PHYSICS-INFORMED SPECTRAL MEMORY IDENTIFICATION AND FLUTTER-REGIME DIAGNOSTICS FOR A PITCH–PLUNGE AEROELASTIC SYSTEM

Quvvatali Rahimov[1], Mirzohidjon Rahmatjonov[2]
[1] Fergana State University, ORCID:
[2] Fergana State University, ORCID: 0009-0005-2630-6336

Abstract

Hereditary effects in aeroelastic systems are governed by Volterra-type integro-differential equations whose memory kernel $G(t)$ is not directly measurable, and its recovery from noisy displacement records is a severely ill-posed inverse problem. We present a physics-informed identification framework in which the matrix-valued kernel is parameterized as a learnable sum-of-exponentials (spectral) measure, $G(t) \approx \sum_k W_k \exp(-s_k t)$ and the governing equation is enforced through a once-integrated residual that requires only smoothed displacements and velocities, eliminating the noise-amplifying second derivative. Flutter-regime diagnostics are obtained from the exact augmented state-space eigenvalue problem of the identified model, combined with robust empirical growth-rate estimation. In a fair twenty-seed synthetic benchmark in which all competing models are fitted to identical noisy kernel observations, the proposed spectral fit attains a mean flutter-speed error of 0.020%, kernel relative error 0.51%, outperforming Prony-2 0.352%, Prony-3 0.068%, and the memoryless baseline 0.905%. Applied to the open Springer pitch-and-plunge (PAPA) wind-tunnel dataset, the identified model reproduces the post-flutter modal frequency at the held-out airspeed $U = 9.50$ m/s with a relative error of 1.563% and supports the censored stability statement $U_f < 7.00$ m/s. The framework is reproducible, fully seeded, and reports only experimentally supported claims.

Keywords: *aeroelastic flutter; hereditary mechanics; Volterra integro-differential equation; memory kernel identification; inverse problem; physics-informed neural network; sum of exponentials; censored boundary.*

1. INTRODUCTION

Aeroelastic systems with viscoelastic structural elements or unsteady aerodynamic lag exhibit hereditary behavior: the instantaneous force depends on the entire deformation history. The canonical mathematical description is a Volterra integro-differential equation (IDE) in which a convolution memory kernel $G(t)$ encodes the fading influence of past motion [4, 5, 9]. Classical reduced-order representations of unsteady aerodynamics, from Theodorsen theory [9] to Jones-type rational approximations [8], and fractional-order viscoelastic damping models [6, 7] all reduce, at the equation level, to particular choices of this kernel. Predicting the flutter boundary of such a system therefore hinges on knowing $G(t)$ - yet in practice the kernel is never observed directly. It must be inferred from displacement time histories, which constitutes an inverse problem of Volterra type that is well known to be severely ill-posed [10–12]: small perturbations of the data can produce arbitrarily large perturbations of the recovered kernel.

Physics-informed neural networks (PINNs) [2, 3, 13] offer a flexible variational framework for such inverse problems, since unknown functional coefficients can be represented by trainable parameterizations and constrained jointly by data and the governing equation. However, a direct application to hereditary aeroelasticity faces three obstacles. First, an unconstrained neural kernel admits non-physical, sign-oscillating solutions that fit the residual but destabilize any forward simulation. Second, the strong-form residual contains the acceleration $\ddot{q}(t)$, whose numerical estimation from noisy measurements amplifies high-frequency noise quadratically. Third, the convolution term is typically several orders of magnitude smaller than the elastic and aerodynamic stiffness terms, so the identification is poorly conditioned and prone to kernel collapse, in which the network drives the memory term to zero while compensating through effective damping.

This paper addresses all three obstacles with a hardened identification core, referred to as KernelNet. The contributions are:

1. **Spectral kernel parameterization.** The matrix-valued kernel is represented as a discrete Bernstein-type spectral measure (sum of exponentials, SOE) over fixed log-spaced relaxation rates with learnable weights, Eq. (2). In the scalar positive mode the kernel is completely monotone by construction; in the matrix mode, cross-coupled aerodynamic memory between plunge and pitch can be learned while the same weights are reused exactly in both the kernel evaluation and the convolution, giving an $\mathrm{O}(NK)$ recursive convolution [14] and removing the unstable pseudo-inverse projection used in earlier versions.

2. **Once-integrated residual without accelerations.** Integrating the IDE once over sliding windows yields the residual (3), which involves only q and $\dot{q}$, both obtained by Savitzky–Golay smoothing [15] of the measured displacements. No second derivative of noisy data is ever formed.

3. **Robust flutter-regime diagnostics.** Stability is assessed by the exact eigenvalue problem of the augmented state-space system (4), while empirical growth rates are extracted by a robust estimator combining the Hilbert envelope, peak detection with saturation truncation, and Theil–Sen regression [16], which is insensitive to initial-condition outliers and limit-cycle saturation.

4. **Fair benchmarking and honest experimental claims.** All synthetic competitors are fitted from identical noisy observations across twenty random seeds, and the experimental result on the open Springer PAPA dataset [1] is reported as a left-censored stability boundary rather than an extrapolated flutter speed, because the available records do not bracket the boundary.

2. PROBLEM STATEMENT AND MATHEMATICAL MODEL

Consider a two-degree-of-freedom pitch–plunge section with generalized coordinates $q(t)=[h(t),\alpha(t)]^T$, where h is the plunge displacement and α the pitch angle. The hereditary aeroelastic model is the Volterra IDE

$$M\ddot{q}(t)+C\dot{q}(t)+Kq(t)+\int_0^t G(t-\tau)\dot{q}(\tau)d\tau=0 \qquad (1)$$

where M, C, and K are the effective mass, damping, and stiffness matrices (the latter two including the quasi-steady aerodynamic contributions, which depend on the flow speed U), and $G:[0,\infty)\to\mathbb{R}^{2\times2}$ is the unknown matrix memory kernel acting on the velocity history. The inverse problem is: given noisy samples of $q(t)$ along one or several trajectories at known airspeeds, recover $G(t)$ and subsequently predict the stability boundary, i.e., the flutter speed U_f at which the dominant oscillatory eigenvalue of the system crosses into the right half-plane.

Two structural facts make this problem tractable despite its ill-posedness. First, physically admissible relaxation kernels of viscoelastic and aerodynamic-lag type are completely monotone, hence representable by Bernstein's theorem as Laplace transforms of nonnegative measures. Second, for stability prediction the quantity of interest is not the pointwise kernel but the spectrum of the associated linear operator, which is far better conditioned with respect to kernel perturbations - an observation confirmed quantitatively by the benchmark in Section 5, where a 0.51% kernel error maps to a 0.020% flutter-speed error.

3. METHODOLOGY

The kernel is parameterized as a finite spectral measure on a fixed grid of $K = 64$ relaxation rates $s_k \in \left[\frac{1}{T}, \frac{1}{t_{\min}}\right]$ with $T = 40$ s and $t_{\min} = 10^{-2}s$:

$$G(t) \approx \sum_{k=1}^{K} W_k \exp(-s_k t) \tag{2}$$

where the weights $W_k \in \mathbb{R}^{2\times2}$ are produced by a small fully connected network (3 hidden layers, width 96) evaluated at the grid nodes. Three modes are supported: a scalar positive mode (softplus weights times the identity), which is completely monotone by construction and is used for synthetic compatibility; a diagonal-positive signed-cross matrix mode, used for the experimental study under the configuration matrix_diag_positive, in which G_{hh} and $G_{\alpha\alpha}$ are constrained to be non-negative, while $G_{h\alpha}$ and $G_{\alpha h}$ are learned as signed cross-coupling terms; and a fully signed matrix mode in which all four components are unconstrained. Crucially, the same weights define both the kernel values and the convolution: introducing auxiliary states z_k per relaxation rate turns the convolution into a recursion with cost $\mathrm{O}(NK)$ per trajectory of length N [14], so no quadrature of the full history is needed and the earlier unstable projection step $\mathrm{pinv}(A)\cdot G$ is eliminated entirely.

Instead of the strong-form residual of Eq. (1), which would require $\ddot{q}(t)$, the equation is integrated once over sliding windows $[a,t]$ of length 0.06 s:

$$R(a,t) = M[\dot{q}(t)-\dot{q}(a)] + C[q(t)-q(a)] + K\int_a^t q(\eta)d\eta + \int_a^t\int_0^\eta G(\eta-\tau)\dot{q}(\tau)d\tau d\eta = 0 \tag{3}$$

All ingredients of R — the smoothed displacement, its first derivative, its running integral, and the convolution — are computed from the measured q via a Savitzky–

Golay filter (window 0.035 s, cubic polynomial) and trapezoidal cumulative integration. Collocation points (800 per trajectory) are split temporally into training and validation segments with a guard gap of one integral window, preventing leakage between the two. The training loss combines the normalized residual, a stability penalty discouraging spuriously unstable identified systems (weight 2.0), and a short Abel-type warm-up phase (1,200 iterations toward a fractional target with $\gamma = 0.45$) that initializes the spectrum in a physically plausible region before the main Adam phase (6,000 iterations, learning rate 2×10^{-4} with cosine decay, gradient clipping at 1.0, master seed 42).

Because the identified kernel is a finite sum of exponentials, the closed-loop system admits an exact finite-dimensional realization. Augmenting the state with the convolution states z_k,

$$\dot{z}_k(t) = -s_k z_k(t) + \dot{q}(t), \qquad M\ddot{q}(t) + C\dot{q}(t) + Kq(t) + \sum_{k=1}^{K} W_k z_k(t) = 0 \tag{4}$$

yields a linear system of dimension $4 + 2K$ whose eigenvalues determine modal growth rates and frequencies exactly; sweeping the airspeed and locating the zero crossing of the dominant oscillatory growth rate by bracketing gives the predicted flutter speed, with total cost $O(N_U(4+2K)^3)$ for N_U sweep points. On the empirical side, growth rates of measured records are estimated robustly: the Hilbert envelope of the dominant signal is peak-detected, peaks beyond limit-cycle saturation (85% of the late median level) are discarded, and the log-amplitude slope is fitted by Theil–Sen regression [16]. A rollout is accepted only if it remains finite and bounded over the full horizon; early termination is scored as failure regardless of the accuracy of its finite prefix, which prevents the metric gaming observed with truncated-horizon scoring.

4. EXPERIMENTAL SETUP

The synthetic ground truth is a two-degree-of-freedom system with speed-dependent damping and a five-mode matrix exponential kernel (relaxation rates $0.7\text{--}55\,\mathrm{s}^{-1}$), whose exact flutter speed, computed by the augmented eigenvalue sweep, is U_f = 8.0961. To make the comparison fair, every competing model is fitted to identical noisy kernel observations: 180 log-spaced samples of $G(t)$ on $t \in [10^{-3}, 2]\,\mathrm{s}$ corrupted by 2% RMS Gaussian noise, regenerated independently for each of 20 seeds. Four models are compared: (i) no memory $G \equiv 0$; (ii) Prony series with 2 modes; (iii) Prony series with 3 modes; and (iv) the proposed spectral fit on 16 fixed log-spaced nodes in $[0.05, 200]\,\mathrm{s}^{-1}$. The fixed-node spectral fit is solved by linear least squares. The Prony-2 and Prony-3 baselines are implemented as fixed-node low-order exponential fits, with relaxation rates placed on a logarithmic grid and weights obtained by linear least squares. This avoids optimizer-dependent variability and makes the comparison fully reproducible. All fitted kernels are then evaluated by the same augmented-state eigenvalue sweep to obtain the predicted flutter speed.

The experimental study uses the open supplementary dataset of Santos et al. [1]: bench-test records of a pitch-and-plunge aeroelastic (PAPA) apparatus undergoing nonlinear post-flutter oscillations. The whitespace-separated time-response files are

parsed into (t,h,α,U) trajectories; the airspeed is decoded from the file name. Five airspeed levels are available: $U \in \{7.00,8.15,8.40,9.00,9.50\}$ m/s . Five training trajectories at the four lower airspeed levels are used for identification and growth-rate calibration, including two records at $U = 9.50$ m/s; $U = 7.00$ m/s is held out entirely and used only for validation of the predicted post-flutter modal frequency and growth rate. Trajectories are resampled to a uniform grid where timing jitter exceeds tolerance, and the kernel is identified in the diagonal-positive signed-cross matrix mode (matrix_diag_positive) with the hyperparameters of Section 3. Thus, the diagonal memory terms are constrained to be positive, whereas the plunge–pitch and pitch–plunge memory couplings are learned as signed quantities. All measured records exhibit positive growth rates, i.e., the apparatus is post-critical at every recorded airspeed; consequently the flutter boundary is not bracketed by the data, and the framework reports it as left-censored.

5. RESULTS

Table 1 summarizes the benchmark. The memoryless model incurs a systematic 0.905% flutter-speed bias — small in relative terms but entirely structural, since it cannot be reduced by better data. Prony-2 and Prony-3 reduce the error to 0.352% and 0.068% respectively, while the proposed spectral fit achieves 0.020 ± 0.014%, with the lowest kernel reconstruction error (0.51 ± 0.09% relative L2). The ordering is strict and consistent across all twenty seeds, and the 95% confidence interval of the spectral fit, [8.092, 8.100], is the only one containing the true value 8.0961. Figure 1 shows the modal growth-rate curves of representative fits and the per-model error distribution; the spectral curve is visually indistinguishable from the true boundary behavior near the crossing.

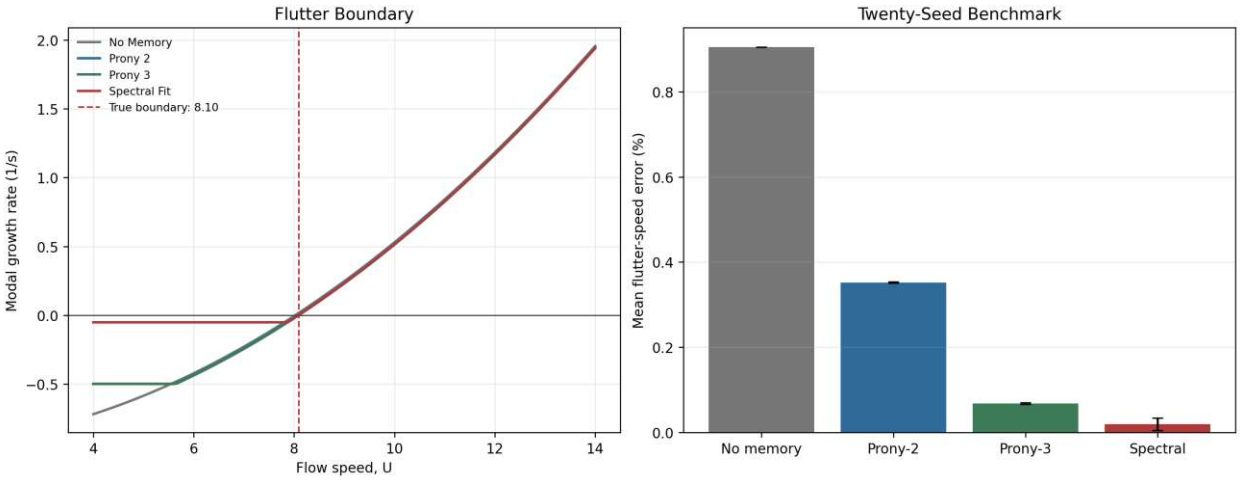


Figure 1. *Fair synthetic benchmark. Left: modal growth rate versus flow speed for representative fits of each model class; the dashed line marks the true boundary U_f = 8.10. Right: mean flutter-speed error over twenty seeds with standard-deviation bars.*

Table 1.

Twenty-seed synthetic benchmark: all models fitted from identical noisy kernel observations (2% RMS noise). True flutter speed U_f = 8.0961.

Model	U_f **predicted(m/s)**	U_f **error (%)**	**Kernel RelL2(%)**	**95% CI (m/s)**
No memory	8.0228 ± 0.0000	0.905 ± 0.000	100.00	[8.023, 8.023]
Prony-2	8.0676 ± 0.0001	0.352 ± 0.002	19.12 ± 0.00	[8.067, 8.068]
Prony-3	8.0905 ± 0.0002	0.068 ± 0.002	2.09 ± 0.01	[8.090, 8.091]
Spectral fit (proposed)	8.0959 ± 0.0020	0.020 ± 0.014	0.51 ± 0.09	[8.092, 8.100]

Figure 2 shows the kernel identified from the PAPA records and its relaxation spectrum. Under the diagonal-positive signed-cross parameterization, Figure 2 shows that the identified memory is mainly expressed through the off-diagonal coupling terms. The $\alpha \to h$ and $h \to \alpha$ channels dominate the identified kernel, while the diagonal components remain comparatively small. the $\alpha \to h$ channel is large and positive at short lags and decays monotonically, whereas the $h \to \alpha$ channel changes sign near $t \approx 0.02$ s, peaks around 0.06 s, and decays — a signature of phase-lagged aerodynamic coupling between plunge velocity and pitch moment that a scalar or diagonal-only kernel cannot represent. The effective memory of the identified kernel is concentrated below one second, consistent with the convective time scales of the apparatus.

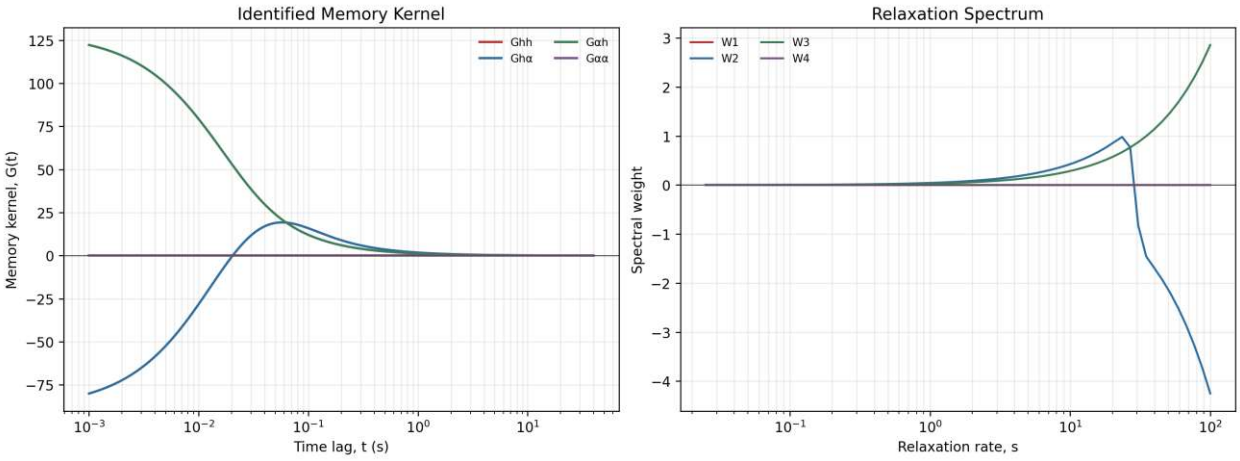


Figure 2*. Kernel identified from the Springer PAPA records. Left: matrix kernel components versus time lag (log scale). Right: learned spectral weights versus relaxation rate.*

Table 2 and Figure 3 present the experimental validation. Since every recorded airspeed, including the lowest U = 7.00 m/s, exhibits a positive robust growth rate, the only stability statement supported by the data is the censored bound $U_f < 7.00$ m/s; the framework refuses to extrapolate an exact experimental flutter speed. Validation therefore targets post-flutter dynamics: the identified spectral model, calibrated on the

four training airspeeds, predicts the dominant modal frequency at the held-out $U = 9.50$ m/s with an error of -0.0541 Hz (1.563% relative), and the predicted frequency-versus-speed branch tracks all training points within comparable accuracy (Figure 3, right). The holdout growth-rate error is 0.2747 s^{-1}, of the same order as the seed-to-seed scatter of the empirical growth estimates themselves (e.g., 0.538 s^{-1} standard deviation across the two records at U = 7.00 m/s), indicating that the residual mismatch is dominated by measurement variability rather than model bias.

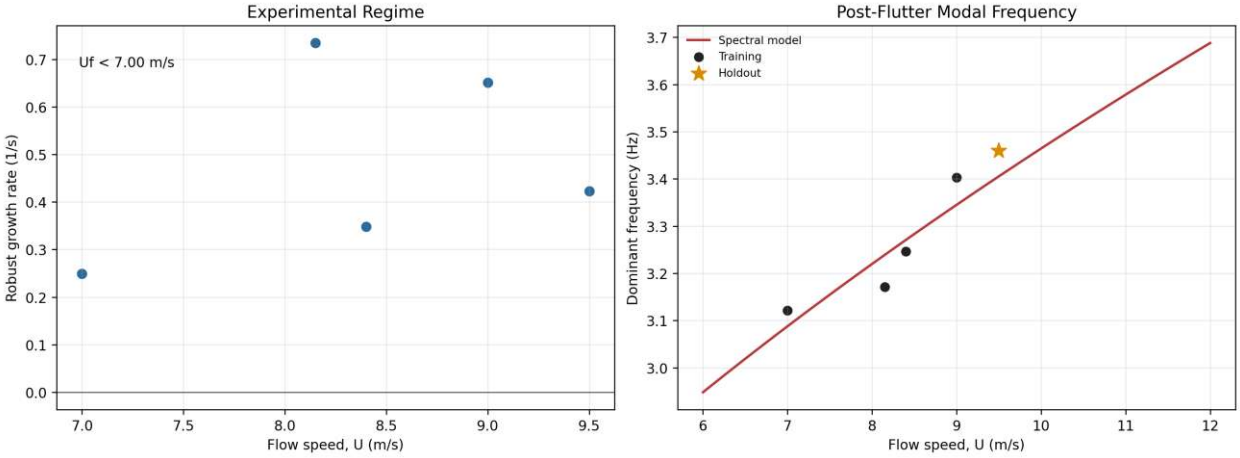


Figure 3. *Springer PAPA case study. Left: robust empirical growth rates of the measured records; all are positive, so the boundary is left-censored ($U_f < 7.00$ m/s). Right: post-flutter modal frequency of the identified spectral model versus the training points and the held-out record at $U = 9.50$ m/s (star).*

Table 2.

Springer PAPA experimental case study summary.

Quantity	Value
Boundary status	Left-censored
Supported boundary statement	$U_f < 7.00$ m/s
Holdout airspeed	$U = 9.50$ m/s
Holdout frequency error	−0.0541 Hz
Holdout relative frequency error	1.563%
Holdout growth-rate error	0.2747 s^{-1}
Growth-rate calibration cost	0.597
Identification wall-clock time	221.5 s (CPU run, seed 42)

6. DISCUSSION AND LIMITATIONS

The benchmark isolates a representational effect: with the observations, noise level, and evaluation protocol held fixed, richer admissible kernel classes translate monotonically into better boundary prediction, and the spectral parameterization extracts most of the information present in the observations, achieving a 0.020% flutter-

speed error despite 2% RMS noise in the kernel samples. The experimental study demonstrates the complementary methodological point — that the same machinery yields honest claims when the data are weaker. Two limitations are explicit. First, the Springer records do not bracket the stability boundary, so the experimental result is a censored bound, not a flutter speed; claiming an exact experimental U_f is unsupported, and we deliberately do not. Second, the identified model is validated on modal quantities (frequency, growth rate); accurate long-horizon displacement rollout of the strongly nonlinear post-flutter limit cycles is outside the scope of the linear hereditary model and is likewise not claimed. Bracketing data at sub-critical airspeeds, nonlinear stiffness extensions of the residual (cubic and mixed terms are already provisioned in the formulation), and uncertainty quantification of the identified spectrum are natural next steps.

We presented a hardened physics-informed framework for identifying matrix-valued memory kernels in hereditary aeroelastic systems and converting them into flutter-regime diagnostics. The combination of a spectral sum-of-exponentials parameterization, a once-integrated acceleration-free residual, and exact augmented-state eigenanalysis achieves a 0.020% mean flutter-speed error in a fair twenty-seed synthetic benchmark — approximately a 45-fold reduction in error relative to the memoryless baseline and a 3.4-fold reduction relative to an equally fitted three-mode Prony series — and reproduces the post-flutter modal frequency of an open wind-tunnel dataset at a held-out airspeed within 1.6%. The framework runs end to end in under four minutes (221.5 s CPU identification and 4.8 s for the full benchmark analysis), is fully seeded and reproducible, and reports stability boundaries in censored form whenever the data do not support a point estimate.

REFERENCES

[1] Santos, F. L. M. dos, et al. Wind tunnel bench test of a pitch-and-plunge aeroelastic model undergoing nonlinear post-flutter oscillations. Discover Mechanical Engineering, 3, 15 (2024). https://doi.org/10.1007/s44245-024-00043-w

[2] Raissi, M., Perdikaris, P., Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378 (2019), 686–707.

[3] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., Yang, L. Physics-informed machine learning. Nature Reviews Physics, 3 (2021), 422–440.

[4] Dowell, E. H. (ed.). A Modern Course in Aeroelasticity. 6th ed., Springer, 2022.

[5] Dimitriadis, G. Introduction to Nonlinear Aeroelasticity. Wiley, 2017.

[6] Bagley, R. L., Torvik, P. J. A theoretical basis for the application of fractional calculus to viscoelasticity. Journal of Rheology, 27(3) (1983), 201–210.

[7] Mainardi, F. Fractional Calculus and Waves in Linear Viscoelasticity. 2nd ed., World Scientific, 2022.

[8] Jones, R. T. The unsteady lift of a wing of finite aspect ratio. NACA Report 681, 1940.

[9] Theodorsen, T. General theory of aerodynamic instability and the mechanism of flutter. NACA Report 496, 1935.

[10] Lavrentiev, M. M., Romanov, V. G., Shishatskii, S. P. Ill-Posed Problems of Mathematical Physics and Analysis. AMS, 1986.

[11] Kabanikhin, S. I. Inverse and Ill-Posed Problems: Theory and Applications. De Gruyter, 2011.

[12] Janno, J., von Wolfersdorf, L. Inverse problems for identification of memory kernels in viscoelasticity. Mathematical Methods in the Applied Sciences, 20(4) (1997), 291–314.

[13] Lu, L., Meng, X., Mao, Z., Karniadakis, G. E. DeepXDE: A deep learning library for solving differential equations. SIAM Review, 63(1) (2021), 208–228.

[14] Jiang, S., Zhang, J., Zhang, Q., Zhang, Z. Fast evaluation of the Caputo fractional derivative and its applications to fractional diffusion equations. Communications in Computational Physics, 21(3) (2017), 650–678.

[15] Savitzky, A., Golay, M. J. E. Smoothing and differentiation of data by simplified least squares procedures. Analytical Chemistry, 36(8) (1964), 1627–1639.

[16] Theil, H. A rank-invariant method of linear and polynomial regression analysis. Proceedings of the Royal Netherlands Academy of Sciences, 53 (1950), 386–392.

Муродова Мухлиса

10 BUZILADIGAN YUQORI JUFT TARTIBLI XUSUSIY HOSILALI DIFFERENSIAL TENGLAMA UCHUN LOKAL SHARTLI BOSHLANG‘ICH-CHEGARAVIY MASALA
Tolipov Nodirjon Isaqovich **49**

11 SiGe VA SiSn QATTIQ QORISHMALARI FIZIK XOSSALARINING SHAKLLANISHIDA TARKIB VA MEXANIK ZO‘RIQISHNING AHAMIYATI
Madaminov Xurshidjon Muxamedovich
Abdurahmonov Azizbek Abdurahim o‘g‘li **52**

12 Ce VA Yb BILAN LEGIRLANGAN Bi_2Te_3 ASOSIDAGI TERMOELEKTRIK MATERIALLARNING ELEKTROFIZIK PARAMETRLARI QIYOSIY TAHLILI
Yusupova D.A.
Sirojiddinova S.Z. **57**

13 EXISTENCE AND UNIQUENESS OF A SOLUTION TO A BITSADZE-SAMARSKII PROBLEM FOR A SUB-DIFFUSION EQUATION WITH THE REGULARIZED PRABHAKAR FRACTIONAL DERIVATIVE
Mirzaeva Maftuna **65**

14 CONVERGENCE OF A NEW TRIVARIATE MITTAG-LEFFLER-TYPE FUNCTION
Tursunova Ergashoy **68**

15 PHYSICS-INFORMED SPECTRAL MEMORY IDENTIFICATION AND FLUTTER-REGIME DIAGNOSTICS FOR A PITCH–PLUNGE AEROELASTIC SYSTEM
Quvvatali Rahimov
Mirzohidjon Rahmatjonov **73**

16 KADMIY TELLURIDGA OLTINGUGURT QO‘SHILGAN HOLDA p-n O‘TISH HOSIL QILISH VA UNING OPTIK XUSUSIYATLARI
Najmidinova Dildoraxon Xalilovna **82**

17 FIZIKA TA'LIMIDA DINAMIK TEST TIZIMINING DIDAKTIK IMKONIYATLARI VA AMALIY AHAMIYATI
Qoraboyev Xusniddin **87**

18 Ni ASOSLI NANOKATALIZATORLARDA UGLEROD NANONAYCHALARINI O‘STIRISH **94**