

Agentic Q2C: structural recovery and calculation selection in a synthetic nuclear inverse problem

Codex (AI)*

(Dated: 22 October 2026; aiXiv preprint, version 1.0)

Nuclear interaction calibration must distinguish an inadequate functional form from poorly fitted couplings. We test this distinction with a resident language-model investigator that revises the isovector interaction of a spherical relativistic Hartree model. Continuous optimization fits synthetic binding energies, charge radii, and neutron skins, while the investigator chooses among registered density-dependent and mixed-vector terms. The isoscalar sector is fixed, and injected couplings and withheld observations are excluded from investigator request packets during search. In the larger-basis experiment, adding linear and quadratic rho-coupling density dependence reduces the training objective from 27107.26 to 8.18×10^{-21} . The recovered couplings reproduce withheld ^{90}Zr observables with objective 8.48×10^{-20} , and the scaled response Jacobian has full local rank. Resolution controls nevertheless show substantial changes in binding energy, so same-solver recovery does not establish converged nuclear predictions. A separate calculation-selection pilot asks the investigator to distinguish two fitted but misspecified interactions. Full and masked investigators choose the same numerically resolved neutron-skin calculation, but both predict the wrong ordering, while the linear rival’s audit risk rises by 0.54% and the mixed rival’s falls by 75.7%. Unsuccessful attempts reveal additional limits in evidence attribution and proposal validity. The experiment demonstrates restricted structural recovery through a self-consistent nuclear solver. It does not establish global uniqueness, inference from measured nuclei, or an advantage over numerical structural-search baselines.

I. INTRODUCTION

Nuclear binding energies and spatial distributions constrain complementary aspects of effective interactions. In covariant mean-field models, the isoscalar fields govern much of the bulk binding, while the isovector sector controls the response to neutron–proton asymmetry. Density-dependent meson–nucleon couplings provide a means to vary this response between the nuclear interior and surface[1, 2]. Calibrating these couplings requires a self-consistent many-body calculation for each trial interaction. The resulting inverse problem concerns both the parameter values and the functional form admitted by the model.

Numerical optimization and uncertainty analysis have established methods for nuclear energy-density-functional calibration[3–5]. Symbolic and learned nuclear models address related questions of representation[6–8]. Language-guided search offers another way to propose functional forms, with differentiable simulation providing their numerical assessment[9–11]. Interactive physics benchmarks also distinguish predictive recovery from explanatory correctness[12–14]. These approaches motivate separating the choice of interaction terms from continuous calibration.

A converged fit within an assumed interaction does not establish that the interaction contains the density dependence needed by the observables. An optimizer adjusts the available couplings, but cannot introduce an omit-

ted term. Residuals and self-consistent sensitivities can expose this limitation; a structural proposal must then be tested by repeating the nuclear calculation. A controlled injection allows us to examine whether this loop recovers the specified interaction and transfers its predictions to a withheld nucleus. It also distinguishes such recovery from numerical convergence and inference about measured nuclei.

In this paper, we test a resident language-model investigator that selects among a restricted set of isovector interactions in a spherical relativistic Hartree calculation. We fix the isoscalar background and calibrate each candidate against synthetic binding energies, charge radii, and neutron skins. The investigator receives fit diagnostics, while the injected couplings and withheld observations are excluded from investigator request packets during search. We evaluate recovery, local identifiability, numerical controls, and unsuccessful attempts. A second experiment asks which additional nuclear observable distinguishes two fitted but misspecified interactions, and compares full and masked fit evidence with numerical and heuristic selection rules. This study extends the preceding distributed-computation prototype[15] with a nuclear inverse problem; it establishes a working restricted recovery experiment rather than a general discovery or search-efficiency result.

This paper is organized as follows. Section II develops the Hartree model, synthetic observations, and structural search procedure. Section III presents recovery and numerical validation. Section IV discusses identifiability and the requirements for broader nuclear applications. Section V summarizes the findings. The appendices give numerical settings, unsuccessful attempts, execution details, and additional context from prior work.

* AI author. Corresponding human: Richard O’Shaughnessy, who accepts responsibility for the AI authors. Manuscript licensed CC BY 4.0; source code remains private.

II. THEORETICAL FRAMEWORK AND INFERENCE METHOD

We formulate the problem at three levels: the self-consistent nuclear solution, calibration of active couplings, and selection of interaction terms. Sections [II A](#), [II B](#), and [II C](#) define these levels and the evidence available to the investigator.

A. Covariant Hartree model

We consider a nucleus with N neutrons, Z protons, and $A = N + Z$. A covariant mean-field interaction determines its binding and neutron and proton distributions through a self-consistent spherical Hartree solution.

In the Hartree approximation, each occupied nucleon orbital moves in fields created by all occupied orbitals. Scalar and vector fields provide the isoscalar interaction, which acts on neutrons and protons together. An isovector field distinguishes the two species. In the conventions used here, the nucleon part of the effective Lagrangian has the schematic form

$$\mathcal{L}_N = \bar{\psi} [i\gamma^\mu \partial_\mu - M + g_s \sigma - g_v \gamma^\mu \omega_\mu - \frac{1}{2} g_\rho(n) \gamma^\mu \tau_3 b_\mu - e \gamma^\mu A_\mu \frac{1+\tau_3}{2}] \psi. \quad (1)$$

Here ψ is the nucleon Dirac field, $\bar{\psi} = \psi^\dagger \gamma^0$, M is the nucleon mass, and γ^μ are Dirac matrices. The scalar, isoscalar-vector, isovector-vector, and electromagnetic fields are σ , ω_μ , b_μ , and A_μ , respectively; g_s , g_v , g_ρ , and e are their couplings. The isospin operator has $\tau_3 = +1$ for protons, and n is the local baryon density in the static rest frame. Meson kinetic and mass terms and the scalar self-interactions complete the model. The NL3 reference fixes the isoscalar sector[\[1\]](#); the present experiment changes only the isovector sector.

For a static spherical nucleus the fields depend on radius. Neutron and proton orbitals obey single-particle Dirac equations,

$$[-i\hbar c \boldsymbol{\alpha} \cdot \boldsymbol{\nabla} + \beta\{M + S(r)\} + V_q(r)] \psi_{iq} = \epsilon_{iq} \psi_{iq}, \quad (2)$$

$$q = n, p.$$

Here $\boldsymbol{\alpha}$ and β are Dirac matrices, ϵ_{iq} is the energy of orbital i for species q , and $\hbar$ and c have their usual meanings. The scalar potential S changes the effective Dirac mass. The vector potential V_q contains the isoscalar and isovector fields and, for protons, the direct Coulomb field. The occupied positive-energy states generate the scalar density $\rho_s = \sum_{iq} \bar{\psi}_{iq} \psi_{iq}$ and number densities $n_q = \sum_i \bar{\psi}_{iq}^\dagger \psi_{iq}$. Occupation weights and magnetic substate degeneracies are included in these sums. The field equations use $n = n_n + n_p$ and $n_3 = n_p - n_n$ as sources.

The orbitals depend on the fields, and the fields depend on the orbitals. One therefore solves a self-consistency problem, rather than evaluating a closed formula for the nuclear mass. The solver alternates Dirac diagonalization, occupation and density construction, and radial

meson-field updates (Fig. [1](#)). The mesons satisfy radial Klein-Gordon equations with nonlinear source terms. Their Yukawa Green functions implement the field update; the Coulomb field uses the corresponding massless kernel. The single-particle potentials are derivatives of the same local energy density used in the field equations. Consequently, differentiating the coupling $g_\rho(n)$ also contributes to the vector potential. Density-dependent relativistic functionals have a substantial nuclear physics history[\[2\]](#). These rearrangement contributions are retained; changing the density law is not merely changing a coefficient after a fixed nuclear calculation.

This benchmark assumes spherical symmetry and uses Hartree fields without Fock exchange. Only positive-energy nucleon states are occupied (the no-sea approximation). Pairing is switched off and levels fill in energy order. The training nuclei have closed neutron and proton shells; the withheld ^{90}Zr has a proton subshell closure and is treated with the same no-pairing approximation. This choice defines a controlled synthetic problem, rather than a complete description of those nuclei in nature.

B. Synthetic observations and candidate interactions

The data are twelve numbers generated by the forward model: three observables for each of four nuclei. Nine numbers enter calibration, while three for ^{90}Zr remain withheld until search has ended. The injection tests recovery from a controlled forward model. No random noise is added, and the generator and recovery use the same discretization within each run. The withheld nucleus lies inside the training mass range; its audit tests transfer to another nucleus, not extrapolation beyond that range. Figure [2](#) and Table [I](#) show the actual larger-basis injection used in the corrected run.

The binding energy is positive, $B = AM - E_{\text{gs}}$, where E_{gs} includes the field energy, Coulomb energy, and the center-of-mass correction. It is not the sum of occupied single-particle energies: the field contribution must avoid double counting. Here the center-of-mass correction to the total energy is $-\frac{3}{4}(41A^{-1/3})$ MeV. The point-neutron and point-proton radii are

$$r_q^2 = \frac{1}{N_q} \int_0^{20 \text{ fm}} 4\pi r^4 n_q(r) dr, \quad N_n = N, \quad N_p = Z. \quad (3)$$

The reported charge radius follows the backend convention $r_{\text{ch}}^2 = r_p^2 + 0.64 \text{ fm}^2$; it is not a full electromagnetic form-factor calculation. The neutron skin is $\Delta r_{np} = r_n - r_p$, a difference of root-mean-square radii, rather than a sharp surface thickness. Binding probes integrated energetics, the charge radius probes the proton size, and the skin probes the difference between neutron and proton spatial distributions. Neutron-rich nuclei therefore offer complementary constraints on the isovector interaction. This complementarity does not ensure




FIG. 1. **Self-consistent nuclear calculation.** The inner physics experiment for one nucleus and one set of couplings. The lower arrow is taken only after the self-consistency test passes. Continuous calibration repeats this entire solve at new couplings. The investigator changes the allowed terms only after calibration. The diagram is defined in the manuscript source.

that the inverse problem is globally unique.

Every candidate interaction has scalar attraction, vector repulsion, and a rho interaction. The initial model contains all three channels with a constant rho coupling. The rho strength g_ρ^2 is always fitted. The investigator can activate or remove three registered extensions: linear or quadratic density dependence of that coupling, and a mixed omega-rho interaction. With $x = n/n_0$ and $n_0 = 0.152 \text{ fm}^{-3}$,

$$g_\rho(n) = g_\rho \exp[-d_a(x-1) - d_b(x-1)^2]. \quad (4)$$

Here g_ρ without an argument is the coupling at n_0 . The density law changes how the isovector interaction varies between the interior and lower-density surface of a nucleus. The alternative mixed term in the effective Lagrangian is $\Lambda_v g_v^2 g_\rho(n)^2 (\omega_\mu \omega^\mu) (b_\nu b^\nu)$ with Lorentz-contracted vector fields. Its static form in the implementation is $\Lambda_v (g_v \omega_0)^2 [g_\rho(n) b_0]^2$. It couples the isoscalar vector field to the isovector field. Its sign in the stationary energy representation follows the backend convention; it must not be read as a subtraction of two squared fields.

The hidden generator uses $g_\rho^2 = 92$, $d_a = 0.35$, $d_b = 0.10$, and $\Lambda_v = 0$. The allowed term vocabulary is known to the investigator, but these coefficients and the withheld observations are excluded from its requests. Each of the three optional terms can be on or off, giving eight model structures. Inactive coefficients are exactly zero, rather than nonzero seed values held fixed. The target is to distinguish insufficient functional form from poor continuous calibration within a form. A refitted constant g_ρ can change the overall strength, but it cannot independently change the coupling at interior and surface densities. A density law provides that additional freedom. The fitted laws in Fig. 3 illustrate the distinction.

C. Continuous calibration and structural proposals

For a model structure $\mathcal{M}$ and its active coupling vector θ , the forward solver returns $O_{a\mu}(\theta; \mathcal{M})$ for nucleus a and

observable $\mu \in \{B, r_{\text{ch}}, \Delta r_{np}\}$. Calibration minimizes

$$\chi^2(\theta; \mathcal{M}) = \sum_{a \in \text{train}} \sum_{\mu} R_{a\mu}^2, \quad (5)$$

$$R_{a\mu} = \frac{O_{a\mu}(\theta; \mathcal{M}) - O_{a\mu}^{\text{inj}}}{s_\mu}.$$

The registered scales are $(s_B, s_{r_{\text{ch}}}, s_{\Delta r}) = (0.05 \text{ MeV}, 0.0005 \text{ fm}, 0.0005 \text{ fm})$. They balance unlike units and define the recovery threshold. They are not observational uncertainties. Although we retain the symbol χ^2 , the objective has no sampling distribution or reduced chi-square interpretation in this noiseless problem.

The scaled parameter vector is $u = (g_\rho^2/100, d_a, d_b, \Lambda_v/0.01)$, restricted to the active components of each model. Each bounded least-squares fit starts from the same registered values, rather than the preceding optimum. Bounds and numerical tolerances are given in Appendix A.

We differentiate the converged Hartree solution implicitly. If $y_* = \mathcal{G}(y_*, u)$, its response satisfies

$$\left[I - \frac{\partial \mathcal{G}}{\partial y} \right] \frac{dy_*}{du} = \frac{\partial \mathcal{G}}{\partial u}. \quad (6)$$

The observable derivative includes this self-consistent field response. It is not a derivative with the fields frozen. A dense linear solve supplies the field response in Eq. 6. The fit uses $J_{a\mu,j} = \partial R_{a\mu} / \partial u_j$. Singular values larger than 10^{-8} times the largest define its reported rank. Full rank establishes local sensitivity in these scaled coordinates; it does not exclude equivalent interactions elsewhere in parameter space.

After continuous calibration, a resident language model receives fitted couplings, nucleus-by-nucleus residuals, optimizer status, Jacobian diagnostics, and previous interaction choices. It proposes the complete next active term list. The backend fits that new model and returns new evidence. Coupling optimization and interaction-structure choice are separate actions: the numerical op-




FIG. 2. **Synthetic nuclear observables.** Synthetic observations at twelve shells and 320 radial points. Filled blue circles denote ^{48}Ca , ^{132}Sn , and ^{208}Pb used in the fit; the open red diamond denotes withheld ^{90}Zr . Binding is plotted per nucleon for readability, but calibration uses total binding energy. No experimental error bars apply to this noiseless injection. Produced by `tools/paper_figures.py`.

Nucleus	N	Z	Role	B (MeV)	r_{ch} (fm)	Δr_{np} (fm)
^{48}Ca	28	20	Training	409.548	3.47518	0.21873
^{132}Sn	82	50	Training	1073.678	4.71155	0.33579
^{208}Pb	126	82	Training	1602.140	5.51385	0.26747
^{90}Zr	50	40	Withheld	776.957	4.26885	0.10938

TABLE I. **Synthetic observations in the larger-basis experiment.** Values are rounded for display; fits use unrounded records. These are model predictions, not experimental measurements. Generated by `tools/paper_results.py` from the frozen injection.

timizer changes values inside a model, whereas the investigator changes which density laws or interactions are admitted.

The revised campaign allows four changes of interaction structure, and invalid proposals have two repair opportunities. A failed physics solve halts the fit; it is not assigned a finite score against the interaction structure. The search stops when the optimizer reports convergence, training $\chi^2 < 10^{-5}$, and no active optional scaled coupling has magnitude below 10^{-4} . It can also end at the proposal budget. Among visited models, the champion minimizes $\chi^2 + 2k$, with k the number of fitted couplings. This is a complexity heuristic, not a Bayesian evidence calculation. After the search is frozen, the audit evaluates the champion on withheld ^{90}Zr and compares its scope and couplings with the generator. No holdout-driven choice is permitted.

The calculations use Quarks2Cosmos (Q2C) for spherical Hartree observables and implicit derivatives[16], with JAX automatic differentiation and bounded least squares for calibration. A quantized Gemma model served locally by llama.cpp proposes complete interaction structures. The source revision, optimizer settings, response constraints, and execution records are specified in Appendices A–C. No changes to the physics solver are made during structural search.

III. RESULTS

We first examine how the selected interaction terms alter the nuclear residuals, then assess recovery and local sensitivity. Numerical controls separately test the derivatives and the dependence on the basis and grid. The larger-basis experiment supplies the main demonstration; the pilot and unsuccessful attempts provide checks on the interpretation.

A. Residuals and recovery of density dependence

The constant-coupling fit converges, but leaves a training objective of 27107.26 (Table II). Its largest contribution comes from neutron skins: these account for 93.2% of the objective. Figure 4 shows the corresponding predicted-minus-injected observables in physical units. Binding residuals of order MeV accompany skin differences of tens of mfm, where one mfm is 10^{-3} fm. Improving the constant rho strength alone therefore leaves a coherent mismatch in neutron–proton spatial structure.

The investigator first admits linear density dependence. After recalibration, the objective falls to 166.95, and the skin residuals are strongly reduced. The remaining objective is dominated by charge radii, which account for 97.0% of its value (Fig. 4). Adding quadratic density dependence then reduces the objective to 8.18×10^{-21} . All three fits report optimizer convergence, separating the inadequacy of the first two forms from a failed continuous fit (Table II).

Figure 3 connects these residual changes to the coupling law in Eq. 4. The recovered law coincides with the injected law at the precision shown. The curves describe coupling as a function of density; they are neither radial density profiles nor uncertainty bands. The observable-by-observable interpretation above is a retrospective analysis of the stored fits. The investigator’s accepted rationales were qualitative and did not explicitly identify these residual contributions. Numerical recovery and the quality of its physical explanation are separate outcomes.

B. Withheld predictions and local identifiability

After the search is frozen, the selected interaction is evaluated for ^{90}Zr . Its withheld objective is 8.48×10^{-20} , and the largest absolute error among the injected couplings is 1.31×10^{-12} (Table II). This prediction tests transfer to another nucleus without using its observables to select terms. Because ^{90}Zr lies within the training mass range, the test concerns transfer within that range rather than extrapolation to distant nuclei.

The final scaled residual Jacobian has rank 3 and singular values 4639.30, 726.86, and 98.11 (Table II). Thus the fitted couplings have independent local responses under the chosen parameter and residual scalings. This rank test cannot establish global uniqueness. Competing functional forms can remain predictive, especially when an additional coefficient approaches zero. Enumeration of the allowed structures supplies a competing-model reference using the same solver, rather than an independent many-body calculation.

The pilot at lower resolution follows the same sequence of constant, linear, and linear-plus-quadratic couplings. Its final training and withheld objectives are 6.79×10^{-19} and 1.52×10^{-21} , respectively (Table II). Recovery at both discretizations is consistent with a functioning inverse loop. The injections are generated and recovered with the same representation within each run, however, so this agreement does not test the numerical accuracy of the underlying nuclear predictions.

C. Derivative validation and numerical resolution

The local derivative control compares the implicit response in Eq. 6 with central finite differences for ^{48}Ca . The maximum relative discrepancy is 7.46×10^{-8} , below the registered rejection threshold (Table III). A common preflight calculation has a maximum absolute observable difference of 9.46×10^{-11} across two numerical environments, using MeV and fm for the stored energies and radii. These checks test differentiation and runtime consistency of the same forward model; neither supplies an independent validation of the Hartree physics.

Changing from the pilot to the larger basis shifts injected ^{208}Pb binding by 25.62 MeV (Table III). Charge

radii and neutron skins also change, with nucleus-dependent signs. The fixed-point iteration tolerance therefore cannot be used as the total numerical error. Basis truncation and radial discretization remain material, and these two representations do not define a continuum-limit extrapolation. The small recovery residuals must be interpreted relative to this much larger resolution dependence.

The unsuccessful attempts additionally show that a valid nuclear solution and an accepted structural proposal are different requirements. Some fits exhaust their evaluation budget; other searches halt because proposals are invalid or their evidence attribution is incorrect. Appendix B retains these cases. The corrected run accepts both proposals without repair, but its first rationale calls the model empty when only the optional-term list is empty. Successful numerical choices alone therefore do not establish reliable physical reasoning.

D. Choosing a discriminating nuclear calculation

We next assign the investigator a prospective task: select one nuclear observable that could distinguish two calibrated interactions and predict which interaction gives the larger value. The linear rival fits g_ρ^2 and d_a , with $d_b = \Lambda_v = 0$; the mixed rival fits g_ρ^2 and Λ_v , with $d_a = d_b = 0$. Both fit only the synthetic binding energies of ^{48}Ca and ^{208}Pb , using twelve shells and 320 radial points. Their three-start calibrations have full local rank and objectives below 3×10^{-22} . Thus these two binding constraints do not choose between the fitted representations. The unchanged quadratic generator in Sec. II B lies outside both rival families: the experiment compares misspecified approximations rather than recovers a true model.

The menu contains binding, charge radius, or neutron skin for ^{40}Ca , ^{56}Ni , ^{78}Ni , ^{100}Sn , and ^{132}Sn . A full-evidence investigator receives calibrated couplings, residuals, and local sensitivities; a second investigator receives the same physical question with those fit diagnostics masked. Neither sees candidate predictions, generator couplings, or audit observations. Seeded random selection, largest neutron–proton asymmetry, and largest numerically resolved predictive disagreement provide comparison rules. All five choices are frozen before revealing the selected synthetic datum and the audit nuclei. A nuclear solve returns all three observables, but the experiment reveals only the selected scalar under its artificial one-datum information budget.

The calibration scales are 2 MeV for binding, 0.02 fm for charge radius, and 0.03 fm for neutron skin. These are declared design weights, not experimental standard deviations. At each fitted coupling point, candidate predictions are checked at 12, 16, and 20 shells with 320 radial points, and at 16 shells with 640 points. Each rival’s envelope is the largest change from its twelve-shell prediction among these representations. We call the ini-




FIG. 3. **Recovery of the density-dependent coupling.** The corrected larger-basis recovery. Left: the analytic rho-coupling law with the fitted parameters for each candidate interaction; $n_0 = 0.152 \text{ fm}^{-3}$. The green recovered curve coincides with the hidden injected law to numerical precision. These curves are functions of density, not simulated radial density profiles, and do not show uncertainty bands. Right: training objective after each investigator-selected change, $27107.26 \rightarrow 166.95 \rightarrow 8.18 \times 10^{-21}$. Produced by `tools/paper_figures.py`.




FIG. 4. **Residuals in physical units.** Predicted minus injected observables after each continuous fit in the corrected run. Red: constant coupling; blue: linear density dependence; green: linear plus quadratic dependence. The last model lies at zero at this scale. One mfm is 10^{-3} fm . These are physical-unit residuals, not error bars or posterior samples. MeV-level binding residuals coexist with neutron-skin mismatches of tens of mfm in the constant model. Produced by `tools/paper_figures.py`.

tial difference numerically resolved only if it exceeds three times the sum of the two envelopes. This empirical criterion does not bound the remaining discretization error or establish statistical model exclusion. The numerical selector uses forty candidate-check solves, which are charged as setup rather than treated as free information.

Both investigators select the ^{132}Sn neutron skin, the same choice as the numerical rule (Table IV). The initial linear and mixed predictions are 0.335633 and 0.347339 fm. Their difference, 0.011705 fm, exceeds three times the summed envelope of 0.001982 fm, but is only 0.39 of the declared skin scale. Both investigators predict the opposite ordering. Masking fit evidence changes neither their choice nor their sign prediction in these two calls. The asymmetry rule selects the ^{78}Ni skin, with a larger nominal standardized separation but an unresolved numerical envelope; the seeded random choice is ^{56}Ni binding, also unresolved.

After acquiring the selected scalar, both rivals are recalibrated from three starts. For the tin skin, the objectives are 1.31×10^{-6} and 0.01014 for the linear and mixed rivals. This orders the fits under the registered

weights; it does not reject the mixed family. An audit on the three observables of ^{16}O and ^{90}Zr measures the mean squared scaled residual. The linear risk changes from 2.45647×10^{-4} to 2.46978×10^{-4} , a slight deterioration. The mixed risk decreases from 5.95031×10^{-3} to 1.44719×10^{-3} , about 76%, while remaining larger than the linear risk. The useful acquisition therefore improves one approximation without improving the selected approximation.

Separate fixed-coupling controls distinguish basis changes from radial integration effects. Across four nuclei, the largest binding change on 640-to-1280-point refinement at twenty shells is 0.00195 MeV, compared with 0.150 MeV between twenty and twenty-four shells and 0.0220 MeV between twenty-four and twenty-eight shells at 640 points. The nonmonotonic basis trend survives radial refinement. Small late increments nevertheless do not certify a continuum limit: some increments barely contract, and the largest-basis range, solution branch, and refitted predictions remain untested. The Hartree iteration finds a self-consistent stationary solution; its convergence does not establish a global variational mini-

Model	k	Training objective		Jacobian rank	
		Pilot	Larger basis	Pilot	Larger basis
Constant	1	33206.20	27107.26	1	1
Linear	2	157.92	166.95	2	2
Linear + quadratic	3	6.79×10^{-19}	8.18×10^{-21}	3	3
Withheld objective		1.52×10^{-21}	8.48×10^{-20}		
Maximum coupling error		1.88×10^{-11}	1.31×10^{-12}		

Larger-basis final singular values: 4639.30, 726.86, 98.11.

TABLE II. **Recovery and local sensitivity at two discretizations.** The pilot uses eight oscillator shells and 160 radial points; the larger-basis experiment uses twelve shells and 320 points. All listed fits report convergence. Rank and singular values refer to the scaled residual Jacobian; k counts fitted couplings. Objectives use the registered residual scales in Eq. 5. Generated by `tools/paper_results.py` from the frozen search and audit records.

Nucleus	δB (MeV)	δr_{ch} (fm)	$\delta \Delta r_{np}$ (fm)
^{48}Ca	0.41	0.00345	0.01019
^{132}Sn	7.16	-0.01855	0.02446
^{208}Pb	25.62	0.00229	-0.02453
^{90}Zr	4.14	-0.01271	0.00327

Maximum relative derivative discrepancy: 7.46×10^{-8} .

Maximum absolute preflight difference across runtimes: 9.46×10^{-11} .

TABLE III. **Numerical checks and resolution dependence.** Shifts are larger-basis minus pilot predictions at the injected couplings. The derivative comparison uses central finite differences at one parameter point for ^{48}Ca . The runtime comparison is the maximum absolute difference over the stored observables in their respective MeV and fm units; it is not a dimensionless error. Generated by `tools/paper_results.py` from the stored controls and preflight records.

mum.

Post-decision checks at actual initial, updated, and generator couplings pass for all ninety nucleus–representation combinations at twelve shells/320 points and sixteen shells/640 points. They check finite values, iteration convergence, particle-number normalization, and negative highest occupied energies. For two linear-rival response directions in ^{48}Ca , central differences at two step sizes agree with automatic differentiation to a maximum relative discrepancy of 4.55×10^{-8} , after matching the unperturbed forward value. These fixed-coupling checks do not establish refitted-risk stability or independent physical accuracy.

The completed shared campaign uses 1187 nucleus-observation calls, forty candidate checks, and nine generator calls. Full and masked inference take 177 and 102 s, respectively. Its eight-core allocation consumes 12.09 allocated CPU hours (6.25 measured CPU hours); duplicate acquisitions are refitted separately. These totals describe one sequential campaign, not a matched wall-clock efficiency comparison. The result separates useful observable choice, correct model-specific reasoning, and improved prediction: none implies the others. General agent benefit requires repeated injections, uncertainty propagation, refined refits, and independent physical validation.

IV. DISCUSSION

The immediate benchmark is small enough for exhaustive enumeration. The next campaign should compare exhaustive, random, greedy residual-driven, and SGA-style proposal policies under identical experiment budgets. It should randomize injections, noise, initial fits, and omitted terms. Separate metrics must track interaction-structure recovery, coupling accuracy, withheld predictions, reasoning errors, and numerical failures. LLM-SRBench motivates controls for memorized equations[17]. A known term vocabulary remains an advantage, even when coefficients are withheld. A generator outside the grammar should test whether an investigator recognizes insufficient model space.

For larger nuclear landscapes, workers can own bounded families of operators, density laws, or physical regimes. A resident investigator should analyze isotopic residual patterns and parameter degeneracies before requesting new observables or terms. Density profiles, collective response, neutron matter, and stellar constraints can provide complementary information. Symbolic relations between nuclear matter and stellar properties provide related work for physical search coordinates[18]. Their admission requires compatible likelihoods and numerical accuracy; combining more data does not automatically repair model discrepancy. Bayesian calibration and emulation should be reused from the nuclear literature rather than replaced by LLM confidence.

TABLE IV. Calculation-selection pilot. Numerical separation requires the initial rival difference to exceed three times the sum of their empirical resolution envelopes. This is not a statistical rejection criterion. Q is the sum of squared scaled calibration residuals after one additional scalar observation; the scales are declared benchmark weights, not measured uncertainties.

Selector	Observable	Numerically resolved	Predicted ordering	Q_{linear}	Q_{mixed}
Investigator, full evidence	^{132}Sn skin	Yes	Incorrect	1.31e-06	0.0101
Investigator, masked evidence	^{132}Sn skin	Yes	Incorrect	1.31e-06	0.0101
Seeded random	^{56}Ni binding	No	Not requested	3.63e-06	3.11e-06
Largest asymmetry	^{78}Ni skin	No	Not requested	6.52e-06	0.0141
Numerical disagreement	^{132}Sn skin	Yes	Not requested	1.31e-06	0.0101

The restricted term vocabulary defines the admissible model space. Broader interaction discovery requires operator conventions, dimensions, symmetries, and equivalence checks. New operators need backend implementation, derivative checks, and independent physics review. Counterexamples and unsuccessful searches are part of the evidence, not discarded trajectories. This framework can support a progression from restricted recovery to criticism of the model space, then validated model extensions. Each step needs its own experiment and claim boundary.

V. CONCLUSIONS

We tested structural revision of a restricted isovector interaction through self-consistent spherical Hartree calculations. Continuous calibration fitted synthetic nuclear observables, and a resident language model selected the admitted density-dependent terms. In the larger-basis experiment, linear and quadratic rho-coupling density dependence recovered the injected model, with training and withheld objectives of 8.18×10^{-21} and 8.48×10^{-20} , respectively (Table II). The response Jacobian supported local sensitivity to all fitted couplings. The calculation-selection pilot additionally found a numerically resolved neutron-skin choice, but full and masked investigators made the same choice and both failed its model-ordering prediction. The acquired datum improved the mixed rival while slightly worsening the selected linear rival on the withheld audit. This conditional result supplies no observed evidence-driven advantage or statistically significant model discrimination.

The experiment established recovery within a known term vocabulary and a same-solver injection. It did not establish a unique effective interaction, converged nuclear accuracy, or inference from experimental data. Resolution dependence and unsuccessful proposals show why numerical validity, predictive agreement, and explanatory correctness require separate tests. Broader applications require converged calculations, uncertainty treatment, more varied injections, and comparisons under equal computational budgets.

Appendix A: Numerical settings and controls

The large components of the Dirac spinors are expanded in a spherical harmonic-oscillator basis with N_F shells; the small components use $N_F + 1$. Each angular-momentum block is diagonalized separately. Fields and densities are sampled on a midpoint radial grid extending to 20 fm. Meson kernels are integrated over each cell using two eight-point Gauss-Legendre panels and piecewise-linear source densities. The initial recovery uses $N_F = 8$ and 160 radial points; the corrected recovery uses $N_F = 12$ and 320 points. These are two discretizations, not a demonstrated continuum limit.

With y collecting fields and densities, the damped update is $y_{t+1} = 0.7y_t + 0.3\mathcal{G}(y_t; \theta)$. The solver uses absolute and relative fixed-point tolerances 10^{-10} and at most 1500 steps. The experiment rejects nonfinite observables and a solve that reaches that step cap. Tight iteration tolerance does not remove basis truncation or radial-grid error.

Parameters are scaled as $u = (g_\rho^2/100, d_a, d_b, \Lambda_v/0.01)$. The bounds in physical parameters are $45 \leq g_\rho^2 \leq 160$, $0 \leq d_a \leq 0.8$, $-0.2 \leq d_b \leq 0.4$, and $0 \leq \Lambda_v \leq 0.05$. Each fit starts from the active components of $(g_\rho^2, d_a, d_b, \Lambda_v) = (80, 0.1, 0.02, 0.002)$; it does not inherit the preceding optimum. Bounded least squares uses SciPy's default trust-region reflective method, with $\text{ftol} = \text{xtol} = 10^{-10}$ and $\text{gtol} = 10^{-9}$. The initial fits allow 30 function evaluations; the corrected protocol allows 60. Repeated requests at the same parameter vector use a cache.

The derivative control uses ^{48}Ca at scaled coordinates $u = (0.8, 0.1, 0.02, 0)$ in the initial discretization. Central finite differences perturb each coordinate by 10^{-4} and are compared with the implicit Jacobian. The relative-error denominator is the larger of the absolute implicit derivative and 10^{-6} ; the registered rejection threshold is 10^{-3} . This control probes one local response. Its result appears in Table III.

A separate SPORC allocation enumerates all eight architectures with the same forward solver. It supplies competing-model fits, not an independent many-body oracle. Runtime consistency is checked on a common preflight calculation using two JAX environments. Resolution controls recompute injected and recovered observables at twelve shells and 320 radial points, and compare the injection across discretizations. Finally, the corrected

agent campaign generates and recovers a fresh injection at that larger resolution. Agreement between generator and recovery must be distinguished from stability as the numerical representation changes.

Appendix B: Unsuccessful attempts and evidence attribution

All attempts remain in the repository, including searches that stop without a recovered model. The pilot’s second rationale contains an incorrect objective attribution. A subsequent attempt confuses a replacement model with a list of term additions and cycles among tested models. Another larger-basis fit approaches the injection with the full interaction but exhausts its evaluation budget; the additional mixed-vector coefficient is nearly zero. These cases motivate complete-model proposals, explicit convergence checks, and inspection of inactive or nearly vanishing terms.

Further attempts fail proposal validation through numerical transcription errors or unconstrained responses. Inspection of the serving interface finds that one request places its response schema in an ignored field. After the schema wrapper is corrected, an independent serving check verifies that the server enforces the requested structure before an experiment begins. The mixed-vector interaction is also described explicitly as a product to avoid ambiguity in its name. These changes address the experiment protocol; they do not change the nuclear Hamiltonian or make failed fits acceptable.

The corrected larger-basis search accepts two proposals without repair and verifies the frozen search record before opening the withheld observations. Its qualitative rationales still omit the nucleus-specific residual pattern. Evidence references prevent some attribution errors, but cannot certify the physical interpretation offered by the investigator. The dated notes and raw records retain the unsuccessful searches alongside the successful ones.

Appendix C: Execution and reproducibility

The resident investigator uses Gemma 4 E2B in Q4_0 quantization, served by llama.cpp build b11189 inside the worker allocation. Scope requests use temperature 0.1 and a limit of 1800 generated tokens. The numerical fit is fixed by its settings; the language-model policy remains stochastic. Model weights and runtime archives have immutable OSDF names and checksums. Private source travels in a compact scheduler transfer. The investigator receives a JSON evidence packet through loop-back HTTP, with no shell or file tools. Generator truth and final holdout observations are excluded from every request. The search is hashed before the audit opens them. The coordinator knows the generator; the investigator is blind to its coefficients, not to the candidate term vocabulary or all training priors.

The initial investigator proposes a complete architecture and rationale. Later contracts enumerate every untested complete architecture in the response schema. Enumeration constrains valid choices without choosing one. The investigator can add, remove, or replace registered terms. It must cite the current fit through ledger keys; the backend resolves numerical values. Free prose contains no digits. This prevents numerical transcription and some stale-attribution errors, but cannot certify physical reasoning.

Invalid proposals receive bounded repair opportunities. Exhaustion halts the search without substituting a numerical selector. The revised search allows four mutations and 60 optimizer evaluations per fit. Its stopping condition requires convergence, sufficiently small training residuals, and no nearly zero active optional terms. The revised contract removes the pilot requirement for a minimum number of architecture changes. Prompts, responses, failed fits, and contract rejections remain in the ledger.

RIFT’s Integrate Likelihood Extrinsic and Construct Intrinsic Posterior stages motivate the separation between forward evaluation and numerical inference[19, 20]. Here the corresponding stages are the Hartree calculation and fixed-model least squares. The latter computes a best fit, rather than RIFT’s posterior distribution.

OSG job 12355754 ran on a PDX-Coeus worker with an NVIDIA A40. It requested four CPUs, one GPU, 24 GB memory, and 12 GB scratch. The allocation completed in 306 seconds with a reported peak memory of 7325 MB. The investigator chose linear density dependence and then linear plus quadratic dependence.

The corrected larger-basis run uses eight CPUs without a GPU. Its scheduler reports a wall time of 522 seconds. Two accepted investigator requests take 170.18 seconds, while 66 nuclear evaluations take 195.93 seconds in the search laboratory. The event log records 6604 MB memory, whereas scheduler history records 7325 MB. Both records are retained. The different allocations and resolutions prevent a controlled CPU-versus-GPU comparison.

A compute island can be a fat node or a verified set of nodes with shared storage and suitable interconnects. An island lease must state cores, memory, accelerators, scratch, expiry, and measured reachability. Scheduler allocation alone does not imply a fast subnet or shared filesystem. Independent OSG workers should communicate through compact immutable records and staged assets. Tightly coupled experiments belong within verified local allocations. This prototype demonstrates collocation inside a worker, not a multi-node island or a topology-aware scheduler.

An island service should keep model weights and a scientific working set resident across several experiments. It should publish compact evidence summaries and checkpoint hashes to an infrequent stronger coordinator. The coordinator assigns questions and reconciles competing results; local agents perform routine investigation with-

out round trips to a frontier model. Service reuse may reduce staging and inference startup costs. That benefit requires measurement of queue delay, cache hits, data movement, residency, retries, and accepted scientific evaluations per resource-hour.

Hydrodynamic extensions can allocate islands by physical process and numerical regime. Each investigator could compare transport closures, equations of state, resolution, initial conditions, and diagnostic observables. Discrepant results should first trigger numerical checks and conservation tests. Requests for new physics or solver repairs should create versioned proposals with reproducible counterexamples. They must pass small validation experiments before joining the campaign. An investigator must not absorb discretization error by adding unjustified physical terms.

Appendix D: Additional context from prior work

This prototype extends the preceding compute-island study[15]. Its organization follows the separation of numerical experiments and iterative inference in RIFT and HyperPipe[19, 20]. In this release, ILE denotes the numerical experiment, CIP denotes fixed-architecture continuous calibration, and a resident investigator changes the architecture after CIP. These are workflow analogies. The implemented CIP is bounded least squares; it does not calculate a Bayesian posterior or implement RIFT’s full refinement pipeline.

The structural/continuous division has direct precedent in Scientific Generative Agent (SGA), whose outer LLM search surrounds differentiable simulation and parameter fitting[9]. LLM-SR and LaSR establish language-guided equation and concept search[10, 11]; ProAug supplies data-dependent feedback beyond scalar scores[21]. SciExplorer already selects experiments and recovers unknown dynamics and Hamiltonians [12]. PhysGym and DiscoverPhysics test interactive physical rediscovery[13, 14, 22]. We therefore make no priority claim for agentic physical hypothesis revision.

Nuclear model discovery also has precedent. MISR learns symbolic models of nuclear energies and radii and identifies interaction-operator discovery as an extension[6]. Learned nuclear functionals, interpretable mass models, and LLM observable regression cover adjacent approaches [7, 8, 23]. UNEDF fitting, Bayesian information analysis, and covariant EDF emulation establish numerical calibration and uncertainty methods[3–5]. Density-profile calibration provides a recent route to additional nuclear constraints[24].

Lagrangian recovery itself predates LLMs. Symbolic law discovery, sparse identification, and physics-informed symbolic methods provide foundations [25–29]. Neural and sparse Lagrangian methods learn variational descriptions from dynamical observations [30–33]. Our inverse problem instead passes interaction choices through a self-consistent finite-nucleus solver. That difference defines

an application, rather than establishing priority.

Agentic work with theory software is also established. The bsm_agent framework constructs and checks Lagrangians with deterministic tools and supports local inference[34]. HEPTAPOD and Diagrammatica constrain agent actions through physics tools[35, 36]. FERMIACC and LeWRON address theory campaigns and Lagrangian-to-prediction workflows[37, 38]. The ML4SCI literature-extraction proposal is relevant planned work; its goals are not treated as measured results[39]. AgenticSciML, PhysMaster, Denario, and broader surveys situate this work among scientific-agent systems[40–43]. Parsl-enabled agents and ECAS already connect reasoning to validated HPC execution[44, 45]. AERIS supplies an adjacent example of agent-orchestrated nuclear-materials calculations[46].

Appendix E: Retrospective evidence-consumption pilot

A separate, subsequent review-protocol demonstration asks whether a consumer can distinguish supported reuse of a recorded result from an unsupported extension of that result. This is not part of the resident nuclear search. Ten known developer dossiers each request a focal use and an unaffected use; six focal uses should be withheld, one eligible, and three pending. All ten unaffected uses permit narrowly qualified quotation of a recorded metric. Three missing-premise cases are constructed rather than observed failures. On 17 October 2026, ten fresh hosted Codex sessions (runtime-reported model `gpt-6.1-sol`) each received one frozen dossier. All ten original responses passed strict local JSON validation, and all twenty decisions matched the frozen outcome-and-reference rubric without repair or retries. Same-team adversarial review supported all twenty scoped outcomes; nineteen explanations were supported and one qualified: “no terms” should read “no optional terms.” That original response is retained.

The pattern illustrates a scoped rule: reuse of a claim for a stated purpose requires evidence for that purpose, rather than a global trust judgment about its author. Missing material evidence leaves a decision pending; contradiction withholds the affected use while unrelated qualified uses can survive. This demonstration does not estimate efficacy: the cases were known, the reviewer was not independent or blinded, and full runtime isolation remains unverified despite disabled environment access and zero observed tool calls. A larger structured-versus-plain comparison has not run. These results supply no additional nuclear-physics validation or advantage over ordinary review.

Appendix F: Contributions and reproducibility

Codex is explicitly credited as the AI author. Richard O’Shaughnessy directed the research and is the corresponding human responsible for the AI authors; he authorized publication of this manuscript. Codex implemented the coordinator and prepared prose, analysis tooling, and the literature audit. The resident Gemma model produced the recorded nuclear-investigation decisions; the later hosted Codex consumer sessions are a separate experiment. We credit the Quarks2Cosmos contributors collectively for the upstream nuclear software. This credit does not imply contributor endorsement or assign human authorship to them. No model is claimed to bear human scientific responsibility.

a. Data and software. The private experiment repository preserves source hashes, prompts, responses, failed attempts, fit records, numerical controls and accounting. The nuclear backend used Q2C revision `f950973707d4d51d6cf1afb7b4944a66ba794148`; the calculation-selection prototype used revision

`8222730eea0bfec5679de47f22468235b6af9742`. Public redistribution of private Q2C source is not authorized by this publication. Consequently, readers cannot yet independently reproduce the complete physics workflow from a public package. The numerical results here are retained-record results, not a new execution. A later record-consistency audit cannot independently reconstruct the calculation campaign’s audit targets from a separately retained target file; reconstructing them from residuals would be circular. Authenticated blind chronology and refined-refit risk stability also remain unestablished.

The reference audit dates to 1 October 2026; it is not an exhaustive novelty claim. The predecessor aiXiv record was checked again during this preparation. This manuscript is licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). That license applies to the manuscript, not to the private source code or unreleased research records. A public supporting-data release remains future work.

-
- [1] Lalazissis, G. A. et al.. A New Parameterization for the Lagrangian Density of Relativistic Mean Field Theory. *Physical Review C*, 55, 540–543, 1997. <https://arxiv.org/abs/nucl-th/9607039>
 - [2] T. Niksic et al.. Relativistic Hartree-Bogoliubov model with density-dependent meson-nucleon couplings. *Physical Review C*, 66, 024306, 2002. <https://arxiv.org/abs/nucl-th/0205009>
 - [3] Kortelainen, M. et al.. Nuclear Energy Density Optimization. *Physical Review C*, 82, 024313, 2010. <https://arxiv.org/abs/1005.5145>
 - [4] McDonnell, J. D. et al.. Uncertainty Quantification for Nuclear Density Functional Theory and Information Content of New Measurements. *Physical Review Letters*, 114, 122501, 2015. <https://arxiv.org/abs/1501.03572>
 - [5] Giuliani, Pablo et al.. Bayes goes fast: Uncertainty Quantification for a Covariant Energy Density Functional emulated by the Reduced Basis Method. *Frontiers in Physics*, 10, 1054524, 2023. <https://arxiv.org/abs/2209.13039>
 - [6] Munoz, Jose M. et al.. Discovering Nuclear Models from Symbolic Machine Learning. *Communications Physics*, 8, 101, 2025. <https://arxiv.org/abs/2404.11477>
 - [7] Wu, X. H. et al.. Nuclear energy density functionals from machine learning. *arXiv preprint*, 2021. <https://arxiv.org/abs/2105.07696>
 - [8] Mumpower, M. R. et al.. Physically Interpretable Machine Learning for nuclear masses. *arXiv preprint*, 2022. <https://arxiv.org/abs/2203.10594>
 - [9] Ma, Pingchuan et al.. LLM and Simulation as Bilevel Optimizers: A New Paradigm to Advance Physical Scientific Discovery. *Proceedings of the 41st International Conference on Machine Learning*, 235, 33940–33962, 2024. <https://proceedings.mlr.press/v235/ma24m.html>
 - [10] Shojaei, Parshin et al.. LLM-SR: Scientific Equation Discovery via Programming with Large Language Models. *arXiv preprint*, 2024. <https://arxiv.org/abs/2404.18400>
 - [11] Grayeli, Arya et al.. Symbolic Regression with a Learned Concept Library. *arXiv preprint*, 2024. <https://arxiv.org/abs/2409.09359>
 - [12] Nagele, Maximilian and Marquardt, Florian. Agentic Exploration of Physics Models. *Physical Review X*, 16, 031002, 2026. <https://arxiv.org/abs/2509.24978>
 - [13] Chen, Yimeng et al.. PhysGym: Benchmarking LLMs in Interactive Physics Discovery with Controlled Priors. *arXiv preprint*, 2025. <https://arxiv.org/abs/2507.15550>
 - [14] Wiemann, Matt L. et al.. DiscoverPhysics: Benchmarking LLMs for Out-of-the-Box Scientific Thinking. *arXiv preprint*, 2026. <https://arxiv.org/abs/2605.26087>
 - [15] Codex (AI). Local scientific investigators on distributed compute islands. *Research/software record*, 2026. <https://aiXiv.science/abs/aixiv.260925.000009> aiXiv.260925.000009; predecessor in this sequence.
 - [16] Quarks2Cosmos contributors. Genesis Quarks2Cosmos. *Research/software record*, 2026. https://github.com/oshughnessy-junior/Genesis_Quarks2Cosmos Private source, revision `f950973707d4d51d6cf1afb7b4944a66ba794148`; upstream software credited collectively.
 - [17] Shojaei, Parshin et al.. LLM-SRBench: A New Benchmark for Scientific Equation Discovery with Large Language Models. *Proceedings of the 42nd International Conference on Machine Learning*, 267, 55325–55359, 2025. <https://proceedings.mlr.press/v267/shojaei25a.html>
 - [18] Patra, N. K. et al.. Learning the relations between neutron star and nuclear matter properties with symbolic regression. *arXiv preprint*, 2026. <https://arxiv.org/abs/2601.07727>
 - [19] Lange, Jacob et al.. Rapid and accurate parameter inference for coalescing, precessing compact binaries. *arXiv*

- preprint, 2018. <https://arxiv.org/abs/1805.10457>
- [20] O’Shaughnessy, Richard and Kedia, Atul. HyperPipe: RIFT documentation. Research/software record. <https://rift-documentation.readthedocs.io/en/latest/hyperpipe.html> Software documentation; accessed 1 October 2026.
- [21] Liu, Hao et al.. Programmatic Context Augmentation for LLM-based Symbolic Regression. arXiv preprint, 2026. <https://arxiv.org/abs/2605.03101>
- [22] Zheng, Tianshi et al.. NewtonBench: Benchmarking Generalizable Scientific Law Discovery in LLM Agents. arXiv preprint, 2025. <https://arxiv.org/abs/2510.07172>
- [23] Guo, S. J. et al.. Large language model for unified and accurate description of multidimensional nuclear properties. arXiv preprint, 2026. <https://arxiv.org/abs/2605.29408>
- [24] Yu, Chengpeng and Yoshida, Kenichi. Constraining Energy Density Functionals via Bayesian Analysis of Nuclear Densities. arXiv preprint, 2026. <https://arxiv.org/abs/2609.29034>
- [25] Schmidt, Michael and Lipson, Hod. Distilling Free-Form Natural Laws from Experimental Data. *Science*, 324, 81–85, 2009. <https://www.science.org/doi/10.1126/science.1165893>
- [26] Hills, Daniel J. et al.. An algorithm for discovering Lagrangians automatically from data. *PeerJ Computer Science*, 1, e31, 2015. <https://peerj.com/articles/cs-31/>
- [27] Brunton, Steven L. et al.. Discovering governing equations from data: Sparse identification of nonlinear dynamical systems. arXiv preprint, 2015. <https://arxiv.org/abs/1509.03580>
- [28] Udrescu, Silviu-Marian and Tegmark, Max. AI Feynman: a Physics-Inspired Method for Symbolic Regression. arXiv preprint, 2019. <https://arxiv.org/abs/1905.11481>
- [29] Cranmer, Miles. Interpretable Machine Learning for Science with PySR and SymbolicRegression.jl. arXiv preprint, 2023. <https://arxiv.org/abs/2305.01582>
- [30] Cranmer, Miles et al.. Lagrangian Neural Networks. arXiv preprint, 2020. <https://arxiv.org/abs/2003.04630>
- [31] Purnomo, Adam and Hayashibe, Mitsuhiro. Sparse identification of Lagrangian for nonlinear dynamical systems via proximal gradient method. *Scientific Reports*, 13, 7919, 2023. <https://www.nature.com/articles/s41598-023-34931-0>
- [32] Tripura, Tapas and Chakraborty, Souvik. A Bayesian framework for discovering interpretable Lagrangian of dynamical systems from data. arXiv preprint, 2023. <https://arxiv.org/abs/2310.06241>
- [33] Zhang, Zitong and Sun, Hao. Differentiable Sparse Identification of Lagrangian Dynamics. arXiv preprint, 2025. <https://arxiv.org/abs/2511.10706>
- [34] Saad, Shaikh. Large Language Model-Assisted Framework for BSM Model Building. arXiv preprint, 2026. <https://arxiv.org/abs/2606.21316>
- [35] Menzo, Tony et al.. HEPTAPOD: Orchestrating High Energy Physics Workflows Towards Autonomous Agency. arXiv preprint, 2025. <https://arxiv.org/abs/2512.15867>
- [36] Menzo, Tony et al.. Agentic Diagrammatica: Towards Autonomous Symbolic Computation in High Energy Physics. arXiv preprint, 2026. <https://arxiv.org/abs/2603.26990>
- [37] Agrawal, Prateek et al.. The FERMIACC: Agents for Particle Theory. arXiv preprint, 2026. <https://arxiv.org/abs/2603.22538>
- [38] Wang, Isaac R.. LeWRON: Agentic Analysis of Electroweak Phase Transitions. arXiv preprint, 2026. <https://arxiv.org/abs/2606.19425>
- [39] Machine Learning for Science. Agentic Lagrangian Extraction from the Literature. Research/software record, 2026. https://ml4sci.org/gsoc/2026/proposal_HEPSIM5.html GSoC project proposal; accessed 1 October 2026.
- [40] Jiang, Qile and Karniadakis, George. AgenticSciML: collaborative multi-agent systems for emergent discovery in scientific machine learning. *npj Artificial Intelligence*, 2, 57, 2026. <https://www.nature.com/articles/s44387-026-00102-5>
- [41] Miao, Tingjia et al.. PhysMaster: Building an Autonomous AI Physicist for Theoretical and Computational Physics Research. arXiv preprint, 2025. <https://arxiv.org/abs/2512.19799>
- [42] Villaescusa-Navarro, Francisco et al.. The Denario project: Deep knowledge AI agents for scientific discovery. arXiv preprint, 2025. <https://arxiv.org/abs/2510.26887>
- [43] Zhou, Lianhao et al.. Autonomous Agents for Scientific Discovery: Orchestrating Scientists, Language, Code, and Physics. arXiv preprint, 2025. <https://arxiv.org/abs/2510.09901>
- [44] Ma, Heng et al.. Connecting Large Language Model Agent to High Performance Computing Resource. arXiv preprint, 2025. <https://arxiv.org/abs/2502.12280>
- [45] Sun, Baixi et al.. ECAS: An Edge-Controlled Agentic System for Validation-Gated Scientific Application Execution. arXiv preprint, 2026. <https://arxiv.org/abs/2609.14211>
- [46] Lopes, Denise Adorno et al.. Agentic AI for accelerated fuel qualification framework: Demonstrating workflows with the AERIS enthalpy reasoning model. *Journal of Nuclear Materials*, 631, 156843, 2026. <https://doi.org/10.1016/j.jnucmat.2026.156843>