AIMS Mathematics, 11(11): 38771–38817.
DOI: 10.3934/math.20261496
Received: 04 August 2026
Revised: 15 September 2026
Accepted: 19 October 2026
Published: 20 November 2026

https://www.aimspress.com/journal/Math

Review

Fractional physics-informed neural networks for fractional PDEs: an operator-aware review of error, stability, cost, and reliability

Muhammad Asim Khan[1], **Ayesha Ayub**[2], **Naglaa F. Soliman**[3,*] **and Abeer D. Algarni**[3]

[1] School of Mathematical Sciences, Universiti Sains Malaysia, Penang 11800, Malaysia
[2] Advanced Telecommunication Research Center (ATRC), Faculty of Electrical and Electronic Engineering, Universiti Tun Hussein Onn Malaysia, Parit Raja 86400, Malaysia
[3] Department of Information Technology, College of Computer and Information Sciences, Princess Nourah Bint Abdulrahman University, Riyadh 11671, Saudi Arabia

* **Correspondence:** Email: nfsoliman@pnu.edu.sa.

Abstract: Fractional physics-informed neural networks (fPINNs) combine neural approximation with numerical or analytical treatments of fractional operators, making their reliability dependent on both network training and operator construction. This Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020-adapted review synthesizes 37 primary studies and classifies them by their dominant residual-construction mechanism: discretization, quadrature or integral, spectral or basis, Monte Carlo, analytical-prior, variable- or distributed-order, and operator-learning methods. The audit identifies formal convergence theory in seven studies, direct classical-solver comparisons in 11, operator-level ablations in 10, detailed cost reporting in eight, and full reproducibility packages in nine. Evidence is strongest for L1/L2 Caputo formulations, whereas Monte Carlo, variable- and distributed-order, and operator-learning approaches require stronger analyses of uncertainty, identifiability, and generalization. Across classes, operator accuracy, stability, boundary treatment, computational cost, and reproducibility remain decisive. Current evidence therefore supports fPINNs primarily as complementary hybrid solvers for inverse, data-constrained, high-dimensional, and repeated-query fractional partial differential equation problems rather than universal replacements for classical methods.

Keywords: fractional partial differential equations; fractional physics-informed neural networks; operator-aware taxonomy; systematic review; fractional operator approximation; error analysis; computational efficiency; reproducibility
Mathematics Subject Classification: 35R11, 35R30, 65M12, 65M15, 68T07

1. Introduction

Fractional partial differential equations (FPDEs) extend classical differential models to systems with memory, nonlocal transport, anomalous diffusion, viscoelastic relaxation, heterogeneous media, and multiscale dynamics [1]. Their numerical treatment is difficult because a fractional derivative at one point may depend on the complete temporal history, an extended spatial domain, or an integral over fractional orders. Time-fractional Caputo models therefore inherit weakly singular kernels and reduced regularity near the initial time, whereas Riesz and fractional-Laplacian models require singular integration, exterior-domain conditions, or compatible spectral representations [2–4]. Recent analytical studies illustrate the continuing breadth of fractional theory and applications. A Riemann-Liouville time-fractional thin-film model has been examined through symmetry reduction and exact-solution construction, and fractional linear Korteweg-deVries (KdV)-type models have been studied through nonlinear dynamics and group analysis [5, 6].

Physics-informed neural networks (PINNs) incorporate governing equations, data, and initial or boundary conditions into a learning objective [7]. For integer-order PDEs, automatic differentiation usually evaluates the differential operator directly. Fractional physics-informed neural networks (fPINNs) differ in a fundamental way: Their residuals ordinarily contain an additional numerical, stochastic, spectral, or analytical approximation of the fractional operator [8]. An fPINN is therefore not a purely neural solver; it is a hybrid analytical-numerical-learning method whose reliability depends jointly on network approximation, optimization, sampling, and the embedded operator treatment.

Furthermore, fPINN research has evolved beyond conventional discretization schemes to encompass phase-field models, high-dimensional Monte Carlo methods, and variable-order formulations [9–11]. More recent advances include Laplace-transform-based methods, approaches for high-dimensional fractional and tempered PDEs, and, among other notable developments, physicsinformed operator learning [12–14]. These studies are frequently grouped under the common label "fPINN," even though their residuals inherit different truncation errors, stability mechanisms, computational bottlenecks, regularity assumptions, and boundary requirements. Comparisons based only on network architecture or final relative error can therefore obscure the numerical method that the network is actually optimizing.

1.1. Knowledge gap and position relative to existing reviews

Fractional-calculus and FPDE surveys explain operators, kernels, applications, and classical discretizations, whereas reviews of PINNs and scientific machine learning primarily emphasize architectures, optimization, software, and integer-order PDE applications. The closest fPINN-focused review provides valuable guidance on implementation and efficiency, particularly for Caputo formulations [15]. However, a cross-family synthesis that treats the embedded fractional-operator approximation as the unit of comparison remains lacking, especially one that quantitatively evaluates theory, classical baselines, computational cost, operator ablation, uncertainty, and reproducibility. Table 1 positions the present review within these related literatures. Accordingly, the contribution lies not in the taxonomic labels alone, but in the operator-aware stratification that connects each approximation mechanism to its distinct reliability requirements. History discretization, singular quadrature, spectral projection, Monte Carlo estimation, analytical transforms, order-space integration,

and operator learning introduce different sources of error, stability constraints, and computational demands and therefore require different forms of validation. This framework also enables evidence claims to be systematically audited rather than inferred from isolated benchmark errors.

Table 1. Positioning of the present review relative to broader review literatures.

Review stream	Representative focus	Limitation for fPINNs	Role in the present review
Fractional-calculus and FPDE references	Fractional operators, anomalous diffusion, Mittag-Leffler functions, memory kernels, and nonlocal models [16–21]	Usually do not analyze neural residual construction or training dynamics	Provides mathematical background for why fPINN residuals are nonlocal and approximation-dependent.
Classical fractional numerical methods	L1/graded meshes, convolution quadrature, finite elements, discontinuous Galerkin methods, compact high-order schemes, fractional Laplacians, and fast memory methods [2–4, 22–29]	Provides scheme-specific stability, convergence, and error analysis, but usually does not address neural approximation, optimization, or sampling error	Supplies both the numerical baseline and the standard of mathematical error control against which discretization-, quadrature-, spectral-, and memory-aware fPINNs should be assessed.
PINN and SciML reviews	PINN architectures, loss functions, optimization, software, and applications [30–34]	Usually emphasizes integer-order PDEs or general PINN methodology	Helps interpret optimization and reproducibility issues, but not the fractional-operator approximation itself.
fPINN-focused implementation review	Caputo-based fPINN implementation, accuracy, and efficiency considerations [15]	Primarily focused on Caputo-based fPINN implementations, accuracy, and efficiency rather than a comparative evidence taxonomy spanning multiple fractional-operator families	Provides the closest review-level comparator and motivates explicit differentiation of the present contribution.
Present review	Operator-aware systematic evidence synthesis of fPINNs for FPDEs	–	Uses the embedded fractional operator as the analytical unit, quantifies reporting and reproducibility gaps, and derives class-specific reliability and validation requirements.

1.2. Review questions, scope, and contributions

The review addresses four questions:
(RQ1) Which fractional operators and approximation mechanisms are embedded in current fPINN and related fractional neural-solver residuals?
(RQ2) What theoretical and empirical evidence is reported for their consistency, stability, convergence, boundary treatment, and error interpretation?
(RQ3) How adequately do studies report classical baselines, computational cost, operator ablation, uncertainty, and reproducibility?
(RQ4) Under which problem characteristics are fPINNs defensible complements to established FPDE solvers?

The primary scope is FPDE-oriented neural solvers published from 2019 to 2026. Earlier fractional-calculus and numerical-analysis references are used only to interpret the embedded residual schemes. The review does not propose a new neural architecture or fractional discretization, and it does not

assume that fPINNs are generally superior to classical solvers.

The main contributions are as follows:

(C1) An operator-aware taxonomy based on the dominant approximation used to evaluate the fractional term in the neural residual;

(C2) A study-level quantitative audit of formal theory, classical-solver comparison, public code, computational-cost reporting, operator ablation, uncertainty reporting, and reproducibility;

(C3) A cross-class synthesis showing that error mechanisms, cost drivers, boundary requirements, and validation standards vary systematically with the embedded operator treatment;

(C4) A reliability framework that separates operator approximation from neural approximation, optimization, sampling, and boundary- and initial-condition enforcement errors; and

(C5) Operator-specific recommendations for benchmarking, cost-error reporting, uncertainty quantification, reproducibility, and future research.

Figure 1 summarizes the analytical logic of the review.

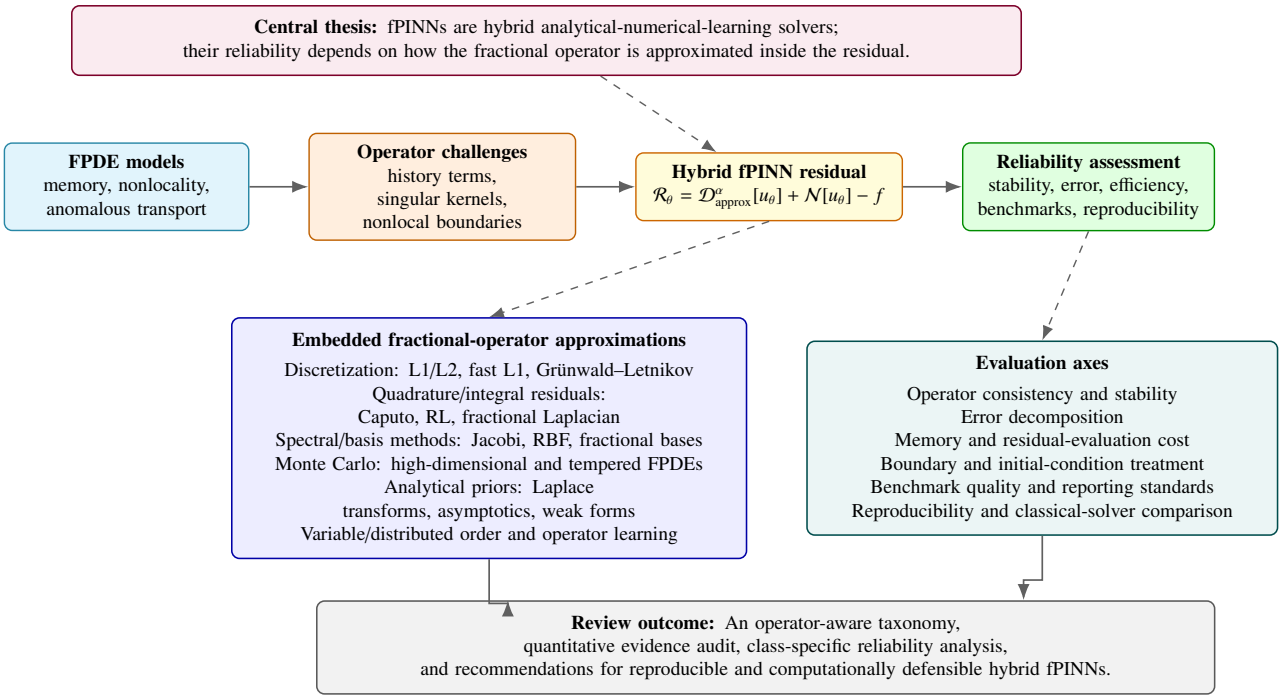

Figure 1. Operator-aware roadmap of the review. An fPINN residual combines neural approximation with an embedded numerical, stochastic, spectral, or analytical approximation of a fractional operator. The review therefore evaluates method families through operator consistency, stability, error, cost, boundary treatment, benchmark quality, and reproducibility.

The remainder of the paper proceeds from the review protocol and corpus audit to the mathematical structure of fractional residuals; analytical and numerical approximation mechanisms; the operator-aware taxonomy, reliability, and evidence maturity; computational efficiency; benchmark standards; and research priorities.

2. Review protocol and evidence corpus

2.1. Review scope and eligibility

This review follows an adapted PRISMA 2020 workflow for heterogeneous methodological literature in applied mathematics, numerical analysis, and scientific machine learning. Eligible primary records were required to solve or identify an FPDE with an fPINN, PINN, DeepONet, a neural operator, or a closely related neural solver; explicitly describe the fractional operator or its approximation; or provide explicitly relevant theory for an FPDE-oriented fractional neural method. Classical numerical, general PINN, and neural-operator references were retained only to interpret the core evidence and were excluded from the primary-study audit.

The eligibility criteria are given in Table 2. Preprints and early-access records had to satisfy the same topical and methodological criteria as peer-reviewed records and the additional status-specific threshold in Table 2. Their publication status was recorded explicitly and treated separately in the evidence-maturity assessment. Non-English records were excluded; this might introduce language bias, but the exclusion was necessary to allow consistent operator definitions, residual formulas, and numbers to be checked.

Table 2. Inclusion and exclusion criteria used for paper screening.

Category	Criteria
Include	Papers that solve FPDEs using fPINNs, PINNs, DeepONet, neural operators, or related neural solvers; papers that explicitly approximate Caputo, Riemann-Liouville, Riesz, Grünwald-Letnikov, fractional Laplacian, tempered, variable-order, or distributed-order operators; papers that discuss stability, convergence, error analysis, computational cost, memory reduction, high-dimensional fractional operators, inverse problems, or reproducibility; and classical fractional numerical-analysis papers need to interpret operator approximation inside fPINNs.
Conditionally include	Preprints, arXiv papers, or early-access papers only when they introduce a technically important method, such as convergence theory, operator-learning fPINNs, high-dimensional Monte Carlo fPINNs, or a new fractional residual formulation. Their bibliographic status must be verifiable, the extraction matrix must label them as preprints or early-access records, and they are not treated as equivalent to independently validated peer-reviewed evidence.
Exclude	Papers limited to integer-order PDEs without fractional operators; papers limited to fractional ODEs unless the operator approximation is transferable to FPDEs; application-only papers that do not specify how the fractional derivative is evaluated; generic machine-learning papers without relevance to FPDEs or nonlocal operators; non-English papers; and papers without accessible bibliographic verification.

Decision rule for fractional-ODE records

Fractional ordinary differential equation (FODE) records were screened using the hierarchical rule in Table 3. The distinction concerns transferability of the operator-residual mechanism, not a general assumption that performance on an ODE extends to an FPDE. An FODE-only record had to define the fractional operator and its residual approximation mathematically. Its coordinatewise approximation also had to couple to spatial differential or nonlocal terms without changing the approximation formula

or consistency argument. In addition, the record had to supply operator-specific validation through a consistency or error result, a controlled resolution study, or a matched classical comparison. Finally, its conclusions had to concern the operator-residual mechanism rather than infer FPDE-level accuracy from an ODE benchmark. Records meeting all four conditions were eligible for the methodological synthesis and quantitative audit, with their FODE scope identified explicitly. A record failing any condition was excluded from the counted corpus; it could be cited only as an external scope-boundary example and could not support prevalence or maturity claims. Records containing an explicit FPDE problem, FPDE benchmark, or FPDE-specific theorem proceeded directly to the remaining primary-study eligibility checks.

Table 3. Operational decision rules for FPDE and fractional-ODE records.

Step	Screening question	Decision
1	Does the record formulate, solve, or analyze an explicit FPDE?	Yes: Assess for the primary-study set under the remaining eligibility criteria. No: Proceed to Step 2.
2	Does the FODE-only record state both the fractional operator and the approximation embedded in the neural residual?	No: Exclude. Yes: Proceed to Step 3.
3	Can the approximation be coupled to spatial differential or nonlocal terms without changing its formula or consistency argument?	No: Exclude from the counted corpus. Yes: Proceed to Step 4.
4	Does the study provide operator-specific validation (a consistency/error result, controlled resolution study, or matched classical comparison) without treating ODE performance as proof of FPDE performance?	No: Exclude from the counted corpus; cite only as an external scope-boundary example if methodologically useful. Yes: Classify as transferable and assess under the remaining primary-study criteria; label its FODE scope explicitly.

2.2. Search strategy and selection process

The principal search window was 2019–2026, beginning with the foundational fPINN formulation [8]. Searches were conducted in Web of Science, Scopus, Google Scholar, ScienceDirect, SpringerLink, SIAM Publications, Wiley Online Library, Taylor & Francis Online, AIP Publishing, arXiv, Crossref, and Semantic Scholar. Additional records were identified through ResearchRabbit-assisted literature mapping, backward citation searching, forward citation searching, and targeted bibliographic verification.

The primary database, platform, and citation search was completed on 24 May 2026. A separate bibliographic-status audit of 2025–2026 and preprint records was completed on 15 July 2026 to verify author names, titles, publication status, journal metadata, and persistent identifiers. This was a metadata and publication-status update rather than a new search or screening round: It did not alter the eligibility criteria, add or remove records, change the 67-record review framework, or change the quantitative coding rules. Supplementary Table S2 summarizes the search concepts and representative keyword groups used to construct the source-specific queries. Records were imported into Zotero and deduplicated using DOI matching, normalized-title matching, and author-year comparison. Potential

duplicates without matching DOIs were checked manually.

Records were screened by title and abstract, followed by full-text or metadata assessment where relevance was uncertain. The detailed PRISMA selection counts are provided in Supplementary Table S3, while Supplementary Table S4 summarizes the reasons for excluding 27 records during the screening and eligibility stages. Figure 2 summarizes the selection process.

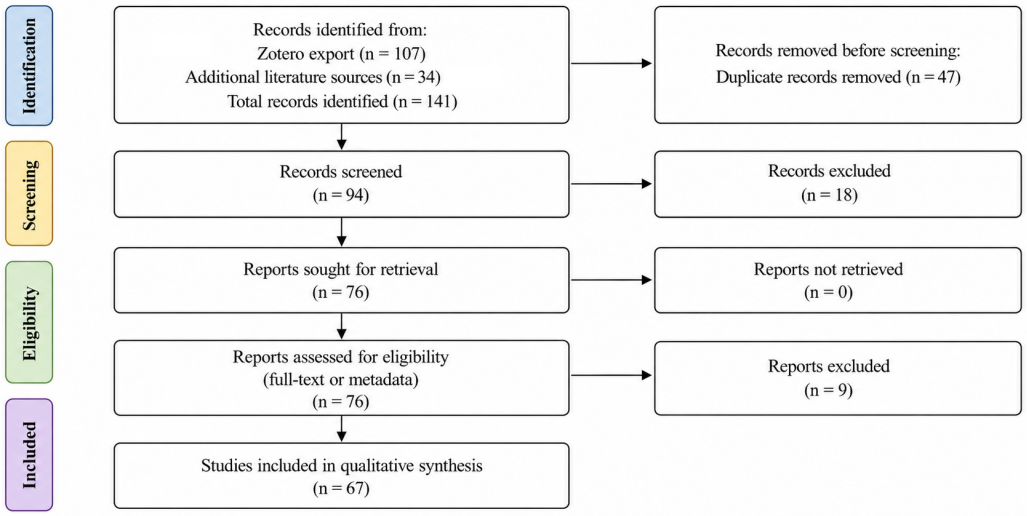

Figure 2. PRISMA 2020-adapted literature-selection flow. A total of 67 records were retained after eligibility assessment. Of these, 38 formed the core evidence set, comprising 37 primary fractional neural-solver studies and one fPINN-focused positioning review. The remaining 29 records formed a contextual supporting layer and were not included in the quantitative evidence-frequency calculations.

2.3. Corpus roles and counting rules

The 67 records retained within the review framework were assigned to two mutually exclusive roles. The core evidence set contains 38 records: 37 primary fPINN or fractional neural-solver studies and one fPINN-focused review used for positioning. The contextual supporting layer contains 29 references used for fractional-calculus background, classical numerical comparison, PINN/SciML context, neural-operator positioning, and reproducibility interpretation. Only the 37 primary studies entered the quantitative audit; the positioning review and contextual sources were excluded. Eight studies suggested during peer review were added post hoc as contextual references and were not retrospectively inserted into the documented search corpus. Three classical numerical studies illustrate scheme-specific stability, convergence, and error analysis in Table 1 [27–29]; two recent studies provide fractional-theory and application context in the Introduction [5, 6]; and three structural-topology-optimization studies support the engineering outlook without providing evidence about fPINN performance [35–37]. None changes the 67-record review framework, the 37-primary-study denominator, or any review-flow or quantitative-audit count. Supplementary Table S5 presents the publication-year distribution and reference-role composition of the 67 retained records, distinguishing the 37 primary studies from the positioning review and 29 contextual supporting references.

At the 15 July 2026 audit cutoff, two records in the 38-record core set were preprints. The

fPINN-focused positioning review [15] was excluded from the quantitative denominator. One primary convergence study [38] was included in the 37-study denominator. An eligible primary preprint was coded in the raw descriptive audit using the same indicator definitions as a peer-reviewed primary study, although publication status remained visible in the extraction matrix. Figure 5 illustrates class-level evidence coverage based on the audited indicators rather than a validated measure of methodological maturity. These descriptive scores should not be interpreted as evidence of independent peer-reviewed validation.

Each primary study was assigned one dominant operator class for class-stratified analysis, and secondary tags such as adaptive activation, loss weighting, multistage training, and multiprecision were coded separately. Operator, approximation, architecture, and thematic tags are otherwise non-exclusive; their totals are not expected to sum to the corpus size.

2.4. Data extraction and methodological appraisal

For each core record, the extraction matrix records bibliographic status, FPDE model, fractional operator, residual approximation, neural framework, problem type, theoretical or empirical stability/error evidence, cost reporting, benchmark characteristics, classical comparison, reproducibility, contribution, and limitation. The complete matrix appears as Supplementary Table S1.

The methodological appraisal considered operator transparency, connection to an established numerical or analytical scheme, stability or convergence evidence, benchmark credibility, classical-solver comparison, computational reporting, reproducibility, and publication status. The original categorical maturity labels combined these dimensions without a reproducible numerical mapping. We therefore replaced them with the continuous, domain-specific evidence-coverage scores in Table 4. These scores are descriptive transformations of the audited indicators, not a validated quality scale, and no overall score is formed across domains.

Table 4. Transparent aggregation rules used to generate the revised class-level evidence-coverage heatmap. For class c, N_c is its number of primary studies.

Heatmap domain	Score	Audited input and interpretation
Theory/error analysis	$S_{c,T} = T_c/N_c$	T_c is the number of studies coded Yes for formal convergence theory.
Comparison/attribution coverage	$S_{c,B} = (B_c + A_c)/(2N_c)$	B_c and A_c are the numbers with a direct classical baseline and a fractional-operator ablation. Equal weights are used because the two indicators represent complementary comparison and attribution evidence; neither was prespecified as more important.
Cost reporting	$S_{c,C} = D_c/N_c$	D_c is the number with Detailed computational-cost reporting. Basic cost information is not counted as complete coverage.
Reproducibility	$S_{c,R} = R_c/N_c$	R_c is the number coded Full for reproducibility. Partial, Limited, and Unclear are not positive numerators.
Display and missing evidence	Continuous scale from 0 to 1	Cells display the percentage and its underlying count(s). Unclear records remain in N_c and not in a positive numerator. No High/Moderate/Limited thresholds and no weighted overall maturity score are used.

Quantitative-audit coding protocol

The categorical audit used the operational definitions in Table 5. Coding was based on the full article together with available supplementary material and publicly accessible code repositories at the audit cutoff. Positive evidence had to be stated or demonstrated in these sources; it was not inferred from method names, citations to external theory, or general statements of availability. When the evidence was insufficient to assign a category defensibly, the entry was coded "Unclear". Unclear entries remained in the denominator but were not counted in a positive numerator. The fPINN-focused positioning review was excluded from all primary-study denominators.

Table 5. Operational codebook for the study-level quantitative audit.

Indicator	Operational coding rule
Formal convergence theory	Yes only when the study states a study-specific theorem, proof, or rigorous convergence/error bound for the neural method. Empirical convergence plots, numerical rates without a proof, or reliance only on the known order of an embedded classical scheme were coded No.
Direct classical-solver comparison	Yes only for a quantitative same-problem comparison with an identified non-neural fractional solver under compatible equations, domains, parameters, and error measures. Using a classical solution only to generate reference values was coded No.
Public code	Yes only when a functioning public repository or executable implementation could be located at the audit cutoff. An unlinked availability statement or an inaccessible repository was coded Unclear; an explicit statement that code was unavailable was coded No.
Computational-cost reporting	None: No meaningful cost information. Basic: At least one of wall-clock time, hardware, memory, asymptotic complexity, or residual-evaluation cost. Detailed: Wall-clock time and hardware together with at least one additional item (memory, fractional-residual evaluation time, inference time, scaling with operator resolution, or a cost–error comparison).
Fractional-operator ablation	Yes only when the fractional approximation or its numerical resolution (for example, h, Q, M, N_s, or Q_α) was varied under fixed or reasonably matched network architecture, sampling, optimization, and training budget. Changes only to architecture, activation, or loss weights were coded No.
Quantitative uncertainty	Applicable to stochastic residuals, noisy inverse problems, Bayesian/UQ methods, and analogous settings. Yes required variance, repeated-seed dispersion, confidence/credible intervals, posterior summaries, or an equivalent quantitative measure; point estimates alone were coded No. Other studies were coded N/A.
Reproducibility	Full: Public code plus sufficiently specified operator approximation, architecture, optimizer and training settings, sampling, and reference-solver construction. Partial: Substantial methodological detail was available, but code or at least one implementation-critical component was missing. Limited: Multiple critical details were absent, preventing faithful reconstruction. Unclear: Available material did not support a defensible classification.

In practical terms, "sufficiently specified" for a Full rating meant that the article, supplement, or public code/configuration files jointly identified the fractional-operator implementation and resolution parameters; the exact values of implementation-critical hyperparameters, including network depth, width, activation, optimizer, learning rate or schedule, training duration, and loss weights; collocation or data-sampling distributions and sample sizes; initial/boundary-condition enforcement; and the reference-solver construction and reported error metric. A random seed or deterministic seed policy was also required when a reported conclusion depended on stochastic sampling, Monte-Carlo estimation, or repeated-run variability; otherwise, this item was recorded as not applicable. Hardware and software-version details were recorded when available but were not mandatory for functional reconstruction and hence did not by themselves determine the Full rating. They were, however, required where appropriate for the separate detailed computational-cost classification and for interpreting runtime comparisons.

The audit was not independently double-coded, and the project records contain only the final study-level classifications reported in Supplementary Table S7, not separate preconsensus ratings from two coders. Consequently, a chance-corrected intercoder coefficient cannot be reconstructed and is not reported. Creating nominally independent ratings after the consensus table had already been seen would not measure the reliability of the original coding process. This limitation is mitigated, but not removed, by publishing the operational codebook, retaining ambiguous cases as unclear, and providing every study-level classification in Supplementary Table S7 so that readers can audit the aggregate counts.

2.5. Synthesis strategy and quantitative audit

The synthesis is methodological rather than chronological. Studies are compared through the approximate residual operator, allowing operator consistency, stability, error, memory and residual-evaluation cost, boundary treatment, benchmark quality, and reproducibility to be examined using common criteria. No pooled accuracy or effect-size meta-analysis was conducted because the governing equations, operators, domains, reference solutions, error metrics, and computational settings were too heterogeneous for meaningful quantitative aggregation.

Table 6 presents the descriptive audit of the 37 primary studies, with the study-level coding reported in Supplementary Table S7. The findings indicate that the principal weakness of the literature is not a shortage of reported benchmark errors but the limited evidence available for attributing, comparing, and reproducing those errors. Formal convergence theory was reported in $7/37$ studies, operator-level ablation in $10/37$, direct classical comparison in $11/37$, publicly available code in $10/37$, and full reproducibility packages in $9/37$. Although $22/37$ studies reported some computational information, only $8/37$ provided detailed cost evidence. Among the 16 studies for which quantitative uncertainty assessment was considered applicable, eight reported a variance estimate, confidence or credible interval, repeated-run dispersion, posterior measure, or an equivalent uncertainty statistic.

Table 6. Quantitative reporting characteristics of the primary fPINN evidence base. Percentages use 37 primary studies unless an applicability-specific denominator is shown.

Indicator	Studies reporting	Percentage	Operational interpretation
Formal convergence theory	7/37	18.9%	Study-specific theorem, proof, or rigorous convergence/error bound
Direct classical-solver comparison	11/37	29.7%	Matched comparison with an appropriate non-neural fractional solver
Publicly available code	10/37	27.0%	Accessible repository or executable implementation at the audit date
Any computational-cost reporting	22/37	59.5%	At least basic runtime, hardware, memory, or computational-cost information
Detailed computational-cost reporting	8/37	21.6%	Runtime and hardware together with memory, residual cost, inference time, or cost-error analysis
Fractional-operator ablation	10/37	27.0%	Comparison of alternative operator approximations or resolutions under controlled learning settings
Quantitative uncertainty reporting	$8/16_{\text{applicable}}$	50.0%	Variance, confidence/credible intervals, repeated-seed dispersion, or equivalent uncertainty evidence
Full reproducibility package	9/37	24.3%	Public code plus the implementation-complete settings defined in the audit codebook
Full or partial reproducibility	35/37	94.6%	At least substantial methodological information sufficient for partial replication

Note: Percentages were calculated conservatively from the study-level audit in Supplementary Table S7.

The 14 records initially identified as potential classical-solver comparisons were further assessed using the explicit eligibility criterion in Table 5, of which 11 satisfied the direct-comparison requirement. The three excluded records did not provide an eligible head-to-head comparison: One used quasi-Monte Carlo/Bayesian optimization (QMC/BO) outputs only as reference data, one provided only an asymptotic-complexity comparison, and one used the finite-difference method (FDM) solely to generate reference values. These records were therefore coded as not meeting the direct-comparison criterion in Supplementary Table S7.

2.6. Narrative synthesis of classical-solver comparisons

Differences in the governing equations, error norms, numerical resolutions, hardware configurations, and stopping criteria preclude pooling the 11 eligible comparisons into a common effect estimate or using vote counting to claim the superiority of one solver family. Nevertheless, these studies support a structured narrative synthesis. Table 7 distinguishes the comparative outcomes

reported by the individual studies from the qualifications required for their interpretation. The corresponding study-level extraction, including the three ineligible comparison designs, is provided in Supplementary Table S10.

Table 7. Narrative synthesis of the 11 eligible direct classical-solver comparisons. "Favorable" and "mixed" describe the results reported in the source studies, not independently standardized effect estimates.

Dominant operator class	n	Reported comparative pattern	Qualification and observed underperformance
Discretization-based	6	Lower errors were reported for selected phase-field and Huxley benchmarks and for the transferable ψ-Caputo FODE study; the foundational fPINN, HNS, and deep-network results were context-dependent.	Pang et al. reported optimization-error saturation under refinement and generally higher training costs than FDM. HNS of order three was favorable at coarse temporal resolution, whereas order five became unstable on fine intervals. Fang et al. reported accuracy similar to a spectral method at one resolution without a matched cost comparison. The ψ-Caputo result is FODE-only and is not evidence of FPDE-wide competitiveness.
Variable/distributed-order	2	Both comparisons were mixed rather than uniformly favorable.	A refined classical scheme achieved lower error than the variable-order PINN at greater runtime; dPINN and FDM rankings changed with order and resolution, and neural optimization was slower.
Spectral/basis	1	RBF-fPINN reported lower error than both the plain fPINN and FDM as the tested solution became sharper.	The comparison matched point counts, not computational budgets, and one study cannot establish class-level consistency.
Quadrature/integral	1	The neural Ritz method reported lower error than FDM at the same degrees of freedom.	Adam optimization was described as expensive, and computational complexity was not matched, so accuracy competitiveness did not establish cost competitiveness.
Analytical-prior	1	DSC-PINN was usually more accurate for weakly singular solutions and used less time and memory at a comparable-accuracy setting.	For $\alpha = 0.75$, FDM achieved a smaller L^∞ error on the finest reported grid, whereas DSC-PINN retained a slightly smaller relative L^2 error, demonstrating norm- and resolution-dependent performance.

The most repeatable favorable pattern is narrow. Several discretization-based variants were competitive on smooth, author-selected time-fractional benchmarks at coarse or sparse resolution. Singularity-aware or basis-enriched constructions were also favorable in the specific regimes they were designed to address [9, 39–41]. The counter-pattern is equally important. Refinement could expose neural optimization error, classical solvers could recover lower error, higher-order neural residuals could become unstable, and favorable accuracy could coexist with greater training cost [8, 11, 42–44]. No Monte Carlo or operator-learning study met the direct-comparison rule. Thus, no FPDE class can yet be described as consistently superior to classical solvers. The defensible conclusion is regime-specific competitiveness, with the strongest practical rationale remaining sparse-data, inverse, irregular, high-dimensional, or repeated-query settings that were seldom isolated in matched head-to-head tests.

2.7. Protocol limitations

The review is subject to terminology, indexing, language, publication-status , and coding-reliability limitations. Relevant work may be described as an fPINN, fractional PINN, neural solver, model-driven neural network, DeepONet, or neural operator. Some transferable methods concern fractional ODEs

or differential-algebraic systems rather than FPDEs and were therefore assessed using the explicit decision rule in Table 3 . Recent records may change status after the July 2026 metadata audit. The assignment of a dominant operator class necessarily simplifies hybrid methods, and several quantitative audit entries remain marked Unclear in Supplementary Table S7. Because independent preconsensus coder ratings were not retained, intercoder agreement could not be estimated retrospectively; the categorical frequencies should therefore be interpreted as a transparent, auditable evidence map rather than as measurements with established coding reliability. Finally, the review is a critical methodological synthesis rather than a statistical meta-analysis; the reported frequencies measure documentation and evidence availability, not comparative solver superiority.

3. Fractional operators and hybrid residual construction

The location of nonlocality (in time, space, fractional order, or several of these variables) determines the residual construction and the conditions under which it is meaningful. This section fixes the operator definitions and notation used in the subsequent analysis. For $0 < \alpha < 1$, the Caputo derivative is commonly written as

$$ {}_0^C D_t^\alpha u(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t (t-s)^{-\alpha} u'(s)\, ds. \tag{3.1} $$

It is widely used in time-fractional diffusion, reaction-diffusion, phase-field, Black-Scholes, and inverse problems because it allows classical-style initial conditions [9, 14, 39, 45]. Its weakly singular memory kernel creates history-dependent residuals and motivates L1/L2, fast-L1, graded-mesh, Hermite, and Laplace-transform fPINN variants [12, 44, 46].

The left-sided Riemann-Liouville derivative is defined by

$$ {}_0^{RL} D_t^\alpha u(t) = \frac{d}{dt}\left[\frac{1}{\Gamma(1-\alpha)} \int_0^t (t-s)^{-\alpha} u(s)\, ds \right]. \tag{3.2} $$

Compared with the Caputo derivative, the placement of the integer derivative changes the interpretation of initial conditions and may affect loss construction in inverse problems. Therefore, fPINN studies should report not only the fractional order but also the exact derivative definition used in the residual.

Grünwald-Letnikov approximations connect fractional differentiation with weighted finite differences:

$$ \begin{aligned} {}^{GL}D^\alpha u(x) &\approx h^{-\alpha} \sum_{k=0}^{N} (-1)^k \binom{\alpha}{k} u(x-kh), \\ \binom{\alpha}{k} &= \frac{\Gamma(\alpha+1)}{\Gamma(k+1)\Gamma(\alpha-k+1)}. \end{aligned} \tag{3.3} $$

This approximation was central to the original fPINN framework [8] and remains important in space–time fractional models such as fractional Black–Scholes equations [47]. Its main residual-level limitation is the appearance of dense or wide nonlocal stencils.

Space-fractional diffusion is often represented by Riesz derivatives or fractional Laplacian operators. For the integral fractional Laplacian in $\mathbb{R}^d$, one common definition is

$$ (-\Delta)^s u(x) = C_{d,s}\, \mathrm{PV} \int_{\mathbb{R}^d} \frac{u(x) - u(y)}{|x-y|^{d+2s}}\, dy, \qquad 0 < s < 1, \tag{3.4} $$

and the spectral characterization is

$$\mathcal{F}[(-\Delta)^s u](\xi) = |\xi|^{2s}\mathcal{F}[u](\xi). \tag{3.5}$$

Integral fractional-Laplacian residuals require appropriate treatment of singular spatial kernels and nonlocal exterior conditions, with domain truncation where applicable. Spectral formulations instead depend on the adopted operator definition, compatible boundary conditions, and spectral approximation. This motivates quadrature, Monte Carlo, spectral, and variational neural formulations [10, 13, 43, 48].

Variable-order and distributed-order models introduce an additional approximation layer because the order is variable, learned, or integrated over. A distributed-order Caputo operator may be written as

$$\mathcal{D}_t^w u(t) = \int_{\alpha_{\min}}^{\alpha_{\max}} w(\alpha)\, {}_0^C D_t^\alpha u(t)\, d\alpha. \tag{3.6}$$

Such models are useful for heterogeneous or multiscale memory processes but introduce identifiability, order-space quadrature, and regularization challenges [11, 42, 49, 50]. Conformable derivatives, by contrast, can often be evaluated through automatic differentiation with a multiplicative factor, but they do not represent the same nonlocal memory as Caputo or Riemann-Liouville derivatives [51]. Accordingly, conformable-derivative studies are used here only as scope-boundary comparisons and are not used to establish the maturity of nonlocal Caputo-, Riemann-Liouville-, Riesz-, or fractional-Laplacian fPINNs.

Representative FPDEs include time-fractional diffusion,

$${}_0^C D_t^\alpha u(x,t) = \kappa\Delta u(x,t) + f(x,t), \qquad 0 < \alpha < 1; \tag{3.7}$$

space-fractional diffusion,

$$u_t(x,t) = -\kappa(-\Delta)^s u(x,t) + f(x,t), \qquad 0 < s < 1; \tag{3.8}$$

space-time fractional advection-diffusion,

$${}_0^C D_t^\alpha u(x,t) + \boldsymbol{v}\cdot\nabla u(x,t) = \kappa\mathcal{D}_x^\beta u(x,t) + f(x,t), \quad 0 < \alpha < 1, \quad 1 < \beta \le 2; \tag{3.9}$$

and distributed-order diffusion,

$$\int_{\alpha_{\min}}^{\alpha_{\max}} w(\alpha)\, {}_0^C D_t^\alpha u(x,t)\, d\alpha = \kappa\Delta u(x,t) + f(x,t). \tag{3.10}$$

These models differ in residual structure: (3.7) requires temporal-history approximation, (3.8) requires spatial nonlocal integration or spectral evaluation, (3.9) combines temporal memory and spatial nonlocality, and (3.10) adds quadrature over fractional order.

Let $u_\theta(x,t)$ denote a neural approximation. A generic fPINN residual can be written as

$$\mathcal{R}_\theta(x,t) = \mathcal{A}_{h,q,M,N}^{\alpha,\beta}[u_\theta](x,t) + \mathcal{N}[u_\theta](x,t) - f(x,t), \tag{3.11}$$

where $\mathcal{A}_{h,q,M,N}^{\alpha,\beta}$ denotes an approximate fractional operator depending on mesh width h, quadrature parameters q, Monte Carlo sample size M, spectral truncation N, memory length, or related numerical parameters. The corresponding loss is typically

$$\mathcal{L}(\theta) = \lambda_r\mathcal{L}_r(\theta) + \lambda_i\mathcal{L}_i(\theta) + \lambda_b\mathcal{L}_b(\theta) + \lambda_d\mathcal{L}_d(\theta). \tag{3.12}$$

For fPINNs, $\mathcal{L}_r$ may contain L1 sums, Grünwald weights, quadrature nodes, Monte Carlo samples, spectral coefficients, inverse-transform approximations, or auxiliary outputs. Thus, the mathematical operator and its approximation directly shape the optimization landscape, computational graph, and total error.

The next two sections examine the analytical and numerical mechanisms used to construct and validate this residual.

4. Analytical structure for residual design and validation

Analytical methods remain important in fPINN research because they provide exact or semi-analytical reference solutions, regularity information, asymptotic behavior, compatible trial spaces, and diagnostic benchmarks for fractional models [16, 18, 19, 21]. Their role is not merely theoretical: Analytical structure can guide residual construction, reduce fractional-operator evaluation cost, identify weak initial singularities, and help separate neural approximation error from operator-approximation or optimization error [12, 43, 45, 52]. However, analytical priors are problem-dependent and must match the chosen fractional operator, boundary setting, and solution regularity [3].

4.1. Transform, spectral, and asymptotic priors

Laplace and Fourier transforms provide exact or semianalytical structures for linear FPDEs. For example, for $0 < \alpha < 1$,

$$\mathcal{L}[{}_0^C D_t^\alpha u(t)](s) = s^\alpha \widehat{u}(s) - s^{\alpha-1} u(0), \tag{4.1}$$

which supports analytical benchmarks and motivates Laplace-fPINNs that shift part of the fractional residual evaluation into the transform domain [12, 46]. Similarly, Fourier and eigenfunction representations support spectral treatment of Riesz derivatives and fractional Laplacians. If $-\Delta\phi_k = \lambda_k\phi_k$, then

$$(-\Delta)^s u = \sum_{k=1}^{\infty} \lambda_k^s u_k \phi_k, \qquad u_k = \langle u, \phi_k \rangle. \tag{4.2}$$

Equation (4.2) corresponds to the spectral definition of the fractional operator on a compatible eigenbasis; it should not be treated as automatically equivalent to the integral fractional Laplacian on bounded domains without specifying boundary and exterior-domain conditions. This structure motivates spectral and basis-enhanced fPINNs based on Jacobi polynomials, radial basis functions, fractional polynomial bases, or compatible eigenfunctions [40, 50, 52, 53].

Mittag-Leffler functions and weak initial singularities are especially relevant for time-fractional diffusion. Because many Caputo solutions are not smooth near $t = 0$, standard fully connected networks may under-resolve early-time behavior. Asymptotic priors, graded sampling, singularity decomposition, and correction terms can therefore improve accuracy and robustness [41, 45].

4.2. Weak, variational, and Green's-function priors

Weak and variational formulations are important when pointwise residuals are unstable or expensive because of singular kernels. For fractional Laplacian problems, a typical energy form is

$$a(u, v) = \frac{C_{d,s}}{2} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \frac{(u(x) - u(y))\,(v(x) - v(y))}{|x - y|^{d+2s}}\, dx\, dy. \tag{4.3}$$

Such formulations connect fPINNs with neural Ritz and finite-element-style error analysis, especially in fractional Sobolev or energy norms [43]. Green's functions and fundamental solutions can also generate reference data, identify convolution structures, and motivate kernel-informed or integral residual formulations, although their direct use in nonlinear and bounded-domain fPINNs remains limited.

4.3. Analytical methods as benchmark generators

Analytical tools are important for benchmark construction because a small residual loss does not necessarily imply a small solution error. Manufactured solutions, Laplace/Fourier solutions, eigenfunction expansions, Green's functions, Mittag-Leffler decay profiles, and high-accuracy weak-form references can help distinguish neural approximation error from fractional-operator approximation, quadrature error, Monte Carlo variance, or optimization error. Table 8 summarizes how analytical information can support residual design, trial-space construction, stability analysis, boundary treatment, and benchmark generation.

Table 8. Analytical tools, potential fPINN implementations, and limitations in hybrid analytical-numerical fPINNs.

Analytical tool or information	**Role in fPINNs**	**Possible implementation**	**Limitation or open issue**
Laplace transform	Time-fractional exact solutions, transform-domain residuals, inverse problems, and Laplace-fPINNs	Laplace-domain residuals with numerical inverse transforms	Particularly convenient for linear and suitably structured semilinear problems; numerical inverse Laplace transforms can introduce additional approximation error.
Fourier transform	Riesz derivatives, fractional Laplacians, periodic benchmarks, and spectral residuals	Fourier/spectral representation of nonlocal spatial operators	Applicability depends on the operator definition, domain geometry, and boundary conditions; direct Fourier representations are most natural in whole-space or periodic settings.
Mittag-Leffler decay	Fractional relaxation, asymptotic decay, and early-/long-time behavior in Caputo models	Enriched trial functions, output transforms, or asymptotic correction terms	Difficult to generalize to nonlinear, high-dimensional, or irregular-domain FPDEs.
Weak initial singularity	Regularity-aware treatment of time-fractional solutions near $t = 0$	Graded sampling, singularity-aware decomposition, and correction terms	Requires prior knowledge of solution regularity and singular structure.
Green's functions	Benchmark generation, convolution structure, and kernel-informed priors	Kernel-informed layers, integral representations, or reference-solution construction	Constructing and evaluating suitable Green's functions can be difficult for nonlinear problems, complex geometries, and nonlocal boundary conditions.

Continued on next page

Analytical tool or information	Role in fPINNs	Possible implementation	Limitation or open issue
Eigenfunction expansions	Spectral fractional operators, reduced basis design, and compatible hard constraints	Spectral trial spaces, basis-enriched networks, or coefficient-space residuals	Sensitive to smoothness, boundary compatibility, and basis truncation.
Weak or variational formulations	Energy residuals, fractional Laplacian stability theory, and Ritz-type neural solvers	Variational or Ritz-type fPINN losses based on energy functionals	Requires suitable test spaces and may increase quadrature complexity.
Known boundary behavior	Nonlocal boundary consistency, especially for fractional Laplacian problems	Boundary-aware trial functions or hard/soft exterior-condition enforcement	Difficult to construct for complex geometries and nonlocal exterior conditions.
Manufactured solutions	Controlled benchmark construction and error attribution	Exact or semianalytical reference solutions for residual and solution-error testing	May be too smooth or artificial compared with realistic FPDE solutions.

Analytical structure is most useful when treated as a verifiable component of residual design or benchmark construction. Regularity, asymptotic decay, function-space compatibility, and nonlocal boundary behavior can then guide the trial space and supply independent diagnostic checks. The next section examines the numerical approximations that complete the hybrid residual.

5. Embedded numerical approximation of fractional operators

Building on the operator definitions in Section 3, this section compares the numerical and semianalytical approximations used to evaluate fractional terms inside fPINN residuals. Each approximation introduces its own resolution parameters and error mechanism in addition to the neural and sampling choices. Let u_θ denote the neural approximation, and let $\mathcal{D}^\alpha$ be a fractional operator. In practice, fPINNs minimize a residual involving an approximate operator,

$$\mathcal{D}^\alpha u_\theta \approx \mathcal{D}^\alpha_{h,q,M,N} u_\theta, \tag{5.1}$$

where h may denote a mesh or auxiliary-point spacing, q quadrature parameters, M Monte Carlo samples, and N spectral or memory truncation parameters. Classical analyses of discretized fractional calculus, convolution quadrature, graded-mesh L1 schemes, finite elements, fast Caputo evaluation, and nonlocal diffusion provide the relevant consistency and stability standards [2–4, 22–26, 54].

5.1. L1 approximation for Caputo derivatives

The L1 approximation is one of the most common numerical schemes for Caputo time-fractional derivatives. On a uniform grid $t_n = n\tau$, for $0 < \alpha < 1$, it has the form

$$\begin{aligned} {}_0^C D_t^\alpha u(t_n) &\approx \frac{\tau^{-\alpha}}{\Gamma(2-\alpha)} \sum_{j=0}^{n-1} a_j^{(\alpha)} [u(t_{n-j}) - u(t_{n-j-1})], \\ a_j^{(\alpha)} &= (j+1)^{1-\alpha} - j^{1-\alpha}. \end{aligned} \tag{5.2}$$

When embedded into fPINNs, (5.2) makes the residual at t_n depend on previous temporal nodes or auxiliary points. The method is therefore not fully mesh-free; it inherits the memory structure of classical fractional time stepping.

L1-based fPINNs are attractive because they are simple, interpretable, and directly connected to classical numerical analysis. They have been used for time-fractional Huxley equations, reaction-diffusion equations, phase-field models, and other Caputo-type FPDEs [39, 45, 55]. However, the L1 scheme may lose accuracy when the solution has weak initial singularities. This motivates graded meshes, correction terms, fast L1 approximations, and singularity-aware fPINN architectures.

5.2. Higher-order, fast, and corrected Caputo schemes

Higher-order approximations such as L2, L2-1_σ, Hermite interpolation, and corrected L1-type schemes can improve accuracy for sufficiently smooth solutions. For example, L2-type approximations have been used in operator-learning fPINNs for multiterm time-fractional, mixed diffusion-wave equations [14]. Hermite neural solvers replace low-order time-fractional approximations with high-order Hermite explicit schemes [44]. Fast L1 and graded-mesh fPINNs address the combined challenge of memory cost and weak initial singularity [45].

The main limitation of higher-order schemes is that their formal accuracy often assumes smoothness that fractional diffusion solutions may not possess near $t = 0$. Therefore, higher-order fPINNs should be interpreted cautiously unless the paper reports regularity handling, graded sampling, correction terms, or singularity-aware decomposition. This is particularly important in time-fractional diffusion, where early-time behavior can dominate the total error.

5.3. Operational-matrix approximations and scope boundary

Operational-matrix methods assemble the weights of a discrete fractional operator before optimization and evaluate the fractional contribution to the neural residual through matrix-vector products. For Caputo derivatives, this separates matrix construction from training and can improve vectorization and constant-factor computational efficiency, but it does not remove temporal nonlocality. A full lower-triangular matrix remains history dependent and, if stored and applied without sparsity, compression, or other exploitable structure, requires $O(N_t^2)$ work and storage when the derivative is evaluated at all N_t time levels. Taheri et al. [56] proposed a nonuniform, nodal operational matrix for Caputo derivatives, mainly for $0 < \alpha < 1$, and evaluated it on fractional delay differential equations and a fractional differential-algebraic system, together with nonfractional delay and differential-algebraic examples. The study reports feasible matrix-based, fractional-residual evaluation, but it does not provide an FPDE benchmark, a convergence or operator-error analysis, or a matched cost-error comparison with a direct history-summation fPINN. Under the FPDE-specific counting rule adopted in this review, it is therefore retained in the supporting set as a transferable scope-boundary reference and is not included in the 38-record core FPDE/fPINN evidence set. An FPDE extension should specify whether the matrix acts along temporal trajectories or approximates a spatial fractional operator and should report consistency, space-time scaling, peak memory, batching and vectorization, initial-condition enforcement, local or exterior boundary treatment, and coupling with multidimensional spatial operators.

5.4. Grünwald-Letnikov and shifted Grünwald approximations

Grünwald-Letnikov approximations, introduced in (3.3), express fractional operators through weighted finite-difference series. Shifted Grünwald and vector Grünwald formulas are often used

to improve stability or represent multidimensional directional fractional operators. In fPINNs, these approximations make the residual stencil-dependent, mesh-dependent, and potentially dense.

The original fPINN paper used Grünwald-Letnikov-type approximations for fractional operators, making this class foundational for fPINNs [8]. Recent time-space fractional financial models also use matrix formulations of Grünwald-Letnikov derivatives inside PINN losses [47]. The advantage of Grünwald-type approximations is their direct connection with fractional finite differences. However, in high-dimensional or long-time simulations, these weighted sums may become computationally expensive unless combined with truncation, acceleration, or structured matrix techniques.

5.5. Quadrature approximations

Quadrature methods approximate fractional derivatives or fractional integrals directly from their integral definitions. For a generic weakly singular fractional integral,

$$I^{\alpha}u(t) = \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} u(s)\, ds, \tag{5.3}$$

a quadrature approximation takes the form

$$I^{\alpha}u(t) \approx \sum_{m=1}^{Q} \omega_m(t) u(s_m), \tag{5.4}$$

where $\{s_m\}_{m=1}^{Q}$ are quadrature nodes, and $\{\omega_m\}_{m=1}^{Q}$ are kernel-dependent weights.

Quadrature-based fPINNs are natural for Caputo, Riemann-Liouville, distributed-order, and fractional Laplacian operators. In fractional Laplacian problems, quadrature rules must handle both kernel singularity and domain truncation. Modified trapezoidal rules have been used for nonlocal Fokker-Planck equations with fractional Laplacians [48]. Distributed-order PINNs require quadrature not only in time but also over the fractional-order variable [42, 49]. Auxiliary-output fPINNs for fractional partial integro-differential equations further show that quadrature may be coupled with network outputs representing integral terms [57].

The main challenge of quadrature-based fPINNs is that quadrature accuracy, node placement, and singularity treatment directly affect the residual. If these details are not reported, the fPINN result is difficult to reproduce and interpret.

5.6. Spectral approximations

Spectral methods approximate the solution using global basis functions, such as Fourier modes, Jacobi polynomials, Chebyshev polynomials, or problem-adapted eigenfunctions. A generic spectral approximation can be written as

$$u_N(x,t) = \sum_{k=0}^{N} c_k(t)\phi_k(x), \tag{5.5}$$

where $\{\phi_k\}$ is a chosen basis. If the fractional operator acts conveniently on the basis, then residual evaluation can be transformed into operations on spectral coefficients.

Spectral-fPINNs use this idea to reduce the need for dense auxiliary-point constructions and to improve accuracy for smooth problems. Spectral collocation has been used in time-fractional phase-field fPINNs [9]. Jacobi-polynomial spectral-fPINNs provide a more explicit operator-approximation

framework and include error-analysis considerations [52]. Radial basis function (RBF)-fPINNs and fractional spectrally adapted PINNs extend this idea by modifying the approximation space itself [40, 53].

The advantage of spectral approximations is high accuracy for smooth solutions and potential efficiency gains. The limitation is sensitivity to solution regularity, boundary conditions, and domain geometry. For FPDEs with weak singularities, nonsmooth data, or complex boundaries, spectral fPINNs may require enrichment, domain decomposition, or adaptive basis selection.

5.7. Monte Carlo and stochastic approximations

Monte Carlo approximations are attractive for high-dimensional nonlocal fractional operators. For an integral operator of the form

$$\mathcal{I}[u](x) = \int_{\Omega} K(x, y)(u(x) - u(y))\, dy, \tag{5.6}$$

a Monte Carlo estimator may be written as

$$\mathcal{I}[u](x) \approx \frac{|\Omega|}{M} \sum_{m=1}^{M} K(x, Y_m)\,(u(x) - u(Y_m)), \tag{5.7}$$

where Y_m are random samples in Ω. This approximation avoids tensor-product quadrature and is therefore useful for high-dimensional fractional Laplacians and tempered fractional PDEs.

Monte Carlo fPINNs were introduced for high-dimensional forward and inverse FPDEs involving nonlocal fractional operators [10]. Quadrature-enhanced Monte Carlo fPINNs have also been proposed for fractional and tempered fractional PDEs to reduce variance and improve high-dimensional scaling [13]. Other recent studies couple Monte Carlo PINNs with optimization strategies for fractional PDEs [58].

The central limitation is variance. Monte Carlo noise enters the residual and can destabilize gradient-based training. Therefore, variance reduction, quasi-Monte Carlo sampling, Gaussian quadrature corrections, control variates, and adaptive sampling are important directions for future fPINN research.

5.8. Laplace-transform and analytical-prior approximations

Laplace-transform fPINNs approximate time-fractional dynamics by transforming the governing equation into the Laplace domain, solving or constraining the transformed problem, and then using numerical inverse Laplace transforms. This approach can reduce the need for explicit history summations in the time domain. Laplace-fPINNs are therefore a distinct class of analytical-prior fPINNs [12, 46].

A key advantage of the Laplace-transform approach is that fractional derivatives with respect to time can be represented algebraically in the transform domain, thereby avoiding direct numerical discretization of the fractional operator. The main numerical challenge is that recovering the physical-domain solution requires numerical inverse Laplace transformation, which introduces additional approximation and conditioning errors. Moreover, because the Laplace transform is linear, nonlinear terms generally do not retain a simple algebraic form and may require additional treatment or

linearization. Complex spatial geometries, boundary conditions, and the incorporation of sparse or noisy data may also complicate the transformed formulation. From the perspective of this review, Laplace-based fPINNs are important because they demonstrate that fractional residuals can be hybridized not only with numerical discretization but also with analytical integral transforms.

5.9. Fast convolution, sum-of-exponentials, and memory compression

A direct Caputo residual over N_t time levels can require $O(N_t^2)$ cumulative work when the derivative is evaluated at all time points. Classical fractional solvers reduce this cost by approximating the weakly singular kernel with compressed representations. A typical sum-of-exponentials approximation has the form

$$(t-s)^{-\alpha} \approx \sum_{\ell=1}^{L} \omega_\ell e^{-\eta_\ell (t-s)}, \qquad 0 < \alpha < 1, \tag{5.8}$$

where ω_ℓ and η_ℓ are positive weights and decay rates, respectively. This representation allows recursive history updates and can reduce the cost of long-time memory evaluation from quadratic to nearly linear or logarithmic complexity, depending on the implementation.

In fPINNs, fast-memory ideas are attractive because the fractional residual is evaluated repeatedly during training. However, direct transfer from classical time-stepping is not straightforward. PINN collocation points may be scattered, adaptively resampled, or batched, whereas fast convolution is usually designed for structured temporal grids. A practical hybrid strategy is therefore to use structured auxiliary history grids for the fractional residual while retaining mesh-flexible spatial collocation for the neural approximation. Future work should compare direct L1, fast L1, short-memory, and sum-of-exponentials fPINNs under the same architecture and benchmark.

5.10. Operator approximation as a source of residual error

Fractional residual construction introduces an operator-approximation error before neural training begins. Its magnitude depends on the discretization, quadrature resolution, Monte Carlo sample size, spectral truncation, memory approximation, or transform inversion used in the residual. The complete error-decomposition framework and its stability implications are developed in Section 7; Table 9 summarizes the residual-level error mechanisms associated with the principal operator families.

Table 9. Fractional operators, embedded fPINN approximations, and residual-level challenges.

Operator class	Representative FPDE setting	Embedded fPINN approximation	Main residual-level challenge
Caputo derivative	Time-fractional diffusion, diffusion-wave, phase-field, reaction-diffusion, Black-Scholes	L1/L2 schemes, fast L1, graded meshes, Hermite interpolation, Laplace-domain treatments	Long memory, weak initial singularity, and auxiliary temporal history points
Riemann-Liouville derivative	Integral and analytical FPDE formulations	Quadrature, finite differences, and transform-based approximations	Singular-kernel treatment and consistency of initial-condition formulation
Grünwald-Letnikov derivative	Space-time fractional advection–diffusion and fractional financial models	Weighted finite differences, shifted/vector Grünwald formulas, and matrix forms	Dense nonlocal stencils, truncation error, and mesh dependence
Operational-matrix treatment	Caputo-type fractional systems; transferable scope-boundary mechanism	Precomputed derivative matrix and matrix-vector residual evaluation [56]	FPDE applicability, matched cost-error performance, and space-time scaling require validation beyond the reported test setting.
Riesz derivative / fractional Laplacian	Space-fractional diffusion, fractional Poisson, nonlocal Fokker-Planck, high-dimensional FPDEs	Shifted differences, Fourier/spectral representations, integral quadrature, Monte Carlo, and Ritz/variational formulations	Singular spatial integration, exterior-boundary treatment, and high-dimensional nonlocal cost
Tempered fractional operator	Tempered diffusion and truncated anomalous transport	Tempered Grünwald forms, tempered Monte Carlo, and quadrature-enhanced residuals	Kernel-parameter sensitivity and variance–cost tradeoff
Variable-order derivative	Heterogeneous media and variable-order advection-diffusion models	Power-series expansion, adaptive approximation, and learned fractional order	Identifiability of $\alpha(x,t)$ and dynamic operator dependence
Distributed-order derivative	Multiscale memory and ultraslow diffusion models	Quadrature over the order variable and multioutput PINN formulations	Coupled time-order quadrature and parameter identifiability
Conformable derivative	Alternative fractional diffusion models	Automatic differentiation with fractional scaling factors	Local derivative structure; interpretation differs from genuinely nonlocal Caputo/RL models
Transform-based fractional treatment	Time-fractional diffusion and inverse problems with usable analytical structure	Laplace-domain residuals and numerical inverse transforms	Transform-inversion error and limited applicability to nonlinear or irregular-domain settings

These approximation mechanisms define the principal methodological classes used in the operator-aware taxonomy. They also determine which refinement studies, stability arguments, cost measures, and boundary checks are required for a defensible evaluation.

6. Operator-aware taxonomy of hybrid fPINNs

The taxonomy uses a prespecified component-level hierarchy to identify the dominant mechanism used to evaluate the fractional term in the residual. Every numerical approximation is first linked to the operator component on which it acts; architecture , activation, loss weighting, precision, optimization, and approximations of purely integer-order terms are retained as secondary attributes unless they alter

the fractional-operator evaluation. This prevents a new training strategy, or a spatial discretization applied only to an integer-order term, from being conflated with a new fractional-operator treatment. Figure 3 provides the visual classification.

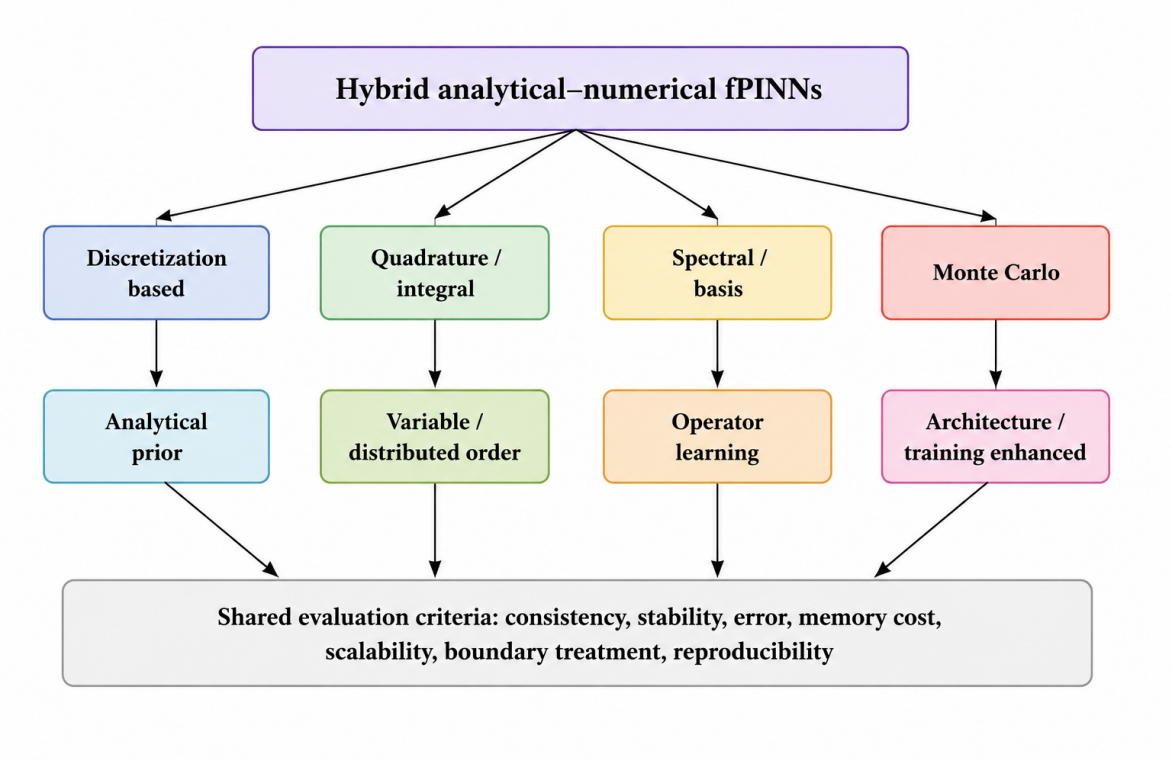

Figure 3. Compact taxonomy of hybrid analytical-numerical fPINNs according to the dominant fractional-operator approximation embedded in the residual.

6.1. Classification rule

Let u_θ be the neural approximation. A generic residual is

$$\mathcal{R}_\theta(x,t) = \mathcal{D}^{\alpha}_{\text{approx}}[u_\theta](x,t) + \mathcal{N}[u_\theta](x,t) - f(x,t), \tag{6.1}$$

where $\mathcal{D}^{\alpha}_{\text{approx}}$ may be a history discretization, finite-difference stencil, operational matrix, quadrature rule, spectral projection, Monte Carlo estimator, transform-domain representation, order-space quadrature, or a physics-informed operator-learning construction. Hybrid studies were coded using the decision rules in Table 10. All mechanisms were retained as tags, but exactly one primary class was used for the mutually exclusive denominators in Table 12.

Table 10. Decision rules for assigning the dominant operator class in hybrid studies.

Decision stage	Operational rule	Consequence/example
1. Component map	Record each approximation and the equation component on which it acts; distinguish fractional from integer-order operators.	An L1 formula for a Caputo derivative and a Fourier method for an integer-order spatial Laplacian are not codominant fractional mechanisms.
2. Direct fractional mechanism	If only one mechanism evaluates a fractional operator, that mechanism determines the primary class.	The L1-Fourier example is classified as discretization-based; Fourier remains a secondary spatial tag. If Fourier instead evaluates a fractional Laplacian, the study proceeds to stage 3.
3. Multiple fractional mechanisms	Select the mechanism that constitutes the principal fractional-operator contribution and is examined in the study's principal analysis, theorem, resolution study, or experiment. If these signals conflict, use the mechanism reported truncation, stochastic, transform-inversion, or order-quadrature error.	An L1 time derivative combined with a spectral fractional-Laplacian approximation is coded by the focal mechanism; the other is recorded as a secondary mechanism and as a plausible alternative assignment.
4. Transform, stochastic, and task exceptions	A transform is analytical-prior when it reformulates or eliminates the nonlocal operator; numerical inversion is secondary unless inversion is the focal approximation. Monte Carlo is primary when the stochastic estimator evaluates the fractional operator. Variable/distributed-order and operator-learning are primary only when order treatment or learning a family-to-solution map is the study's principal task.	These rules distinguish Laplace-plus-quadrature, quadrature-enhanced Monte Carlo, distributed-order-plus-time-difference, and DeepONet-plus-L2 combinations.
5. Residual ambiguity	If stages 1–4 do not yield a unique assignment, retain both tags, mark the alternative explicitly, and include the study in the alternative-assignment stress test; no architecture-based tie-break is used.	The study-level primary/secondary crosswalk and the prespecified alternatives are reported in Supplementary Table S8 and the machine-readable coding file.

This rule changes one assignment relative to the original table. Wang et al. [9] used a second-order fractional backward-difference formula for the Caputo time derivative and spectral collocation for the integer-order spatial direction. The paper is therefore classified as discretization-based under stages 1–2, although spectral collocation is retained as a secondary tag. A novelty-first rule would instead place it in the spectral/basis class, so that reassignment is included in the sensitivity analysis below.

6.2. Discretization-based residuals

Discretization-based fPINNs embed L1/L2, Grünwald–Letnikov, finite-difference, matrix, Hermite, or fast-memory approximations in the loss [8, 39, 44, 45, 47]. Their main advantage is a direct connection to classical fractional schemes; their main limitation is that mesh dependence, temporal history, truncation error, regularity assumptions, and stability restrictions remain present even though the neural representation is continuous. Operational-matrix acceleration is adjacent to this class, but the cited study without an FPDE benchmark is retained only as a transferable scope-boundary reference [56].

6.3. Quadrature and integral residuals

Quadrature-based methods approximate the defining integral of a Caputo, Riemann-Liouville, distributed-order, integro-differential, or fractional-Laplacian operator:

$$\mathcal{D}^{\alpha} u_{\theta}(x,t) \approx \sum_{m=1}^{Q} \omega_m(x,t)\, u_{\theta}(\xi_m, \tau_m). \tag{6.2}$$

Quadrature-based methods apply the integral approximation in (5.4) to Caputo, Riemann-Liouville, distributed-order, integro-differential, or fractional-Laplacian operators [42, 48, 49, 57]. Reliability depends on singular-kernel treatment, node placement, domain truncation, order-space resolution, and explicit nonlocal boundary conditions.

6.4. Spectral and basis-enhanced residuals

The spectral and basis constructions introduced in Section 4 and Section 5 can achieve high accuracy with fewer degrees of freedom when the solution, operator, geometry, and basis are compatible [9, 40, 50, 52, 53]. Their class-level evidence must nevertheless be interpreted through regularity, truncation, conditioning, and boundary compatibility; smooth manufactured solutions alone do not establish robustness to initial or boundary singularities.

6.5. Monte Carlo residuals

Monte Carlo fPINNs use stochastic estimators such as (5.7) for high-dimensional fractional and tempered fractional Laplacians [10, 13, 58]. Their dimensional scalability is obtained at the price of estimator and gradient variance. Sample size, sampling law, random seeds, variance reduction, and confidence intervals are therefore part of the numerical method, not optional implementation details.

6.6. Analytical-prior residuals

Analytical-prior methods incorporate Laplace transforms, inverse transforms, asymptotic structure, singularity decomposition, or compatible basis information [12,41,46]. Such priors can reduce explicit history summation or improve approximation of weak singularities, but introduce transform-inversion, prior-mismatch, or problem-specific modeling errors. Their value is strongest when the analytical structure is stated, independently validated, and compared with a time-domain residual.

6.7. Variable- and distributed-order residuals

Variable-order methods approximate or learn an order $\alpha(x,t)$, whereas distributed-order methods combine temporal history with quadrature over an order interval [11,42,49,50]. In addition to operator approximation, these methods face identifiability, regularization, and order-space discretization errors. Forward accuracy does not by itself establish that an inferred order or order distribution is uniquely or robustly identifiable.

6.8. Operator-learning formulations

Physics-informed operator-learning methods target mappings from forcing functions, coefficients, initial data, or fractional orders to solution fields rather than a single solution. DeepONet provides

the general foundation, and fPINN-DeepONet couples it with an L2 Caputo approximation for multi-term diffusion-wave equations [14, 59]. Amortized inference is potentially valuable in repeated-query settings, but the embedded fractional approximation remains part of the loss, and generalization must be evaluated across both input functions and operator parameters.

6.9. Architecture and training as secondary attributes

Recent advances in fPINNs include auxiliary-output formulations with adaptive loss weighting [57], improved network architectures, adaptive activation functions, and gradient-flow optimization strategies [60–62]. Multiprecision computing has also been explored to enhance the numerical performance of multistage fPINNs [63]. They do not, however, remove the approximation error of the fractional residual. A credible architecture comparison must therefore hold the operator approximation fixed or report a joint ablation.

6.10. Taxonomy summary and cross-class findings

Table 11 summarizes the operator treatment, strength, and best-supported use of each family.

Table 11. Taxonomy of fPINNs and related neural fractional solvers.

Class	Representative references	Operator treatment	Strength	Limitation / application
Discretization-based	[8, 39, 44, 45, 47]	L1/L2, fast L1, Grünwald-Letnikov, Hermite, or finite-difference matrices	Close to classical schemes	Mesh and memory costs; diffusion, reaction-diffusion, and financial FPDEs
Quadrature/integral	[42, 48, 49, 57]	Quadrature of fractional integrals, Laplacians, distributed-order terms, or FPIDEs	Flexible for nonlocal models	Kernel singularities and quadrature error; integral and distributed-order problems
Spectral/basis	[9, 40, 50, 52, 53]	Spectral collocation, Jacobi bases, RBFs, or fractional polynomials	Potentially high accuracy for sufficiently smooth solutions	Basis and geometry dependent; smooth benchmark FPDEs
Monte Carlo	[10, 13, 58]	Monte Carlo or stochastic quadrature for fractional and tempered operators	Offers potential scalability to high-dimensional problems	Sampling variance and noisy gradients; high-dimensional FPDEs
Analytical-prior	[12, 41, 46]	Laplace transforms, inverse transforms, or singularity decomposition	Uses known solution structure	Problem specific; inverse, singular, and time-fractional problems
Variable/distributed order	[11, 42, 49, 50]	Variable-order discretization, order-space quadrature, or adaptive bases	Models heterogeneous memory	Identifiability and order-quadrature error; multiscale-memory models
Operator learning	[14, 59]	DeepONet or neural operators with fractional physics constraints	Efficient repeated inference after training within the learned problem family	Needs representative training data; parametric and repeated-query FPDEs
Training enhanced	[57, 60–63]	Adaptive activations, loss weighting, auxiliary outputs, multistage or multiprecision training	Improves optimization	Does not remove operator-approximation error; stiff training regimes

To move beyond descriptive classification, the primary studies were stratified by dominant operator

class. Table 12 compares formal theory, direct classical baselines, detailed cost evidence, operator ablation, and full reproducibility. The resulting pattern is evaluated together with the alternative-assignment stress test in Table 13; class percentages for small groups are treated as descriptive rather than inferential.

Table 12. Class-stratified evidence characteristics under the operator-aware taxonomy. Values are reported as n/N; architecture and training modifications are treated as secondary attributes.

Dominant operator class	N	Formal theory	Classical baseline	Detailed cost	Operator ablation	Full reprod.	Principal evidence limitation
Discretization-based	15	1/15	6/15	4/15	7/15	5/15	Long-memory cost, initial singularities, mesh dependence, and insufficient controlled comparison of alternative fractional schemes.
Quadrature/integral	4	1/4	1/4	0/4	0/4	1/4	Singular-kernel resolution, domain truncation, quadrature sensitivity, and nonlocal boundary reporting.
Spectral/basis	7	4/7	1/7	0/7	1/7	0/7	Dependence on solution smoothness, basis compatibility, geometry, and truncation resolution.
Monte Carlo	3	0/3	0/3	1/3	1/3	2/3	Residual and gradient variance, limited confidence-interval reporting, and absence of matched deterministic cost–accuracy comparisons.
Analytical-prior	3	0/3	1/3	1/3	0/3	0/3	Transform-inversion or prior-mismatch error and challenges in extending analytical-prior formulations to less structured problems.
Variable/distributed-order	3	1/3	2/3	2/3	1/3	0/3	Identifiability, order-space discretization, regularization, and uncertainty quantification.
Operator-learning	1	0/1	0/1	0/1	0/1	1/1	Limited out-of-distribution evaluation, a one-study evidence base, and the absence of standardized parametric FPDE benchmarks.
Scope-boundary conformable derivative	1	0/1	0/1	0/1	0/1	0/1	The derivative is local and chain-rule compatible and is therefore not directly comparable with nonlocal Caputo, Riemann-Liouville, Riesz, or fractional-Laplacian formulations.

Table 13. Sensitivity of the class-stratified findings to simultaneous alternative assignment of all 10 hybrid studies with a plausible secondary class. Each evidence vector is ordered as formal theory/classical baseline/detailed cost/operator ablation/full reproducibility. A dash indicates that no study remains in that class under the stress assignment.

Class	Primary N	Primary evidence vector	Stress-test N	Stress-test evidence vector
Discretization-based	15	1/15; 6/15; 4/15; 7/15; 5/15	13	0/13; 5/13; 2/13; 6/13; 6/13
Quadrature/integral	4	1/4; 1/4; 0/4; 0/4; 1/4	8	1/8; 1/8; 2/8; 1/8; 2/8
Spectral/basis	7	4/7; 1/7; 0/7; 1/7; 0/7	8	3/8; 2/8; 2/8; 1/8; 0/8
Monte Carlo	3	0/3; 0/3; 1/3; 1/3; 2/3	2	0/2; 0/2; 0/2; 0/2; 1/2
Analytical-prior	3	0/3; 1/3; 1/3; 0/3; 0/3	3	1/3; 2/3; 1/3; 1/3; 0/3
Variable/distributed-order	3	1/3; 2/3; 2/3; 1/3; 0/3	2	2/2; 1/2; 1/2; 1/2; 0/2
Operator-learning	1	0/1; 0/1; 0/1; 0/1; 1/1	0	–
Scope-boundary conformable derivative	1	0/1; 0/1; 0/1; 0/1; 0/1	1	0/1; 0/1; 0/1; 0/1; 0/1

Supplementary Table S8 details the primary and secondary operator mechanisms and alternative classifications of the 10 hybrid studies considered in the sensitivity analysis. These assignments provide the basis for assessing the robustness of the class-level findings reported in Table 13. The stress test deliberately applies all plausible secondary assignments simultaneously and therefore represents a conservative perturbation rather than a competing preferred taxonomy. The main absolute-count findings are stable: Discretization-based studies remain the largest class and retain the most classical baselines and operator ablations, whereas spectral/basis studies retain the largest number of formal-theory contributions and still have no fully reproducible study. Proportions in the two- and three-study classes change materially, and the singleton operator-learning class disappears when its embedded L2 approximation is made primary. We therefore interpret small-class percentages cautiously, retain the multilabel crosswalk, and do not use Table 12 to infer an architecture-independent causal effect.

Discretization-based methods form the largest class and provide the greatest number of classical comparisons and operator ablations, but still lack general end-to-end convergence and consistent long-memory cost reporting. Spectral/basis methods contain a comparatively high concentration of formal approximation analysis but no full reproducibility and limited evidence for nonsmooth solutions. Monte Carlo studies provide stronger code availability within a small evidence base but no formal theory or matched deterministic baselines in the present coding. Variable/distributed-order studies report relatively more detailed cost evidence, although identifiability and uncertainty remain unresolved. The operator-learning class is represented by one primary study and is not preserved under the alternative-assignment stress test; its apparent reproducibility cannot substitute for independent replication or out-of-distribution validation. These class-specific differences support operator-dependent evaluation standards rather than a single generic fPINN checklist, but the sensitivity results caution against over interpreting rates from sparse classes.

7. Reliability, evidence maturity, and practical interpretation

This section integrates stability and error analysis with the quantitative evidence audit. The purpose is not to rank architectures, but to identify what must be controlled before an fPINN result can be interpreted as evidence for the intended FPDE rather than only for an approximated residual problem.

7.1. Error decomposition and residual perturbation

Let u be the FPDE solution and u_θ the trained neural approximation. A useful conceptual bound is

$$\|u - u_\theta\|_X \lesssim C_{\mathrm{stab}} \left(\varepsilon_{\mathrm{op}} + \varepsilon_{\mathrm{res}} + \varepsilon_{\mathrm{bc/ic}}\right) + \varepsilon_{\mathrm{NN}} + \varepsilon_{\mathrm{opt}} + \varepsilon_{\mathrm{samp}} + \varepsilon_{\mathrm{round}}, \tag{7.1}$$

where X is a problem-appropriate norm, and C_{stab} denotes a stability factor. This factor is determined by the well-posedness and continuous-dependence estimate for the particular FPDE; it may depend on the fractional order, coefficients, domain, boundary conditions, final time, and chosen norms. It can become large near an ill-conditioned or weakly stable regime and need not exist for an ill-posed formulation. The terms represent operator approximation, evaluated residual, boundary/initial-condition enforcement, neural approximation, optimization, sampling/generalization, and round-off or implementation error. Equation (7.1) is therefore an organizing inequality rather than a universal theorem . A computable stability factor or compatible residual-to-solution estimate is rarely available

for the FPDEs and trained fPINNs considered here, so a small right-hand-side component should not be interpreted as a certified solution-error bound. The precise inequality must instead be established for the specific operator, function space, and well-posedness setting using an available FPDE stability estimate .

The defining additional term for fPINNs is ε_{op}. In practice,

$$\mathcal{D}^{\alpha} u_{\theta} = \mathcal{D}^{\alpha}_{h,q,M,N} u_{\theta} + \epsilon_{\text{op}}, \tag{7.2}$$

where h, q, M, and N may represent a history-grid spacing, quadrature rule, Monte Carlo sample size, or spectral/memory truncation. Minimizing a residual based on $\mathcal{D}^{\alpha}_{h,q,M,N}$ can at best solve a perturbed residual equation unless the approximate operator converges in a norm compatible with the stability of the original FPDE. A small training residual is therefore insufficient when the operator approximation is unresolved. These components are not independently observable in a trained model and need not contribute additively in practice. Refining the operator can leave the total error unchanged when network capacity, optimization, or sampling already sets an error floor. Conversely, increasing network capacity cannot remove the bias of an unresolved operator. The coupling is two-way because changing h, q, M, or N changes the residual landscape and its gradients, whereas the regularity and frequency content of u_{θ} affect the error made when the fractional operator is applied to it. Empirical separation therefore requires controlled, nested ablations. An operator-resolution sweep should vary only h, q, M, or N while holding the architecture, collocation design, optimization budget, loss weights, and preferably random seeds fixed; the resulting change in error would be a conditional proxy for ε_{op}, and a plateau would indicate that another component may dominate. A complementary capacity sweep should vary network width, depth, or basis richness with a fixed, independently validated operator approximation to probe ε_{NN}, although repeated seeds would be needed to limit confounding by ε_{opt}. A two-factor grid crossing operator resolution with network capacity is more informative than either one-factor sweep because nonparallel convergence curves reveal interaction. In every case, solution error should be measured against an independent reference solution, and the residual should be re-evaluated with a finer or otherwise independent operator, as evaluation with the same operator used in training can conceal operator bias. These diagnostics estimate attribution; they do not identify the individual terms exactly without a problem-specific stability and approximation theory. Figure 4 summarizes the resulting reliability pathway.

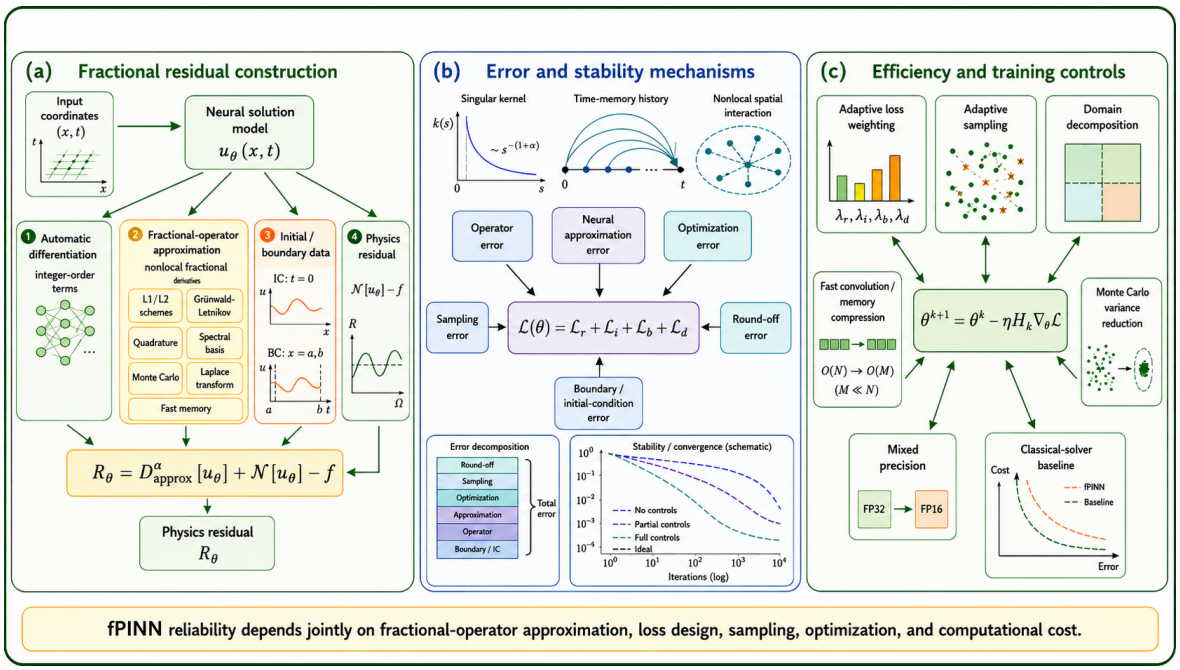

Figure 4. Reliability pipeline for hybrid fPINNs, showing how fractional residual construction, error mechanisms, sampling, optimization, and efficiency controls jointly affect robustness and reproducibility.

7.2. Class-specific stability mechanisms

Stability mechanisms differ across fractional-operator classes. History-discretization methods inherit temporal-memory accumulation, reduced initial regularity, and potential stability restrictions from their embedded numerical schemes. Quadrature-based methods may encounter difficulties associated with singular kernels, nonlocal boundary conditions, and numerical integration. Spectral and basis-enhanced methods depend on solution regularity, basis compatibility, truncation, and coefficient conditioning. Monte Carlo residuals introduce stochastic fluctuations into residual estimates and training gradients. Analytical-prior methods may incur errors arising from numerical inverse transforms or mismatched asymptotic assumptions. Variable- and distributed-order inverse problems introduce additional parameter-identifiability and regularization challenges, whereas operator-learning methods require careful assessment of generalization beyond the training distribution.

These effects are compounded by PINN training challenges, including imbalanced loss convergence, inefficient collocation sampling, and difficulties in adaptive loss weighting [64–66]. Variational formulations and domain decomposition offer alternative approaches [67], although optimization improvements cannot compensate for inconsistencies in the underlying fractional operator.

7.3. Evidence coverage and defensible present use

The class differences above should be interpreted as evidence coverage rather than performance rankings. Linking that coverage to defensible present use requires keeping the quantitative indicators separate from application-oriented judgment. Figure 5 presents the reproducible class-

level coverage calculation, and Table 14 provides a separate, explicitly interpretive account of defensible present uses and unresolved limitations. The figure uses the corrected dominant classes from Table 12; it does not split discretization subclasses that were not separately coded in the quantitative audit. Supplementary Table S9 provides the class-level data and evidence-coverage scores underlying Figure 5, including study counts and indicators of theory, comparison, cost reporting, and reproducibility. These scores were calculated using the aggregation rules in Table 4 and the study-level audit in Supplementary Table S7, ensuring transparency and verifiability.

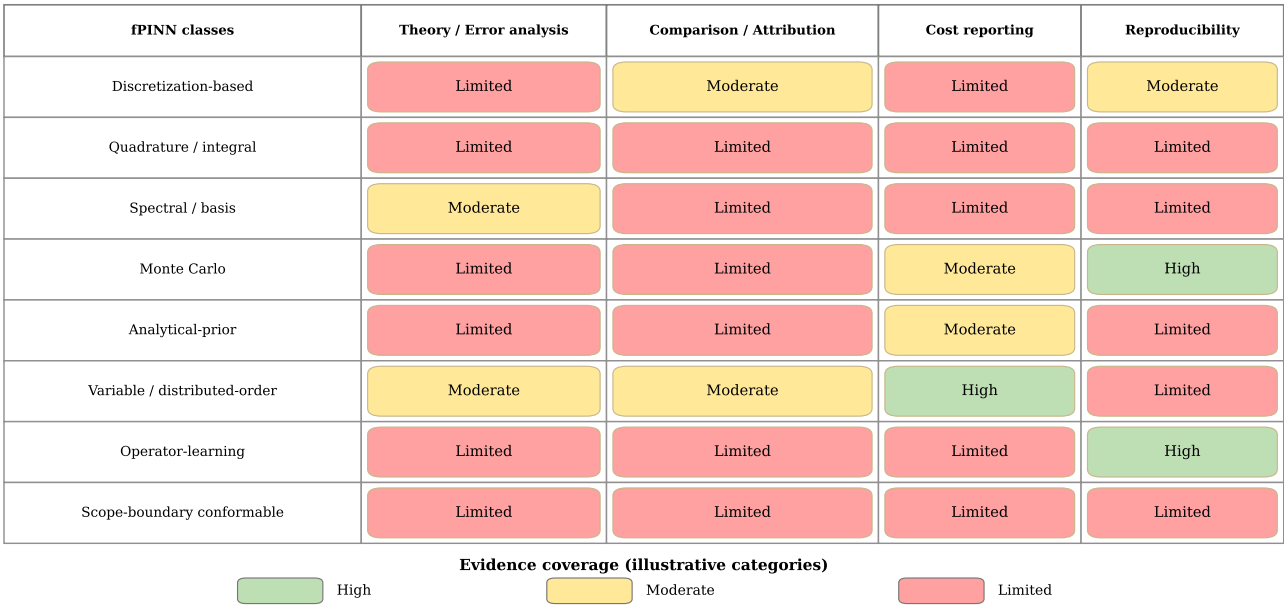

Figure 5. Class-level evidence coverage calculated from the study-level audit ($N = 37$). Theory/error-analysis coverage is T_c/N_c, where T_c is the number of studies reporting formal convergence theory. Comparison-and-attribution coverage is the equally weighted mean of direct-classical-baseline and fractional-operator-ablation coverage, $(B_c + A_c)/(2N_c)$, where B_c and A_c denote the corresponding numbers of studies. Cost-reporting coverage is D_c/N_c, where D_c is the number of studies providing detailed computational-cost evidence, and reproducibility coverage is R_c/N_c, where R_c is the number meeting the full-reproducibility criteria. Each cell reports the percentage and its underlying count or counts. The continuous color scale is descriptive and does not represent categorical maturity levels or comparative method performance.

Table 14. Interpretive readiness and limits of major fPINN classes. This table is a qualitative application-oriented synthesis and is not used to calculate Figure 5.

Class	Evidence interpretation	Defensible present use	Main unresolved issue
L1/L2 Caputo discretization	Relatively well-established discretization approaches within the reviewed literature, although evidence coverage remains incomplete	Controlled forward/inverse benchmarks when regularity, history grid, and operator formula are reported	No general end-to-end convergence guarantee; long-memory cost and initial singularities
Grünwald/matrix discretization	Sparse class-specific reporting within the broader discretization evidence base	Low-dimensional space-time fractional benchmarks with explicit stencil and boundary treatment	Dense stencils, truncation sensitivity, and limited operator ablation
Quadrature, spectral, and analytical-prior methods	Heterogeneous and strongly problem-dependent evidence	Smooth or structured problems aligned with the quadrature, basis, transform, or variational formulation	Singular kernels, geometry dependence, transform inversion, and weak-regularity behavior
Monte Carlo fPINNs	Small evidence base emphasizing high-dimensional feasibility	High-dimensional feasibility studies where conventional deterministic quadrature may be computationally expensive	Residual and gradient variance, confidence intervals, and cost–accuracy comparisons
Variable/distributed-order fPINNs	Small evidence base emphasizing proof-of-concept forward and inverse studies	Proof-of-concept studies with explicit order quadrature and regularization	Identifiability, order-space discretization error, and uncertainty quantification
Operator-learning fPINNs	Single-study primary evidence base in the present corpus	Repeated-query problems within the reported training distribution	Few standardized FPDE datasets and insufficient out-of-distribution testing

The coverage scores should not be interpreted as universal rankings or as estimates of method performance. They measure only the prevalence of prespecified evidence items within each class. The 100% reproducibility cell for operator learning, for example, represents $1/1$ study and therefore does not establish mature or independently replicated evidence. Even the comparatively best-supported settings lack a general convergence theory that simultaneously controls operator approximation, network approximation, optimization, sampling, and condition enforcement.

7.4. Operator ablation and convergence evidence

Only 10/37 primary studies in the quantitative audit varied the fractional approximation or its resolution. This is a central attribution problem: Varying the depth, activation, loss weights, or optimizer while fixing the operator cannot establish whether an accuracy gain is numerical or neural. Informative comparisons include L1 versus L2, uniform versus graded history grids, direct versus fast memory, deterministic quadrature versus Monte Carlo, finite-difference versus spectral residuals, and time-domain versus transform-domain formulations under matched learning settings.

Formal convergence evidence is similarly sparse. The foundational fPINN work separates several error sources, neural Ritz formulations provide energy-norm analysis for integral fractional Laplacians, spectral-fPINNs connect basis approximation with error behavior, and recent work studies PINN convergence for bounded-domain fractional diffusion [8, 38, 43, 52]. A general theory remains open because the residual operator, neural approximation, optimization, and sampling errors are coupled. Table 15 consolidates the main cross-cutting weaknesses.

Table 15. Evidence-weighted gaps in the current fPINN literature.

Gap	Evidence in current literature	Consequence	Priority
Limited operator-ablation studies	Many studies fix one fractional approximation while varying the network or optimizer	Accuracy gains are difficult to attribute	High
Weak cost reporting	Residual-evaluation time, memory, hardware, and auxiliary-point cost are often omitted	Efficiency claims are hard to verify	High
Insufficient uncertainty reporting in Monte Carlo fPINNs	Variance and confidence intervals are often absent	Stochastic residual reliability is unclear	High
Limited boundary analysis for fractional Laplacians	Exterior-domain treatment is often underreported	Reported errors may hide boundary inconsistency	High
Sparse convergence theory	Most studies are empirical and problem-specific	Reliability cannot be certified across operators and geometries	High
Limited operator-learning benchmarks	Few parametric FPDE datasets are standardized	Generalization claims remain difficult to compare	Medium

7.5. When classical solvers remain preferable

For a low-dimensional forward problem with fixed parameters, regular geometry, and a credible mesh-based formulation, finite-difference, finite-element, spectral, convolution-quadrature, or fractional-Laplacian methods may be cheaper, more accurate, and easier to verify [3, 25, 26, 54, 68]. fPINNs are most defensible when they exploit information or amortization that classical solvers do not naturally provide: sparse or irregular observations, inverse identification, uncertain coefficients, high-dimensional nonlocal operators, variable/distributed order, or repeated parametric queries. The relevant comparison is therefore problem-specific accuracy and cost, not the generic claim that a neural method is mesh-free or universally superior.

8. Computational efficiency and scaling

The quantitative audit yielded some computational information in 22/37 primary studies and detailed cost information in 8/37 primary studies. The fractional residual is an important consideration because, in an fPINN, it often incurs a greater cost than the network itself. Epoch counts and numbers of trainable parameters are therefore insufficient efficiency measures without reporting the history, quadrature, stochastic, spectral, or order-space workload.

8.1. Cost model and residual scaling

Let N_{epoch} denote the number of training epochs, N_r the number of residual collocation points, N_t the number of temporal history points, Q the number of deterministic quadrature nodes, M the number of Monte Carlo samples, N_s the number of spectral modes, and P the number of trainable network parameters. A schematic training cost is

$$C_{\text{train}} \approx N_{\text{epoch}}\Big[C_{\text{AD}}(P, N_r) + C_{\text{frac}}(N_r, N_t, Q, M, N_s) + C_{\text{opt}}(P)\Big], \tag{8.1}$$

where automatic differentiation, fractional residual evaluation, and optimization are separated. The decomposition makes clear that a modest network can still be expensive when C_{frac} is large.

Table 16 gives indicative residual-level scaling. The estimates require implementation assumptions and should not be interpreted as hardware-independent runtime guarantees.

Table 16. Indicative asymptotic costs for fractional residual evaluation in fPINNs. Actual runtime depends on implementation, batching, vectorization, hardware, and caching.

Residual approximation	Indicative arithmetic cost	Main cost driver	Suggested reporting items
Caputo L1/L2 history summation	$O(N_r N_h)$ per residual sweep; $O(N_t^2)$ for a complete N_t-level trajectory	Long temporal memory and growth of the number of history terms with evaluation time	Number of time levels, maximum and average history length, time step, history storage, and uniform versus graded mesh
Fast-memory Caputo approximation	$O(LN_t)$ per structured trajectory; $O(N_t \log N_t)$ when $L = O(\log N_t)$	Recursive updates of L compressed-kernel or sum-of-exponentials states	Compression tolerance, number L of exponentials or memory states, grid structure, state-reuse strategy, storage, and kernel-approximation error
Fractional-Laplacian deterministic quadrature	$O(N_r Q)$	Domain-wide nonlocal integration using Q quadrature nodes per residual point	Quadrature rule, node count, singularity treatment, domain truncation, spatial dimension, and exterior-condition handling
Monte Carlo fractional residual	$O(N_r M)$	Stochastic operator evaluation using M samples per residual point	Sample size, sampling law, seed policy, variance-reduction method, repeated-run dispersion, and confidence intervals
Spectral fractional approximation	$O(N_s \log N_s)$ per FFT-based transform on an N_s-point grid; $O(N_s^2)$ for direct dense evaluation of N_s retained modes at N_r residual points	Spectral transforms, dense operator matrices, and basis-projection operations	Basis type, number of modes, dimensional convention for N_s, transform or matrix implementation, truncation criterion, and conditioning
Distributed-order residual	$O(N_r Q_\alpha N_h)$ for direct temporal-history evaluation	Coupled quadrature over fractional order and temporal history	Number Q_α of order nodes, order-quadrature rule, history length, temporal discretization, weight function, and use of direct or compressed memory

Notation: N_r is the number of residual points; N_t is the number of temporal levels; $N_h \le N_t$ is the maximum number of history nodes used for one residual evaluation; Q is the number of deterministic quadrature nodes; M is the number of Monte Carlo samples; N_s is the number of spectral modes; Q_α is the number of order-space quadrature nodes; and L is the number of compressed-kernel states. Fast-memory scaling assumes that recursive history states can be reused on a structured temporal grid.

8.2. Temporal memory and spatial nonlocality

Because the cumulative history work described in Section 5 is repeated throughout training, it can dominate automatic differentiation. Fast L1, graded meshes, Hermite approximations, and Laplace-domain formulations pursue different strategies: memory compression, improved accuracy per history point, and analytical transformation [12, 44–46]. Classical sum-of-exponentials and fast-convolution methods are promising, but their transfer to scattered or adaptively resampled collocation sets requires structured auxiliary history grids or matrix-free implementations. For fractional Laplacians and Riesz operators, direct quadrature depends on the number of residual points and integration points and requires the treatment of singular kernels and exterior domains. Dimension-dependent growth in quadrature cost can be avoided by using Monte Carlo estimators, with the trade-off that they may lead to higher estimator variance, which can either increase the optimization time or decrease the accuracy [10, 13]. A useful efficiency comparison should therefore measure errors relative to the total computational budget, rather than relative to the speed of residual evaluation.

8.3. Order-space, basis, and operator-learning trade-offs

Distributed-order residuals combine temporal history with quadrature over the fractional order,

$$\int_{\alpha_{\min}}^{\alpha_{\max}} w(\alpha)\, {}_0^C D_t^{\alpha} u(x,t)\, d\alpha, \tag{8.2}$$

so order resolution multiplies the cost of the time-fractional approximation [42, 49]. Spectral and basis-enhanced methods can reduce the number of degrees of freedom or auxiliary evaluations when the fractional operator and solution are well-represented in the selected basis. Their efficiency and accuracy nevertheless depend on solution regularity, basis compatibility, truncation resolution, and domain geometry [40, 52].

Operator learning changes the economic question from the cost of one solve to the cost of training plus repeated inference. Amortized inference is meaningful only when the training distribution is representative, the cost of generating any training data is disclosed, and both in-distribution and out-of-distribution errors are reported [14, 69, 70].

8.4. Precision, hardware, and minimum cost evidence

Long history sums, slowly decaying weights, singular kernels, and spectral transforms can be sensitive to numerical precision [63]. GPU acceleration benefits batched network evaluation but may not accelerate irregular history or quadrature operations unless they are vectorized. Future studies should therefore separate network and operator costs. Reports should specify residual, boundary, initial, data, auxiliary, quadrature, or Monte Carlo point counts, together with approximation parameters, architecture, and numerical precision. They should also provide wall-clock training and inference times, peak memory, hardware and software, repeated-run variability, and error-versus-cost curves against an appropriate classical solver.

9. Benchmarks, reproducibility, and reporting standards

Cross-study comparison is difficult because nominally similar FPDEs can be implemented with different fractional definitions, auxiliary grids, quadrature rules, exterior conditions, sampling

distributions, reference solutions, and computational budgets. Benchmarking must therefore evaluate both the learned solution and the numerical approximation embedded in the residual.

9.1. Benchmark design, tiering, and taxonomy alignment

A useful benchmark suite should combine three levels. Tier 1 verifies implementation with smooth exact or manufactured solutions. Tier 2 introduces characteristic fractional difficulties such as weak initial singularities, singular spatial kernels, nonlocal exterior conditions, or order-space quadrature. Tier 3 tests settings in which fPINNs may offer a substantive advantage: noisy inverse problems, irregular observations, high-dimensional nonlocal operators, variable/distributed order, and repeated-query operator learning.

Each benchmark must match the claimed method. A smooth time-fractional diffusion problem can compare L1, L2, fast-memory, Hermite, spectral, or Laplace-domain residuals, but cannot by itself validate a high-dimensional Monte Carlo fractional-Laplacian solver. Likewise, a single manufactured solution is insufficient for claims about singularity robustness or out-of-distribution generalization. Table 17 maps benchmark families to their main diagnostic purpose.

Table 17. Recommended benchmark families for future fPINN evaluation. The table proposes benchmark categories rather than reporting a standardized public benchmark suite.

Benchmark family	Main fPINN class evaluated	Main diagnostic purpose
Time-fractional diffusion with smooth exact solution	Discretization-based, Laplace-transform, spectral fPINNs	Verifies Caputo residual implementation and baseline accuracy
Time-fractional diffusion with weak initial singularity	Fast-L1, graded-mesh, singularity-aware, analytical-prior fPINNs	Tests robustness to limited regularity near $t = 0$
Space-fractional advection-diffusion	Grünwald-Letnikov, Riesz, spectral fPINNs	Evaluates nonlocal spatial stencils and mesh dependence
Fractional Laplacian on bounded domain	Quadrature, Monte Carlo, variational/Ritz, spectral fPINNs	Tests singular spatial integration and exterior boundary treatment
High-dimensional fractional or tempered FPDE	Monte Carlo and quadrature-enhanced Monte Carlo fPINNs	Tests scalability, variance control, and curse-of-dimensionality claims
Fractional phase-field or reaction-diffusion model	Spectral, L1/L2, adaptive-loss, architecture-enhanced fPINNs	Tests nonlinear stability, stiffness, and long-time behavior
Variable-order FPDE	Variable-order fPINNs and inverse fPINNs	Tests identifiability and order-dependent residual construction
Distributed-order diffusion	Distributed-order and multioutput fPINNs	Tests order-space quadrature and multiscale memory effects
Parametric fractional diffusion-wave equation	fPINN-DeepONet and neural-operator methods	Tests solution-operator learning and out-of-distribution generalization
Inverse fractional PDE benchmark	All fPINN classes with data loss	Tests parameter recovery, noise robustness, and regularization

9.2. Reference solutions and error metrics

Reference solutions should be exact, manufactured, transform-based, or generated by a converged classical method. Numerical references must report mesh or modal resolution, time step, quadrature, tolerance, and convergence verification; otherwise, the reported neural error may partly measure reference-solver error. At minimum, solution accuracy should include

$$\mathrm{Rel}\, L^2 = \frac{\|u_\theta - u_{\mathrm{ref}}\|_2}{\|u_{\mathrm{ref}}\|_2}, \tag{9.1}$$

$$\mathrm{Rel}\, L^\infty = \frac{\|u_\theta - u_{\mathrm{ref}}\|_\infty}{\|u_{\mathrm{ref}}\|_\infty}, \tag{9.2}$$

$$\mathrm{Err}_p = \frac{|p_\theta - p_{\mathrm{true}}|}{|p_{\mathrm{true}}|}, \tag{9.3}$$

where p is an inferred quantity such as a fractional order or diffusion coefficient. A separate test-set residual measure is

$$\mathcal{R}_{\mathrm{err}} = \left(\frac{1}{N_{\mathrm{test}}} \sum_{j=1}^{N_{\mathrm{test}}} |\mathcal{R}_\theta(x_j, t_j)|^2 \right)^{1/2}. \tag{9.4}$$

Residual error should be evaluated with an independently refined operator whenever possible. A small residual computed by the same under-resolved approximation used in training does not establish a small solution error. Energy norms are appropriate for variational fractional-Laplacian methods, whereas stochastic residuals require means, dispersion, and intervals over repeated runs.

9.3. Minimum reporting and reproducibility standard

A compilation of the information needed to distinguish operator, neural, optimization, and reference-solver effects is provided in Table 18. It is important to distinguish training, validation, in-distribution test, and out-of-distribution test sets in operator-learning studies. For Monte Carlo methods, the following should be included: sampling law, sample size, seed policy, variance reduction, and intervals. If this standard were applied consistently, it would be possible to compare accuracy and cost without conflating better neural optimization with a more accurate or more expensive fractional residual.

Table 18. Minimum reporting checklist for future fPINN papers.

Item	Required information
Problem definition	FPDE, domain, final time, boundary conditions, initial conditions, source term, exact/reference solution
Fractional operator	Operator type, fractional order, kernel, normalization constant, memory interval, exterior-domain treatment, boundary interpretation
Operator approximation	L1/L2/Grünwald/quadrature/spectral/Monte Carlo/Laplace/Hermite details, step sizes, quadrature nodes, spectral modes, stencil length, Monte Carlo samples
Neural network	Architecture, number of layers, neurons per layer, activation function, initialization, parameter count, hard/soft constraint design
Training	Optimizer, learning rate, epochs, batch size, stopping criterion, loss weights, adaptive weighting strategy, precision
Sampling	Residual points, boundary points, initial points, data points, auxiliary fractional points, quadrature points, sampling distribution, adaptive strategy
Error metrics	Relative L^2, relative L^∞, residual error, parameter error, energy norm where appropriate, confidence intervals for stochastic methods
Computational efficiency	Residual-evaluation time, total training time, inference time, memory use, hardware, CPU/GPU type, software framework
Baseline comparison	Classical finite difference, finite element, spectral, quadrature, Monte Carlo, analytical, or semianalytical reference method
Ablation study	Effect of operator approximation, loss weights, sampling density, auxiliary points, network size, and fractional order
Reproducibility	Code, data, random seeds, reference-solver settings, hyperparameters, trained model if applicable, and version information

10. Research priorities and future directions

The evidence synthesis identifies four priorities that are both technically important and directly connected to the documented gaps.

10.1. Operator-consistent and memory-efficient residuals

Fast-memory L1 schemes, sum-of-exponentials compression, matrix-free structured operators, weak or Ritz formulations, singularity-aware trial spaces, and adaptive quadrature should be evaluated under matched architectures and computational budgets. The key question is not whether these mechanisms can be embedded in a network, but whether they preserve operator consistency while reducing end-to-end cost. Long-time Caputo problems and bounded-domain fractional Laplacians are the most immediate test cases.

10.2. Coupled convergence, uncertainty, and identifiability theory

A general fPINN theory must connect the consistency of the approximate fractional operator to stability of the FPDE and then account for network approximation, optimization, sampling, and condition enforcement. Monte Carlo residuals require nonasymptotic variance and gradient-error analysis; variable/distributed-order inverse problems require identifiability, regularization, and uncertainty bounds. Reporting a point estimate for a fractional order is not sufficient when multiple order functions or kernels can explain the same sparse data.

10.3. Operator learning and open benchmark infrastructure

Physics-informed operator learning is promising for repeated-query FPDEs, but the present evidence base is too small for broad claims. Open parametric datasets should cover forcing functions, coefficients, geometries, initial conditions, fractional orders, and operator definitions, with explicit in- and out-of-distribution partitions. Shared implementations of classical baselines and residual operators would allow the cost of training to be compared with the break-even number of repeated solves.

10.4. Application-specific model validation and engineering design

Most retained studies concern canonical diffusion, phase-field, financial, or high-dimensional nonlocal benchmarks. Proposed applications in cyber-security, cyber-physical monitoring, communications, or other data-rich domains remain exploratory. Such applications should first justify why a particular fractional operator represents the observed memory or nonlocality, establish parameter identifiability, and compare against established stochastic-process, time-series, communication-system, or domain-specific baselines. Application novelty cannot substitute for validation of the fractional model itself.

Engineering design is a plausible extension, particularly when the governing physics contains hereditary material response, anomalous transport, nonlocal damage, or another mechanism for which a fractional model is independently justified. Structural topology optimization provides a concrete example. Its current computational workflows include phased or two-stage physics-informed refinement, parameterized level-set representations, and asymptotic-concentration schemes for driving intermediate material densities toward discrete designs [35–37]. These developments provide potential interfaces for fPINNs rather than direct evidence of their effectiveness. An fPINN could serve as a state or adjoint surrogate inside a staged optimization loop; infer fractional constitutive or transport parameters from sparse measurements; or amortize repeated FPDE solves across loads, geometries, and material parameters through operator learning. Parameterized level-set variables could supply a compact geometry input, and an asymptotic-concentration or projection step could remain responsible for manufacturable near-discrete material layouts.

Such integration would require stricter evidence than a forward-solution benchmark. The complete design loop should be compared with an established finite-element and adjoint-based topology optimizer using matched constraints and stopping rules. Evaluation should report feasibility; compliance or the relevant engineering objective; stress and manufacturing constraints; design sensitivity accuracy; operator and mesh refinement; robustness to unseen loads and geometries; and total cost, including training and reference-data generation. For safety-critical structures, final candidates should be independently verified with a high-fidelity solver. Until these tests are available,

topology optimization should be presented as an application prospect for fPINNs, not as an established capability of the reviewed evidence base.

11. Conclusions

This review treats the embedded fractional-operator approximation as the analytical unit of fPINN research. That perspective changes the interpretation of the field: fPINNs are not a homogeneous class of mesh-free neural solvers, but a collection of hybrid analytical-numerical-learning methods whose residuals inherit different consistency errors, stability requirements, boundary conditions, and computational bottlenecks.

The quantitative audit shows that the literature remains stronger in benchmark accuracy reporting than in evidence attribution and reproducibility: Formal theory appeared in seven studies, operator ablation and public code in 10 each, direct classical comparison in 11, detailed cost in eight, and full reproducibility in nine. Direct comparisons showed regime-specific competitiveness rather than a class-wide advantage. Favorable reported accuracy was often confined to coarse or sparse grids, smooth manufactured solutions, or a method's target singularity, whereas refinement, optimization cost, and high-order instability supplied explicit counterexamples.

The most defensible evidence concerns controlled Caputo problems with transparent L1/L2-type residuals and credible reference solutions, although even this setting lacks a general end-to-end convergence theory. Spectral, quadrature, and analytical-prior methods can be compelling when regularity, geometry, and operator definitions align with their assumptions. Monte Carlo, variable/distributed-order, and operator-learning methods address problems for which conventional discretizations may be difficult or repeatedly expensive, but require stronger variance, identifiability, and generalization evidence.

The practical implication is not that fPINNs should replace classical FPDE methods. Their strongest role is complementary: combining governing equations with sparse or irregular data, solving inverse problems, treating high-dimensional nonlocal operators, or amortizing repeated parametric queries. Progress will depend on operator-level ablations, independently refined residual checks, matched classical baselines, cost-error curves, uncertainty reporting, and open reproducible implementations.

Author contributions

Muhammad Asim Khan: Conceptualization, methodology, investigation, data curation, formal analysis, and writing–original draft; Ayesha Ayub: Methodology, validation, formal analysis, and writing–review and editing; Naglaa F. Soliman: Project administration and validation; Abeer D. Algarni: Validation and resources. All authors reviewed and approved the final manuscript.

Acknowledgments

The authors extend their appreciation to the Deanship of Scientific Research and Libraries in Princess Nourah bint Abdulrahman University for funding this research work through the Supporting Publication in Review Article Program, Grant No. (RAP-2026-96)

Funding Statement

This work was funded by the Deanship of Scientific Research and Libraries in Princess Nourah bint Abdulrahman University for funding this research work through the Supporting Publication in Review Article Program, Grant No. (RAP-2026-96)

Data availability

The literature-extraction matrix, search documentation, corpus statistics, quality criteria, study-level quantitative evidence audit , dominant-class crosswalk, and Figure 5 class-level data are provided in the supplementary material. The executable notebooks reproduce the class-stratified sensitivity analysis and evidence-coverage heatmap. No new experimental FPDE dataset or solver implementation was generated.

Use of Generative-AI tools declaration

Generative artificial intelligence tools were used to assist with language editing and proofreading of the manuscript. The authors independently reviewed and verified the resulting text, references, mathematical content, and scientific interpretations.

References

1. H. Sun, Y. Zhang, D. Baleanu, W. Chen, Y. Chen, A new collection of real world applications of fractional calculus in science and engineering, *Commun. Nonlinear Sci. Numer. Simul.*, **64** (2018), 213–231. https://doi.org/10.1016/j.cnsns.2018.04.019
2. M. Stynes, E. O'Riordan, J. L. Gracia, Error analysis of a finite difference method on graded meshes for a time-fractional diffusion equation, *SIAM J. Numer. Anal.*, **55** (2017), 1057–1079. https://doi.org/10.1137/16M1082329
3. G. Acosta, J. P. Borthagaray, A fractional Laplace equation: regularity of solutions and finite element approximations, *SIAM J. Numer. Anal.*, **55** (2017), 472–495. https://doi.org/10.1137/15M1033952
4. S. Jiang, J. Zhang, Q. Zhang, Z. Zhang, Fast evaluation of the Caputo fractional derivative and its applications to fractional diffusion equations, *Commun. Comput. Phys.*, **21** (2017), 650–678. https://doi.org/10.4208/cicp.OA-2016-0136
5. J. G. Liu, K. L. Tian, Formulation, symmetry analysis and solutions for the time fractional thin film-type equation, *J. Geom. Phys.*, **228** (2026), 105923. https://doi.org/10.1016/j.geomphys.2026.105923
6. J. G. Liu, Y. F. Zhang, J. Q. Wang, Investigation of nonlinear dynamics and group analysis for fractional linear KdV-type models, *Commun. Nonlinear Sci. Numer. Simul.*, **152** (2026), 109449. https://doi.org/10.1016/j.cnsns.2025.109449

7. M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.*, **378** (2019), 686–707. https://doi.org/10.1016/j.jcp.2018.10.045
8. G. Pang, L. Lu, G. E. Karniadakis, fPINNs: Fractional physics-informed neural networks, *SIAM J. Sci. Comput.*, **41** (2019), A2603–A2626. https://doi.org/10.1137/18M1229845
9. S. Wang, H. Zhang, X. Jiang, Fractional physics-informed neural networks for time-fractional phase field models, *Nonlinear Dyn.*, **110** (2022), 2715–2739. https://doi.org/10.1007/s11071-022-07746-3
10. L. Guo, H. Wu, X. Yu, T. Zhou, Monte Carlo fPINNs: deep learning method for forward and inverse problems involving high dimensional fractional partial differential equations, *Comput. Methods Appl. Mech. Eng.*, **400** (2022), 115523. https://doi.org/10.1016/j.cma.2022.115523
11. S. Wang, H. Zhang, X. Jiang, Physics-informed neural network algorithm for solving forward and inverse problems of variable-order space-fractional advection-diffusion equations, *Neurocomputing*, **535** (2023), 64–82. https://doi.org/10.1016/j.neucom.2023.03.032
12. X. B. Yan, Z. Q. J. Xu, Z. Ma, Laplace-fPINNs: Laplace-based fractional physics-informed neural networks for solving forward and inverse problems of a time fractional equation, *East Asian J. Appl. Math.*, **14** (2024), 657–674. https://doi.org/10.4208/eajam.2023-197.171223
13. Z. Hu, K. Kawaguchi, Z. Zhang, G. E. Karniadakis, Tackling the curse of dimensionality in fractional and tempered fractional PDEs with physics-informed neural networks, *Comput. Methods Appl. Mech. Eng.*, **432** (2024), 117448. https://doi.org/10.1016/j.cma.2024.117448
14. B. Lu, Z. Hao, C. Moya, G. Lin, FPINN-deeponet: a physics-informed operator learning framework for multi-term time-fractional mixed diffusion-wave equations, *J. Comput. Phys.*, **538** (2025), 114184. https://doi.org/10.1016/j.jcp.2025.114184
15. D. Dabiri, J. DaRosa, M. Saadat, D. Mangal, S. Jamali, A detailed and comprehensive account of fractional physics-informed neural networks: from implementation to efficiency, *Artif. Intell. Eng.*, **2** (2026), 251–264. https://doi.org/10.1049/aie2.70024
16. I. Podlubny, *Fractional differential equations*, San Diego: Academic Press, 1999.
17. A. A. Kilbas, H. M. Srivastava, J. J. Trujillo, *Theory and applications of fractional differential equations*, Amsterdam: Elsevier, 2006.
18. K. Diethelm, *The analysis of fractional differential equations*, Berlin: Springer-Verlag, 2010. https://doi.org/10.1007/978-3-642-14574-2
19. R. Gorenflo, A. A. Kilbas, F. Mainardi, S. V. Rogosin, *Mittag-Leffler functions, related topics and applications*, Berlin: Springer-Verlag, 2014. https://doi.org/10.1007/978-3-662-61550-8
20. R. Metzler, J. Klafter, The random walk's guide to anomalous diffusion: a fractional dynamics approach, *Phys. Rep.*, **339** (2000), 1–77. https://doi.org/10.1016/S0370-1573(00)00070-3
21. F. Mainardi, *Fractional calculus and waves in linear viscoelasticity*, London: Imperial College Press, 2010.
22. C. Lubich, Discretized fractional calculus, *SIAM J. Math. Anal.*, **17** (1986), 704–719. https://doi.org/10.1137/0517050

23. C. Lubich, Convolution quadrature and discretized operational calculus. I, *Numer. Math.*, **52** (1988), 129–145. https://doi.org/10.1007/BF01398686

24. B. Jin, R. Lazarov, Z. Zhou, Error estimates for a semidiscrete finite element method for fractional order parabolic equations, *SIAM J. Numer. Anal.*, **51** (2013), 445–466. https://doi.org/10.1137/120873984

25. A. Bonito, J. E. Pasciak, Numerical approximation of fractional powers of elliptic operators, *Math. Comp.*, **84** (2015), 2083–2110. https://doi.org/10.1090/S0025-5718-2015-02937-8

26. B. Jin, R. Lazarov, Z. Zhou, Numerical methods for time-fractional evolution equations with nonsmooth data: a concise overview, *Comput. Methods Appl. Mech. Eng.*, **346** (2019), 332–358. https://doi.org/10.1016/j.cma.2018.12.011

27. C. Ran, X. Xu, C. Hou, X. Zhang, Numerical approximation of fourth-order fractional diffusion-wave systems using finite difference and discontinuous Galerkin method, *Netw. Heterog. Media*, **20** (2025), 1346–1366. https://doi.org/10.3934/nhm.2025058

28. F. Safari, M. A. Zaky, M. Fardi, C. Hong, Fast and accurate high-order discretization of nonlinear fractional integro-differential equations in multi-dimensional settings, *Int. Commun. Heat Mass Transf.*, **175** (2026), 111204. https://doi.org/10.1016/j.icheatmasstransfer.2026.111204

29. X. Zhang, Y. Feng, L. Wei, High-precision numerical simulation for fractional convection-diffusion-reaction equations, *Math. Methods Appl. Sci.*, **49** (2026), 13928–13945. https://doi.org/10.1002/mma.70756

30. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nat. Rev. Phys.*, **3** (2021), 422–440. https://doi.org/10.1038/s42254-021-00314-5

31. S. Cuomo, V. Schiano Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics-informed neural networks: where we are and what's next, *J. Sci. Comput.*, **92** (2022), 88. https://doi.org/10.1007/s10915-022-01939-z

32. C. Meng, S. Griesemer, D. Cao, S. Seo, Y. Liu, When physics meets machine learning: a survey of physics-informed machine learning, *Mach. Learn. Comput. Sci. Eng.*, **1** (2025), 20. https://doi.org/10.1007/s44379-025-00016-0

33. A. A. Sophiya, A. K. Nair, S. Maleki, S. K. Krishnababu, A comprehensive analysis of PINNs: variants, applications, and challenges, *arXiv*, 2025. https://doi.org/10.48550/arXiv.2505.22761

34. L. Lu, X. Meng, Z. Mao, G. E. Karniadakis, DeepXDE: a deep learning library for solving differential equations, *SIAM Rev.*, **63** (2021), 208–228. https://doi.org/10.1137/19M1274067

35. D. Wang, Y. Ning, C. Xiang, A. Chen, A two-stage network framework for topology optimization incorporating deep learning and physical information, *Eng. Appl. Artif. Intell.*, **133** (2024), 108185. https://doi.org/10.1016/j.engappai.2024.108185

36. Y. Liu, Z. Li, P. Wei, W. Wang, Parameterized level-set based topology optimization method considering symmetry and pattern repetition constraints, *Comput. Methods Appl. Mech. Eng.*, **340** (2018), 1079–1101. https://doi.org/10.1016/j.cma.2018.04.034

37. M. Cui, W. Li, G. Li, X. Wang, The asymptotic concentration approach combined with isogeometric analysis for topology optimization of two-dimensional linear elasticity structures, *Electron. Res. Arch.*, **31** (2023), 3848–3878. https://doi.org/10.3934/era.2023196

38. E. Abdo, L. Chai, R. Hu, X. Yang, Convergence analysis of PINNs for fractional diffusion equations in bounded domains, *arXiv*, 2026. https://doi.org/10.48550/arXiv.2601.01462

39. J. Shi, X. Yang, X. Liu, A novel fractional physics-informed neural networks method for solving the time-fractional Huxley equation, *Neural Comput. Appl.*, **36** (2024), 19097–19119. https://doi.org/10.1007/s00521-024-10177-3

40. M. Mohammadi, R. Mokhtari, M. Ramezani, RBF-fPINNs: radial basis function-enhanced fractional physics-informed neural networks, *Eng. Comput.*, **42** (2026), 42. https://doi.org/10.1007/s00366-025-02258-1

41. H. Li, Z. Tan, DSC-PINN: a decomposed singularity-capturing PINN for time-fractional PDEs involving initial singularities, *J. Comput. Phys.*, **561** (2026), 114965. https://doi.org/10.1016/j.jcp.2026.114965

42. M. Ramezani, M. Mohammadi, R. Mokhtari, dPINNs: a physics-informed framework for forward and inverse problems governed by distributed-order derivatives, *Eng. Anal. Bound. Elem.*, **179** (2025), 106418. https://doi.org/10.1016/j.enganabound.2025.106418

43. Z. Hao, M. Park, Z. Cai, Neural network method for integral fractional Laplace equations, *East Asian J. Appl. Math.*, **13** (2023), 95–118. https://doi.org/10.4208/eajam.010122.210722

44. Z. Hao, M. Park, Z. Cai, HNS: an efficient Hermite neural solver for solving time-fractional partial differential equations, *Chaos Solitons Fract.*, **181** (2024), 114637. https://doi.org/10.1016/j.chaos.2024.114637

45. J. Shi, X. Liu, X. Yang, Data-driven solutions and parameter estimation of the high-dimensional time-fractional reaction-diffusion equations using an improved fPINN method, *Nonlinear Dyn.*, **113** (2025), 9577–9604. https://doi.org/10.1007/s11071-024-10786-6

46. R. Zhang, F. Li, J. Liu, LT-PINNs: physics-informed neural networks based on Laplace transform for solving Caputo-type fractional partial differential equations, *Chin. Phys. B*, **35** (2026), 030201. https://doi.org/10.1088/1674-1056/adf827

47. S. M. Nuugulu, K. C. Patidar, D. T. Tarla, An efficient physics-informed neural network solution to the time-space fractional Black-Scholes equation, *Oper. Res. Forum*, **6** (2025), 172. https://doi.org/10.1007/s43069-025-00570-6

48. S. Jiang, X. Li, Solving the non-local Fokker-Planck equations by deep learning, *Chaos*, **33** (2023), 043107. https://doi.org/10.1063/5.0128935

49. W. Liu, Y. Liu, H. Li, Y. Yang, Multi-output physics-informed neural network for one- and two-dimensional nonlinear time distributed-order models, *Netw. Heterog. Media*, **18** (2023), 1899–1918. https://doi.org/10.3934/nhm.2023080

50. H. Momeni, A. Yazdani Cherati, A. Valinejad, AFP-PINN: adaptive fractional polynomial-based physics-informed neural network for solving distributed-order time-fractional diffusion equations and its error analysis, *Phys. Scr.*, **100** (2025), 076026. https://doi.org/10.1088/1402-4896/adeaf8

51. Y. Ye, H. Fan, Y. Li, X. Liu, H. Zhang, Deep neural network methods for solving forward and inverse problems of time fractional diffusion equations with conformable derivative, *Neurocomputing*, **509** (2022), 177–192. https://doi.org/10.1016/j.neucom.2022.08.030
52. T. Zhang, D. Zhang, S. Shi, Z. Guo, Spectral-fPINNs: spectral method based fractional physics-informed neural networks for solving fractional partial differential equations, *Nonlinear Dyn.*, **113** (2025), 12565–12588. https://doi.org/10.1007/s11071-025-11067-6
53. Q. Wu, F. Ma, W. Wo, FS-PINNs: fractional spectrally adapted physics-informed neural networks for fractional partial differential equations, *Math. Methods Appl. Sci.*, **48** (2025), 13882–13889. https://doi.org/10.1002/mma.11149
54. R. H. Nochetto, E. Otárola, A. J. Salgado, A PDE approach to fractional diffusion in general domains: a priori error analysis, *Found. Comput. Math.*, **15** (2015), 733–791. https://doi.org/10.1007/s10208-014-9208-x
55. X. Kang, Y. Li, Y. Li, J. Hu, K. Zheng, Solving the fractional Allen-Cahn equation and the fractional Cahn-Hilliard equation with the fractional physics-informed neural networks, *Fractal Fract.*, **9** (2025), 773. https://doi.org/10.3390/fractalfract9120773
56. T. Taheri, A. Afzal Aghaei, K. Parand, Accelerating fractional PINNs using operational matrices of derivative, *arXiv*, 2024. https://doi.org/10.48550/arXiv.2401.14081
57. J. Zhang, Y. Zhao, Y. Tang, Adaptive loss weighting auxiliary output fPINNs for solving fractional partial integro-differential equations, *Physica D*, **460** (2024), 134066. https://doi.org/10.1016/j.physd.2024.134066
58. T. Ahmad, M. Sulaiman, D. Bassir, F. S. Alshammari, G. Laouini, Enhanced numerical solutions for fractional PDEs using Monte Carlo PINNs coupled with cuckoo search optimization, *Fractal Fract.*, **9** (2025), 225. https://doi.org/10.3390/fractalfract9040225
59. L. Lu, P. Jin, G. Pang, Z. Zhang, G. E. Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.*, **3** (2021), 218–229. https://doi.org/10.1038/s42256-021-00302-5
60. H. Ren, X. Meng, R. Liu, J. Hou, Y. Yu, A class of improved fractional physics informed neural networks, *Neurocomputing*, **562** (2023), 126890. https://doi.org/10.1016/j.neucom.2023.126890
61. A. D. Jagtap, K. Kawaguchi, G. E. Karniadakis, Adaptive activation functions accelerate convergence in deep and physics-informed neural networks, *J. Comput. Phys.*, **404** (2020), 109136. https://doi.org/10.1016/j.jcp.2019.109136
62. S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, *SIAM J. Sci. Comput.*, **43** (2021), A3055–A3081. https://doi.org/10.1137/20M1318043
63. N. Xue, M. Chen, Multiprecision computing for multistage fractional physics-informed neural networks, *Chaos*, **35** (2025), 113109. https://doi.org/10.1063/5.0287809
64. S. Wang, X. Yu, P. Perdikaris, When and why PINNs fail to train: a neural tangent kernel perspective, *J. Comput. Phys.*, **449** (2022), 110768. https://doi.org/10.1016/j.jcp.2021.110768

65. M. A. Nabian, R. J. Gladstone, H. Meidani, Efficient training of physics-informed neural networks via importance sampling, *Comput. Aided Civ. Infrastruct. Eng.*, **36** (2021), 962–977. https://doi.org/10.1111/mice.12685

66. L. D. McClenny, U. M. Braga-Neto, Self-adaptive physics-informed neural networks, *J. Comput. Phys.*, **474** (2023), 111722. https://doi.org/10.1016/j.jcp.2022.111722

67. E. Kharazmi, Z. Zhang, G. E. Karniadakis, hp-VPINNs: variational physics-informed neural networks with domain decomposition, *Comput. Methods Appl. Mech. Eng.*, **374** (2021), 113547. https://doi.org/10.1016/j.cma.2020.113547

68. C. Li, F. Zeng, *Numerical methods for fractional calculus*, Boca Raton: CRC Press, Taylor & Francis Group, 2015.

69. Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, A. Anandkumar, Fourier neural operator for parametric partial differential equations, *International Conference on Learning Representations*, 2021.

70. N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, et al., Neural operator: learning maps between function spaces with applications to PDEs, *J. Mach. Learn. Res.*, **24** (2023), 1–97.



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)