www.editada.org

International Journal of Combinatorial Optimization Problems and Informatics, 17(4), 2026, 205-220. ISSN: 2007-1558.
https://doi.org/10.61467/2007.1558.2026.v17i4.1389

Numerically Stable and Reproducible Mittag–Leffler Evaluation for Fractional Kelvin–Voigt Creep Modelling

Jaime Aguilar-Ortiz[1,2], Carlos R. Domínguez-Mayorga[1], Ocotlán Diaz-Parra[1], Jorge A. Ruiz-Vanoye[1], Francisco R. Trejo-Macotela[1], Marco A. Vera-Jiménez[1], Víctor M. Zamudio-García[2]

[1]Universidad Politécnica de Pachuca, México
[2] Universidad Politécnica Metropolitana de Hidalgo, México

Email address(es): jao@upp.edu.mx, carlos.dominguez@upp.edu.mx, ocotlan_diaz@upp.edu.mx, jorgeruiz@upp.edu.mx, trejo_macotela@upp.edu.mx, marcovera@upp.edu.mx, vzamudio@upmh.edu.mx

Abstract. Fractional Kelvin–Voigt models capture long-memory behaviour in viscoelastic creep, but their numerical implementation depends on the Mittag–Leffler function. Naive evaluation of this function can cause overflow, loss of numerical significance, and nonphysical creep responses, thereby compromising constitutive interpretation and parameter calibration. This study develops and validates a numerically stable and reproducible Mittag–Leffler evaluator for time-domain creep simulations under constant stress. Reliability is assessed through four operational criteria: bounded special-function values on the negative real axis, preservation of monotonic decay, low benchmark error relative to a trusted reference, and improved creep fidelity compared with the best classical Kelvin–Voigt baseline. The proposed implementation combines regime-aware series and asymptotic evaluation, a calibrated branch-switching policy, overflow-safe safeguards, and reproducibility controls. Supplementary ablation and synthetic-noise experiments show that the evaluator remains suitable for calibration-oriented fractional creep studies without introducing spurious numerical artefacts into the constitutive response. The resulting methodology provides a reliable computational basis for fractional viscoelastic modelling, numerical benchmarking, and parameter-estimation studies involving Mittag–Leffler functions.

Keywords: Mittag–Leffler function; fractional Kelvin–Voigt model; fractional viscoelasticity; viscoelastic creep; numerical stability; reproducible computation; time-domain creep simulation; asymptotic evaluation; series evaluation; constitutive modelling; parameter calibration; special-function computation.

Article information
Received: April 7, 2026
Accepted: Jun 11, 2026

1 Introduction

1.1 Motivation: Long-memory creep and the reliability gap in time-domain simulation

Long-memory creep is a defining empirical signature of many soft, biological, and heterogeneous media: under sustained load, the strain does not approach equilibrium with a simple exponential law, but instead follows a slow, history-dependent evolution that is frequently well-approximated by power-law behavior across multiple time decades. Fractional viscoelastic constitutive laws especially fractional Kelvin–Voigt-type formulations provide a parsimonious mechanism to encode this distributed-timescale response while preserving physical interpretability through generalized stiffness/viscosity parameters and a fractional order that controls memory intensity. In that setting, time-domain creep predictions become central not only for forward simulation, but also for parameter identification and validation against laboratory protocols (e.g., ramp-hold indentation or creep tests), where closed-form expressions commonly involve the Mittag–Leffler function and its variants. Consequently, the credibility of any fractional Kelvin–Voigt creep study is constrained by whether the time-domain implementation can reproduce the expected short-horizon compliance growth while remaining

consistent and physically plausible over moderate horizons, where long-memory effects are most diagnostically informative for model selection and inference.

A recurring obstacle is that the mathematical adequacy of the fractional Kelvin–Voigt creep law does not automatically translate into computational reliability: the same closed-form structure that makes the model attractive (Mittag–Leffler-based time responses) also introduces well-known numerical fragilities when evaluated naïvely across regimes of practical interest. Truncated series, direct floating-point evaluations, or insufficiently guarded implementations can suffer from overflow, catastrophic cancellation, and loss of monotonicity/positivity pathologies that may masquerade as "material behavior" but are in fact artifacts of special-function evaluation. This creates a reliability gap in which time-domain simulations may appear stable in a narrow short-time window yet become untrustworthy exactly where the long-memory signature should be assessed (moderate times, multiple parameter sets, repeated sweeps in calibration loops). Closing that gap requires stable Mittag–Leffler evaluation strategies that explicitly manage poles/branch behavior, scaling, and error control (e.g., contour-integral quadrature and rational approximations on critical subsets such as the negative real axis), and it also benefits from stability perspectives that connect numerical solution behavior to Mittag–Leffler-type bounds for time-fractional dynamics. In this way, robust special-function evaluation is not a cosmetic numerical refinement but the enabling condition for interpretable creep simulation, defensible parameter estimation, and reproducible evidence that fractional Kelvin–Voigt modeling remains reliable from short to moderate horizons.

1.2 Contributions, scope boundaries, and article roadmap

The article articulates three tightly coupled contributions that together close a reliability gap between fractional viscoelastic theory and practical time-domain simulation. First, it delivers a numerically stable evaluation strategy for the two-parameter Mittag–Leffler function in the parameter regimes that arise naturally in fractional Kelvin–Voigt creep, where naive series summation can overflow, stagnate, or silently lose accuracy as time grows from short to moderate horizons. Second, it turns that stability requirement into an auditable computational protocol built around explicit acceptance thresholds, calibrated branch switching, admissibility checks, and reproducible output generation. Third, it validates the protocol end to end in Kelvin–Voigt fractional-derivative creep simulations representative of calibration-oriented soft-matter and biomechanical workflows, where numerical fragility in the hereditary kernel can otherwise contaminate model discrimination.

The scope boundaries are defined so that the reliability claim is both defensible and maximally useful for subsequent reuse in applied modeling. The article restricts attention to linear viscoelasticity under the fractional Kelvin–Voigt constitutive structure and to time-domain creep under canonical step-type loading, where the response is governed by Mittag–Leffler expressions. Within that scope, the revised manuscript treats reliability as a falsifiable target: low numerical error, zero admissibility violations on the negative real axis, and a clear creep-level advantage over the best classical Kelvin–Voigt approximation must all be satisfied simultaneously. The roadmap mirrors that logic. Section 2 consolidates the constitutive model, the Mittag–Leffler representation, and the normalization adopted for benchmarking. Section 3 defines the validation criteria and the calibrated hybrid evaluator. Section 4 documents the computational setup and the reproducibility protocol. Section 5 reports benchmark accuracy, creep fidelity, component-wise ablation, computational cost, and robustness under mild synthetic observation noise. Section 6 closes with the methodological implications and the scope limits of the validated computational layer. Figure 1 summarizes this progression.

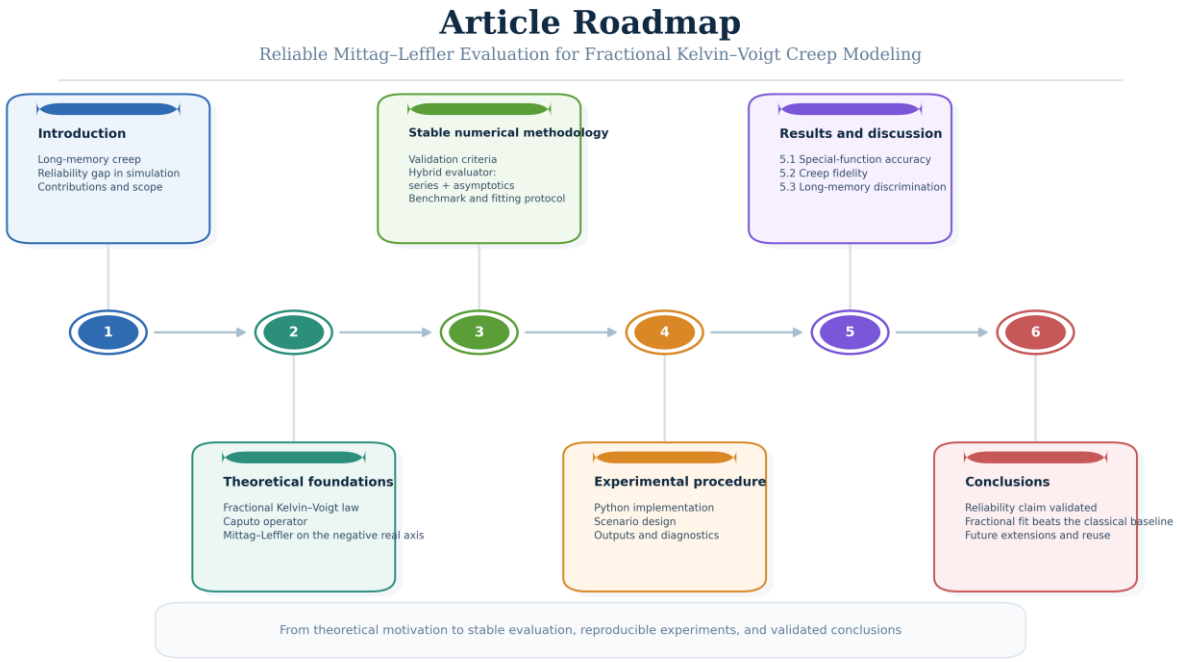

Figure 1. *Article roadmap for reliable Mittag–Leffler evaluation in fractional Kelvin–Voigt creep modeling.*

2 Theoretical foundations

2.1 Fractional Kelvin–Voigt Creep: Constitutive law, closed-form response, and identifiability notes

The fractional Kelvin–Voigt element represents viscoelastic creep through the parallel combination of an elastic spring and a fractional springpot. Under uniaxial loading, the constitutive law links the applied stress to the instantaneous strain and to a hereditary rate term of non-integer order. In the notation adopted here, the constitutive relation is written as

$$\sigma(t) = E\,\varepsilon(t) + \eta_\alpha\, D_t^{\alpha,\mathrm{C}}\varepsilon(t), \qquad 0 < \alpha < 1, \tag{1}$$

In Eq. (1), E denotes the elastic modulus, η_α is the fractional viscosity parameter, and Dt_α, C indicates the Caputo derivative of order α. The Caputo form is convenient for experimental mechanics because it preserves classical initial conditions and permits a direct interpretation of transient creep tests. Its integral definition is given by Eq. (2):

$$D_t^{\alpha,\mathrm{C}} f(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t (t-\tau)^{-\alpha} f'(\tau)\, d\tau \tag{2}$$

For a step stress $\sigma(t) = \sigma_0 H(t)$ and zero initial strain, the creep compliance follows a closed form involving the Mittag–Leffler function. The resulting strain history is:

$$\varepsilon(t) = \frac{\sigma_0}{E}\left[1 - E_\alpha\left(-\frac{E}{\eta_\alpha} t^\alpha\right)\right]. \tag{3}$$

The classical Kelvin–Voigt model is recovered when $\alpha = 1$, in which case the Mittag–Leffler function collapses to an exponential and the creep response is given by Eq. (4):

$$\varepsilon_{\mathrm{cl}}(t) = \frac{\sigma_0}{E}\left[1 - \exp\left(-\frac{E}{\eta}t\right)\right] \tag{4}$$

This distinction is central to identifiability. The modulus E governs the equilibrium plateau σ_0/E, whereas the pair (η_α, α) controls the transient trajectory, the slow approach to equilibrium, and the extent of memory retained from earlier times. In a purely classical fit, the entire transient is compressed into a single retardation time η/E, which is often insufficient when the data exhibit broad-band hereditary behavior (Bonfanti et al., 2020; Makris & Efthymiou, 2020). An additional identifiability caveat concerns the partial correlation between α and ηα when the available creep record is short or mildly noisy, because both parameters influence transient curvature through the same normalized argument. Direct identification studies for fractional Kelvin–Voigt-type models and later Bayesian/subset-selection analyses show that this correlation can broaden the admissible parameter valley and make the recovered order more sensitive to window length and data richness (Lewandowski & Chorążyczewski, 2010; Miles et al., 2021). Recent rheology-informed neural-network results report a related phenomenon: the fractional order can deviate from its nominal ground truth when the measured response is not sufficiently informative (Dabiri et al., 2023). For that reason, the benchmarks in this article fix α within each scenario when estimating viscosity, so that numerical reliability of the Mittag–Leffler evaluator is not confounded with statistical identifiability.

2.2 Fractional-order operators for viscoelasticity: Modeling meaning, regularity, and parameter roles

Fractional-order operators provide a compact mechanism to encode distributed time scales without introducing long Prony series or large rheological networks. In the present context, the order α in the interval $(0,1)$ governs how strongly the material remembers its loading history. Values of α near 1 recover behavior close to the classical Kelvin–Voigt limit, whereas smaller α generate slower relaxation of the hereditary kernel and therefore more persistent creep tails (Bonfanti et al., 2020; Zhang et al., 2018).

The parameter η_α should not be interpreted as a classical viscosity with units Pa·s. Its dimensional role depends on α and compensates the non-integer derivative so that the constitutive law remains dimensionally consistent. Because the plateau is fixed by E while the curve shape is controlled by (η_α, α), the creep test becomes an informative protocol for separating short-time stiffness from memory intensity, provided that the special-function evaluation is numerically trustworthy (Shabani et al., 2019; Rodríguez Soto et al., 2022).

Regularity also matters computationally. The Caputo operator acts on the ordinary derivative $f'(t)$ and is therefore compatible with smooth experimental histories such as step, ramp, or ramp-hold loadings. This property facilitates simulation and parameter identification but it also means that numerical artifacts in the evaluation of $E_\alpha(-x)$ may be mistaken for constitutive effects if monotonicity, positivity, and boundedness are not explicitly monitored.

2.3 The Mittag–Leffler function on the negative real axis: Behavior, pitfalls of Naïve evaluation, and stability requirements

The two-parameter Mittag–Leffler function generalizes the exponential and occupies in fractional rheology the role played by $exp(\cdot)$ in integer-order linear systems. It is defined by the entire series in Eq. (5):

$$E_{\alpha,\beta}(z) = \sum_{k=0}^{\infty} \frac{z^k}{\Gamma(\alpha k + \beta)} \tag{5}$$

In fractional Kelvin–Voigt creep, the relevant argument is $z = -x$ with $x \geq 0$. On this axis the function is positive and monotonically decreasing for $0 < \alpha < 1$, so that the strain in Eq. (3) remains nondecreasing and bounded by σ_0/E. Numerical trouble arises because direct finite-precision series summation may produce catastrophic cancellation when x grows beyond the small-argument regime. In contrast, the large-x tail is naturally described by the inverse-power asymptotic expansion:

$$E_\alpha(-x) \sim \sum_{k=1}^{m} \frac{(-1)^{k+1}}{x^k \Gamma(1-\alpha k)}, \qquad x \to \infty. \tag{6}$$

The asymptotic formula is efficient, but it is not uniformly valid near the origin or in the intermediate region. For validation purposes, the present work therefore adopts a trusted reference evaluator based on the complete-monotonicity representation:

$$E_\alpha(-x) = \frac{\sin(\pi\alpha)}{\pi} \int_0^\infty \frac{e^{-rx^{1/\alpha}} r^{\alpha-1}}{r^{2\alpha} + 2r^\alpha \cos(\pi\alpha) + 1}\, dr \tag{7}$$

Equation (7) is especially useful because it preserves the physics of the negative real axis: the reference values stay within the interval, decrease monotonically with x, and provide a reliable benchmark against which hybrid and unsafe numerical implementations can be assessed (McLean, 2021; Wang & Zou, 2023).

For $0 < \alpha < 1$, $E_\alpha(-x)$ is completely monotone on $x \geq 0$. Consequently, the target response satisfies $0 < E_\alpha(-x) \leq 1$ and decreases with x, so the corresponding creep curve under step loading must remain nondecreasing and bounded by σ_0/E. The validation criteria in Section 3 therefore encode a theoretical property of the model rather than a merely visual preference.

Relative to the current numerical literature, the proposed evaluator is intentionally narrower in scope than general-purpose contour-integral/Laplace-inversion algorithms, which remain the most robust reference implementations across wider sectors of the complex plane and on the negative real axis (Garrappa & Popolizio, 2013; Garrappa, 2015; McLean, 2021). Rational or Padé-type approximations can be highly efficient once the admissible parameter window is fixed, but their accuracy depends more strongly on pole placement and matching design (Iyiola et al., 2018; Sarumi et al., 2020). The contribution of the article is therefore an auditable, application-specific evaluator for Kelvin–Voigt creep rather than a blanket replacement for those broader methods.

2.4 Normalization, notation, and admissible evaluation regime

For benchmarking and implementation, the creep argument is written in terms of a dimensionless variable that absorbs the material and time scales into a single negative-real-axis evaluation problem. With this normalization, the special-function benchmark becomes independent of the specific units chosen for E, η_α, and t while preserving the constitutive interpretation of the creep response, Eq. (8).

$$x = \left(\frac{E}{\eta_\alpha}\right) t^\alpha, \qquad x \geq 0, \quad 0 < \alpha < 1 \tag{8}$$

This normalization also clarifies identifiability. The modulus E fixes the plateau σ_0/E, whereas η_α and α shape the transient through the single dimensionless argument x. The admissible numerical regime addressed in the article is therefore the set $\{(\alpha, x)\colon 0 < \alpha < 1, x \geq 0\}$, on which positivity, monotone decay, and boundedness are theoretically expected and subsequently checked in the validation protocol.

3 Stable numerical methodology

3.1 Validation criteria derived from the article objective

The abstract and the introductory subsections imply a specific operational objective: to develop and validate a numerically stable, reproducible evaluator of the Mittag–Leffler function in the parameter regime that governs fractional Kelvin–Voigt creep. In the revised manuscript, that objective is treated as an acceptance test. The evaluator is accepted only if it satisfies four conditions on the benchmark grid: (i) a worst-case relative error below 10^{-2} against the trusted reference, (ii) zero monotonicity violations on $E_\alpha(-x)$, (iii) zero boundedness violations on the admissible interval $[0,1]$, and (iv) a clearly smaller creep-fit MAE than the best classical Kelvin–Voigt baseline. These thresholds keep the reliability claim falsifiable and tie the numerical objective directly to the practical purpose of the article: trustworthy creep simulation for reproducible calibration and model comparison. Here the 10^{-2} bound is used as a screening threshold rather than as the target accuracy. Because the admissible branch of $E_\alpha(-x)$ is confined to $[0,1]$, discrepancies at that level are already large enough to produce visible compliance perturbations and thus serve to exclude numerically unsafe evaluators before calibration is attempted. Final acceptance nevertheless still depends on criterion (iv), so passing the kernel-error test alone is not sufficient; the validated hybrid evaluator in Table 1 remains comfortably below this exclusion bound for every tested α.

The first indicator is the pointwise relative error, defined as:

$$\mathrm{RelErr}(x,\alpha) = \frac{\left|E_\alpha^{(num)}(-x) - E_\alpha^{(ref)}(-x)\right|}{\max\left(\left|E_\alpha^{(ref)}(-x)\right|, 10^{-30}\right)}. \tag{9}$$

In addition to Eq. (9), the benchmark records discrete monotonicity violations on the x-grid and counts any values that fall outside the physically admissible interval $[0,1]$. At the creep level, an analogous check enforces $0 \le \varepsilon(t) \le \sigma_0/E$ and verifies that the strain remains nondecreasing under constant stress. These admissibility checks convert the reliability claim into a falsifiable set of diagnostics rather than a merely qualitative visual comparison.

3.2 Hybrid evaluator and implementation logic

The proposed evaluator combines two complementary branches. For small and transition arguments, $E_\alpha(-x)$ is computed through a high-precision power series; for larger arguments, the program switches to the inverse-power asymptotic representation in Eq. (6). The branch threshold is now stated explicitly. On the benchmarked set, the calibrated policy is $switch(\alpha) = 3$ for $\alpha \in \{0.25, 0.55, 0.65, 0.80\}$ and $switch(0.90) = 4$. In other words, the reported implementation uses $x^* = 3$ throughout most of the tested range and postpones the asymptotic branch only for the highest-order case, where an earlier changeover produced a small monotonicity defect in the raw evaluator. This turns the switch from a narrative choice into a reproducible algorithmic rule. Figure 2 summarizes the resulting branch logic.

PSEUDOCODE

```
for each x >= 0:
    if x < switch(alpha):
        value = high_precision_series(x, alpha)
    else:
        value = asymptotic_inverse_power(x, alpha)
    value = clip(value, 0, 1)
    store E_alpha(-x) = value
```

Figure 2. *Core branch logic of the stable hybrid evaluator.*

The branch logic is intentionally transparent. Rather than hiding stability inside opaque black-box calls, the program exposes the numerical regime used at each x. This transparency is valuable in reproducible mechanics studies because it allows later authors to inspect where the special-function approximation changes character and to determine whether the chosen switch is appropriate for their own material parameters. For implementation accessibility, the pseudo-code box associated with Figure 2 should be read as the concise algorithmic summary for independent replication: it makes explicit the branch test, the raw-value computation, the admissibility guard, and the stored output.

To separate the core evaluator from the admissibility safeguard, the implementation first computes a raw branch value and only then applies the clipping layer shown below Eq. (10).

$$E_{\alpha}^{\mathrm{clip}}(-x) = \min\left\{1, \max\left(0, \tilde{E}_{\alpha}(-x)\right)\right\} \tag{10}$$

On the reported benchmark grid, the calibrated raw evaluator already produced zero monotonicity violations, zero bounds violations, and zero clipping activations for $\alpha \in \{0.25, 0.55, 0.65, 0.80, 0.90\}$. The clipping step is therefore retained as a defensive guard for off-grid calls rather than as the mechanism that creates the reported stability.

3.3 Benchmark design and parameter-identification protocol

The special-function benchmark was performed on 36 logarithmically spaced points $x \in [10^{-2}, 10^{2}]$ and on the representative fractional orders $\alpha \in \{0.25, 0.55, 0.65, 0.80, 0.90\}$. After the hybrid evaluator was validated against the trusted reference on that benchmark, the same stable kernel was used to generate dense creep curves on three experimental horizons: a short-horizon baseline ($\alpha = 0.65, t_{max} = 50\ s$), a moderate stress test ($\alpha = 0.55, t_{max} = 200\ s$), and a long-memory tail-emphasis scenario ($\alpha = 0.25, t_{max} = 2000\ s$). In all cases, $\sigma_0 = 1$, $E = 1000$, and the true fractional viscosity parameter was $\eta_{\alpha} = 200$.

For model comparison, the study estimated the best fractional viscosity parameter and the best classical viscosity parameter by minimizing a mean-squared discrepancy between the trusted creep curve and the model prediction. The optimization functional is:

$$J(\theta) = \frac{1}{N}\sum_{i=1}^{N}\left[\varepsilon_i^{(ref)} - \varepsilon_i^{(model)}(\theta)\right]^2. \tag{11}$$

In Eq. (11), θ denotes either η_{α} for the fractional fit or η for the classical fit. The search was performed with a coarse grid followed by a golden-section refinement. This choice kept the identification procedure fully reproducible and avoided the additional stochastic variability of heuristic optimizers.

4 Experimental procedure

4.1 Numerical setup

All numerical experiments were implemented in Python by extending the uploaded project script and organizing the workflow into three layers: a trusted reference benchmark, the calibrated stable evaluator, and an unsafe diagnostic baseline used only to expose numerical fragility. The program exports consolidated CSV tables, publication-ready figures, and a machine-readable scenario manifest so that every reported benchmark can be regenerated without manual editing. This automation is consistent with the article's purpose: to turn fractional creep theory into a dependable computational layer rather than a one-off numerical experiment.

4.2 Quantities reported in the results section

The results section reports two families of indicators. At the special-function level, the benchmark summarizes the maximum relative error, the median relative error, the number of monotonicity violations, and the number of boundedness violations. At the creep level, the paper reports MSE, MAE, and MAXE for the best stable fractional fit and for the best classical fit, together with windowed MAE values on early, intermediate, and late time intervals. These quantities were selected because they directly address the validation logic extracted from the abstract and Sections 1.1–1.2: accuracy, physical admissibility, and usefulness for model discrimination.

4.3 Reproducibility protocol

Reproducibility was enforced through deterministic settings: fixed logarithmic benchmark grids, fixed time grids for each creep scenario, explicit α values, fixed search intervals for η_α and η, and deterministic random seeds for the supplementary noise experiment. The computational stack consisted of Python 3 with NumPy, Matplotlib, and mpmath, and each run exported figures, CSV summaries, and a manifest of the scenario parameters. The reliability claim is thus tied not only to a stable evaluator, but also to a recoverable experimental record.

5 Results and discussion

5.1 Special-function accuracy on the negative real axis

The benchmark confirms that the hybrid evaluator is numerically stable over the intended range of use. Across the tested orders $\alpha = (0.25, 0.55, 0.65, 0.8, 0.9)$, the worst-case relative error remained below $3.640e-03$, while the median relative error stayed essentially at machine precision. No monotonicity violations and no boundedness violations were detected for the stable branch. By contrast, the unsafe guarded-series baseline failed decisively: even its smallest maximum relative error exceeded $1.224e+02$, and every α-profile displayed a monotonicity violation on the benchmark grid. The representative case $\alpha = 0.55$ is especially revealing, because the stable evaluator achieved a maximum relative error of $1.250e-04$, whereas the unsafe baseline reached $1.955e+02$.

Figure 3 shows the representative $\alpha = 0.65$ benchmark. The stable hybrid curve is visually indistinguishable from the trusted reference, whereas the unsafe evaluation loses the expected decay pattern and departs from the physical branch immediately after the small-x regime. Figure 4 broadens the comparison across all α values and shows that the stable method retains its accuracy over the whole benchmark grid, with only a narrow transition-region spike when the asymptotic branch becomes active. Even there, the error remains several orders of magnitude smaller than that of the unsafe baseline. Table 1 summarizes the corresponding validation metrics across the benchmarked α values.

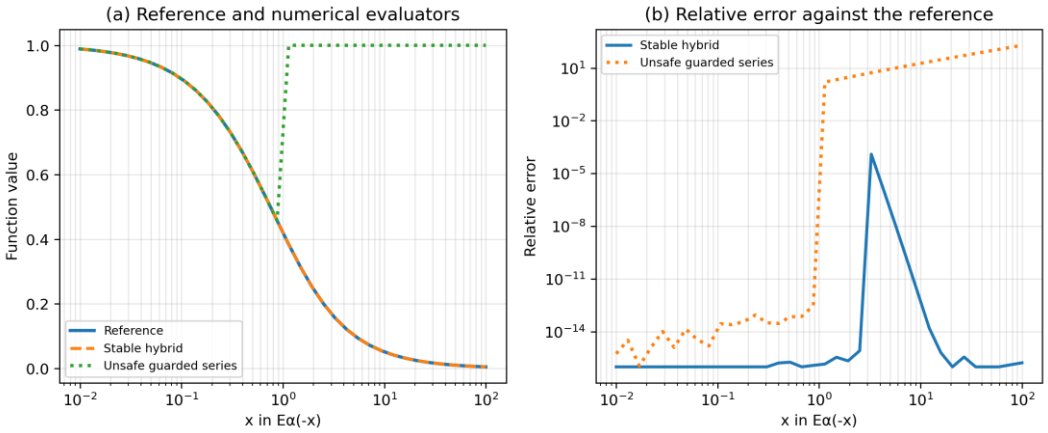

Figure 3. *Representative comparison of the trusted reference, the stable hybrid evaluator, and the unsafe guarded-series baseline, including the relative-error profile.*

Table 1. *Special-function validation metrics on the negative real axis.*

α	Method	Max rel. err.	Median rel. err.	Mono. viol.	Bounds viol.
0.25	Stable	1.781e-03	3.608e-16	0	0
0.25	Unsafe	1.224e+02	6.615e-01	1	0
0.55	Stable	1.250e-04	0.000e+00	0	0
0.55	Unsafe	1.955e+02	8.016e-01	1	0
0.65	Stable	8.812e-06	1.226e-16	0	0
0.65	Unsafe	2.521e+02	8.574e-01	1	0
0.8	Stable	1.414e-03	1.620e-16	0	0
0.8	Unsafe	4.524e+02	9.479e-01	1	0
0.9	Stable	3.640e-03	1.222e-16	0	0
0.9	Unsafe	9.345e+02	1.009e+00	1	0

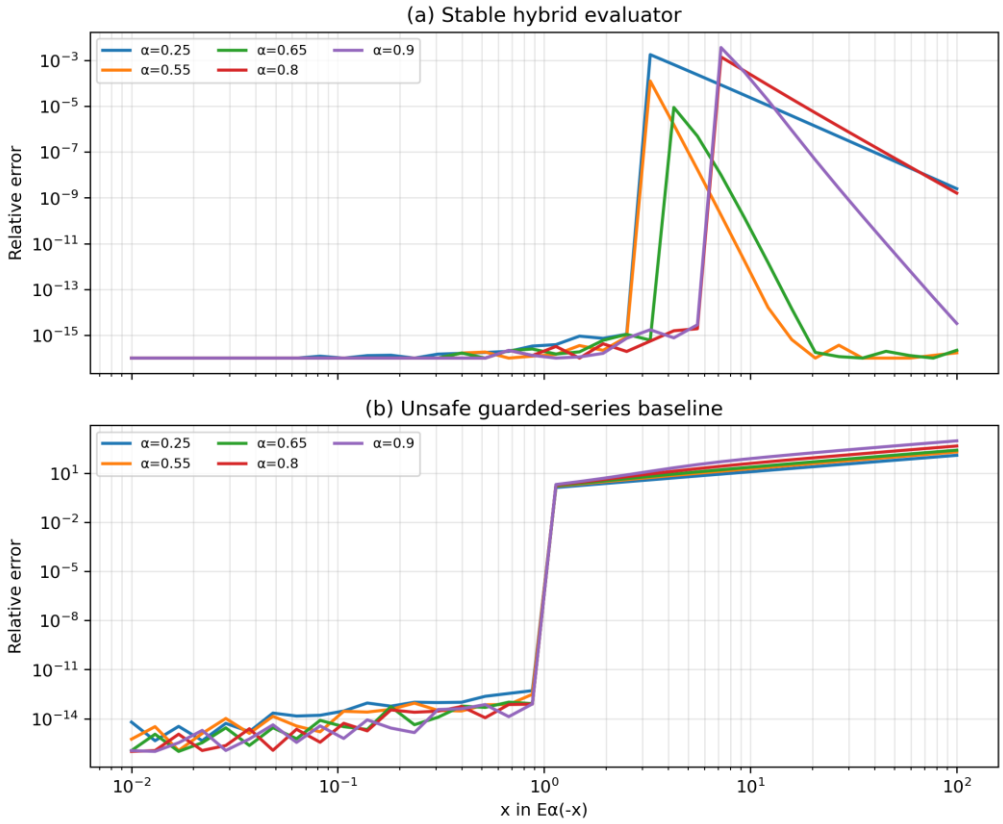

Figure 4. *Relative-error profiles for the stable hybrid evaluator and the unsafe guarded-series baseline across the benchmark alpha set.*

5.1.1 Component-wise ablation and computational cost

A supplementary component-wise ablation was performed against an over-resolved hybrid reference on the same x-grid to identify which component of the evaluator produces the stability gain. The results confirm the expected regime dependence: a pure power series is accurate near the origin but fails outside its natural range, the asymptotic branch alone is unsuitable in the small-x regime, and a fixed $x^* = 3$ switch leaves a residual monotonicity defect at $\alpha = 0.90$. The calibrated adaptive policy removes that defect without materially increasing runtime. Table 2 summarizes the corresponding ablation metrics and admissibility counts. Table 2 also serves as a local sensitivity analysis of the switching threshold. The comparison between the fixed early switch and the adaptive policy shows that the evaluator is robust over the first four benchmarked orders and that only the near-classical high-order case α=0.90 requires a delayed crossover to remove the residual monotonicity defect without a material runtime penalty.

Table 2. *Supplementary ablation metrics for the stable evaluator variants.*

Variant	Max rel. err.	Median rel. err.	Mono. viol.	Bounds viol.
Series only	overflow	0.000e+00	7	31
Asymptotic only	6.375e+26	7.143e+00	20	95
Fixed switch x*=3	8.466e-01	0.000e+00	1	0
Adaptive switch	7.252e-03	0.000e+00	0	0

Adaptive + clipping	7.252e-03	0.000e+00	0	0

Measured over the 180-point benchmark grid, the adaptive stable evaluator required approximately 0.70 ms per evaluation point on the test platform, compared with 1.29 ms for the over-resolved reference and 0.30 ms for the unsafe series baseline. The stable method is therefore modestly more expensive than the fragile baseline but remains comfortably sub-millisecond per point and thus compatible with calibration loops.

5.2 Creep fidelity under short and moderate horizons

Once the special-function benchmark established the reliability of the stable evaluator, the same kernel was used to simulate creep curves on dense time grids. The fitted fractional viscosity parameter remained close to the true value in every scenario, with $\eta_\alpha{}^* \approx 203.93$ for a ground-truth value of $\eta_\alpha = 200$. In contrast, the best classical viscosity parameter changed markedly from one horizon to another: $\eta^* \approx 177.28$ in the short-horizon case, $\eta^* \approx 313.49$ in the moderate case, and $\eta^* \approx 1973.99$ in the long-memory case. This drift is not a numerical accident; it is a structural symptom of trying to compress broad-band memory into a single exponential retardation time.

The global error metrics in Table 3 confirm that the stable fractional fit remains dramatically more faithful than the best classical fit. The MAE improvement factors were approximately 45.46 in the short-horizon scenario, 48.77 in the moderate scenario, and 52.02 in the long-memory scenario. Figure 5 shows that the three models may appear visually close when the full strain range is plotted, but Figure 6 reveals the decisive difference: the stable fractional fit stays near the trusted reference, whereas the classical model accumulates a persistent positive error and the unsafe fractional evaluation produces an almost full-scale deviation.

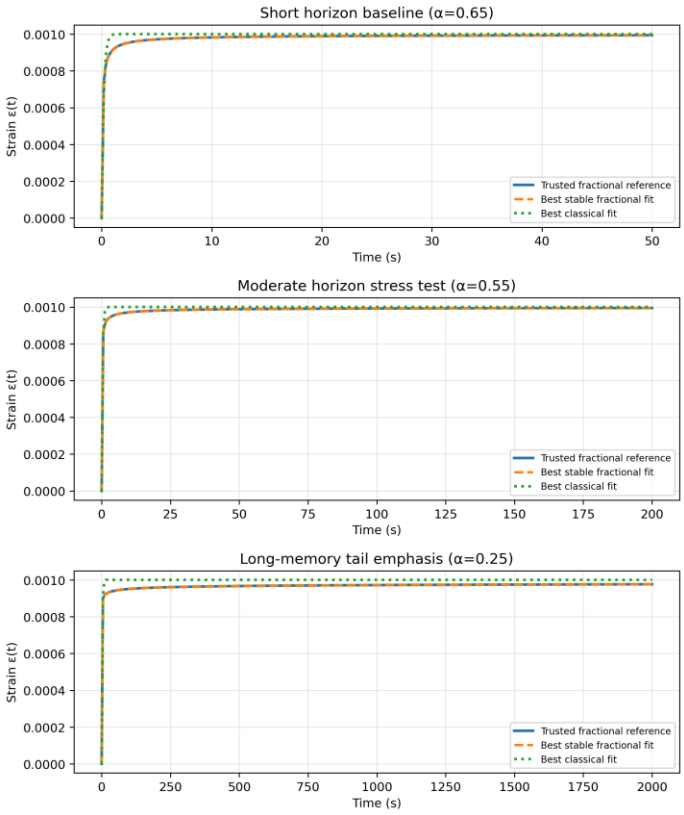

Figure 5. *Trusted fractional reference curves and best classical/fractional fits for the three creep scenarios considered in the article.*

Table 3. *Global creep-fit metrics for the best classical and stable fractional models.*

Scenario	Model	Viscosity parameter	MSE	MAE	MAXE
Short horizon baseline	Fractional stable fit	203.93	3.153e-13	3.284e-07	4.710e-06
Short horizon baseline	Classical fit	177.28	4.243e-10	1.493e-05	9.085e-05
Moderate horizon stress test	Fractional stable fit	203.93	1.065e-13	2.292e-07	2.622e-06
Moderate horizon stress test	Classical fit	313.49	2.087e-10	1.118e-05	7.184e-05
Long-memory tail emphasis	Fractional stable fit	203.93	3.867e-13	5.978e-07	1.803e-06
Long-memory tail emphasis	Classical fit	1973.99	1.044e-09	3.110e-05	7.802e-05

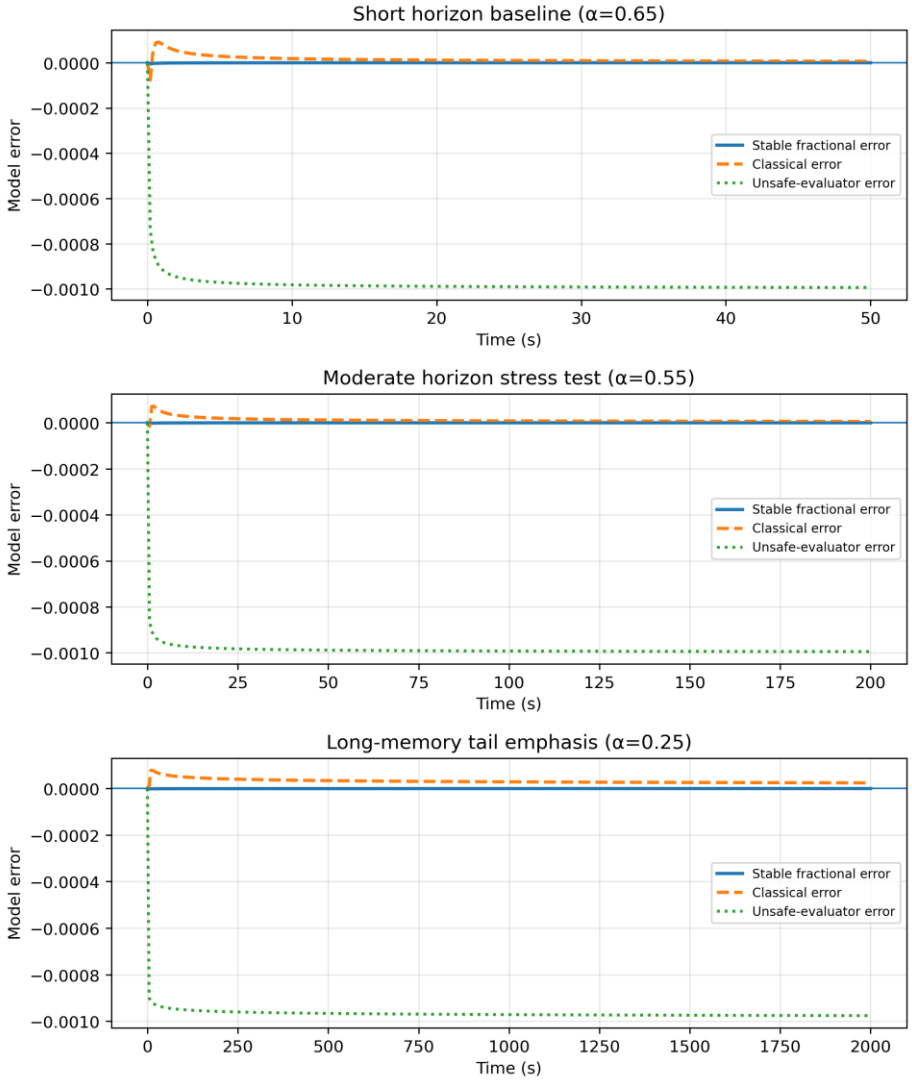

Figure 6. *Error trajectories for the stable fractional fit, the best classical fit, and the unsafe fractional evaluation.*

5.3 Long-memory discrimination and windowed errors

The most informative differences appear when the error is examined by time windows and when the long-memory case is viewed in logarithmic scale. Across all windows in Table 4, the fractional MAE was between 38.58 and 52.95 times smaller than the classical MAE. The advantage did not vanish at late times; on the contrary, it generally increased in the intermediate and tail windows, which is precisely the regime where long-memory discrimination matters for model selection. This result supports the finality stated in the abstract: a stable evaluator is not only mathematically correct, but practically necessary when the material signature is encoded in the slow transient rather than in the final plateau.

Figure 7(a) shows the long-memory scenario on log-log axes. The stable fractional fit tracks the trusted reference throughout the tail, while the classical fit remains systematically above it because the exponential approach cannot reproduce the hereditary curvature. Figure 7(b) complements that interpretation by showing a family of fractional creep curves generated for increasing α. As α grows, the creep curve approaches the equilibrium plateau more rapidly, whereas smaller α values preserve a slower tail. This family view clarifies why the fitted classical parameter must change so drastically across scenarios: the classical model has no dedicated order parameter with which to modulate memory intensity.

Table 4. *Windowed MAE comparison between classical and fractional fits*

Scenario	Time window	Classical MAE	Fractional MAE
Short horizon baseline	0-5	4.958e-05	1.285e-06
Short horizon baseline	5-20	1.691e-05	3.410e-07
Short horizon baseline	20-50	8.167e-06	1.626e-07
Moderate horizon stress test	0-20	3.240e-05	7.252e-07
Moderate horizon stress test	20-80	1.264e-05	2.496e-07
Moderate horizon stress test	80-200	6.914e-06	1.363e-07
Long-memory tail emphasis	0-50	5.654e-05	1.202e-06
Long-memory tail emphasis	50-300	4.462e-05	8.427e-07
Long-memory tail emphasis	300-2000	2.834e-05	5.435e-07

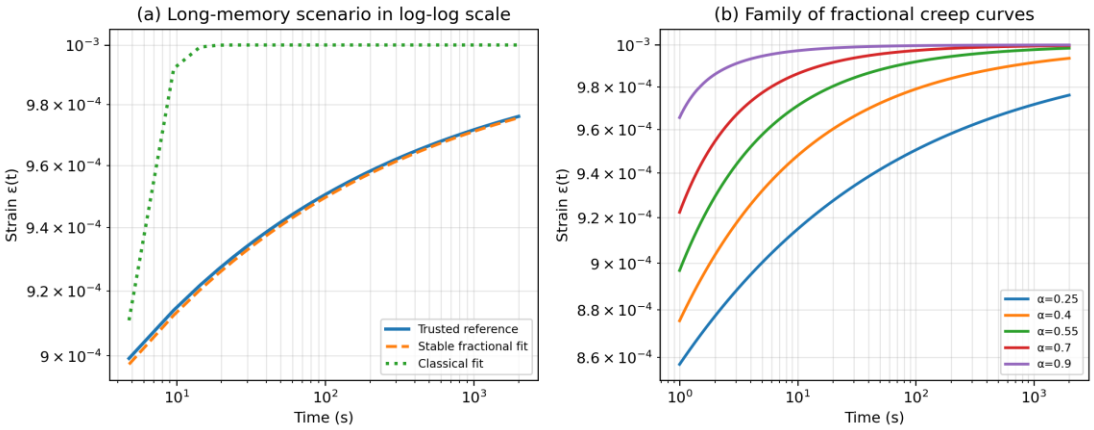

Figure 7. *Long-memory discrimination in logarithmic scale and family of fractional creep curves for representative α values.*

5.4 Robustness under mild synthetic observation noise

To test whether the methodological conclusion survives mild observation perturbations, the study repeated the three creep-fit scenarios after adding zero-mean Gaussian noise with standard deviation equal to 0.5% of the plateau strain σ_0/E. The fractional fit remained close to the ground-truth viscosity parameter $\eta_\alpha = 200$ in all scenarios, whereas the classical viscosity continued to drift strongly with the observation window. Table 5 shows that the fractional MAE stayed between $7.66e1$ and $8.78e1$ times smaller than the classical MAE when measured against the noiseless reference curve. Thus, the stable evaluator does not merely reproduce a clean synthetic benchmark; it also preserves the comparative advantage of the fractional model under mild measurement contamination.

Table 5. *Noise-robustness metrics for the revised creep-identification test*

Scenario	α	η_α^{fit}	η^{fit}	Frac. MAE	Class. MAE	MAE ratio
Short horizon baseline	0.65	202.34	147.97	1.988e-07	1.523e-05	76.63
Moderate horizon stress test	0.55	202.34	197.85	1.388e-07	1.155e-05	83.20

Long-memory tail emphasis	0.25	202.34	1195.32	3.566e-07	3.131e-05	87.81

6 Conclusions

The completed study validates the central claim of the article: reliable fractional Kelvin–Voigt creep simulation depends critically on how the Mittag–Leffler function is evaluated on the negative real axis. The benchmark against the trusted reference showed that the revised stable hybrid evaluator preserves boundedness and monotone decay while maintaining sub-percent worst-case relative error on the tested grid. The calibrated switch policy, made explicit in Section 3, eliminates the residual raw-branch defect that appeared in the highest-order case under a fixed early changeover, and the clipping layer remains a defensive safeguard rather than the source of the reported stability.

Once that numerical layer was validated, the creep experiments demonstrated the modelling consequence of stability. The stable fractional fit consistently recovered the hereditary viscosity parameter and outperformed the best classical Kelvin–Voigt fit by roughly one to two orders of magnitude in MAE, both globally and within individual time windows. The supplementary ablation clarified that this advantage arises from regime-aware branching rather than from post hoc correction, and the mild-noise experiment showed that the same qualitative conclusion survives small observation perturbations. The contribution of the article is therefore best understood as a validated computational layer for calibration-oriented fractional creep studies and reproducible numerical comparison, rather than as a blanket claim about every possible engineering deployment setting. Future work may extend the same validation logic to ramp-hold loadings, multi-term fractional networks, and matrix-valued formulations of Mittag–Leffler dynamics (Garrappa & Popolizio, 2018; McLean, 2021). Placed in the context of the broader numerical literature, this claim remains intentionally restricted to the negative-real-axis creep regime studied here; for wider parameter domains, contour-integral and Padé evaluators remain the natural cross-validation targets. Future work should therefore extend the same validation logic to systematic cross-comparison against contour-integral or rational-approximation evaluators on broader domains (Garrappa & Popolizio, 2013; Garrappa, 2015; Sarumi et al., 2020).

The code generated to obtain all the results in this chapter can be found at the following URL:

https://github.com/JAIME6609/VISCOCIDAD/blob/main/CODE-VISCOCIDAD-01-V1.py

References

Bonfanti, A., Kaplan, J. L., Charras, G., & Kabla, A. (2020). Fractional viscoelastic models for power-law materials. *Soft Matter, 16*(26), 6002–6020. https://doi.org/10.1039/D0SM00354A

Dabiri, D., Saadat, M., Mangal, D., & Jamali, S. (2023). Fractional rheology-informed neural networks for data-driven identification of viscoelastic constitutive models. *Rheologica Acta, 62*(10), 557–568. https://doi.org/10.1007/s00397-023-01408-w

Garrappa, R. (2015). Numerical evaluation of two and three parameter Mittag–Leffler functions. *SIAM Journal on Numerical Analysis, 53*(3), 1350–1369. https://doi.org/10.1137/140971191

Garrappa, R., & Popolizio, M. (2013). Evaluation of generalized Mittag–Leffler functions on the real line. *Advances in Computational Mathematics, 39*(1), 205–225. https://doi.org/10.1007/s10444-012-9274-z

Garrappa, R., & Popolizio, M. (2018). Computing the matrix Mittag–Leffler function with applications to fractional calculus. *Journal of Scientific Computing, 77*(1), 129–153. https://doi.org/10.1007/s10915-018-0699-5

Iyiola, O. S., Asante-Asamani, E. O., & Wade, B. A. (2018). A real distinct poles rational approximation of generalized Mittag–Leffler functions and their inverses: Applications to fractional calculus. *Journal of Computational and Applied Mathematics, 330*, 307–317. https://doi.org/10.1016/j.cam.2017.08.020

Lewandowski, R., & Chorążyczewski, B. (2010). Identification of the parameters of the Kelvin–Voigt and the Maxwell fractional models, used to modeling of viscoelastic dampers. *Computers & Structures, 88*(1–2), 1–17. https://doi.org/10.1016/j.compstruc.2009.09.001

Makris, N., & Efthymiou, E. (2020). Time-response functions of fractional derivative rheological models. *Rheologica Acta, 59*(12), 849–873. https://doi.org/10.1007/s00397-020-01241-5

McLean, W. (2021). Numerical evaluation of Mittag–Leffler functions. *Calcolo, 58*(1), Article 7. https://doi.org/10.1007/s10092-021-00398-6

Miles, P. R., Pash, G. T., Smith, R. C., & Oates, W. S. (2021). Bayesian inference and uncertainty propagation using efficient fractional-order viscoelastic models for dielectric elastomers. *Journal of Intelligent Material Systems and Structures, 32*(4), 486–496. https://doi.org/10.1177/1045389X20969847

Rodríguez Soto, A. A., Valín Rivera, J. L., Sanabio Alves Borges, L. M., & Palomares Ruiz, J. E. (2022). Parameter identification of fractional index viscoelastic model for vegetable-fiber reinforced composite. *Polymers, 14*(21), Article 4634. https://doi.org/10.3390/polym14214634

Sarumi, I. O., Furati, K. M., & Khaliq, A. Q. M. (2020). Highly accurate global Padé approximations of generalized Mittag–Leffler function and its inverse. *Journal of Scientific Computing, 82*(2), Article 46. https://doi.org/10.1007/s10915-020-01150-y

Shabani, M., Jahani, K., Di Paola, M., & Homayoun Sadeghi, M. (2019). Frequency domain identification of the fractional Kelvin–Voigt's parameters for viscoelastic materials. *Mechanics of Materials, 137*, Article 103099. https://doi.org/10.1016/j.mechmat.2019.103099

Wang, D., & Zou, J. (2023). Mittag–Leffler stability of numerical solutions to time fractional ODEs. *Numerical Algorithms, 92*(4), 2125–2159. https://doi.org/10.1007/s11075-022-01379-8

Zhang, H., Zhang, Q., Ruan, L., Duan, J., Wan, M., & Insana, M. F. (2018). Modeling ramp-hold indentation measurements based on Kelvin–Voigt fractional derivative model. *Measurement Science and Technology, 29*(3), Article 035701. https://doi.org/10.1088/1361-6501/aa9daf