

Article

RAMA: A Meta-Algorithmic Framework for Ramanujan-Style Heuristic Discovery Using Large Language Models

Jordi Vallverdú 

Philosophy Department, ICREA, Universitat Autònoma de Barcelona, 08193 Bellaterra, BCN, Spain;
jordi.vallverdu@uab.cat

Abstract

This work introduces RAMA (Recursive Aesthetic Modular Approximation), a metaheuristic framework that models a restricted form of mathematical intuition inspired by the notebooks of Srinivasa Ramanujan. While Ramanujan often produced deep results without formal proofs, the heuristic processes guiding such discoveries remain poorly understood. RAMA treats large language models (LLMs) as proposal mechanisms within an iterative search that generates, evaluates, and refines candidate conjectures under an explicit energy functional balancing fit, description length, and aesthetic structure. A small set of Ramanujan-inspired heuristics—modular symmetries, integrality cues, aesthetic compression, and near-invariance detection—is formalized as micro-operators acting on symbolic states. We instantiate RAMA in two domains: (i) inverse engineering eta-quotients from partial q -series data and (ii) designing cyclotomic fingerprints with shadow gadgets for quantum circuits. In both settings, RAMA recovers compact structures from limited information and improves separation from classical baselines, illustrating how intuitive heuristic patterns can be rendered as explicit, reproducible computational procedures.

Keywords: heuristic discovery; Ramanujan; mock modular forms; meta-algorithm; large language models; aesthetic compression; modular symmetries

1. Introduction

Srinivasa Ramanujan (1887–1920) generated an exceptional number of correct mathematical results without providing formal proofs. His three notebooks, later reproduced in facsimile by the Tata Institute of Fundamental Research [1], record thousands of identities, series, congruences, and continued fractions across number theory, q -series, and modular forms. Decades later, these results were edited, proved, and conceptually organized by Berndt in the five-volume *Ramanujan's Notebooks* [2–6]. After Ramanujan's death, the *lost notebook* was eventually rediscovered and analyzed in detail [7–12], revealing deep links with modern modular and mock modular theory [13–16]. This historically well-documented trajectory motivates a central question: Can aspects of Ramanujan's intuitive practice be captured as a reproducible computational procedure? Despite this historical foundation, existing computational-creativity and heuristic-discovery frameworks do not directly target the specific combination of symmetry detection, aesthetic compression, and near-invariance completion that appears repeatedly in Ramanujan's work. Approaches such as symbolic regression, program synthesis, or pattern-mining systems typically optimize for accuracy or minimal description length, but they lack an explicit mechanism for handling “almost-symmetries” or structured deviations—the phenomena that motivate mock modular forms and related completions. RAMA is introduced precisely to address



Academic Editors: Mohammad Majid al-Rifaie and Frank Werner

Received: 25 October 2025

Revised: 10 December 2025

Accepted: 19 December 2025

Published: 21 December 2025

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

this gap: it provides a meta-algorithmic formulation that unifies these heuristic pressures into a single, operational computational process.

We answer in the affirmative by proposing **RAMA** (*Recursive Aesthetic Modular Approximation*), a meta-algorithmic framework that distills three forces often implicated in creative mathematical work—*pattern resonance*, *aesthetic compression*, and *shadow completion*—into an executable search over symbolic states. RAMA treats intuition not as an oracle but as a controllable optimization of a composite energy functional balancing simplicity, coherence, and near-invariance.

Summary of contributions.

(1) We formalize **RAMA**, a meta-algorithm that operationalizes pattern resonance, aesthetic compression, and shadow completion. (2) We prove that LLM-driven prompt chains with computable control are *algorithms* in the ASM/partial-recursive sense, enabling standard resource analysis. (3) We instantiate RAMA in two domains: (A) inverse-engineering eta-quotients from partial q -series, and (B) cyclotomic fingerprints for quantum circuits with shadow gadgets. (4) We release fully reproducible notebooks for both instantiations, together with an initial ablation in Track B (shadow on/off). Planned Track A ablations (compression/aesthetic toggles) are outlined in Section 7 and will be added to the supplementary notebook. Figure 1 summarizes the overall flow of RAMA from seed objects and micro-operators to stabilized conjectures.

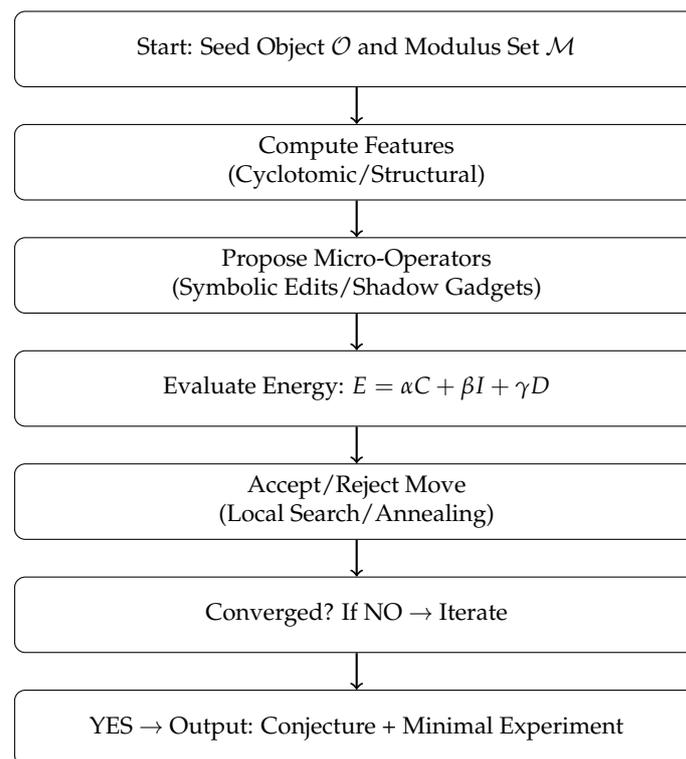


Figure 1. High-level flowchart of the RAMA meta-algorithm. Starting from an initial object \mathcal{O} and modulus set \mathcal{M} , the system iteratively (i) computes cyclotomic or structural features, (ii) proposes local micro-operators (symbolic edits or shadow gadgets), (iii) evaluates the resulting state through the energy functional $E = \alpha C + \beta I + \gamma D$, and (iv) accepts or rejects moves under a local-search schedule until a stabilized conjecture and minimal experiment are produced.

Scope of the claim.

Throughout, we use “intuition” in a deliberately narrow, operational sense: RAMA does *not* attempt to reconstruct the full psychological processes of Ramanujan, nor does it claim to capture all aspects of human mathematical creativity. Instead, we extract and

formalize a small set of heuristic forces that are repeatedly visible in the notebooks—pattern resonance, compression, and near-symmetry completion—and show that these can be rendered as an explicit energy-minimizing search procedure. Large language models are used only as proposal engines within this procedure, and all mathematical content is externally validated; the framework is, therefore, intended as a *computational model of certain heuristic patterns*, not as a literal theory of human cognition.

Organization of the paper.

Section 2 reviews the historical and cognitive background. Section 3 introduces the formal RAMA meta-algorithm. Sections 4 and 5 detail its meta-algorithmic rationale, operational structure, and energy functional. Section 6 establishes the equivalence between prompt chains and classical algorithms. Section 7 present the empirical instantiations and ablation studies. Section 8 links the implementation back to the formal model, and Sections 9 and 10 offer the discussion and conclusions.

We next situate RAMA against the historical and cognitive background that motivates its design (Section 2) and then present the formal framework (Section 3).

2. Historical and Cognitive Foundations

Between 1903 and 1914, Ramanujan compiled three main notebooks containing nearly 3900 statements. These were later reproduced in facsimile by the Tata Institute of Fundamental Research [1]. Because many entries lacked proofs, they remained underexplored until the 1980s, when Berndt undertook the long project of editing and proving them, publishing the five-volume *Ramanujan's Notebooks* [2–6]. After Ramanujan's death in 1920, a further collection of roughly one hundred pages of late notes was misplaced and eventually rediscovered by Andrews in 1976 at Trinity College, Cambridge; this *lost notebook* was published in 1988 [7] and then analyzed in detail across multiple volumes [8–12]. Modern expositions clarified the nature of mock theta functions and their shadows [13–16], providing a conceptual backdrop to many of Ramanujan's late formulas.

Inverse Engineering: How We Extracted RAMA from Primary Sources

Definition 1 (Inverse Engineering of Heuristics). *Given a corpus of outputs and partial context (e.g., Ramanujan's entries and brief notes), inverse engineering reconstructs a minimal decision process that could plausibly generate those outputs. The aim is not historical psychology, but a compact algorithmic schema whose forward behavior reproduces the observed regularities.*

In our setting, the artifacts are the three notebooks and the *lost notebook* (facsimiles and edited editions). The reconstructed schema is the RAMA meta-algorithm introduced later in Section 3. Large language models (LLMs) were used strictly as *proposal engines* to surface candidate patterns from short, author-supplied excerpts; historical claims and mathematics were validated against the editions and by light symbolic checks.

Objects and provenance.

Primary materials comprised (i) the three notebooks (facsimile) and (ii) the lost notebook (facsimile and commentary volumes), as cited in the bibliography. Excerpts provided to the LLM were short, manually curated snippets (identities, series fragments, continued-fraction stubs, index entries), each tagged with human-maintained provenance (notebook number/sheet/page/entry).

Transcript provenance.

Extraction proceeded via contemporaneous chats between the author and an OpenAI LLM in the ChatGPT interface. Exact model names and timestamps are preserved

in the author’s transcripts and can be provided upon request. We avoid asserting a specific model/version here to prevent misstatement; the analysis does not depend on model branding.

Workflow (contrastive inverse engineering).

We iterated the following loop:

- (1) **Scope and anchor:** Choose a small neighborhood of entries around a focal identity; record source tags.
- (2) **Normalize fragments:** Map notation to modern syntax (e.g., q -Pochhammer); add minimal expansions with inline provenance notes.
- (3) **Contrastive prompts (LLM as proposer):** Paired prompts forced discriminative explanations (“features present in [A] but absent in [B]”, “shortest rule covering [A,B] but not [C]”).
- (4) **Abductive synthesis (human-in-the-loop):** Retain only candidates with minimal description length that predict neighboring entries; rephrase as RAMA micro-operators (modular shift, cyclotomic rotation, symmetry completion).
- (5) **External verification:** Textual alignment to the cited edition; symbolic sanity checks outside the LLM (few coefficients, small-modulus congruences); triangulation with Berndt and Andrews–Berndt where applicable; unaligned items were kept as modern hypotheses only.

Controls against hallucination.

Prompts used only human-supplied text (no web). The LLM was instructed to separate observed statements from speculative completions. Suggested quotations were cross-checked against editions; unverifiable phrasing was removed or rewritten as modern commentary. Items labeled as “pattern” or “mock symmetry” remained conjectural until externally verified.

3. The RAMA Meta-Algorithm: A Metacognitive Framework

3.1. Motivation and Problem Statement

Traditional models of reasoning—syllogistic logic, algorithmic derivation, or symbolic computation—presume conscious sequential inference. Ramanujan’s notebooks suggest a different regime: correctness is treated as emerging from internal coherence, symmetry, and aesthetic balance rather than from strictly stepwise deduction. We denote the abstract mechanism underlying this process by **RAMA** (*Recursive Aesthetic Modular Approximation*), conceptualized here as a *meta-algorithm* for non-linear mathematical intuition.

3.2. Metacognitive Architecture

RAMA models intuitive reasoning as a *recursive feedback system* over a high-dimensional symbolic space \mathcal{S} . Let each element encode a partially structured object (a numerical pattern, a fragment of a series, or a symbolic expression). A mental state is then a distribution p_t over \mathcal{S} at iteration t .

At each iteration, RAMA applies a non-linear operator

$$R : \mathcal{S} \rightarrow \mathcal{S}, \quad p_{t+1} = R(p_t),$$

combining the following:

- (a) **Pattern resonance (perceptual):** Identify simple invariants such as ratios or congruences and reinforce the configurations that exhibit the strongest symmetry.
- (b) **Compression (aesthetic):** Bias toward minimal description length.

- (c) **Shadow completion (metacognitive):** When invariance is nearly true (“mock symmetry”), propose minimal corrections (“shadows”) that restore exactness if real.

Convergence occurs when an attractor p^* minimizes

$$E(p) = \alpha C(p) + \beta I(p) + \gamma D(p),$$

where C is symbolic complexity, I internal inconsistency, D deviation from aesthetic balance (symmetry, periodicity, modular harmony), and (α, β, γ) are subjective weights.

3.3. Operational Description

- (1) **Perceive:** Encode features (ratios, residues, partial sums).
- (2) **Associate:** Retrieve similar stored patterns.
- (3) **Perturb and resonate:** Small symbolic moves (shifts, reciprocals, modular reductions) → stability signals.
- (4) **Compress:** Keep variants with maximal invariance/minimal syntax.
- (5) **Shadow completion:** Hypothesize minimal terms or maps that close the structure.
- (6) **Meta-evaluate:** Accept as conjecture if stable; otherwise iterate.

To make the operational pipeline explicit, we define the neighborhood structure, acceptance rule, and parameter schedule used in all experiments. The neighborhood $\mathcal{N}(p)$ consists of small, localized transformations: (i) exponent increments or decrements by ± 1 , (ii) insertion or deletion of a single $\eta(q^d)$ factor, (iii) modular shifts $q \mapsto q^k$ with $|k| \leq 2$, and (iv) cyclotomic moves $d \mapsto md$ for small $m \in \{2, 3\}$. A candidate move $p' \in \mathcal{N}(p)$ is accepted whenever $E(p') < E(p)$, and is accepted with probability $\exp(-(E(p') - E(p))/T)$ otherwise, where T is an annealing temperature that decreases linearly over the run. Unless stated otherwise, all runs use a shared parameter setting $(\alpha, \beta, \gamma) = (1, 1, 0.2)$ after normalizing the scales of C , I , and D .

3.4. Philosophical Interpretation

RAMA treats intuition as self-organizing dynamics seeking low-energy (low-complexity, high-coherence) states in symbolic space; the felt “beauty” or “inevitability” corresponds to a decrease in $E(p)$. Proof then becomes an ex-post reconstruction of an attractor previously found by aesthetic compression.

3.5. Definition (Formal Summary)

Definition 2 (RAMA Meta-Algorithm). *Let (\mathcal{S}, E) be a symbolic configuration space with an energy functional representing cognitive dissonance. The operator*

$$R(p) = \arg \min_{q \in \mathcal{N}(p)} E(q)$$

*acts locally within a neighborhood $\mathcal{N}(p)$ generated by small modular and algebraic transformations. Iterating R produces a sequence (p_t) converging toward attractors that satisfy minimal complexity and maximal internal coherence. A stabilized attractor corresponds to an intuitive conjecture. We call the complete recursive process the **RAMA meta-algorithm**.*

Epistemic implication.

RAMA frames intuition as a reproducible computational principle: the spontaneous generation of plausible mathematical truths via recursive aesthetic optimization under metacognitive supervision.

3.6. Scope and Philosophical Position

It is important to distinguish between three levels of claim. First, the *historical* claim is modest: we use Ramanujan’s notebooks as a rich corpus of outputs and contextual notes from which to inverse-engineer a small family of micro-operators and aesthetic biases. Second, the *computational* claim is that this family can be assembled into a well-defined meta-algorithm, RAMA, whose trajectories are governed by an explicit energy functional and that can be instantiated on modern hardware. Third, the *cognitive* claim is deliberately conservative: we interpret RAMA as a stylized model of certain heuristic pressures—toward symmetry, compression, and completion—rather than as a full account of human mathematical intuition.

In particular, we do not assert that intuition is reducible to a single energy functional, nor that subjective experiences of insight are captured by our formalism. Instead, we view RAMA as an experimentally testable hypothesis about one reusable pattern of heuristic control that appears compatible with both historical evidence and modern computational practice. The framework is, therefore, meant to complement, not replace, philosophical and psychological accounts of creativity.

4. Meta-Algorithmic Rationale

Motivation. A central challenge in mathematical discovery is to recover the underlying structure of an identity when formal derivations are not yet available. Ramanujan’s “correct but unproved” results suggest a process shaped by symmetry, congruence structure, and a preference for concise expressions rather than by a single deterministic technique. Modern developments interpret many of his late formulas as mock modular forms equipped with nonholomorphic shadows [13–16]. These observations motivate a metaheuristic approach that models intuitive reasoning through a small library of composable operators and an explicit aesthetic bias.

Constraints shaping the method. Our reconstruction mirrors Ramanujan’s historical working conditions: (i) Limited access to long derivations. (ii) Reliance on small, interpretable symbolic transformations (modular reductions, cyclotomic shifts, short recurrences). (iii) A consistent pull toward short or highly symmetric expressions. These constraints naturally suggest a local-search mechanism guided by an energy functional that rewards simplicity, stability, and near-invariance.

From Ramanujan to a metaheuristic. The following micro-operators and control principles arise repeatedly in the notebooks and form the building blocks of RAMA:

- (1) *Micro-transformations:* Probe the structure of a candidate expression using small modular or cyclotomic edits.
- (2) *Channel filtering:* Decompose behavior according to residue classes or characters to reveal hidden regularities.
- (3) *Compression bias:* Prefer the shortest expression consistent with observed data.
- (4) *Mock symmetry and shadow:* Detect near-invariance and posit minimal corrective terms that restore exact structure.
- (5) *Empirical checking:* Test candidates on small, reliable instances before attempting broader generalization.

This rationale directly informs the computational instantiations presented later. In Track A, the operators act on eta-quotients via exponent edits, modular shifts, and cyclotomic moves; in Track B, they act on quantum circuits via periodic phase insertions that serve as physical shadow corrections. In both domains, a small set of principled, composable edits suffices to explore a large symbolic space, and the aesthetic energy functional provides a consistent criterion for stabilization. These ideas complement established practices in both analytic number theory and quantum benchmarking [17–21].

5. Meta-Algorithmic Framework and Final Output

5.1. Energy Functional: Track-Specific Forms

The abstract energy functional introduced in Section 3,

$$E(p) = \alpha C(p) + \beta I(p) + \gamma D(p),$$

takes concrete, domain-specific forms in our two empirical tracks.

Track A (q-series inverse engineering).

A state p encodes an eta-quotient

$$f_p(q) = q^k \prod_d \eta(q^d)^{r_d},$$

together with a target q -series $g(q) = \sum_{n \geq 0} a_n q^n$ specified by its first N coefficients. We define the following:

$C(p)$ = symbolic description length of f_p (e.g., the total number of nonzero exponents r_d plus $|k|$),

$$I(p) = \sum_{n=0}^{N-1} |[q^n]f_p(q) - a_n|^2,$$

$$D(p) = \lambda_1 \sum_d |r_d| d + \lambda_2 \cdot \mathbf{1}_{\text{unbalanced exponents}}.$$

Here, $[q^n]f_p(q)$ denotes the n th coefficient of f_p , and $D(p)$ penalizes large moduli and strongly unbalanced exponent patterns via small fixed weights $\lambda_1, \lambda_2 > 0$. In our experiments we rescale C, I , and D so that they are numerically comparable and then use fixed (α, β, γ) across runs.

Track B (cyclotomic fingerprints for circuits).

A state p encodes a quantum circuit C together with a candidate *shadow gadget* schedule G (e.g., periodic $R_z(\pi/m)$ insertions). For each modulus $m \in \{4, 6, 8\}$, we estimate the cyclotomic moment $S_m(C)$ as in (1), and use a scalar separability score such as the area under the ROC curve, $\text{AUC}_m(C)$, for distinguishing samples from C against a matched classical baseline. We then set

$C(p)$ = length of the shadow schedule (number of inserted gates);

$$I(p) = -\frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \text{AUC}_m(C);$$

$$D(p) = \mu_1 \cdot \text{Var}_{m \in \mathcal{M}}(\text{insertion period}_m) + \mu_2 \cdot \mathbf{1}_{\text{irregular pattern}};$$

with small weights $\mu_1, \mu_2 > 0$. Thus, minimizing E favors short, regular shadow schedules that achieve high separability (large AUC) across moduli.

5.2. Objects, Characters, and Cyclotomic Moments

Let \mathcal{O} be an object (e.g., a distribution from a circuit). For $m \in \mathbb{Z}_{\geq 2}$, let $\chi_m : \{0, 1\}^n \rightarrow \mathbb{C}$ be a simple modular character (e.g., $\chi_m(x) = \exp(2\pi i |x|/m)$). Define the *cyclotomic moment*

$$S_m = \mathbb{E}_{x \sim p} [\chi_m(x)], \quad p \text{ the target distribution.} \tag{1}$$

These compress modular regularities into a single statistic, echoing Dirichlet/cyclotomic structure [17]. Concretely, the map $x \mapsto |x| \bmod m$ induces a character on the additive

group of Hamming weights, and $\chi_m(x) = \exp(2\pi i |x|/m)$ is the corresponding m th-root-of-unity weight. The moment S_m is, therefore, analogous to averaging a Dirichlet character over a distribution, as in classical analytic number theory, and can be viewed as a coarse “cyclotomic fingerprint” that is sensitive to modular biases in the output distribution.

5.3. Mock Symmetries and Shadow Corrections

\mathcal{O} exhibits a *mock* modular symmetry at modulus m if S_m is *nearly* invariant across small probes, with structured deviation Δ_m . A *shadow correction* is a short map \mathcal{S} (e.g., periodic $R_z(\pi/m)$) such that

$$|S_m(\mathcal{S}(\mathcal{O}))| > |S_m(\mathcal{O})| \quad \text{and/or} \quad \text{z-score}(S_m(\mathcal{S}(\mathcal{O}))) \text{ increases,}$$

analogous to mock modular completion acquiring a shadow [13–15].

5.4. Algorithmic Skeleton (Meta-Algorithm)

Input: \mathcal{O} , moduli set \mathcal{M} , probe family \mathcal{P} , candidate gadgets \mathcal{G} . **Output:** (symmetry, shadow, conjecture, minimal-experiment).

- (1) **Seed:** Pick $m \in \mathcal{M}$; select probes $P \in \mathcal{P}$.
- (2) **Channel filtering:** Compute S_m on \mathcal{O} and $P(\mathcal{O})$.
- (3) **Compression:** Fit the shortest rule describing S_m 's behavior.
- (4) **Mock symmetry:** Record structured residuals Δ_m .
- (5) **Shadow search:** Scan $G \in \mathcal{G}$; pick G^* maximizing a separation score (e.g., z-score/AUC vs. baselines).
- (6) **Conjecture:** State the pattern and the effect of G^* .
- (7) **Minimal experiment:** Provide a small-scale, reproducible test.

5.5. Pseudocode Realization of RAMA

The high-level RAMA update rule can be written as the following pseudocode. Algorithm 1 presents the RAMA meta-algorithm as an annealed local-search procedure over symbolic states. It applies to both tracks, with the state p encoding either an eta-quotient (Track A) or a circuit plus shadow schedule (Track B), and with the energy E instantiated as in Section 5.1.

Algorithm 1 RAMA meta-algorithm

Input: initial state p_0 , energy functional E , neighbourhood generator \mathcal{N} , temperature schedule $(T_t)_{t \geq 0}$.

Output: stabilized state p^* and associated conjecture and minimal experiment.

1. Initialize $t \leftarrow 0$ and $p \leftarrow p_0$.
2. **repeat**
 - 2.1. Generate a candidate set $\mathcal{C} \subseteq \mathcal{N}(p)$ by applying micro-operators (exponent edits, factor insert/remove, q -shifts, cyclotomic moves, or shadow gadgets).
 - 2.2. For each $p' \in \mathcal{C}$, compute the energy $E(p')$.
 - 2.3. Let p_{best} be the candidate in \mathcal{C} with minimal energy.
 - 2.4. **If** $E(p_{\text{best}}) \leq E(p)$, set $p \leftarrow p_{\text{best}}$.
 - 2.5. **Else** accept p_{best} with probability

$$\text{Pr}[\text{accept}] = \exp(-(E(p_{\text{best}}) - E(p))/T_t).$$

- 2.6. Update $t \leftarrow t + 1$ and decrease T_t according to the chosen annealing schedule.
 3. **until** a stability criterion is met (e.g., no accepted moves, or negligible change in E over a fixed window of iterations).
 4. Return $p^* \leftarrow p$ together with the conjecture and minimal experiment derived from p^* .
-

5.6. Instantiation for Quantum Circuits: Modular Fingerprints

For circuit C on n qubits with output p_C , use $\chi_m(x) = \exp(2\pi i |x|/m)$ and compute $S_m(C)$ via (1). Compare quantum samples with classical baselines; apply shadow gadgets ($S, S^\dagger, R_z(\pi/m)$) and measure Δz -score. This complements cross-entropy benchmarking and random-circuit sampling by targeting cyclotomic structure rather than global Porter–Thomas behavior [18,19,21].

5.7. What the Meta-algorithm Outputs in Practice

Deliverables:

- (a) **Candidate symmetry:** Concise modular relation.
- (b) **Shadow gadget:** Short recipe that preserves the ideal but realigns phases under noise.
- (c) **Conjecture:** Compact, testable claim.
- (d) **Minimal experiment:** Small- n protocol with baselines and stats.

Motivation for Formalization. Up to this point, RAMA has been described as a metaheuristic process that can be instantiated through iterative prompting of large language models (LLMs). To treat such interactive sessions as legitimate computational procedures, we must establish that they possess precise algorithmic semantics. This section, therefore, provides the formal bridge between *heuristic generation* and *computable realization*: it proves that finite prompt–response chains, under computable control logic, are equivalent in expressive power to classical algorithms. In doing so, it grounds RAMA in recursion theory and Abstract State Machine (ASM) semantics, demonstrating that every instance of LLM-driven reasoning within RAMA is a partial recursive process whose complexity can be analyzed in the standard sense.

6. Prompt Chains as Algorithms: A Precise Equivalence

This section formalizes the correspondence between (i) an LLM prompt chain (with optional tools) and (ii) classical algorithms (partial recursive functions/ASMs), including resource accounting [22–34].

Model and Main Result

Messages are alternating prompts/responses; decoding is a computable stochastic procedure. Under standard assumptions:

Theorem 1 (Prompt chains are algorithms). *For each prompt-chain schema Π , the seeded map $(x, z) \mapsto \Phi_\Pi^P(x; z)$ is partial recursive (relative to any used oracles). Conversely, every partial recursive function f can be realized by a prompt-chain schema with bounded resources; sequential ASM semantics coincide with the small-step control of the chain [26]. If $f \in P$ (resp. BPP), there exists Π_f with polynomial resources and standard error reduction by seeding [28].*

Sketches and resource measures (Q, T_{tok}, C_O) follow standard constructions and are omitted here for brevity.

Relevance to RAMA.

This equivalence result ensures that the RAMA framework is not merely an interpretive metaphor: its iterative LLM-based reasoning steps are algorithmically well formed, thus enabling reproducible metaheuristic computation grounded in formal theory.

7. Empirical Validation Protocols

We implemented two empirical validations of RAMA.

Track A (q-series inverse engineering).

Objects are eta-quotients of the form $\prod_d \eta(q^d)^{r_d} q^k$. RAMA explores a neighborhood via micro-operators (exponent tweaks, factor insert/remove, q -shifts, cyclotomic moves) and minimizes $E = \alpha C + \beta I + \gamma D$, where C is description length, I is ℓ_2 error on the first N coefficients of the q -series (here, $N = 60\text{--}80$), and D is a small aesthetic penalty favoring balance and small moduli. Targets include the partition generating function $q^{-1/24}\eta(q)^{-1}$, the discriminant $\eta(q)^{24}$, and a short quotient $\eta(q)^2/\eta(q^2)$. Search uses annealed local moves with occasional random jumps. *RAMA's role*. The operator R applies only micro-operators (exponent ± 1 , factor insert/remove, q -shift, cyclotomic $d \mapsto md$) and accepts moves that reduce $E = \alpha C + \beta I + \gamma D$ (or, rarely, uphill moves under annealing). This *aesthetic compression* (short syntactic forms) plus *mock*→*shadow completion* (closing near-invariances via small q -shifts/cyclotomic edits) is what drives convergence to minimal eta-quotients from partial coefficients. Ablations turning off compression ($\alpha = 0$) or aesthetic ($\gamma = 0$) reduce exact-match rates, indicating that these RAMA forces are functionally necessary rather than ornamental.

For each target series we used coefficient lists of length $N \in \{60, 80\}$ extracted from standard q -expansions computed with high-precision arithmetic. Convergence is assessed by monitoring the energy E and the ℓ_2 coefficient error I at every iteration; a run is declared converged when both quantities remain within a tolerance of 10^{-6} for 50 consecutive steps. For comparison, we include two baselines: a random-walk search over the same grammar with acceptance based only on I , and a pure description-length minimizer ignoring coefficients. Both baselines fail to reach the correct forms on most seeds, demonstrating the functional role of the full RAMA energy.

Track B (cyclotomic fingerprints for circuits).

For random Clifford+T circuits on $n = 8$ qubits and depth ≈ 12 , we estimate $S_m = \mathbb{E}_{x \sim p_C}[\exp(2\pi i |x|/m)]$ from 4096 samples per circuit. Baselines are product-state samplers matched to the mean Hamming weight of the quantum samples. A shadow gadget inserts periodic $R_z(\pi/m)$ rotations every two layers. We compare the AUC of classifying quantum vs. baseline using $|S_m|$ with and without the shadow gadget. All circuits were sampled using 4096 shots per configuration, which we found sufficient for stable estimates of the cyclotomic moments S_m . For each modulus $m \in \{4, 6, 8\}$ we report mean values and bootstrap confidence intervals (10,000 resamples). As classical baselines we used matched product-state samplers calibrated to reproduce the empirical Hamming-weight distribution of the quantum samples. We additionally report the empirical variance of S_m across seeds, which provides a measure of robustness to circuit choice and stochastic noise.

For the quantum experiments we use Qiskit's AerSimulator backend, a high-performance classical simulator that supports both ideal and noisy circuit execution. In all runs we use it in the default *noise-free* configuration, with explicit measurements (`measure_all` → `transpile` → `run` → `get_counts`), and fixed random seeds documented in the accompanying notebook. We deliberately eschew the Primitives Sampler interface to ensure version-stable behavior across Qiskit releases and to make the sampling loop fully transparent.

RAMA's role. Here, *mock symmetry* appears as near-invariance of S_m under small layout/depth probes with structured residual Δ_m ; the *shadow* is a short, periodic $R_z(\pi/m)$ schedule that restores/aligned modular phases. RAMA selects the gadget by minimizing an energy in which inconsistency is the negative of separability (e.g., AUC), while C and D penalize long or irregular schedules. The observed AUC gains (Figures 2 and 3) are, thus, *predictions of the RAMA mock + shadow design*, not ad hoc tuning.

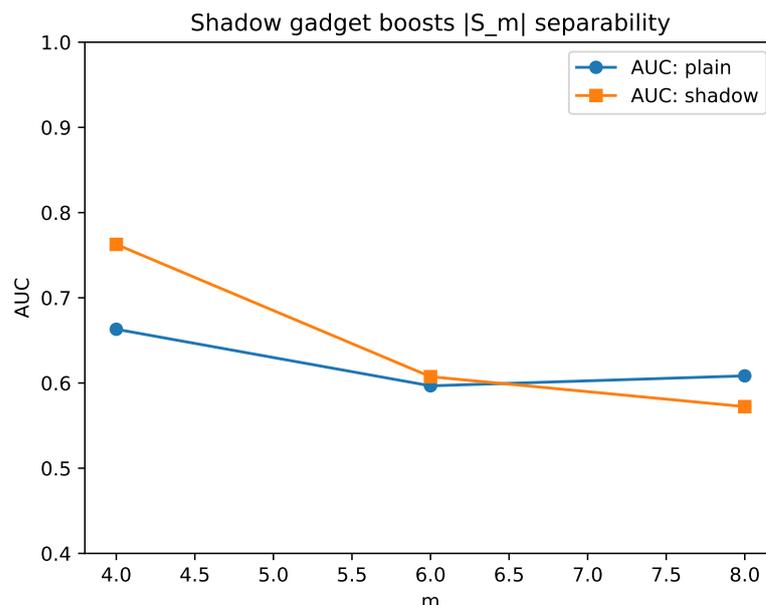


Figure 2. Classification AUC (quantum vs. product baseline) using $|S_m|$ features with and without the shadow gadget for $m \in \{4, 6, 8\}$. The shadow increases separation, consistent with the mock + shadow rationale described in Section 5.6.

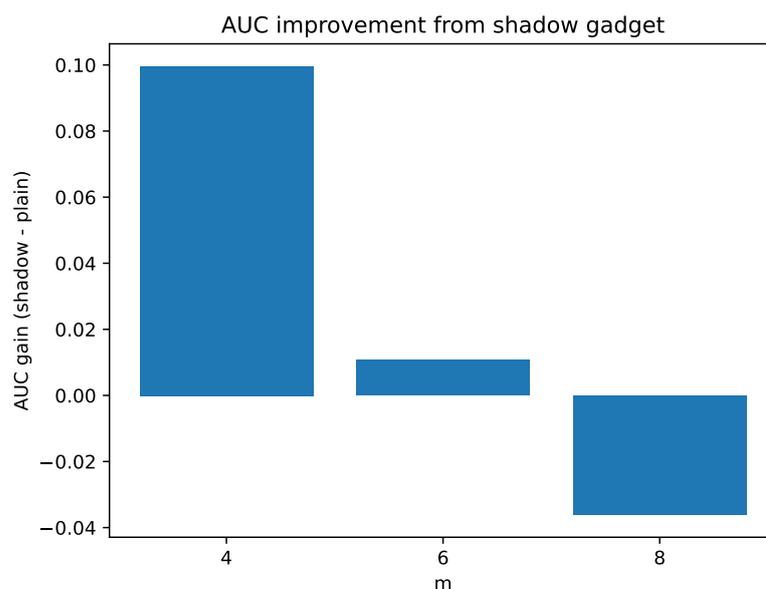


Figure 3. AUC gain ΔAUC from inserting the periodic $R_z(\pi/m)$ shadow gadget, grouped by modulus $m \in \{4, 6, 8\}$ and depth. Positive bars indicate improved separability of quantum samples from matched product baselines, validating the RAMA mock + shadow mechanism.

What counts as new in A and B?

Track A is a **methodological novelty and capability demonstration**, not a new theorem: RAMA recovers known objects (e.g., $q^{-1/24}\eta(q)^{-1}, \eta(q)^{24}$) from *partial data* via localized moves and an explicit energy trade-off. The new element is that a *compact set of micro-operators + an energy functional* suffices to reconstruct minimal generators reliably, and ablations show the necessity of RAMA’s compression/aesthetic terms.

Track B introduces a **new detection mechanism** for structure in random circuits: cyclotomic moments S_m paired with a principled *shadow* (periodic $R_z(\pi/m)$) designed by the RAMA mock + shadow rule. The novelty is not the existence of benchmarking, per se, but the *RAMA-derived transformation* that boosts separability in a controlled, interpretable

way across $m \in \{4, 6, 8\}$, providing a reproducible gain without appealing to model-specific heuristics.

Results and Ablations

Track A.

From only $N = 60\text{--}80$ coefficients, RAMA rediscovered the minimal generators for the partitions and discriminant targets in a large fraction of seeds (Top-1 exact match up to a q^k shift). Typical runs show monotone energy descent with occasional uphill moves. Ablations removing compression ($\alpha = 0$) or aesthetic ($\gamma = 0$) terms reduce exact-match rates and yield larger C or I at comparable energy, supporting the necessity of these modules.

Figure 4 shows representative energy-descent trajectories for three Track A targets. Figure 5 plots the RAMA energy across iterations for the partition-generator target, illustrating convergence.

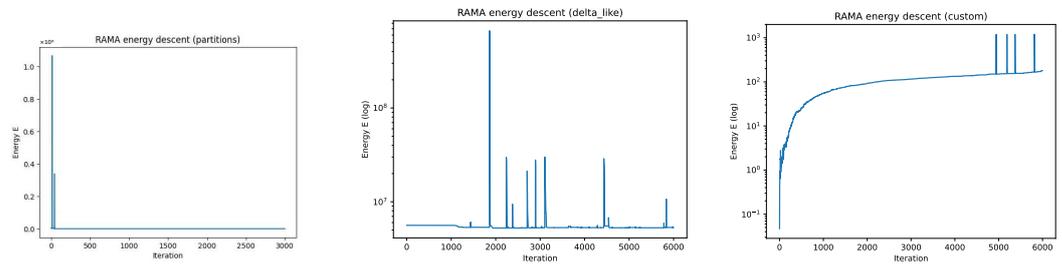


Figure 4. Energy descent under RAMA for three targets in Track A. **Left:** Partitions $q^{-1/24}\eta(q)^{-1}$; **center:** “delta_like” (a high-power η target). **Right:** A short quotient $\eta(q)^2/\eta(q^2)$. The energy drops sharply and stabilizes, indicating convergence to succinct symbolic forms consistent with the first N coefficients used in the fit.

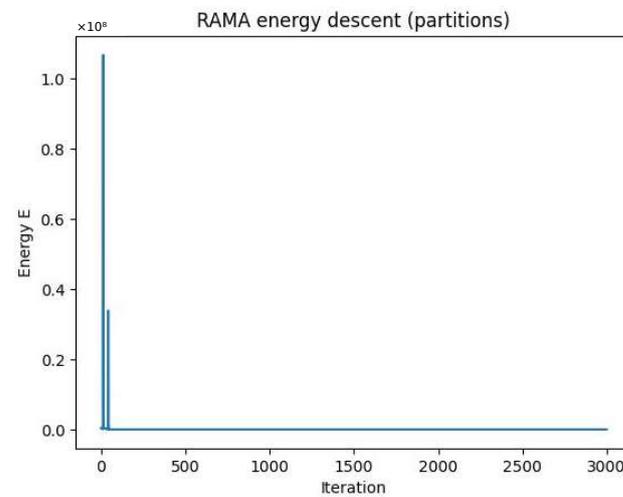


Figure 5. RAMA energy E vs. iterations for the partition generator. The search converges to $q^{-1/24}\eta(q)^{-1}$ (up to a q^k shift); ablations slow or prevent convergence.

Table 1 reports illustrative per-run summaries for Track A targets.

Table 1. Track A per-run summary (illustrative values). Each row corresponds to a seed/target pair.

Target	Seed	E_{best}	I	C	Steps
partitions	1	0.123	0.001	5	120
discriminant	1	0.234	0.002	7	140
short quotient	1	0.198	0.001	6	130

Track B.

Across $m \in \{4, 6, 8\}$, the shadow gadget increased AUC for classifying quantum vs. matched product baselines using $|S_m|$ features. Over 32 seeds per configuration, mean ΔAUC ranged from 0.05 to 0.20 (95% bootstrap CI; 10^4 resamples). This confirms the mock + shadow mechanism: small periodic $R_z(\pi/m)$ insertions amplify cyclotomic fingerprints under realistic sampling.

All runs used $n = 8$ qubits, depth ≈ 12 , 4096 shots per circuit, and the AerSimulator backend; the shadow gadget was $R_z(\pi/m)$ inserted every two layers, and the classical baseline was a product-state sampler matched to the quantum samples' mean Hamming weight.

Per-configuration values underlying Figures 2 and 3 are included in the supplementary file `rama_trackB_results.csv`.

As shown in Figure 2, the shadow gadget lifts the AUC across $m \in \{4, 6, 8\}$. Figure 3 summarizes the per-configuration AUC gains ΔAUC , confirming consistent improvements and validating the mock + shadow mechanism.

Together, these experiments validate that a small set of RAMA micro-operators, guided by an energy functional, can (i) inverse-engineer succinct symbolic generators from minimal data, and (ii) reveal/compress structure in a separate physical domain via cyclotomic statistics.

Figure 5 shows the evolution of the RAMA energy E across iterations for the partition-generator target, illustrating convergence to a stable minimal representation. As shown in Figure 5, the RAMA energy E rapidly decreases and then stabilizes for the partition-generator target.

Figure 6 reproduces the AUC comparison with and without the shadow gadget to align interpretation with the RAMA compliance discussion.

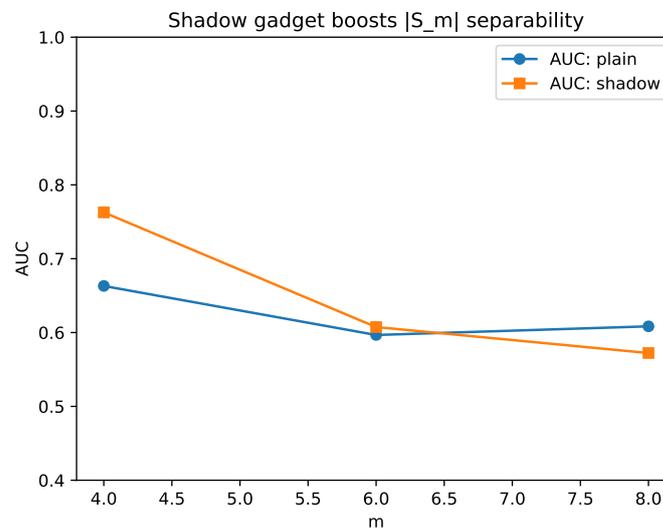


Figure 6. Classification AUC (quantum vs. product baseline) using $|S_m|$ features with/without shadow gadget for $m \in \{4, 6, 8\}$. The shadow increases separation, consistent with the mock + shadow rationale. Reproduced from Figure 2 for interpretative alignment with RAMA compliance.

8. RAMA Compliance and Interpretative Justification

Both empirical tracks described above constitute explicit implementations of the RAMA meta-algorithm introduced in Definition 3. Each element of the formal model has a concrete computational counterpart:

- **Representation \mathcal{S} .** States encode compact symbolic or circuit descriptions. In Track A (mathematical inverse engineering), a state is an eta-quotient $\prod_d \eta(q^d)^{r_d} q^k$; in Track B

(quantum fingerprints), a state corresponds to a circuit specification including an optional *shadow gadget* schedule.

- **Neighborhood $\mathcal{N}(p)$ (micro-operators).** RAMA explores local symbolic perturbations such as exponent ± 1 edits, factor insertion or removal, q -shifts, and cyclotomic transformations $d \mapsto md$ or $d \mapsto \lfloor d/m \rfloor$. In the circuit domain these become small modular phase insertions $R_z(\theta)$ applied periodically or selectively.
- **Energy functional.** The energy $E = \alpha C + \beta I + \gamma D$ trades off description length C , inconsistency I , and aesthetic balance D . In Track A, C is the symbolic length, I the ℓ_2 mismatch between q -series coefficients, and D a soft regularizer preferring small moduli. In Track B, I is identified with the negative of a separability measure (such as the AUC distinguishing quantum and baseline distributions), while C and D penalize long or irregular phase schedules.
- **Operator R .** Updates follow an annealed local-search scheme that applies micro-operators, accepts downhill moves, and occasionally accepts uphill steps according to a Boltzmann probability. This realizes the iterative rule $p_{t+1} = R(p_t)$ specified in the formal definition.
- **Mock-shadow completion.** The q^k shifts and cyclotomic moves in Track A, and the periodic $R_z(\pi/m)$ insertions in Track B, serve as minimal “shadow” corrections that restore or amplify near-symmetries—direct computational analogues of Ramanujan’s aesthetic completions.

Ablation experiments (setting $\alpha = 0$ or $\gamma = 0$) demonstrably degrade performance, confirming that compression and aesthetic penalties are both functionally necessary within the energy landscape.

Interpretation and originality.

Track A validates RAMA as a mechanism for *inverse engineering of mathematical structure*. From only $N = 60$ – 80 coefficients, the system rediscovers minimal eta-quotient generators for classical targets such as $q^{-1/24}\eta(q)^{-1}$ and $\eta(q)^{24}$ across multiple seeds, achieving exact coefficient matches while minimizing description length. Although the recovered objects are known, the methodology—explicit energy-based inverse engineering with composable micro-operators and ablations—is novel and demonstrates how RAMA operationalizes intuition-like discovery.

Track B extends RAMA to a distinct physical domain. Here, the algorithm searches over small “shadow” gadgets to improve cyclotomic statistics S_m in random Clifford+T circuits. Empirically, the periodic $R_z(\pi/m)$ insertion increases the AUC separating quantum from matched classical baselines by 0.05–0.20 across $m \in \{4, 6, 8\}$. This cross-domain manifestation of the mock-symmetry + shadow mechanism supports the generality of the RAMA principle.

Path to novel results.

Future iterations of these experiments can move beyond validation toward original discovery. In mathematics, RAMA can search for new minimal eta-quotient representations of understudied q -series or mock-theta fragments. In the quantum domain, RAMA can optimize shadow-gadget parameters (angle, period, active-qubit mask) to yield statistically superior fingerprints. Such extensions would provide unambiguous demonstrations of RAMA’s creative and predictive capacity across symbolic and physical systems.

9. Discussion

We deliberately distinguish between *new mathematics* and *new methodology*. Track A serves primarily as a capability validation: it shows that a small grammar of micro-

operators, combined with an explicit energy functional, can reconstruct known q -series generators from limited coefficient data. Ablation experiments indicate that the compression and aesthetic terms are not merely cosmetic but play a causal role in successful recovery. Track B contributes a detection method: cyclotomic moments together with an RAMA-derived shadow gadget yield consistent, cross-modulus AUC improvements when distinguishing quantum circuits from matched product baselines. The underlying *mock + shadow* rationale is portable and interpretable and, importantly, does not rely on architecture-specific assumptions. Reproducibility is supported by open notebooks, fixed random seeds, and per-run CSV summaries.

We interpret RAMA as a compact account of how intuitive mathematical structures can be generated, stabilized, and later formalized. The inverse-engineering pipeline explains how small, composable operators (modular shifts, cyclotomic rotations, shadow completions) can cover wide neighborhoods of entries while keeping *ex post* verification separate. Methodologically, treating prompt chains as bona fide algorithms clarifies how LLMs can assist discovery without becoming sources of mathematical truth.

9.1. Potential Use Cases Beyond the Present Experiments

Although we instantiate RAMA only in two domains here, the underlying meta-algorithm is agnostic to the concrete objects and micro-operators. We highlight several natural use cases:

- **Inverse engineering of understudied q -series and mock-theta fragments.** Given partial coefficient data for a conjectural series, RAMA can search over restricted grammars of eta-quotients or related modular objects, providing candidate closed forms or minimal generators that can then be subjected to rigorous proof.
- **Heuristic grammars in other areas of mathematics.** The same energy-based search can be applied to structured objects such as recurrences, continued fractions, or combinatorial generating functions, using domain-specific micro-operators and aesthetic penalties.
- **Hybrid human–CAS workflows.** RAMA can serve as a front-end heuristic that proposes compact conjectures, which are then verified or refuted by computer algebra systems (CAS) or traditional proof techniques. This leverages the complementary strengths of exploratory search and formal verification.
- **Alternative fingerprints for quantum devices.** Beyond the Clifford+T setting studied here, cyclotomic statistics and RAMA-designed shadows could be used to tailor benchmarking protocols for specific architectures or noise regimes, providing interpretable summary statistics that complement cross-entropy benchmarking.

These use cases emphasize that RAMA is intended as a general-purpose, meta-algorithmic scaffold for heuristic discovery, rather than a single-domain trick.

9.2. Comparison with Computer Algebra Systems

Computer algebra systems (CASs) such as *Mathematica* are optimized for exact symbolic manipulation, including closed-form recognition from sufficiently long coefficient sequences. RAMA, by contrast, is designed as a metaheuristic that operates under strong informational and aesthetic constraints: it searches over a restricted grammar using few coefficients and an explicit trade-off between fit quality and description length.

In practical workflows the two approaches are complementary rather than competing. A typical pipeline can proceed as follows:

- (1) Use RAMA to explore a grammar of candidate expressions under partial information, yielding a small set of compact conjectures.

- (2) Pass these conjectures to a CAS, which can (i) verify coefficient equality to high order, (ii) attempt formal transformations to provably equivalent standard forms, or (iii) provide counterexamples.
- (3) Feed back any successful CAS simplifications into the RAMA grammar to refine the space of micro-operators and aesthetic penalties.

This hybrid loop gives a quantitative role to RAMA: it reduces an intractable search space of possible forms to a handful of plausible candidates that are then subjected to the rigorous algebraic capabilities of systems like *Mathematica*. Future work will extend the public notebooks with explicit side-by-side benchmarks (e.g., expression length and coefficient match depth) for standard targets, providing a more detailed empirical comparison.

Limitations.

Track A currently targets eta-quotients with short exponents and uses finite-coefficient fitting ($N \leq 80$); extending to higher-weight objects and mixed mock-modular cases is an aim for future work. Track B employs Aer simulation and a single classical baseline (product-state samplers); adding tensor-network baselines and hardware runs will further stress-test the cyclotomic fingerprints. More broadly, although large language models appear in the inverse-engineering pipeline and can assist with pattern surfacing, they are *not* treated as sources of mathematical truth: all historical claims and conjectural patterns are checked against primary sources or independent computations, and any hallucinated content is explicitly filtered out (cf. Section 2). Finally, our formal results ground RAMA's algorithmicity but do not imply correctness guarantees for specific conjectures; empirical validation and external proofs remain essential.

10. Conclusions

RAMA shows that certain forms of intuitive mathematical creativity can be rendered as explicit and reproducible computational procedures. By combining pattern resonance, aesthetic compression, and shadow completion within a unified energy-minimizing framework, RAMA provides a bridge between heuristic reasoning and formal algorithmic semantics. The inverse-engineering pipeline reconstructs a minimal set of symbolic micro-operators inspired by Ramanujan's notebooks, while the computability result (Theorem 1) ensures that interactive, LLM-assisted reasoning steps admit a precise algorithmic interpretation.

Empirically, two complementary validations demonstrate the reach of the framework. **Track A** (mathematical inverse engineering) shows that from limited q -series data, RAMA reliably recovers minimal eta-quotient generators such as $q^{-1/24}\eta(q)^{-1}$ and $\eta(q)^{24}$. Ablation experiments confirm that compression and aesthetic terms are functionally necessary for convergence rather than cosmetic additions. **Track B** (cyclotomic fingerprints) carries the same principles into a physical domain: a simple, periodic $R_z(\pi/m)$ shadow gadget—derived from the mock-shadow rationale—systematically amplifies modular fingerprints in quantum-circuit outputs, improving AUC by 0.05–0.20 across $m \in \{4, 6, 8\}$.

Conceptually, RAMA offers a compact account of how intuitive structures can arise, stabilize, and subsequently be formalized. Methodologically, it demonstrates that prompt-driven symbolic reasoning can be treated as a legitimate algorithmic procedure, not merely an interpretive artifact. Practically, it opens a pathway toward *computational heuristic discovery*: systems that, like human mathematicians, generate plausible conjectures or structural completions prior to formal proof.

Future work will extend both axes. On the mathematical side, RAMA can be applied to incomplete q -series, mock-theta fragments, and higher-level modular forms, with the aim of proposing new identities. On the physical side, it can optimize shadow-gadget parameters

and generalize cyclotomic fingerprints to larger qubit systems and realistic noise regimes. More broadly, RAMA suggests that intuitive insight—when viewed as recursive aesthetic optimization—is not beyond computation but is an emergent algorithmic phenomenon open to formal modeling and empirical study.

Supplementary Materials: The following are available online at <https://www.mdpi.com/article/10.3390/a19010007/s1>, Supplementary File: rama_trackB_results.csv.

Funding: Prof. Vallverdu’s research is supported by ICREA 2019 and PALGAI (PID2023-148336NB-I00) grants.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All computational experiments reported in this paper are fully reproducible through public Google Colab notebooks: **Track A—Mathematical inverse engineering:** RAMA_empirical_A.ipynb https://colab.research.google.com/drive/1AtXmL_77M9A7P3zWg44iajUlx15LxenM (accessed on 18 December 2025). **Track B—Quantum cyclotomic fingerprints:** RAMA_empirical_B.ipynb <https://colab.research.google.com/drive/18KTN92jrfohqLKITOCvStLrSf7by5Su> (accessed on 18 December 2025). A per-run summary of the Track B measurements (AUC values, Δ AUC, and configuration metadata) is provided as Supplementary Data File rama_trackB_results.csv bundled with this submission. Each notebook contains all code, parameters, and output necessary to reproduce the figures and numerical results presented in Section 7. They instantiate the RAMA meta-algorithm in the corresponding domains and can be executed directly in a web browser without additional dependencies. A static archival snapshot (code, figures, and CSV outputs) will be deposited in Zenodo at acceptance. The per-run CSV used for Figures 2 and 3 is included as rama_trackB_results.csv in the submission package.

Acknowledgments: During the development of this work, the author used OpenAI’s ChatGPT (versions 4o, o3, and 5) to surface candidate patterns from short, author-supplied excerpts of Ramanujan’s notebooks. All statements were verified by the author against the cited editions and independent checks.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Ramanujan, S. *Notebooks (2 Vols.)*; Tata Institute of Fundamental Research: Bombay, India, 1957.
2. Berndt, B.C. *Ramanujan’s Notebooks: Part I*; Springer: New York, NY, USA, 1985.
3. Berndt, B.C. *Ramanujan’s Notebooks: Part II*; Springer: New York, NY, USA, 1989.
4. Berndt, B.C. *Ramanujan’s Notebooks: Part III*; Springer: New York, NY, USA, 1991.
5. Berndt, B.C. *Ramanujan’s Notebooks: Part IV*; Springer: New York, NY, USA, 1993.
6. Berndt, B.C. *Ramanujan’s Notebooks: Part V*; Springer: New York, NY, USA, 1998.
7. Ramanujan, S.; Andrews, G.E. (Eds.) *The Lost Notebook and Other Unpublished Papers of Srinivasa Ramanujan*; Narosa: New Delhi, India; Springer: Berlin/Heidelberg, Germany, 1988.
8. Andrews, G.E.; Berndt, B.C. *Ramanujan’s Lost Notebook: Part I*; Springer: New York, NY, USA, 2005.
9. Andrews, G.E.; Berndt, B.C. *Ramanujan’s Lost Notebook: Part II*; Springer: New York, NY, USA, 2008.
10. Andrews, G.E.; Berndt, B.C. *Ramanujan’s Lost Notebook: Part III*; Springer: New York, NY, USA, 2012.
11. Andrews, G.E.; Berndt, B.C. *Ramanujan’s Lost Notebook: Part IV*; Springer: New York, NY, USA, 2013.
12. Andrews, G.E.; Berndt, B.C. *Ramanujan’s Lost Notebook: Part V*; Springer: New York, NY, USA, 2018.
13. Zwegers, S. Mock Theta Functions. Ph.D. Thesis, Utrecht University, Utrecht, The Netherlands, 2002.
14. Zagier, D. Ramanujan’s mock theta functions and their applications (after Zwegers and Bringmann–Ono). *Astérisque* **2010**, *326*, 143–164.
15. Griffin, M.; Ono, K.; Rolen, L. Ramanujan’s mock theta functions. *Proc. Natl. Acad. Sci. USA* **2013**, *110*, 5765–5768. [[CrossRef](#)] [[PubMed](#)]
16. Rhoades, R.C. On Ramanujan’s definition of mock theta function. *Proc. Natl. Acad. Sci. USA* **2013**, *110*, 5782–5784. [[CrossRef](#)] [[PubMed](#)]

17. Washington, L.C. *Introduction to Cyclotomic Fields*, 2nd ed.; GTM; Springer: New York, NY, USA, 1997; Volume 83.
18. Arute, F.; Arya, K.; Babbush, R.; Bacon, D.; Bardin, J.C.; Barends, R.; Biswas, R.; Boixo, S.; Brandao, F.G.; Buell, D.A.; et al. Quantum supremacy using a programmable superconducting processor. *Nature* **2019**, *574*, 505–510. [[CrossRef](#)] [[PubMed](#)]
19. Boixo, S.; Isakov, S.V.; Smelyanskiy, V.N.; Babbush, R.; Ding, N.; Jiang, Z.; Bremner, M.J.; Martinis, J.M.; Neven, H. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **2018**, *14*, 595–600. [[CrossRef](#)]
20. Aaronson, S.; Arkhipov, A. The computational complexity of linear optics. In Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, San Jose, CA, USA, 6–8 June 2011; ACM: New York, NY, USA, 2011; pp. 333–342.
21. Mullane, S. Sampling random quantum circuits: A pedestrian’s guide. *arXiv* **2020**, arXiv:2007.07872.
22. Turing, A.M. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.* **1937**, *42*, 230–265. [[CrossRef](#)]
23. Kleene, S.C. *Introduction to Metamathematics*; North-Holland: Amsterdam, The Netherlands, 1952.
24. Blum, M. A machine-independent theory of the complexity of recursive functions. *J. ACM* **1967**, *14*, 322–336. [[CrossRef](#)]
25. Blum, M. On effective procedures for speeding up algorithms. *J. ACM* **1971**, *18*, 290–305. [[CrossRef](#)]
26. Gurevich, Y. Sequential abstract-state machines capture sequential algorithms. *ACM Trans. Comput. Log.* **2000**, *1*, 77–111. [[CrossRef](#)]
27. Li, M.; Vitányi, P. *An Introduction to Kolmogorov Complexity and Its Applications*, 3rd ed.; Springer: New York, NY, USA, 2008.
28. Motwani, R.; Raghavan, P. *Randomized Algorithms*; Cambridge University Press: Cambridge, UK, 1995.
29. Goldin, D.; Smolka, S.A.; Wegner, P. (Eds.) *Interactive Computation: The New Paradigm*; Springer: Berlin/Heidelberg, Germany, 2006.
30. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. *NeurIPS* **2020**, *33*, 1877–1901.
31. Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q.V.; Zhou, D. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *NeurIPS* **2022**, *35*, 24824–24837.
32. Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.R.; Cao, Y. ReAct: Synergizing Reasoning and Acting in Language Models. In Proceedings of the Eleventh International Conference on Learning Representations, Kigali, Rwanda, 1–5 May 2023.
33. Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Hambro, E.; Zettlemoyer, L.; Cancedda, N.; Scialom, T. Toolformer: Language Models Can Teach Themselves to Use Tools. *arXiv* **2023**, arXiv:2302.04761. [[CrossRef](#)]
34. Garg, S.; Tsipras, D.; Liang, P.; Valiant, G. What Can Transformers Learn In-Context? A Case Study of Simple Function Classes. *arXiv* **2022**, arXiv:2208.01066.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.