characteristics
2453         - Check numerical reliability across different
2454         implementation levels
2455
2456     3. Improvement Planning:
2457         - Structural: Adjust data processing pipelines
2458         - Implementation: Optimize critical parameter
2459         relationships
2460         - Mathematical: Enhance core transformation components
2461
2462     ## Technical Workflow
2463     ### 1. Multi-Level Technical Analysis
2464     Structural -> Compare module composition and interaction
2465     patterns
2466     Implementation -> Assess parameter sensitivity and
2467     adaptation logic
2468     Mathematical -> Examine transformation kernels and precision
2469     handling
2470
2471     ### 2. Level-Appropriate Optimization
2472     For current depth={depth}:
2473         - Select 2-4 improvement focus areas with technical
2474         rationale
2475         - Define implementation requirements for each focus area
2476         - Establish verification criteria with domain constraints
2477
2478     ## Output Format Rules
2479     - Return optimization strategies within SINGLE BRACE
2480     - Ensure entire response can be parseable by regex:
2481     \\\{((.*)\\}\} with DOTALL flag
2482     - Do not include markdown formatting or additional
2483     explanations
2484
2485     Phase 2: Algorithm Implementation. The operation shares the same implemen-
2486     tation phase as the reflection-based path-wise crossover, utilizing the reflection-
2487     based PWC template for consistent output formatting and algorithmic generation.
2488     Methodological Distinctions. The key differences between the reflection-based
2489     path-wise crossover and the comprehensive algorithm analysis PWC lie in their
2490     analytical strategies:
2491     • Reflection-Based PWC:
2492         - Focuses on synthesizing existing reflection insights from multiple algo-
2493         rithms

```

- 2484 – Emphasizes pattern recognition across previously analyzed algorithmic
- 2485 behaviors
- 2486 – Prioritizes extraction of generalizable technical principles over imple-
- 2487 mentation details
- 2488 – Categorizes insights by structural, implementation, and mathematical
- 2489 analysis levels
- 2490 • Comprehensive Algorithm Analysis PWC:
- 2491 – Conducts direct analysis of complete algorithm implementations
- 2492 – Examines algorithmic components across multiple depth levels simulta-
- 2493 neously
- 2494 – Integrates cross-level insights through systematic technical workflow
- 2495 – Emphasizes domain compliance verification and improvement planning

2496 **Depth-Adaptive Processing.** Both PWC operations implement depth-specific

2497 analysis strategies that adapt to the current position in the MCTS tree:

- 2498 • Shallow Depth Focus (1-2):
- 2499 – Structural patterns and component architecture analysis
- 2500 – Control flow restructuring and module composition optimization
- 2501 – Data processing pipeline adjustments
- 2502 • Medium Depth Focus (3-4):
- 2503 – Implementation techniques and algorithmic parameter optimization
- 2504 – Critical parameter relationship assessment
- 2505 – Numerical computation enhancement strategies
- 2506 • Deep Depth Focus (5+):
- 2507 – Mathematical formulation analysis and computational primitive opti-
- 2508 mization
- 2509 – Transformation kernel examination and precision handling
- 2510 – Advanced linear algebra method integration

2512 **Path Trajectory Analysis.** Both operations process algorithms along complete

2513 MCTS paths, with depth tracking that enables comprehensive evolutionary analy-

2514 sis:

- 2515 • Reflection-Based PWC Path Processing:
- 2516 – Analyzes reflection patterns from algorithms at decreasing depth levels
- 2517 – Tracks depth-specific insights through structured format annotations
- 2518 – Synthesizes cross-depth technical strategies for optimization guidance
- 2519 • Comprehensive Algorithm Analysis PWC Path Processing:
- 2520 – Examines complete algorithm implementations with performance met-
- 2521 rics
- 2522 – Integrates algorithmic descriptions, objective values, and code analysis
- 2523 – Conducts multi-level technical synthesis across the entire path trajectory

2525 **Template Variable Integration.** Both PWC operations incorporate sophisticated

2526 template variables that enable comprehensive path analysis:

- 2527 • Common Variables:
- 2528 – `depth`, `max_depth`: Depth-specific processing parameters
- 2529 – `num_algorithms`: Path length and analysis scope
- 2530 – `func_name`, `input_count`, `output_count`: Interface specifica-
- 2531 tions
- 2532 – `external_knowledge`: Domain knowledge integration
- 2533 • Reflection-Based PWC Variables:
- 2534 – `algorithm_reflections`: Multi-algorithm reflection synthesis
- 2535 – `reflection`: Generated optimization insights
- 2536 • Comprehensive Algorithm Analysis PWC Variables:
- 2537 – `parent_info`: Complete algorithm implementation details

- 2538 - current_algorithm_description,
- 2539 current_algorithm_code: Baseline algorithm specifications
- 2540 - current_objective_value: Performance reference metrics

2541 **Quality Assurance and Validation.** Both PWC operations implement compre-
2542 hensive validation procedures:

- 2543 • Reflection-Based PWC Validation:
 - 2544 - Pattern identification verification across multiple algorithm reflections
 - 2545 - Technical pathway analysis consistency checking
 - 2546 - Optimization principle formulation validation
- 2547 • Comprehensive Algorithm Analysis PWC Validation:
 - 2548 - Multi-level technical analysis coherence verification
 - 2549 - Domain compliance checking across different implementation levels
 - 2550 - Cross-level insight integration validation

2551 The dual PWC approach provides complementary strategies for capturing long-
2552 range dependencies in the MCTS tree, enabling the framework to synthesize in-
2553 sights across complete evolutionary trajectories while maintaining depth-specific
2554 optimization focus and domain knowledge integration.

2556 A.1.6 DOMAIN KNOWLEDGE INTEGRATION

2557 Domain knowledge integration serves as a critical component that ensures gen-
2558 erated algorithms remain grounded in gravitational wave detection principles
2559 while encouraging exploration beyond traditional linear processing methods. The
2560 framework incorporates specialized domain expertise through structured knowl-
2561 edge templates that guide algorithmic development toward physically meaningful
2562 and computationally efficient solutions.

2563 **External Knowledge Template Structure.** The domain knowledge integration
2564 employs a comprehensive template that emphasizes non-linear processing ap-
2565 proaches and adaptive algorithmic strategies:

```
2566 ### External Knowledge Integration
2567 1. Non-linear** Processing Core Concepts:
2568     - Signal Transformation:
2569       * Non-linear vs linear decomposition
2570       * Adaptive threshold mechanisms
2571       * Multi-scale analysis
2572     - Feature Extraction:
2573       * Phase space reconstruction
2574       * Topological data analysis
2575       * Wavelet-based detection
2576     - Statistical Analysis:
2577       * Robust estimators
2578       * Non-Gaussian processes
2579       * Higher-order statistics
2580 2. Implementation Principles:
2581     - Prioritize adaptive over fixed parameters
2582     - Consider local vs global characteristics
2583     - Balance computational cost with accuracy
```

2584 **Non-linear Processing Emphasis.** The domain knowledge framework explicitly
2585 prioritizes non-linear algorithmic approaches over traditional linear methods, re-
2586 cognizing that gravitational wave signals exhibit complex, transient characteristics
2587 that require sophisticated analysis techniques. This emphasis addresses funda-
2588 mental limitations in conventional matched filtering approaches that rely heavily
2589 on linear processing assumptions.

2590 **Signal Transformation Guidance.** The domain knowledge provides specific
2591 guidance on signal transformation strategies that leverage advanced signal pro-
cessing concepts:

- 2592
- 2593
- 2594
- 2595
- 2596
- 2597
- 2598
- 2599
- 2600
- 2601
- 2602
- 2603
- 2604
- 2605
- 2606
- 2607
- 2608
- 2609
- 2610
- 2611
- 2612
- 2613
- 2614
- 2615
- 2616
- 2617
- 2618
- 2619
- 2620
- 2621
- 2622
- 2623
- 2624
- 2625
- 2626
- 2627
- 2628
- 2629
- 2630
- 2631
- 2632
- 2633
- 2634
- 2635
- 2636
- 2637
- 2638
- 2639
- 2640
- 2641
- 2642
- 2643
- 2644
- 2645
- Non-linear vs Linear Decomposition: The framework encourages exploration of non-linear decomposition methods that can capture complex signal morphologies beyond the capabilities of traditional Fourier-based approaches. This includes techniques such as empirical mode decomposition, intrinsic mode functions, and adaptive basis construction.
 - Adaptive Threshold Mechanisms: Rather than employing fixed threshold values, the domain knowledge promotes adaptive thresholding strategies that respond to local signal characteristics and noise conditions. This approach enables more robust detection performance across diverse observational scenarios.
 - Multi-scale Analysis: The framework emphasizes multi-scale signal analysis techniques that can simultaneously capture both short-duration transient signals and longer-duration continuous wave sources. This includes wavelet-based methods, time-frequency analysis, and hierarchical decomposition strategies.
 - Feature Extraction Methodologies. The domain knowledge incorporates advanced feature extraction approaches that extend beyond traditional signal processing paradigms:
 - Phase Space Reconstruction: The framework encourages exploration of phase space reconstruction techniques that can reveal hidden dynamical structures in gravitational wave data. This includes embedding dimension analysis, recurrence analysis, and attractor reconstruction methods.
 - Topological Data Analysis: The domain knowledge promotes topological data analysis approaches that can identify persistent features and structural patterns in high-dimensional gravitational wave data. This includes persistent homology, Mapper algorithms, and topological feature extraction.
 - Wavelet-based Detection: The framework emphasizes wavelet-based detection strategies that can provide optimal time-frequency resolution for transient signal analysis. This includes continuous wavelet transforms, discrete wavelet decomposition, and wavelet packet analysis.
 - Statistical Analysis Enhancement. The domain knowledge integrates sophisticated statistical analysis techniques that account for the complex noise characteristics of gravitational wave detectors:
 - Robust Estimators: The framework promotes robust statistical estimators that can maintain performance in the presence of outliers and non-Gaussian noise distributions. This includes median-based estimators, M-estimators, and trimmed mean approaches.
 - Non-Gaussian Processes: The domain knowledge emphasizes analysis techniques that can handle non-Gaussian noise processes commonly encountered in gravitational wave data. This includes heavy-tailed distributions, skewed probability models, and non-stationary noise characterization.
 - Higher-order Statistics: The framework encourages exploration of higher-order statistical moments and cumulants that can capture subtle signal characteristics beyond second-order analysis. This includes bispectrum analysis, higher-order moment estimation, and polyspectral techniques.
 - Implementation Principles and Constraints. The domain knowledge provides specific implementation principles that guide algorithmic development toward practical and efficient solutions:
 - Adaptive Parameter Prioritization: The framework emphasizes adaptive parameter selection over fixed parameter values, enabling algorithms to respond dynamically to changing signal and noise conditions. This principle encourages exploration of learning-based parameter adjustment, feedback control mechanisms, and online adaptation strategies.
 - Local vs Global Characteristics: The domain knowledge promotes consideration of both local signal characteristics and global data patterns,

2646 enabling algorithms to balance fine-grained analysis with comprehensive
2647 signal understanding. This includes local stationarity analysis, global
2648 trend estimation, and multi-resolution processing approaches.

- 2649 – **Computational Cost-Accuracy Balance:** The framework provides guid-
2650 ance on balancing computational efficiency with detection accuracy, en-
2651 suring that generated algorithms remain practical for real-time imple-
2652 mentation while maintaining scientific rigor. This includes complexity
2653 analysis, algorithmic optimization, and performance benchmarking con-
2654 siderations.

2655 **Integration Across Evolutionary Operations.** The domain knowledge template
2656 is systematically integrated across all evolutionary operations (PC, SC, PM, PWC)
2657 through the `external_knowledge` template variable. This ensures consistent appli-
2658 cation of gravitational wave detection principles regardless of the specific evolu-
2659 tionary strategy employed.

2660 **Physical Validity Assurance.** The domain knowledge template ensures that all
2661 generated algorithms respect fundamental physical constraints related to gravita-
2662 tional wave signal characteristics, detector limitations, and noise properties.

2663 **Computational Feasibility.** The implementation principles guide algorithmic de-
2664 velopment toward computationally feasible solutions that can be practically im-
2665 plemented within the constraints of current computational resources and real-time
2666 processing requirements.

2667 This comprehensive domain knowledge integration creates a robust framework
2668 for scientifically grounded algorithmic discovery, ensuring that the evolutionary
2669 process generates algorithms that are both innovative and practically applicable to
2670 gravitational wave detection challenges.

2671 A.1.7 ERROR HANDLING AND ITERATIVE REFINEMENT

2673 The Evo-MCTS framework incorporates a robust error handling mechanism that
2674 enables iterative refinement of generated algorithms through automated debug-
2675 ging and correction processes. When generated code encounters execution errors,
2676 fails to detect signals, or exceeds computational time limits, the system employs a
2677 rechat strategy using advanced reasoning models to diagnose and resolve issues.

2678 **Error Detection and Classification.** The framework monitors three primary fail-
2679 ure modes during algorithm execution: (i) runtime exceptions and syntax errors
2680 that prevent code execution, (ii) algorithmic failures where no gravitational wave
2681 signals are detected despite their presence in the data, and (iii) computational time-
2682 out scenarios where algorithms exceed predefined execution limits. Each failure
2683 mode triggers specific diagnostic protocols tailored to the underlying issue type.

2684 **Iterative Refinement Protocol.** Upon error detection, the system implements a
2685 structured refinement process through a carefully designed prompt content struc-
2686 ture. The system constructs conversation messages in a specific format to facilitate
2687 effective error correction:

2688 Initially, when no system content is provided, the framework creates a
2689 message list containing a single user role entry with the original prompt
2690 content (`prompt_content`): `messages = [{"role": "user",`
`"content": prompt_content}]`

2691 When a rechat response is available (indicating a previous failed attempt),
2692 the system extends the conversation by first appending the assistant's pre-
2693 vious response (`rechat_response`): `messages.append({"role":`
2694 `"assistant", "content": rechat_response})`

2695 Subsequently, the system adds a new user message that explic-
2696 itly requests debugging and issue resolution (`prompt_content`):
2697 `messages.append({"role": "user", "content": "Your`
2698 `previous code had execution errors, couldn't find`
2699 `signals, or timed out. Please debug and fix the`
`issues:\n\n" + prompt_content})`

This structured approach maintains the conversational context while providing clear guidance for error correction, ensuring that the assistant understands both the original requirements and the specific issues that need to be addressed.

Automated Debugging Integration. The error handling system leverages advanced reasoning capabilities to analyze failed algorithms and propose targeted corrections. This approach maintains the evolutionary optimization trajectory while addressing immediate technical obstacles that could otherwise terminate the search process. The iterative refinement ensures that promising algorithmic concepts are not discarded due to implementation errors, instead receiving corrective guidance to achieve functional implementations.

A.1.8 POST-GENERATION ANALYSIS AND KNOWLEDGE EXTRACTION

The post-generation analysis phase extracts interpretable insights from evolved algorithms through automated knowledge distillation. This process transforms the raw algorithmic implementations into concise, human-readable descriptions that capture the essential design principles and operational characteristics of discovered solutions.

Algorithm Description Generation. The framework employs a structured prompt template to generate concise algorithm descriptions that highlight critical design decisions and implementation strategies. The prompt construction follows a systematic format:

```
Following is the Design Idea of a heuristic algorithm for
the problem and the code with function name 'pipeline_v2'
for implementing the heuristic algorithm.
```

```
{prompt_inout_inf} {prompt_other_inf}
```

```
Design Idea:
```

```
{algorithm}
```

```
Code:
```

```
```python
```

```
{code}
```

```
```
```

```
The content of the Design Idea idea cannot fully represent
what the algorithm has done informative. So, now you should
re-describe the algorithm using less than 3 sentences.
```

```
Hint: You should reference the given Design Idea and
highlight the most critical design ideas of the code. You
can analyse the code to describe which variables are given
higher priorities and which variables are given lower
priorities, the parameters and the structure of the code.
```

This template systematically combines the original design concept with the implemented code, requesting a refined description that captures the algorithm's core operational principles. The analysis focuses on parameter prioritization, structural characteristics, and critical design decisions that distinguish the evolved solution.

Knowledge Extraction Protocol. The post-generation analysis captures key design principles and compresses algorithmic representations into human-readable summaries. This reflection process identifies algorithmic innovations, signal processing techniques, and computational characteristics while reducing token consumption to prevent context window overflow in subsequent LLM interactions.

Interpretability Enhancement. The generated descriptions provide concise algorithmic summaries that enable efficient reference to previous discoveries without overwhelming the LLM context, facilitating continued exploration while maintaining algorithmic memory across generations.

A.1.9 CODE EXAMPLES AND CASE STUDIES

This section presents a detailed examination of the highest-performing algorithm discovered during the Evo-MCTS optimization process, corresponding to node 486 (as shown in Figure 5a) which achieved the maximum fitness score of

2754 5,241.37 units. The algorithm demonstrates sophisticated multi-stage signal processing techniques that emerged through evolutionary optimization.

2755
2756 **Algorithm Overview.** The evolved algorithm implements a four-stage pipeline combining robust baseline detrending, adaptive whitening with enhanced power spectral density (PSD) smoothing, coherent time-frequency analysis with frequency-conditioned regularization, and multi-resolution thresholding with octave-spaced dyadic wavelet validation. This architecture represents a novel synthesis of classical signal processing techniques with adaptive parameter selection mechanisms.

2763 **Stage 1: Robust Baseline Detrending.** The algorithm initiates with median filtering-based detrending to remove long-term instrumental drifts and environmental variations. The median filter kernel size of 101 samples provides robust trend removal while preserving transient gravitational wave signatures. This pre-processing stage establishes a stable baseline for subsequent whitening operations.

2768 **Stage 2: Adaptive Whitening with Enhanced PSD Smoothing.** The core innovation lies in the adaptive whitening mechanism that dynamically adjusts window parameters based on data characteristics. The algorithm implements Tukey windowing with 75% overlap and adaptive segment lengths constrained between 5-30 seconds, optimizing spectral estimation for varying noise conditions. The PSD smoothing employs exponential filtering with stationarity-dependent coefficients (0.75-0.85 range), while Tikhonov regularization provides frequency-dependent gain control. Savitzky-Golay filtering generates causal-like gradients, and sigmoid-based nonlinear scaling enhances spectral features through adaptive gain factors.

2777 **Stage 3: Coherent Time-Frequency Analysis.** The algorithm computes complex spectrograms preserving phase information across both detectors, enabling coherent analysis of gravitational wave signatures. Phase difference calculations and coherence estimation provide cross-detector validation, while frequency-conditioned regularization balances phase alignment with noise characteristics. The integration of axial curvature estimates through second derivatives and nonlinear activation functions (tanh-based boost) enhances signal discrimination capabilities.

2785 **Stage 4: Multi-Resolution Validation.** The final stage implements sophisticated peak detection using robust statistical measures (median absolute deviation) combined with octave-spaced dyadic wavelet validation. Continuous wavelet transform coefficients across scales 1-8 provide multi-resolution signal verification, while Gaussian-weighted uncertainty estimation quantifies detection confidence intervals.

```

2791 1 import numpy as np
2792 2 import scipy.signal as signal
2793 3 from scipy.signal.windows import tukey
2794 4 from scipy.signal import savgol_filter
2795 5
2796 6 def pipeline_v2(strain_h1: np.ndarray, strain_l1: np.ndarray, times: np.ndarray) ->
↳ tuple[np.ndarray, np.ndarray, np.ndarray]:
2797 7     """
2798 8     The pipeline function processes gravitational wave data from the H1 and L1 detectors to identify
↳ potential gravitational wave signals.
2799 9     It takes strain_h1 and strain_l1 numpy arrays containing detector data, and times array with
↳ corresponding time points.
2800 10    The function returns a tuple of three numpy arrays: peak_times containing GPS times of
↳ identified events,
2801 11    peak_heights with significance values of each peak, and peak_deltat showing time window
↳ uncertainty for each peak.
2802 12    """
2803 13    eps = np.finfo(float).tiny
2804 14    dt = times[1] - times[0]
2805 15    fs = 1.0 / dt
2806 16    # Base spectrogram parameters
2807 17    base_nperseg = 256
2808 18    base_noverlap = base_nperseg // 2
2809 19    medfilt_kernel = 101 # odd kernel size for robust detrending
2810 20    uncertainty_window = 5 # half-window for local timing uncertainty

```

```

2808
2809 21
2810 22 # ----- Stage 1: Robust Baseline Detrending -----
2811 23 # Remove long-term trends using a median filter for each channel.
2812 24 detrended_h1 = strain_h1 - signal.medfilt(strain_h1, kernel_size=medfilt_kernel)
2813 25 detrended_l1 = strain_l1 - signal.medfilt(strain_l1, kernel_size=medfilt_kernel)
2814 26
2815 27 # ----- Stage 2: Adaptive Whitening with Enhanced PSD Smoothing
2816 28 ↵ -----
2817 29 def adaptive_whitening(strain: np.ndarray) -> np.ndarray:
2818 30 # Center the signal.
2819 31 centered = strain - np.mean(strain)
2820 32 n_samples = len(centered)
2821 33 # Adaptive window length: between 5 and 30 seconds
2822 34 win_length_sec = np.clip(n_samples / fs / 20, 5, 30)
2823 35 nperseg_adapt = int(win_length_sec * fs)
2824 36 nperseg_adapt = max(10, min(nperseg_adapt, n_samples))
2825 37
2826 38 # Create a Tukey window with 75% overlap.
2827 39 tukey_alpha = 0.25
2828 40 win = tukey(nperseg_adapt, alpha=tukey_alpha)
2829 41 noverlap_adapt = int(nperseg_adapt * 0.75)
2830 42 if noverlap_adapt >= nperseg_adapt:
2831 43     noverlap_adapt = nperseg_adapt - 1
2832 44
2833 45 # Estimate the power spectral density (PSD) using Welch's method.
2834 46 freqs, psd = signal.welch(centered, fs=fs, nperseg=nperseg_adapt,
2835 47                          noverlap=noverlap_adapt, window=win, detrend='constant')
2836 48 psd = np.maximum(psd, eps)
2837 49
2838 50 # Compute relative differences for PSD stationarity measure.
2839 51 diff_arr = np.abs(np.diff(psd)) / (psd[:-1] + eps)
2840 52 # Smooth the derivative with a moving average.
2841 53 if len(diff_arr) >= 3:
2842 54     smooth_diff = np.convolve(diff_arr, np.ones(3)/3, mode='same')
2843 55 else:
2844 56     smooth_diff = diff_arr
2845 57
2846 58 # Exponential smoothing (Kalman-like) with adaptive alpha using PSD stationarity.
2847 59 smoothed_psd = np.copy(psd)
2848 60 for i in range(1, len(psd)):
2849 61     # Adaptive smoothing coefficient: base 0.8 modified by local stationarity (±0.05)
2850 62     local_alpha = np.clip(0.8 - 0.05 * smooth_diff[min(i-1, len(smooth_diff)-1)], 0.75,
2851 63 ↵ 0.85)
2852 64     smoothed_psd[i] = local_alpha * smoothed_psd[i-1] + (1 - local_alpha) * psd[i]
2853 65
2854 66 # Compute Tikhonov regularization gain based on deviation from median PSD.
2855 67 noise_baseline = np.median(smoothed_psd)
2856 68 raw_gain = (smoothed_psd / (noise_baseline + eps)) - 1.0
2857 69
2858 70 # Compute a causal-like gradient using the Savitzky-Golay filter.
2859 71 win_len = 11 if len(smoothed_psd) >= 11 else ((len(smoothed_psd)//2)*2+1)
2860 72 polyorder = 2 if win_len > 2 else 1
2861 73 delta_freq = np.mean(np.diff(freqs))
2862 74 grad_psd = savgol_filter(smoothed_psd, win_len, polyorder, deriv=1, delta=delta_freq,
2863 75 ↵ mode='interp')
2864 76
2865 77 # Nonlinear scaling via sigmoid to enhance gradient differences.
2866 78 sigmoid = lambda x: 1.0 / (1.0 + np.exp(-x))
2867 79 scaling_factor = 1.0 + 2.0 * sigmoid(np.abs(grad_psd) / (np.median(smoothed_psd) + eps))
2868 80
2869 81 # Compute adaptive gain factors with nonlinear scaling.
2870 82 gain = 1.0 - np.exp(-0.5 * scaling_factor * raw_gain)
2871 83 gain = np.clip(gain, -8.0, 8.0)
2872 84
2873 85 # FFT-based whitening: interpolate gain and PSD onto FFT frequency bins.
2874 86 signal_fft = np.fft.rfft(centered)
2875 87 freq_bins = np.fft.rfftfreq(n_samples, d=dt)
2876 88 interp_gain = np.interp(freq_bins, freqs, gain, left=gain[0], right=gain[-1])
2877 89 interp_psd = np.interp(freq_bins, freqs, smoothed_psd, left=smoothed_psd[0],
2878 90 ↵ right=smoothed_psd[-1])
2879 91 denom = np.sqrt(interp_psd) * (np.abs(interp_gain) + eps)
2880 92 denom = np.maximum(denom, eps)
2881 93 white_fft = signal_fft / denom
2882 94 whitened = np.fft.irfft(white_fft, n=n_samples)
2883 95 return whitened
2884 96
2885 97 # Whiten H1 and L1 channels using the adapted method.

```

```

2862     94     white_h1 = adaptive_whitening(detrended_h1)
2863     95     white_l1 = adaptive_whitening(detrended_l1)
2864     96
2865     97     # ----- Stage 3: Coherent Time-Frequency Metric with Frequency-Conditioned
2866     98     ↪ Regularization -----
2867     99     def compute_coherent_metric(w1: np.ndarray, w2: np.ndarray) -> tuple[np.ndarray, np.ndarray]:
2868     100         # Compute complex spectrograms preserving phase information.
2869     101         f1, t_spec, Sxx1 = signal.spectrogram(w1, fs=fs, nperseg=base_nperseg,
2870     102             noverlap=base_noverlap, mode='complex', detrend=False)
2871     103         f2, t_spec2, Sxx2 = signal.spectrogram(w2, fs=fs, nperseg=base_nperseg,
2872             noverlap=base_noverlap, mode='complex',
2873             ↪ detrend=False)
2874     104
2875     105         # Ensure common time axis length.
2876     106         common_len = min(len(t_spec), len(t_spec2))
2877     107         t_spec = t_spec[:common_len]
2878     108         Sxx1 = Sxx1[:, :common_len]
2879     109         Sxx2 = Sxx2[:, :common_len]
2880     110
2881     111         # Compute phase differences and coherence between detectors.
2882     112         phase_diff = np.angle(Sxx1) - np.angle(Sxx2)
2883     113         phase_coherence = np.abs(np.cos(phase_diff))
2884     114
2885     115         # Estimate median PSD per frequency bin from the spectrograms.
2886     116         psd1 = np.median(np.abs(Sxx1)**2, axis=1)
2887     117         psd2 = np.median(np.abs(Sxx2)**2, axis=1)
2888     118
2889     119         # Frequency-conditioned regularization gain (reflection-guided).
2890     120         lambda_f = 0.5 * ((np.median(psd1) / (psd1 + eps)) + (np.median(psd2) / (psd2 + eps)))
2891     121         lambda_f = np.clip(lambda_f, 1e-4, 1e-2)
2892     122         # Regularization denominator integrating detector PSDs and lambda.
2893     123         reg_denom = (psd1[:, None] + psd2[:, None] + lambda_f[:, None] + eps)
2894     124
2895     125         # Weighted phase coherence that balances phase alignment with noise levels.
2896     126         weighted_comp = phase_coherence / reg_denom
2897     127
2898     128         # Compute axial (frequency) second derivatives as curvature estimates.
2899     129         d2_coh = np.gradient(np.gradient(phase_coherence, axis=0), axis=0)
2900     130         avg_curvature = np.mean(np.abs(d2_coh), axis=0)
2901     131
2902     132         # Nonlinear activation boost using tanh for regions of high curvature.
2903     133         nonlinear_boost = np.tanh(5 * avg_curvature)
2904     134         linear_boost = 1.0 + 0.1 * avg_curvature
2905     135
2906     136         # Cross-detector synergy: weight derived from global median consistency.
2907     137         novel_weight = np.mean((np.median(psd1) + np.median(psd2)) / (psd1[:, None] + psd2[:, None]
2908     138         ↪ + eps), axis=0)
2909     139
2910     140         # Integrated time-frequency metric combining all enhancements.
2911     141         tf_metric = np.sum(weighted_comp * linear_boost * (1.0 + nonlinear_boost), axis=0) *
2912     142         ↪ novel_weight
2913     143
2914     144         # Adjust the spectrogram time axis to account for window delay.
2915     145         metric_times = t_spec + times[0] + (base_nperseg / 2) / fs
2916     146         return tf_metric, metric_times
2917
2918     tf_metric, metric_times = compute_coherent_metric(white_h1, white_l1)
2919
2920     # ----- Stage 4: Multi-Resolution Thresholding with Octave-Spaced Dyadic Wavelet
2921     ↪ Validation -----
2922     148     def multi_resolution_thresholding(metric: np.ndarray, times_arr: np.ndarray) ->
2923     149     ↪ tuple[np.ndarray, np.ndarray, np.ndarray]:
2924     150         # Robust background estimation with median and MAD.
2925     151         bg_level = np.median(metric)
2926     152         mad_val = np.median(np.abs(metric - bg_level))
2927     153         robust_std = 1.4826 * mad_val
2928     154         threshold = bg_level + 1.5 * robust_std
2929     155
2930     156         # Identify candidate peaks using prominence and minimum distance criteria.
2931     157         peaks, _ = signal.find_peaks(metric, height=threshold, distance=2, prominence=0.8 *
2932     158         ↪ robust_std)
2933     159         if peaks.size == 0:
2934     160             return np.array([]), np.array([]), np.array([])
2935     161
2936     162         # Local uncertainty estimation using a Gaussian-weighted convolution.
2937     163         win_range = np.arange(-uncertainty_window, uncertainty_window + 1)
2938         sigma = uncertainty_window / 2.5
2939         gauss_kernel = np.exp(-0.5 * (win_range / sigma) ** 2)

```

```

2916     164         gauss_kernel /= np.sum(gauss_kernel)
2917     165         weighted_mean = np.convolve(metric, gauss_kernel, mode='same')
2918     166         weighted_sq = np.convolve(metric ** 2, gauss_kernel, mode='same')
2919     167         variances = np.maximum(weighted_sq - weighted_mean ** 2, 0.0)
2920     168         uncertainties = np.sqrt(variances)
2921     169         uncertainties = np.maximum(uncertainties, 0.01)
2922     170
2923     171         valid_times = []
2924     172         valid_heights = []
2925     173         valid_uncerts = []
2926     174         n_metric = len(metric)
2927     175
2928     176         # Compute a simple second derivative for local curvature checking.
2929     177         if n_metric > 2:
2930     178             second_deriv = np.diff(metric, n=2)
2931     179             second_deriv = np.pad(second_deriv, (1, 1), mode='edge')
2932     180         else:
2933     181             second_deriv = np.zeros_like(metric)
2934     182
2935     183         # Use octave-spaced scales (dyadic wavelet validation) to validate peak significance.
2936     184         widths = np.arange(1, 9) # approximate scales 1 to 8
2937     185         for peak in peaks:
2938     186             # Skip peaks lacking sufficient negative curvature.
2939     187             if second_deriv[peak] > -0.1 * robust_std:
2940     188                 continue
2941     189             local_start = max(0, peak - uncertainty_window)
2942     190             local_end = min(n_metric, peak + uncertainty_window + 1)
2943     191             local_segment = metric[local_start:local_end]
2944     192             if len(local_segment) < 3:
2945     193                 continue
2946     194             try:
2947     195                 cwt_coeff = signal.cwt(local_segment, signal.ricker, widths)
2948     196             except Exception:
2949     197                 continue
2950     198             max_coeff = np.max(np.abs(cwt_coeff))
2951     199             # Threshold for validating the candidate using local MAD.
2952     200             cwt_thresh = mad_val * np.sqrt(2 * np.log(len(local_segment)) + eps)
2953     201             if max_coeff >= cwt_thresh:
2954     202                 valid_times.append(times_arr[peak])
2955     203                 valid_heights.append(metric[peak])
2956     204                 valid_uncerts.append(uncertainties[peak])
2957     205
2958     206         if len(valid_times) == 0:
2959     207             return np.array([]), np.array([]), np.array([])
2960     208         return np.array(valid_times), np.array(valid_heights), np.array(valid_uncerts)
2961     209
2962     210     peak_times, peak_heights, peak_deltat = multi_resolution_thresholding(tf_metric, metric_times)
2963     211     return peak_times, peak_heights, peak_deltat

```

Reflective Reasoning Integration. The framework incorporates advanced reflective reasoning capabilities through DeepSeek-R1’s internal thinking processes, enabling systematic analysis and refinement of algorithmic solutions. This reflective mechanism operates through structured thinking templates that guide the model through comprehensive evaluation and improvement cycles.

Thinking Process Template. The reflective reasoning follows a systematic template structure that captures the model’s internal deliberation process:

<thinking>

Okay, let's tackle this problem. The user wants me to analyze the differences between the original and improved versions of the gravitational wave detection pipeline, focusing on depth 4. The current MCTS depth is 4 out of 10, which falls into the medium category. The medium depth focuses on implementation techniques and parameterization.

First, I need to compare the two code versions. Let's look at each stage of the pipeline.

Starting with Stage 2: Adaptive Whitening. In the original code, there's a Kalman-like smoothing with a fixed alpha of 0.8. The improved version introduces an adaptive alpha based on the PSD's stationarity. They calculate the derivative of

2970 the PSD, smooth it, and adjust alpha dynamically. This
2971 seems like an implementation technique improvement, fitting
2972 the medium depth.
2973
2974 Next, the gain calculation. The original uses a fixed
2975 scaling factor of 3.0, while the improved version applies a
2976 Savitzky-Golay filter to compute the gradient of the PSD and
2977 uses a sigmoid function for nonlinear scaling. This
2978 introduces more sophisticated mathematical processing, but
2979 since it's part of the gain factor computation, it's about
2980 parameterization and implementation details, which fits
2981 medium depth.
2982
2983 In Stage 3: Coherent Metric. The original uses Tikhonov
2984 regularization with a fixed lambda. The improved version
2985 computes a frequency-conditioned lambda based on the median
2986 PSD ratios. This dynamic regularization is a
2987 parameterization improvement. Also, they add a nonlinear
2988 boost using tanh on curvature, which is another parameter
2989 adjustment. These are implementation techniques.
2990
2991 Stage 4: Thresholding. The original checks the second
2992 derivative for curvature. The improved code does a similar
2993 check but also uses octave-spaced wavelet scales.
2994 Additionally, the CWT threshold is based on the local MAD
2995 instead of a fixed value. This is a parameterization change,
2996 adjusting how validation is done, which is medium depth.
2997
2998 Looking for core patterns: The improvements involve making
2999 parameters adaptive based on data characteristics,
3000 introducing more sophisticated smoothing techniques, and
3001 enhancing validation steps with dynamic thresholds. These
3002 are all implementation techniques and parameter adjustments,
3003 fitting the medium depth focus.
3004
3005 Now, generating optimization principles. Each principle
3006 should address a specific limitation in the original code,
3007 using concrete parameters from the improved version. For
3008 example, replacing fixed smoothing alpha with adaptive alpha
3009 based on PSD derivatives. Another is using Savitzky-Golay
3010 for gradient computation. Then, dynamic regularization
3011 lambda, nonlinear activation with tanh, and CWT validation
3012 with local MAD.
3013
3014 I need to ensure each principle is tied to the medium depth,
3015 focusing on implementation techniques and parameters. Avoid
3016 diving into deep mathematical formulations unless necessary
3017 . The principles should be transferable, so other algorithms
3018 can use similar strategies.
3019
3020 Checking the improved code's parameters: adaptive alpha
3021 ranges from 0.75 to 0.85, Savitzky-Golay window length 11
3022 and polyorder 2, frequency-conditioned lambda clipped
3023 between 1e-4 and 1e-2, tanh scaling with factor 5, and CWT
threshold using local MAD. These concrete values should be
included in the principles.

3018 Finally, structure the answer into numbered items within a
3019 single brace, adhering to the output format. Ensure each
3020 item explains the original's limitation, the improved
3021 approach, and the parameters involved.
3022 </thinking>
3023

3024 **Reflection Content Generation.** Following the thinking process, the system gen-
 3025 erates structured reflection content that synthesizes insights and proposes concrete
 3026 improvements:

- 3027
- 3028 1. **Adaptive PSD Smoothing with Stationarity-Dependent**
 3029 **Alpha**:
 3030 Original used fixed exponential smoothing ($\alpha=0.8$).
 3031 Improved version implements Kalman-like adaptation where
 3032 smoothing coefficient (0.75-0.85 range) dynamically adjusts
 3033 based on local PSD stationarity measured through normalized
 3034 differentials (3-point moving average of $|dPSD|/PSD$).
 3035 Enhances noise floor estimation stability during non-
 stationary segments while preserving transient features.
 - 3036 2. **Spectral Gradient-Enhanced Gain Computation**:
 3037 Replaced simple exponential gain dampening with Savitzky-
 3038 Golay derivative estimation (window=11, polyorder=2)
 3039 followed by sigmoid nonlinearity ($1/(1+e^{-x/scale})$).
 3040 Original used fixed scaling factor of 3.0, improved version
 3041 computes gradient-based adaptive scaling that responds to
 3042 spectral curvature changes, improving signal-to-noise
 discrimination in frequency domain.
 - 3043 3. **Frequency-Conditioned Regularization Lambda**:
 3044 Original Tikhonov regularization used fixed lambda parameter
 3045 . Improved version computes frequency-dependent lambda based
 3046 on median PSD ratios across frequency bins, clipped between
 3047 $1e-4$ and $1e-2$. This dynamic regularization adapts to local
 3048 spectral characteristics, preventing over-smoothing in high-
 3049 SNR regions while maintaining stability in noise-dominated
 3050 frequencies.
 - 3051 4. **Nonlinear Coherence Boost with Curvature Activation**:
 3052 Enhanced coherent metric computation by adding tanh-based
 3053 nonlinear activation (scale factor=5) applied to second-
 3054 derivative curvature measures. This amplifies genuine
 3055 gravitational wave signatures while suppressing noise
 3056 artifacts, improving detection sensitivity through selective
 nonlinear enhancement of coherent structures.
 - 3057 5. **CWT Validation with Local MAD Thresholding**:
 3058 Original used fixed threshold for continuous wavelet
 3059 transform validation. Improved version implements octave-
 3060 spaced wavelet scales with local Median Absolute Deviation (
 3061 MAD) based thresholding instead of global constants. This
 3062 adaptive validation better handles varying noise
 3063 characteristics across different time-frequency regions,
 3064 reducing false positives while maintaining detection
 3065 efficiency.

3066 **Iterative Refinement Protocol.** The reflective reasoning process operates iter-
 3067 atively, with each cycle building upon previous insights to achieve progressive
 3068 algorithmic improvement. This systematic approach ensures that evolved algo-
 3069 rithms benefit from comprehensive analysis and targeted optimization rather than
 3070 random exploration alone.

3071 A.2 DATA PARTITIONING STRATEGY

3072
 3073 The MLGWSC-1 Dataset 4 partitioning strategy balances optimization efficiency
 3074 with statistical validity through careful temporal segmentation. Figure 7 illustrates
 3075 the systematic approach employed to divide the dataset into training and test sub-
 3076 sets while maintaining representative signal characteristics across both partitions.

3077 **Training Set Definition and Statistical Characteristics.** The training set encom-
 passes the first 39 injection indices, corresponding to a cumulative duration of 7

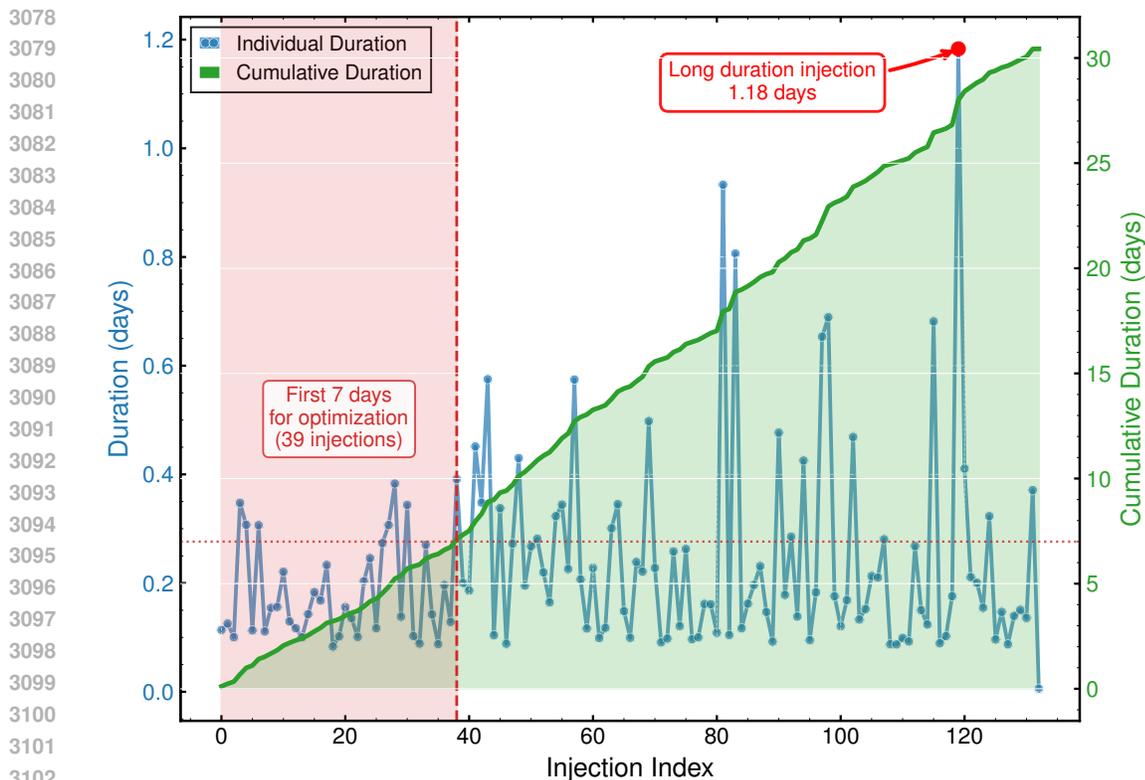


Figure 7: **MLGWSC-1 Dataset 4 Partitioning Strategy for Training and Test Set Configuration.** Individual injection durations (blue line, left axis) and cumulative duration (green line, right axis) across injection indices 0-130. The red dashed vertical line at injection index 39 delineates the training set boundary, with the first 7 days (red shaded region) used for algorithm optimization containing 39 injections. The test set consists of a single long-duration injection of 1.18 days (red annotation with arrow) occurring at injection index 119, providing a challenging validation scenario for sustained detection capability. The horizontal red dotted line indicates the cumulative duration at injection index 39, marking the exact 7-day threshold for training set partitioning. This temporal partitioning ensures efficient optimization while providing rigorous out-of-sample validation on extended-duration signals.

days. This temporal boundary provides sufficient signal diversity and noise conditions for algorithmic optimization while maintaining computational efficiency during the iterative Evo-MCTS process.

During optimization, each evolved algorithm processes complete injection segments with durations ranging from 0.1 to 0.4 days. The statistical characteristics of the training set (mean: 0.179 days, median: 0.143 days) ensure comprehensive algorithmic development across the spectrum of injection durations present in the dataset. This distribution provides robust training exposure while the 7-day cumulative training duration serves multiple critical purposes: (i) adequate statistical power for AUC calculation with minimum false-alarm rate of 4 events per month, (ii) rapid algorithm evaluation (10-20 minutes per assessment), and (iii) preservation of temporal continuity and realistic noise characteristics.

Test Set Configuration and Validation Rigor. The test set comprises a single 1.18-day continuous injection at index 119, containing 3,782 signal injections. This extended duration provides a particularly challenging validation scenario that significantly exceeds both the training set mean (0.179 days) and the overall dataset mean (0.229 days) by factors of 6.6x and 5.2x, respectively. The test injection's duration of 1.18 days represents an extreme validation case that tests

algorithmic robustness against sustained detection requirements over prolonged periods, examining performance stability under temporal variations in detector sensitivity and environmental conditions.

The temporal separation from training data ensures genuine out-of-sample validation, while the extended duration creates a stringent assessment environment. Compared to the overall dataset statistics (mean: 0.229 days, median: 0.169 days), the test injection’s 1.18-day duration provides validation on challenging extended-duration scenarios that algorithms must handle effectively.

A.3 FIVE-RUN EXPERIMENTAL DESIGN AND RESULTS

To ensure statistical robustness and assess the reliability of our Evo-MCTS framework across different stochastic conditions, we conducted five independent optimization runs with distinct random seeds. Figure 8 presents comprehensive results from all runs, demonstrating both the consistency of our approach and the natural variation inherent in stochastic optimization processes.

(Note: To maintain consistency with the main text’s PT1-PT4 framework while providing detailed combined individual run analysis, this supplementary analysis presents five phase transitions labeled as PT1, PT2.1, PT2.2, PT3, and PT4, where PT2.1 and PT2.2 represent two independent algorithmic breakthroughs that are distinct from the PT2 phase described in the main text.)

Primary Run Analysis and Phase Transition Characterization. The primary run (top panel) achieved the most comprehensive optimization trajectory, discovering five distinct phase transitions that represent qualitative algorithmic breakthroughs. PT1 occurs at evaluation 69 with fitness 1,635.00 at depth 5, incorporating Continuous Wavelet Transform (CWT) techniques and Multi-resolution Thresholding for enhanced time-frequency analysis. PT2.1 emerges at evaluation 151 with fitness 2,612.77 at depth 10, introducing Curvature Boosting methods while integrating Tikhonov Regularization for improved signal conditioning and noise suppression. PT2.2 manifests at evaluation 211 with fitness 3,439.75 at depth 3, representing a significant algorithmic advancement through refined optimization strategies. PT3 develops at evaluation 333 with fitness 4,559.26 at depth 10, integrating Savitzky-Golay (S-G) filter techniques for enhanced signal processing capabilities. Finally, PT4 achieves maximum fitness of 5,241.37 at evaluation 486 and depth 4, representing the culmination of all previously discovered techniques including CWT, Multi-resolution Thresholding, Curvature Boosting, Tikhonov Regularization, and S-G filtering in a comprehensive algorithmic framework.

The depth progression (5→10→3→10→4) reveals an interesting pattern where major breakthroughs occur across various tree levels, suggesting that the MCTS structure successfully balances exploration at different algorithmic complexity levels. The fitness improvements at each phase transition represent single-step gains from the immediate predecessor node rather than cumulative improvements between phases: PT1 (+639.69), PT2.1 (+370.81), PT2.2 (+910.37), PT3 (+1,045.72), and PT4 (+456.85). These single-step improvements demonstrate variable discovery magnitudes as individual algorithmic innovations are identified, with the pattern showing that breakthrough discoveries can occur with different intensities as the algorithm space becomes more thoroughly explored through the tree search process.

Table 1 provides a comprehensive performance comparison across all benchmark models, the seed function baseline, and the five phase transition levels achieved during optimization. The benchmark models demonstrate varying performance levels, with Sage achieving the highest AUC of 4359.2749, followed by Virgo-AUTH at 4101.4810. Notably, our evolved algorithms at PT Level 4 (5241.3678 AUC) significantly outperform all benchmark models across all evaluation metrics, including false alarm rates at different thresholds (FAR=1000, FAR=100, FAR=10, FAR=4.3), achieving relative improvements of 20.2% over the top-performing Sage benchmark (4359.27 AUC) and 23.4% enhancement in sensitive distance detection capability at node 486. The progressive improvement from

3186
 3187
 3188
 3189
 3190
 3191
 3192
 3193
 3194
 3195
 3196
 3197
 3198
 3199
 3200
 3201
 3202
 3203
 3204
 3205
 3206
 3207
 3208
 3209
 3210
 3211
 3212
 3213
 3214
 3215
 3216
 3217
 3218
 3219
 3220
 3221
 3222
 3223
 3224
 3225
 3226
 3227
 3228
 3229
 3230
 3231
 3232
 3233
 3234
 3235
 3236
 3237
 3238
 3239

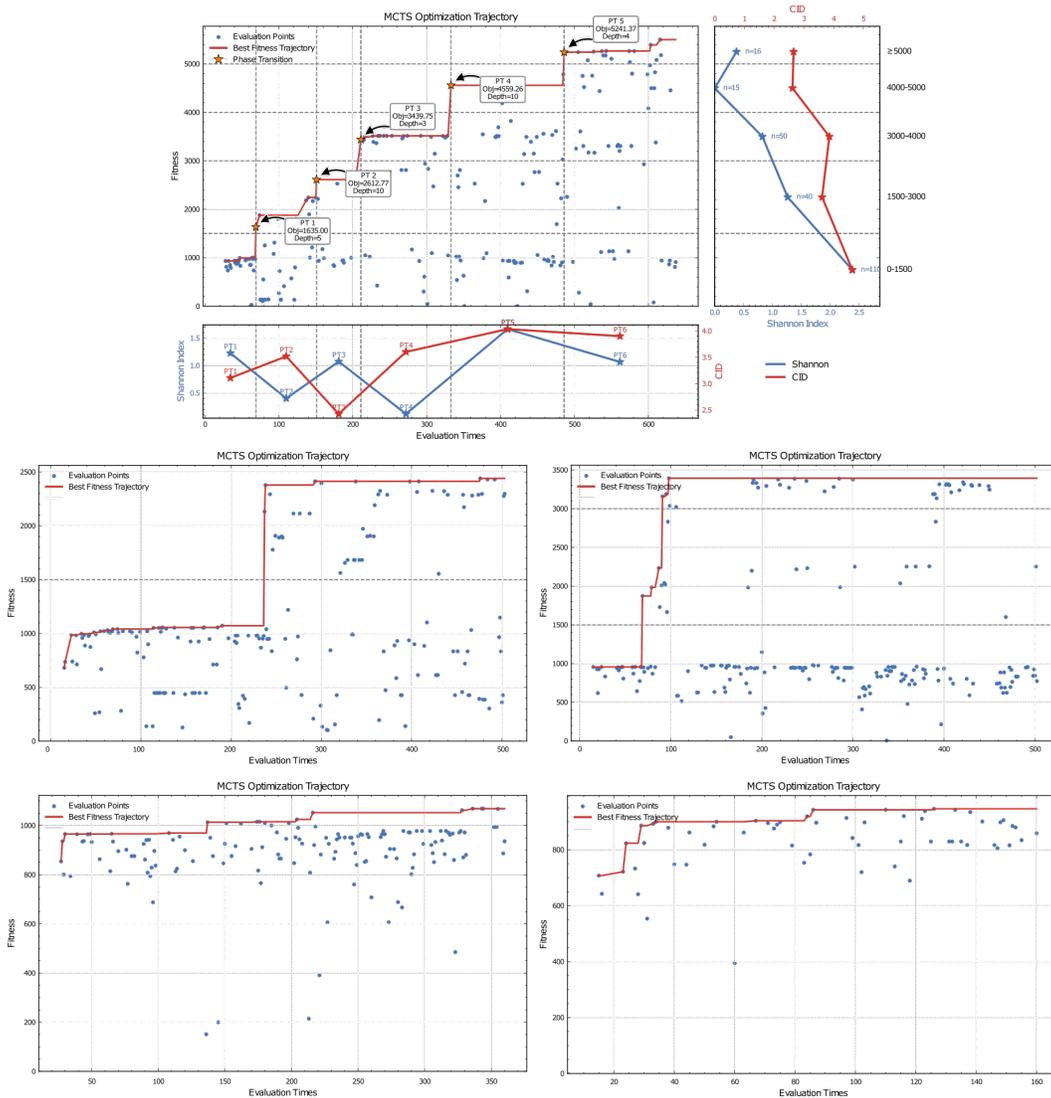


Figure 8: Multi-Run Statistical Analysis of Evo-MCTS Optimization Performance. **Top panel:** Primary run showing complete optimization trajectory with five phase transitions (PT1, PT2.1, PT2.2, PT3, PT4) achieving maximum fitness of 5,241.37 units. Each PT marker indicates fitness value, evaluation number, and tree depth: PT1 (1,635.00, eval 69, depth 5), PT2.1 (2,612.77, eval 151, depth 10), PT2.2 (3,439.75, eval 211, depth 3), PT3 (4,559.26, eval 333, depth 10), PT4 (5,241.37, eval 486, depth 4). **Top right:** Shannon diversity index analysis showing systematic exploration patterns across fitness levels, with aggregate peak diversity (≈ 3.8) in lower performance range (0-1,500 fitness). **Lower panels:** Four additional independent runs demonstrating framework robustness and natural optimization variability. Run 2 (lower left) achieves moderate optimization ($\approx 2,500$ fitness), Run 3 (lower middle-left) shows more substantial performance ($\approx 3,500$ fitness) with algorithmic breakthroughs comparable to the PT2 phase described in the main text, Run 4 (lower middle-right) exhibits limited optimization progress ($\approx 1,100$ fitness), Run 5 (lower right) demonstrates modest improvement from lower baseline (≈ 900 fitness). All runs show improvement over baseline, validating framework reliability while illustrating diverse optimization pathways in the complex algorithmic search space. Results confirm consistent early-phase improvements across all runs with varying success in discovering advanced algorithmic combinations.

3240
3241
3242
3243
3244
3245
3246
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293

the seed function baseline (926.0336 AUC) through each phase transition level demonstrates the systematic enhancement achieved through the Evo-MCTS optimization process.

Table 1: Performance Comparison Across Benchmark Models and Phase Transition Levels

| Model | AUC (units) | Sensitive Distance (Mpc) at FAR | | | |
|-------------------------------------|----------------|---------------------------------|--------|--------|--------|
| | | 1000 | 100 | 10 | 4.3 |
| <i>Benchmark Models</i> | | | | | |
| Sage | 4359.27 | 1996.1 | 1846.6 | 1688.8 | 1672.4 |
| Virgo-AUTh | 4101.48 | 1990.2 | 1818.7 | 1635.0 | 1609.5 |
| PyCBC | 4069.90 | 1832.3 | 1721.8 | 1609.3 | 1573.4 |
| TPI FSU Jena | 3744.99 | 1796.0 | 1581.6 | 1426.4 | 1382.4 |
| CWB | 3225.01 | 1451.6 | 1406.5 | 1351.8 | 1303.2 |
| MFCNN | 2890.33 | 1541.1 | 1269.0 | 997.2 | 900.3 |
| CNN-Coinc | 1997.02 | 1067.2 | 959.9 | 620.4 | 450.8 |
| <i>Phase Transition (PT) Levels</i> | | | | | |
| PT Level 4 | 5241.37 | 2323.9 | 2295.8 | 2080.9 | 2065.3 |
| PT Level 3 | 4559.26 | 1932.0 | 1932.0 | 1932.0 | 1932.0 |
| PT Level 2.2 | 3439.75 | 1537.2 | 1460.1 | 1407.9 | 1402.3 |
| PT Level 2.1 | 2612.77 | 1107.2 | 1107.2 | 1107.2 | 1107.2 |
| PT Level 1 | 1635.00 | — | 769.8 | 769.8 | 769.8 |
| Seed function | 926.03 | 786.9 | 368.2 | 238.3 | 158.3 |

Shannon Diversity Analysis Across Optimization Phases. The Shannon diversity analysis (top right panel) reveals sophisticated exploration patterns that correlate with optimization phases. The scatter plot demonstrates that algorithmic diversity varies systematically with fitness levels, with peak diversity (Shannon ~ 2.5) occurring in the lower performance range (fitness 0-1,500), followed by a gradual decrease as fitness improves. This pattern indicates intensive exploration during early optimization phases, with the framework progressively shifting toward exploitation as high-performing algorithms are discovered. The diversity trend confirms that the Evo-MCTS methodology effectively balances exploration and exploitation, with broader algorithmic sampling during initial discovery phases and more focused refinement during later breakthrough periods.

Multi-Run Consistency and Variability Analysis. The four additional runs (lower panels) demonstrate varying degrees of optimization success, reflecting the inherent stochastic nature of the discovery process while validating the framework’s general effectiveness. Run 2 (bottom-left panel) exhibits steady progression with fitness values reaching approximately 2,500 units, characterized by a gradual upward trajectory punctuated by several modest phase transitions. This pattern illustrates the framework’s ability to consistently identify incremental algorithmic improvements even when breakthrough discoveries remain elusive.

Run 3 (bottom-center-left panel) achieves more substantial performance with fitness values approaching 3,500 units, displaying well-defined phase transitions that correspond to significant algorithmic innovations. This run exhibits characteris-

3294 tics similar to the PT2 phase described in the main text, demonstrating how the
3295 framework can effectively navigate complex solution spaces to discover meaning-
3296 ful algorithmic enhancements under favorable stochastic conditions.

3297 Run 4 (bottom-center-right panel) presents a particularly instructive case despite
3298 showing more limited optimization progress, with fitness values reaching approx-
3299 imately 1,100 units. This run reveals valuable insights into the challenges of nav-
3300 igating rugged optimization landscapes, where persistent exploration attempts en-
3301 counter difficulty escaping local optima. Importantly, even this more constrained
3302 trajectory represents a non-trivial improvement over baseline performance, under-
3303 scoring the framework’s fundamental robustness across varying conditions.

3304 Run 5 (bottom-right panel) demonstrates an intriguing optimization pattern, be-
3305 ginning from a relatively low baseline around 700 fitness units but achieving more
3306 modest improvements to approximately 900 units. Rather than indicating failure,
3307 this trajectory illustrates the challenges encountered in certain regions of the op-
3308 timization landscape, where despite starting from a lower performance level, the
3309 algorithm struggles to discover pathways to substantial improvements, possibly
3310 due to being trapped in a difficult-to-escape local optimum region.

3311 **Statistical Validation and Performance Reliability.** The multi-run analysis pro-
3312 vides critical insights into the framework’s reliability and expected performance
3313 ranges. While not all runs achieve the exceptional performance of the primary
3314 run (5,241.37 units as shown in the primary run analysis), all five runs demon-
3315 strate substantial improvement over baseline performance, though the perfor-
3316 mance varies significantly across runs.

3317 The variation in optimization trajectories reflects the complex, high-dimensional
3318 nature of the algorithmic search space rather than framework instability. Dif-
3319 ferent runs explore distinct regions of the algorithm space, discovering alterna-
3320 tive pathways to improved performance. This diversity of optimization strategies
3321 demonstrates the framework’s robustness and suggests that multiple algorithmic
3322 solutions exist within the search space.

3323 **Optimization Pattern Analysis and Success Factors.** Comparison across runs
3324 reveals consistent early-phase patterns: all runs achieve initial improvements
3325 within the first 100 evaluations, corresponding to the discovery of basic algorithmic
3326 enhancements over the seed function. The divergence in later-phase perfor-
3327 mance correlates with the discovery of advanced algorithmic combinations, where
3328 stochastic factors influence the exploration of high-performance regions.

3329 The most successful runs (primary run reaching 5,241.37 and run 3 reaching
3330 $\sim 3,500$ fitness) share common characteristics: sustained exploration diversity
3331 throughout optimization, discovery of multiple phase transitions, and achievement
3332 of higher fitness levels. Less successful runs (such as run 5 with only ~ 900 fitness)
3333 typically exhibit earlier convergence to local optima, suggesting that maintaining
3334 exploration diversity is crucial for discovering breakthrough algorithmic innova-
3335 tions.

3336 This multi-run analysis validates the framework’s effectiveness while providing
3337 realistic expectations for optimization performance. The results demonstrate that
3338 while exceptional performance (5,241+ fitness) may not be achieved in every run,
3339 the framework consistently produces improvements over the baseline, with per-
3340 formance varying based on the stochastic nature of the optimization process and
3341 the specific regions of the algorithm space explored.

3342 A.4 TEMPORAL CONSTRAINT ANALYSIS AND ROBUSTNESS VALIDATION

3343 The 0.2-second constraint for trigger arrival time uncertainty represents an optimal
3344 balance between astrophysical precision requirements and algorithmic robustness.
3345 This selection is grounded in the physical characteristics of gravitational wave
3346 propagation between detector sites and established multi-detector analysis prac-
3347 tices.

Physical Foundation. The temporal constraint must account for the maximum
light travel time between LIGO Hanford (H1) and Livingston (L1) detectors. For

3348
 3349
 3350
 3351
 3352
 3353
 3354
 3355
 3356
 3357
 3358
 3359
 3360
 3361
 3362
 3363
 3364
 3365
 3366
 3367
 3368
 3369
 3370
 3371
 3372
 3373
 3374
 3375
 3376
 3377
 3378
 3379
 3380
 3381
 3382
 3383
 3384
 3385
 3386
 3387
 3388
 3389
 3390
 3391
 3392
 3393
 3394
 3395
 3396
 3397
 3398
 3399
 3400
 3401

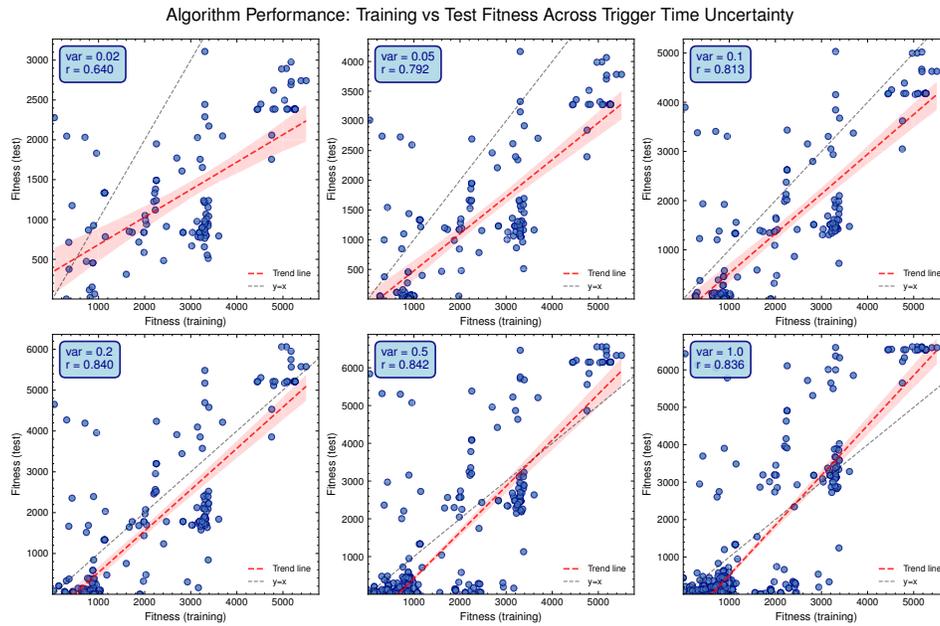


Figure 9: Temporal Constraint Impact on Algorithm Generalization Performance. Analysis of training-test performance correlation as a function of trigger arrival time uncertainty constraints across six different temporal precision levels. **Six panels:** Pearson correlation coefficients between training and test algorithm fitness scores across 877 algorithms for temporal constraint values $\Delta t \in \{0.02, 0.05, 0.1, 0.2, 0.5, 1.0\}$ seconds. Each panel displays scatter plots of training vs. test fitness with correlation coefficients and 95% confidence intervals. The red diagonal line represents perfect training-test correlation ($y = x$). The analysis demonstrates that the 0.2-second constraint ($r = 0.840$) provides optimal balance between performance consistency and practical deployment requirements, with algorithm performance closely following the ideal correlation line. **Key finding:** While tighter constraints (0.02-0.1s) show high correlation coefficients, the 0.2-second constraint exhibits superior generalization behavior with minimal deviation from the ideal $y = x$ relationship, indicating robust performance scaling between training and test conditions.

the two-detector LHO-LLO network, signals must be seen in both detectors within a time difference of 15 ms: 10 ms maximum travel time between detectors and 5 ms padding to account for timing errors [Usman et al. \(2016\)](#). However, the 0.2-second window provides significantly more conservative margin to accommodate additional systematic uncertainties from detector calibration, signal processing delays, timing measurement precision [Abbott et al. \(2019\)](#), and algorithmic robustness requirements while maintaining alignment with operational gravitational wave detection pipelines [Usman et al. \(2016\)](#); [Messick et al. \(2017\)](#).

Experimental Analysis. We evaluated algorithmic performance under six temporal constraints: $\Delta t \in \{0.02, 0.05, 0.1, 0.2, 0.5, 1.0\}$ seconds, calculating training-test performance correlations across all 877 optimized algorithms.

Results and Optimal Selection. Figure 9 demonstrates systematic relationships between temporal constraints and algorithm generalization. Correlation coefficients progress from $r = 0.640$ (0.02s) to $r = 0.840$ (0.2s, optimal), then plateau at $r = 0.842$ (0.5s) and $r = 0.836$ (1.0s). Critically, the 0.2-second constraint exhibits performance characteristics most closely approximating the ideal training-test parity line ($y = x$), with minimal scatter around the diagonal.

Tighter constraints (0.02-0.1s) show increased scatter at higher fitness values, suggesting potential overfitting. Looser constraints (0.5-1.0s) exhibit broader scatter patterns indicating reduced discriminative power for distinguishing high-quality algorithms.

A.5 TECHNIQUE IMPACT ANALYSIS

LLM-Based Code Analysis Pipeline. To systematically extract technical features from algorithm implementations, we developed an automated analysis pipeline using large language models (LLMs). Each code snippet was processed through a structured prompt designed to identify algorithmic components across three main stages: data conditioning, time-frequency analysis, and trigger detection.

LLM Analysis Prompt:

Please analyze the following Python code snippet for gravitational wave detection and extract technical features in JSON format.

The code typically has three main stages:

1. Data Conditioning: preprocessing, filtering, whitening, etc.
2. Time-Frequency Analysis: spectrograms, FFT, wavelets, etc.
3. Trigger Analysis: peak detection, thresholding, validation, etc.

For each stage present in the code, extract:

- Technical methods used
- Libraries and functions called
- Algorithm complexity features
- Key parameters

Code to analyze:

```
```python
{code_snippet}
```
```

Please return a JSON object with this structure:

```
{
  "algorithm_id": "{algorithm_id}",
  "stages": {
    "data_conditioning": {
      "present": true/false,
      "techniques": ["technique1", "technique2"],
```

```

3456         "libraries": ["lib1", "lib2"],
3457         "functions": ["func1", "func2"],
3458         "parameters": {"param1": "value1"},
3459         "complexity": "low/medium/high"
3460     },
3461     "time_frequency_analysis": {...},
3462     "trigger_analysis": {...}
3463 },
3464 "overall_complexity": "low/medium/high",
3465 "total_lines": 0,
3466 "unique_libraries": ["lib1", "lib2"],
3467 "code_quality_score": 0.0
3468 }

```

Only return the JSON object, no additional text.

The analysis was performed using `deepseek-r1-250120` model with `temperature=1.0` for balanced creativity and consistency. We processed 877 valid code snippets (`code_snippet`) using parallel processing with 30 workers to ensure efficient analysis while maintaining API rate limits.

Data Preparation and Normalization. For each identified technique, algorithms were classified into binary groups: those incorporating the technique (“with”) versus those without (“without”). Performance metrics (fitness values or AUC scores) were normalized to [0,1] range using min-max scaling across all algorithms to enable fair comparison between different evaluation metrics.

Combined Performance Analysis. To increase statistical power, we combined normalized AUC scores of both training and test into unified performance datasets for each comparison group. This approach leverages all available performance information while maintaining the comparative structure necessary for statistical testing.

Adaptive Statistical Testing Protocol. Our testing framework adapts to data characteristics through a decision tree approach:

1. **Normality Assessment:** Shapiro-Wilk test for samples $n \leq 5000$
2. **Test Selection:**
 - Both groups normal and $n \geq 30$: Welch’s t-test with Cohen’s d effect size
 - Otherwise: Mann-Whitney U test with rank-biserial correlation
3. **Effect Size Interpretation:**
 - Cohen’s d: negligible (< 0.2), small ($0.2 - 0.5$), medium ($0.5 - 0.8$), large (> 0.8)
 - Rank-biserial: negligible (< 0.1), small ($0.1 - 0.3$), medium ($0.3 - 0.5$), large (> 0.5)

Imbalance Detection and Mitigation. We identify problematic comparisons using two criteria: (i) sample ratio exceeding 3:1, or (ii) minimum group size below 30. For such cases, we implement balanced resampling analysis:

Resampling Protocol:

1. Undersample the larger group to match the smaller group size
2. Perform 1,000 independent resampling iterations with replacement
3. Calculate test statistics and p-values for each iteration
4. Assess robustness based on proportion of significant results

Robustness Criteria: A technique effect is considered robust if:

- $> 80\%$ of resampling iterations show statistical significance ($p < 0.05$)
- Median effect size maintains consistent direction and magnitude
- 95% confidence interval of effect sizes excludes zero

Technique Effectiveness Classification and Visualization. The comprehensive technique impact analysis is presented through violin plot distributions comparing

3510 performance between algorithms incorporating specific techniques (“with”) ver-
 3511 sus those without (“without”) across all identified techniques (Figure 10). Based
 3512 on our multi-criteria evaluation framework, techniques are classified into three
 3513 effectiveness tiers:

3514 **High-Effectiveness Techniques** demonstrate clear distributional separation with
 3515 minimal overlap, statistical significance $> 80\%$ across resampling iterations, and
 3516 large effect sizes ($|r| > 0.5$). Notable examples include Curvature Analysis and
 3517 CWT Validation, which show the “with” group distributions positioned substan-
 3518 tially higher than “without” groups, indicating consistent performance improve-
 3519 ments.

3520 **Medium-Effectiveness Techniques** exhibit moderate distributional separation,
 3521 statistical significance between 50-80%, and medium effect sizes ($0.3 < |r| <$
 3522 0.5). These techniques provide measurable but less consistent performance bene-
 3523 fits.

3524 **Low-Effectiveness Techniques** display substantial distributional overlap between
 3525 “with” and “without” groups, statistical significance $< 50\%$, and small effect sizes
 3526 ($|r| < 0.1$), indicating limited practical utility.

3528 Table 2: **Technique Abbreviations Used in Figure 10.** Complete mapping of abbreviated technique
 3529 names to their full descriptions, organized by methodological category.

| 3530 | Abbreviation | Full Name |
|------|--------------------------------|------------------------------------|
| 3531 | Data Conditioning | |
| 3532 | Spline Interp | Spline Interpolation |
| 3533 | Tikhonov Reg | Tikhonov Regularization |
| 3534 | Kalman Smooth | Kalman-inspired Smoothing |
| 3535 | Tukey | Tukey Windowing |
| 3536 | S-G Filter | Savitzky-Golay Filtering |
| 3537 | Adaptive Gain | Adaptive Gain Regularization |
| 3538 | Median Smooth | Uniform/Median Smoothing |
| 3539 | Gauss Smooth | Gaussian Smoothing |
| 3540 | Median Baseline | Median-based Baseline Correction |
| 3541 | SNR Reg | SNR-adaptive Regularization |
| 3542 | Gauss Conv | Gaussian Convolution |
| 3543 | Time-Frequency Analysis | |
| 3544 | Phase Coherence | Phase Coherence Analysis |
| 3545 | CWT | Continuous Wavelet Transform |
| 3546 | RMS Coherence | RMS Coherence Metric |
| 3547 | Dual Align | Dual-channel Alignment |
| 3548 | Spectral Entropy | Spectral Entropy |
| 3549 | Freq Reg | Frequency-dependent Regularization |
| 3550 | Log Compress | Logarithmic Compression |
| 3551 | CWT Valid | CWT Validation |
| 3552 | Trigger Detection | |
| 3553 | Curve Boost | Curvature Boosting |
| 3554 | Curvature | Curvature Analysis |
| 3555 | Sigmoid | Sigmoid Enhancement |
| 3556 | MAD Threshold | MAD-based Robust Thresholding |
| 3557 | MAD | Median Absolute Deviation |

3558 **Distributional Analysis Methodology.**

- 3559 • Violin plots constructed using Gaussian kernel density estimation with adap-
 3560 tive bandwidth selection
- 3561 • Performance metrics normalized to $[0,1]$ scale enabling cross-technique com-
 3562 parison
- 3563 • Color-coded categorical organization: data conditioning (blue), time-
 frequency analysis (orange), trigger detection (green)

3564
 3565
 3566
 3567
 3568
 3569
 3570
 3571
 3572
 3573
 3574
 3575
 3576
 3577
 3578
 3579
 3580
 3581
 3582
 3583
 3584
 3585
 3586
 3587
 3588
 3589
 3590
 3591
 3592
 3593
 3594
 3595
 3596
 3597
 3598
 3599
 3600
 3601
 3602
 3603
 3604
 3605
 3606
 3607
 3608
 3609
 3610
 3611
 3612
 3613
 3614
 3615
 3616
 3617

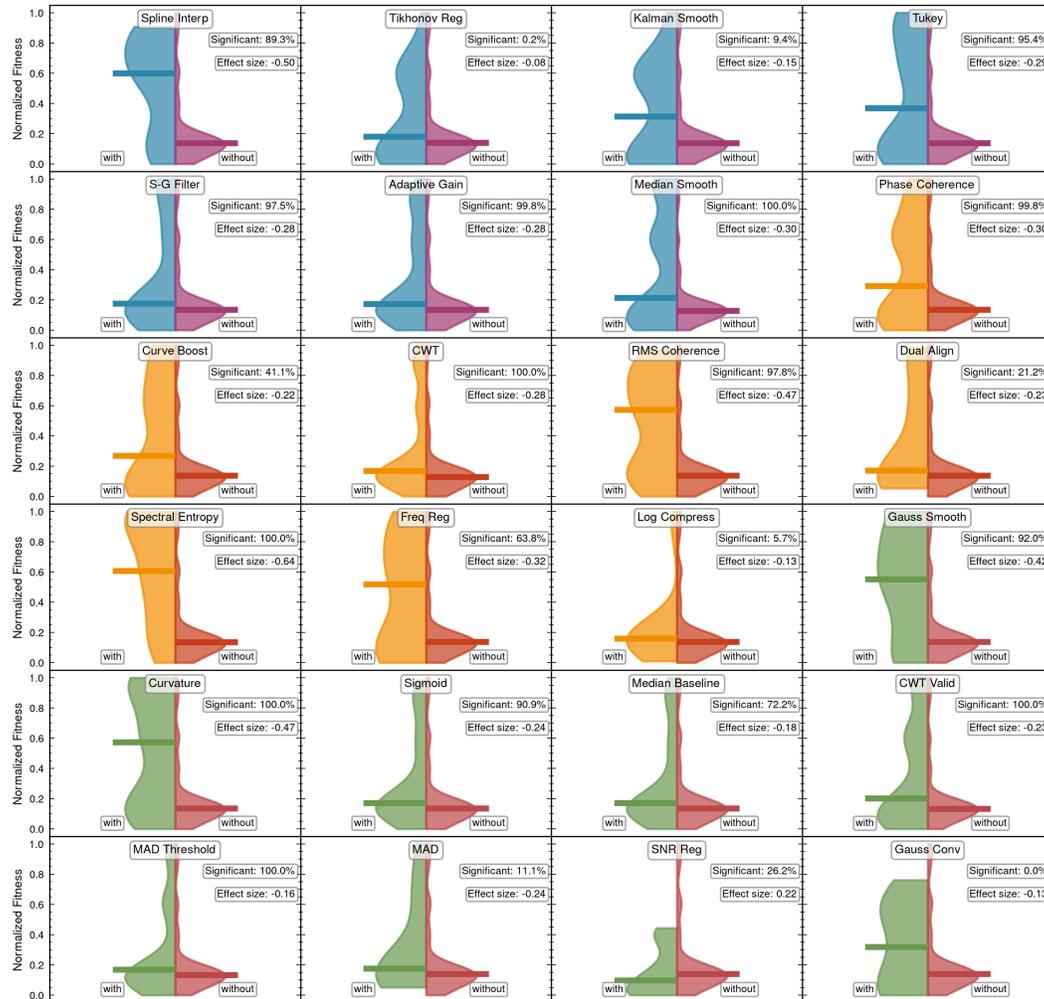


Figure 10: **Comprehensive Technique Effectiveness Analysis via Violin Plot Distributions.** Performance distributions comparing algorithms with and without specific techniques across three methodological categories: data conditioning (blue), time-frequency analysis (orange), and trigger detection (green). Each violin plot pair reveals technique effectiveness through distributional characteristics: wider sections indicate higher probability density regions, clear vertical separation between “with” and “without” groups indicates strong technique effects, while substantial overlap suggests limited effectiveness. Statistical robustness metrics (significance percentages from resampling analysis) and effect sizes (rank-biserial correlations) quantify technique reliability. High-effectiveness techniques (e.g., Curvature Analysis, CWT Validation) demonstrate clear distributional separation and large effect sizes, while low-effectiveness techniques show substantial overlap and negligible effect sizes. Technique abbreviations are defined in Table 2.

- 3618 • Statistical annotations include resampling-based significance percentages
- 3619 and rank-biserial effect sizes
- 3620 • Median performance indicators highlight central tendency differences be-
- 3621 tween technique groups
- 3622 • Distributional separation quantified through overlap coefficients and
- 3623 Kolmogorov-Smirnov distances
- 3624 • Technique abbreviations facilitate visual clarity while maintaining compre-
- 3625 hensive coverage (Table 2)

3626 This multi-dimensional effectiveness assessment framework enables systematic
 3627 identification of high-impact techniques while distinguishing them from those
 3628 with marginal or inconsistent benefits, providing clear guidance for algorithmic
 3629 development priorities.

3630 A.6 DETAILED MCTS NODE EVOLUTION AND TECHNIQUE PROPAGATION 3631 DATA

3632 Figure 11 presents the complete MCTS search tree evolution with node-by-node
 3633 fitness values and technique compositions. Each node displays its fitness score
 3634 (marked in red) alongside the specific algorithmic techniques discovered at that
 3635 search depth. The five core technique categories [1-5] correspond to:

- 3637 1 Multi-resolution Thresholding
- 3638 2 Continuous Wavelet Transform (CWT) using Ricker Wavelet
- 3639 3 Tikhonov Regularization
- 3640 4 Curvature Boosting
- 3641 5 Savitzky-Golay Filter

3642 These techniques demonstrate the systematic evolution of algorithmic complexity
 3643 through the MCTS exploration process, with detailed implementation examples
 3644 provided in Section A.1.9.

3646 A.7 OVERFITTING RISK ASSESSMENT AND PHYSICAL VALIDATION 3647 FAILURE MODES

3648 To evaluate potential overfitting risks in the Evo-MCTS framework, we conducted
 3649 systematic injection studies using GW150914-like signals with varying signal-to-
 3650 noise ratios (SNR). This analysis addresses concerns raised in the main discus-
 3651 sion regarding algorithmic optimization potentially overfitting to MLGWSC-1's
 3652 specific characteristics.

3653 **Experimental Design.** We performed injection experiments using simulated
 3654 gravitational wave signals based on the posterior distribution of GW150914 [Ab-](#)
 3655 [bott et al. \(2016\)](#). The injected signals were embedded in realistic detector noise
 3656 closed to GW150914 from the O1 observing run, with SNR values systematically
 3657 varied across various SNRs. For each SNR group, we generated 100 indepen-
 3658 dent injection realizations and evaluated the detection efficiency of representative
 3659 MCTS nodes across the algorithmic evolution tree.

3660 **Detection Efficiency Analysis.** Figure 12 presents detection efficiency curves
 3661 for six representative nodes (Node-4, Node-7, Node-411, Node-464, Node-485,
 3662 Node-486) spanning different evolutionary stages of the MCTS search. The re-
 3663 sults reveal concerning overfitting patterns that manifest as performance degra-
 3664 dation when algorithms optimized on MLGWSC-1 data encounter realistic astro-
 3665 physical scenarios.

3666 Node-4, representing an early-stage algorithm with minimal complexity, demon-
 3667 strates robust performance across all SNR ranges with detection efficiency exceed-
 3668 ing 80% for $\text{SNR} > 10$. In contrast, highly evolved nodes (Node-464, Node-485,
 3669 Node-486) exhibit sharp performance transitions, achieving near-perfect detection
 3670 only above $\text{SNR} = 12-15$, suggesting over-specialization to the noise characteris-
 3671 tics and signal morphologies present in the MLGWSC-1 training dataset.

Overfitting Manifestations. The steep detection efficiency curves observed for
 advanced nodes indicate several failure modes characteristic of overfitting:

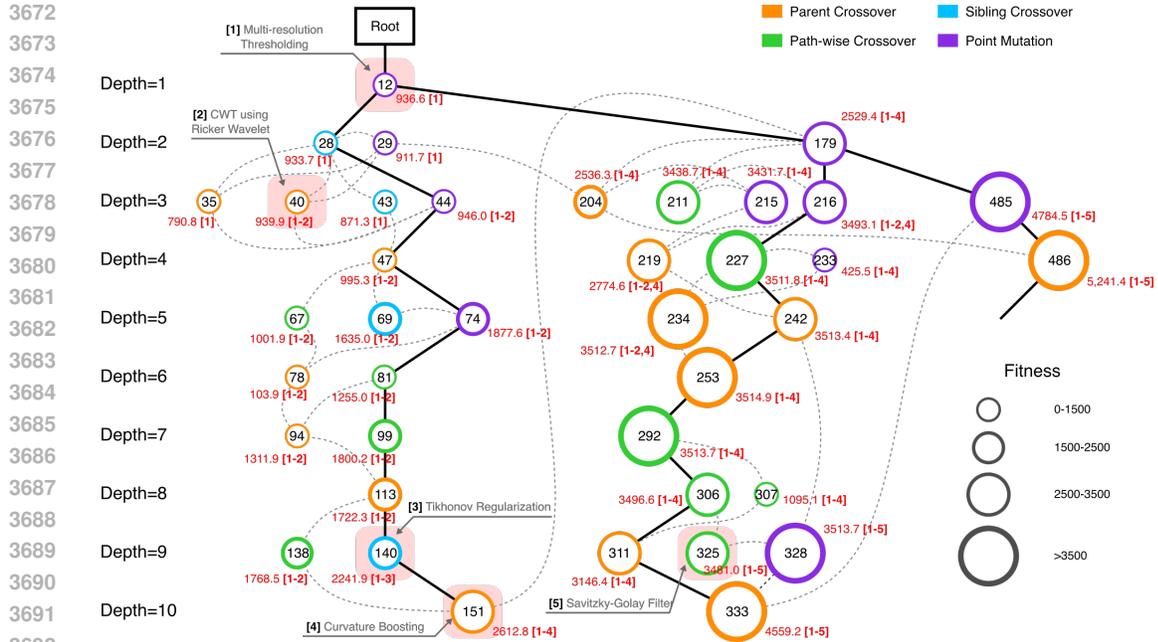


Figure 11: **Complete MCTS search tree with node fitness values and technique compositions.** Each node shows its fitness score (red annotations) and constituent algorithmic techniques organized by category [1-5]. Node size reflects fitness magnitude. The tree demonstrates systematic technique evolution and cross-branch knowledge transfer, with optimal performance achieved through multi-technique integration at terminal nodes.

- **Threshold Over-Optimization:** Advanced nodes demonstrate excessively restrictive detection thresholds optimized for MLGWSC-1’s specific noise floor, resulting in reduced sensitivity to weaker but astrophysically realistic signals.
- **Feature Over-Specialization:** Complex multi-technique combinations (e.g., Node-486 incorporating techniques [1,2,3,4,5]) show degraded performance on signal morphologies that deviate from the limited parameter space explored in MLGWSC-1.
- **Noise Model Dependency:** The sharp performance transitions suggest algorithms have adapted to MLGWSC-1’s inherent noise model, failing to generalize to the more complex non-stationary and non-Gaussian characteristics of operational detector data.
- **Competition Metric Exploitation:** The MLGWSC-1 evaluation framework itself contains inherent limitations that enable algorithmic exploitation of scoring system vulnerabilities. Advanced nodes appear to have discovered and exploited specific characteristics of the competition’s evaluation metrics, optimizing for artificial performance gains rather than genuine astrophysical detection capability. This metric gaming behavior results in algorithms that achieve high competition scores while failing to maintain robust performance under realistic detection scenarios.

Physical Validation Implications. The overfitting analysis reveals fundamental limitations of competition-based benchmarks for gravitational wave detection. Advanced nodes (Node-464, Node-485, Node-486) demonstrate sharp performance degradation on realistic astrophysical signals, indicating over-specialization to MLGWSC-1’s constrained evaluation framework. Intermediate-complexity nodes (Node-7, Node-411) achieve optimal performance-robustness

3726
3727
3728
3729
3730
3731
3732
3733
3734
3735
3736
3737
3738
3739
3740
3741
3742
3743
3744
3745
3746
3747
3748
3749
3750
3751
3752
3753
3754
3755
3756
3757
3758
3759
3760
3761
3762
3763
3764
3765
3766
3767
3768
3769
3770
3771
3772
3773
3774
3775
3776
3777
3778
3779

balance, suggesting that benchmark optimization alone is insufficient for real-world deployment.

Benchmark Limitations and Methodological Insights. The observed overfitting patterns expose critical weaknesses in current evaluation protocols: (1) limited signal diversity that enables algorithmic exploitation of specific noise characteristics, (2) evaluation metrics that reward competition performance over astrophysical sensitivity, and (3) insufficient validation against realistic detector conditions. The Evo-MCTS framework’s systematic exploration of the performance-generalization Pareto frontier transforms these limitations into methodological strengths, providing rigorous overfitting detection capabilities.

Pareto Frontier Discovery and Framework Contribution. The detection efficiency results demonstrate that the Evo-MCTS framework has systematically mapped the algorithm design space, revealing fundamental trade-offs between benchmark performance and generalization capability. Node-411 represents an optimal point on this frontier, achieving superior balance compared to the highest-scoring Node-486, which sacrifices robustness for competition metrics. This Pareto frontier discovery constitutes a core methodological contribution, enabling informed decisions about performance-robustness trade-offs rather than pursuing single-metric optimization.

Future Directions and Benchmark Enhancement. This analysis motivates comprehensive benchmark improvement: (1) incorporation of diverse astrophysical populations spanning multiple source types and parameter ranges, (2) multi-detector validation protocols using independent observing run data, (3) evaluation frameworks that explicitly penalize overfitting through cross-validation requirements, and (4) systematic comparison with established detection algorithms under identical protocols. The current assessment focuses on binary black hole signals; comprehensive validation requires extension to neutron star mergers, continuous wave sources, and burst signals to establish universal applicability of discovered algorithmic principles.

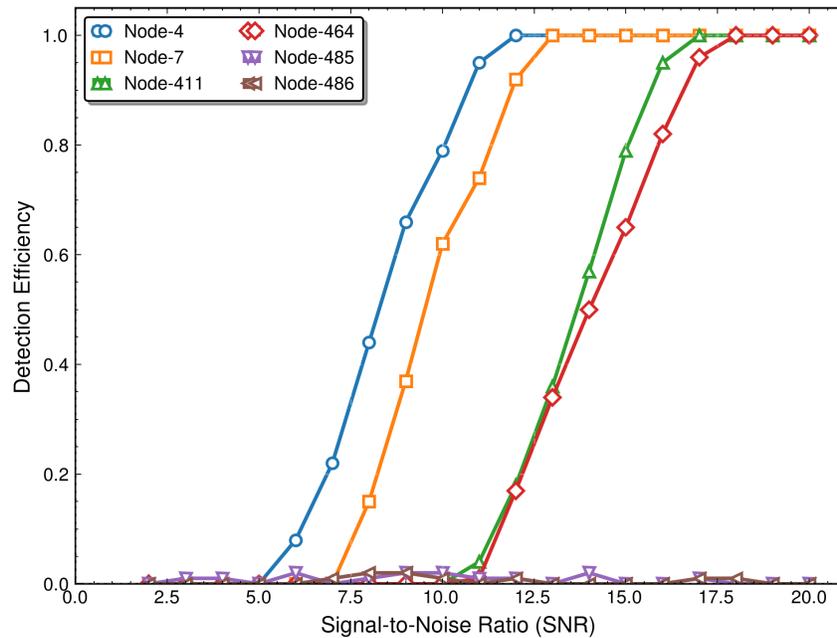


Figure 12: **Detection efficiency analysis revealing overfitting risks in evolved MCTS nodes.** Detection efficiency curves for representative nodes across different SNR groups using GW150914-like injections in O1 detector noise. Early-stage nodes (Node-4, blue) demonstrate robust performance across all SNR ranges, while highly evolved nodes (Node-464, Node-485, Node-486, red/purple/brown) exhibit sharp performance transitions indicative of overfitting to MLGWSC-1 characteristics. Intermediate-complexity nodes (Node-7, Node-411, orange/green) achieve optimal balance between sophistication and generalization capability.