THE UNIVERSITY OF OSAKA / RESEARCH DISCUSSION

GPBiCG(m, ℓ, k) Adaptive strategies and learned initialization

Dr. Ronan Dupont

Nagoya University / Zhang–Sogabe Laboratory

4–6 November 2026

Joint GPBiCG(m, ℓ, k) work with Tomohiro Sogabe and Shao-Liang Zhang.
From static numerical evidence to a KAKENHI research agenda.

名古屋大学 Nagoya University

MOONSHOT RESEARCH & DEVELOPMENT PROGRAM

Can a good solver decision be made before solving?

The numerical problem

Solve $Ax = b$, with A large, sparse and nonsymmetric. Convergence must be checked, not inferred from a recommendation.

1. **Extend the solver:** does a third phase improve GPBiCG(m, ℓ)?
2. **Adapt the triplet:** can live control rescue a poor start?
3. **Learn the choice:** does a predicted triplet repay its own cost?

The research question is not only which solver wins, but how much evidence a reliable decision requires.

Motivation: ICMASE2026 study; MLK serial studies, October 2026; FY2027 KAKENHI research plan.

Research roadmap

I GPBiCG family
Static contribution and eight-matrix evidence

II Live adaptivity
Prescribed-start rescue and blind starts

III Learning
Training matrices and triplet prediction

IV Measured results
Matched curves and complete CPU costs

V Conclusions
Established results and open questions

GPBiCG family

First establish the static contribution. Adaptivity follows from matrix-dependent choices.

A third phase expands the family, not a universal ranking

All phases use the product-type update

$$r_{n+1} = t_n - \eta_n y_n - \zeta_n A t_n.$$

Phase	Number of updates	Mechanism
BiCGSTAB-like	m	one-direction minimization
GPBiCG-like	ℓ	two-direction minimization
LOPBiCG-like	k	orthogonality-based choice

- Repeat these three phases as one complete cycle.
- $k = 0$ recovers GPBiCG(m, ℓ); $k > 0$ adds useful choices.
- More choices also increase the parameter-selection problem.

van der Vorst (1992); Zhang (1997); Fujino (2002); Zhao et al. (2025), DOI 10.1002/nla.70033. Exact coefficients and cycle pseudocode: appendix.

GPBiCG(m, ℓ, k) Cycle Pseudo-Code

1: Initialize $\boldsymbol{x}_0$, $\boldsymbol{r}_0 = \boldsymbol{b} - A\boldsymbol{x}_0$, and a shadow residual $\boldsymbol{r}_0^*$.
2: **while** $\|\boldsymbol{r}_n\|_2 / \|\boldsymbol{b}\|_2 > \varepsilon$ **do**
3: **for** $j = 1, \ldots, m$ **do**
4: Perform one BiCGSTAB-like update.
5: **end for**
6: **for** $j = 1, \ldots, \ell$ **do**
7: Perform one GPBiCG residual-minimizing update.
8: **end for**
9: **for** $j = 1, \ldots, k$ **do**
10: Perform one LOPBiCG orthogonality update.
11: **end for**
12: Evaluate the true relative residual.
13: **end while**

At each update

- build the BiCG direction;
- form the intermediate residual $\boldsymbol{t}_n$;
- select (ζ_n, η_n) from the active phase;
- update $\boldsymbol{x}_{n+1}$ and $\boldsymbol{r}_{n+1}$.

Stopping rule

The experiments report the true residual $\|\boldsymbol{b} - A\boldsymbol{x}_n\|_2 / \|\boldsymbol{b}\|_2$, not only the recursively updated residual norm.

Static Numerical Comparison

Eight nonsymmetric matrices from the SuiteSparse Matrix Collection

Matrix	N	nnz	$\kappa_2(A)$	Matrix	N	nnz	$\kappa_2(A)$
utm300	300	3,155	1.50×10^{6}	memplus	17,758	99,147	1.29×10^{5}
add20	2,395	13,151	1.20×10^{4}	epb2	25,228	175,027	2.62×10^{3}
Goodwin_017	3,317	97,773	1.30×10^{5}	wang4	26,068	177,196	4.03×10^{5}
poisson3Da	13,514	352,762	1.12×10^{3}	crashbasis	160,000	1,750,416	–

N: dimension; nnz: nonzero entries; $\kappa_2(A)$: 2-norm condition number; –: unavailable.

Comparators

- BiCGSTAB and BiCGSTAB2;
- GPBiCG and GPBiCG(m, ℓ).

All solvers use the same matrix, right-hand side, initial guess and stopping rule.

Grid and stopping rule

$(m, \ell, k) \in \{1, \ldots, 7\}^3$ (343 triplets).

Initial guess: $\boldsymbol{x}_0 = \mathbf{0}$.

Right-hand side: $\boldsymbol{b} = A\mathbf{1}$, with $\mathbf{1} = (1, \ldots, 1)^{\mathsf{T}}$ the all-ones vector.

Stopping rule: $\|\boldsymbol{b} - A\boldsymbol{x}_n\|_2 / \|\boldsymbol{b}\|_2 \leq 10^{-12}$.

Settings: none and conventional ILU(0).

Without Preconditioning: Residual Histories

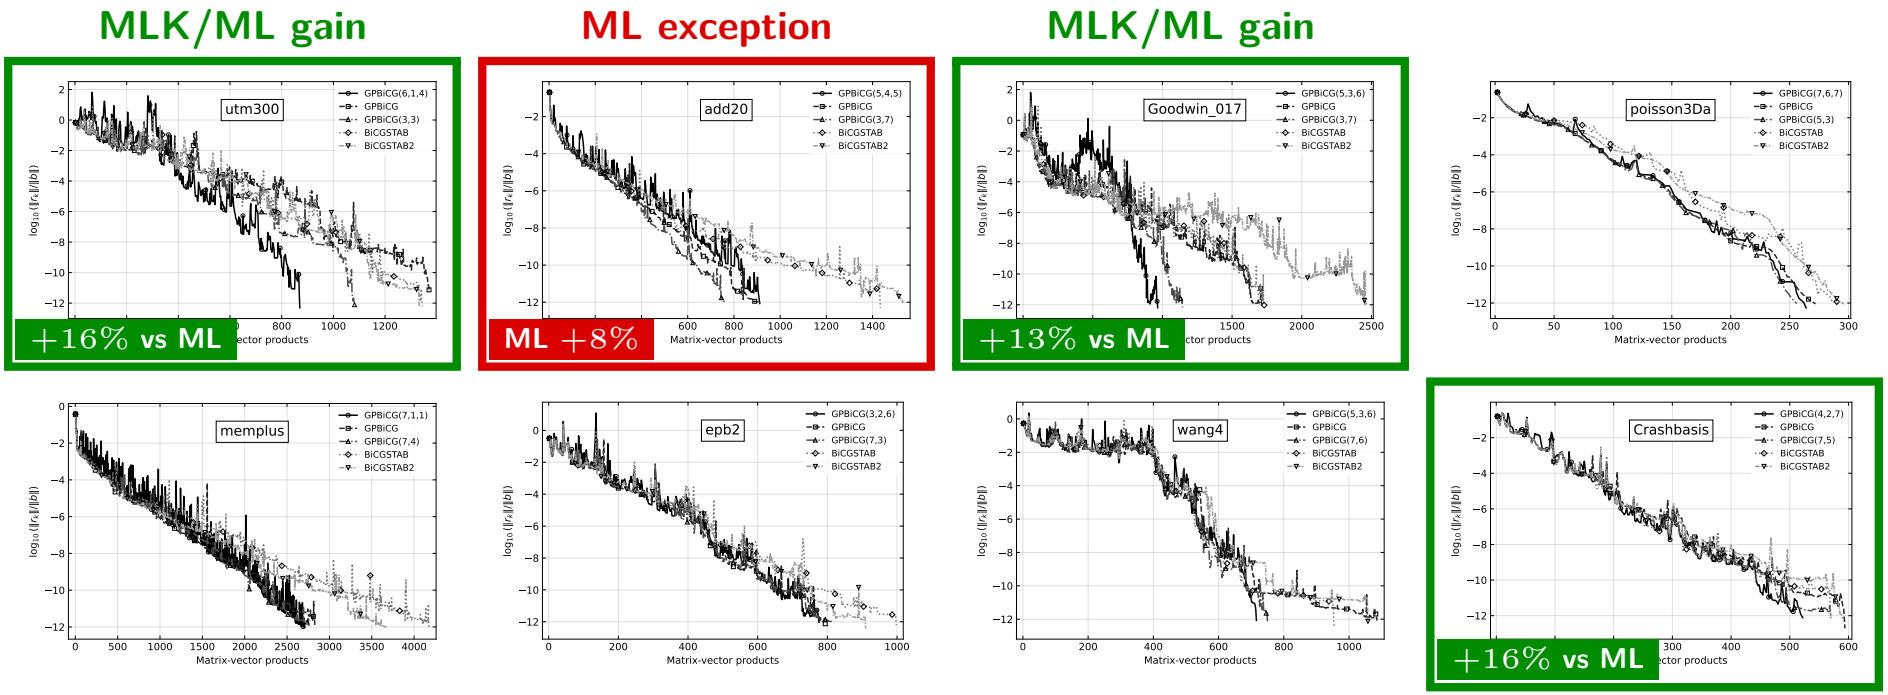

Figure 1: Residual histories without preconditioning. The frames identify the trajectories discussed below; all eight cases remain visible.

- At 10^{-12}, GPBiCG(m, ℓ, k) is faster than GPBiCG(m, ℓ) on 7/8 matrices; only `add20` favours the two-parameter family.
- The largest CPU gains over GPBiCG(m, ℓ) are about 16% on `utm300`, 13% on `Goodwin_017`, and 16% on `Crashbasis`.
- The residual paths are not uniformly ordered: the useful triplet is still matrix-dependent, which motivates the adaptive study.

With ILU(0): Residual Histories

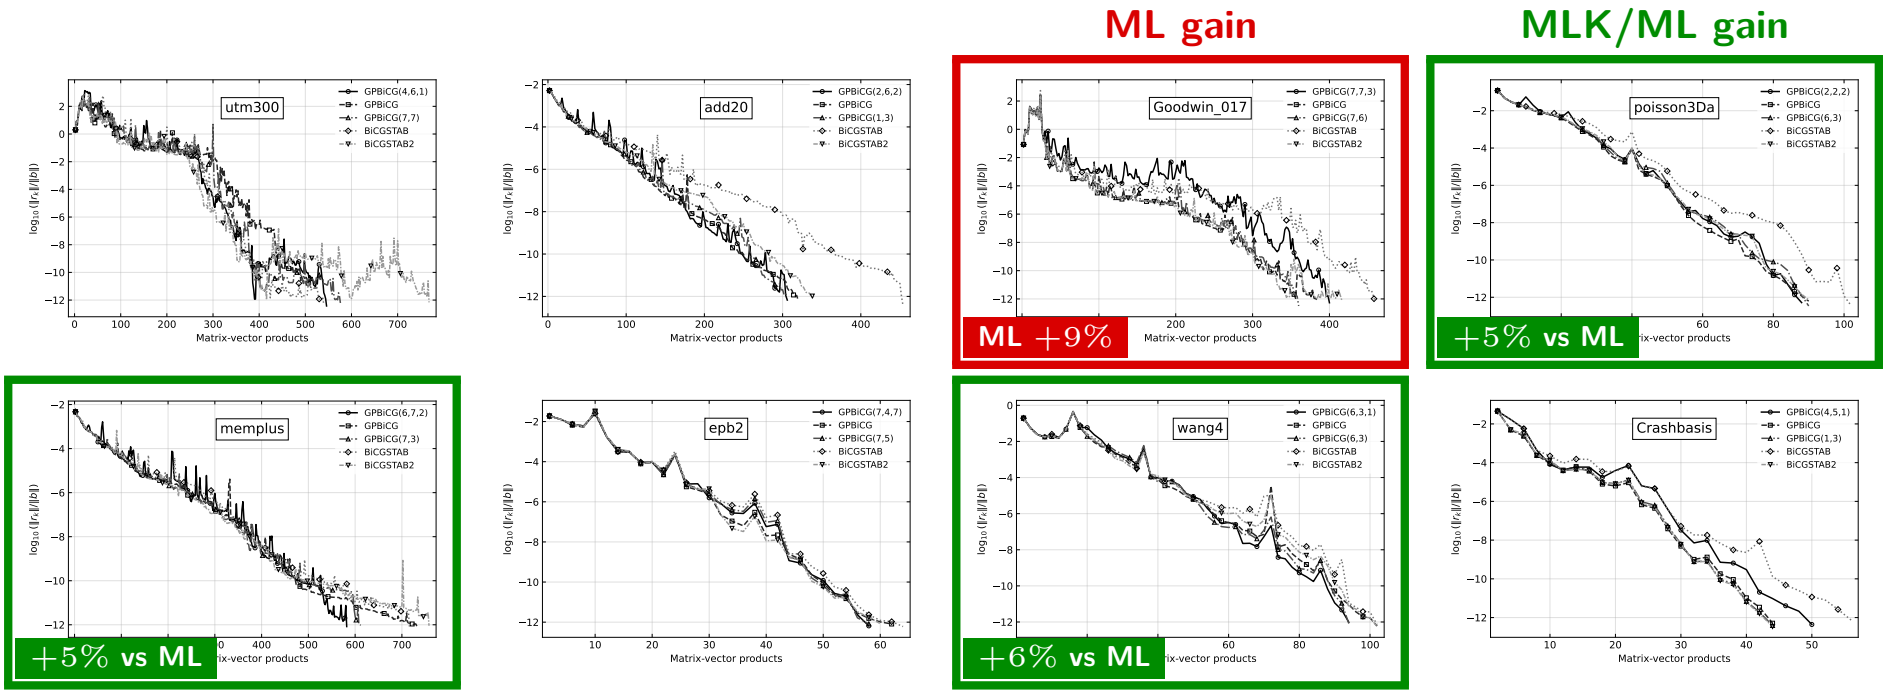

Figure 2: Residual histories with ILU(0). Frames direct attention to the mixed MLK/ML outcome quantified in the CPU table.

- With ILU(0), GPBiCG(m, ℓ, k) is faster than GPBiCG on all six comparable converged rows, with speedups from 1.13 to 1.37.
- Against GPBiCG(m, ℓ) the conclusion is deliberately mixed: 3 wins, 1 near tie, and 3 losses on the converged comparable rows.
- ILU(0) therefore supports the method, but also shows that preconditioning reduces the room for the extra k-phase to help.

Static Results: CPU Synthesis

Without preconditioning

- GPBiCG(m, ℓ, k) is the fastest solver on 6/8 matrices.
- It is faster than GPBiCG(m, ℓ) on 7/8 matrices.
- Largest CPU reductions vs GPBiCG(m, ℓ): about 16% on `utm300` and `Crashbasis`, and 13% on `Goodwin_017`.

With ILU(0)

- GPBiCG(m, ℓ, k) is the fastest converged method on 4/7 comparable rows.
- It is faster than plain GPBiCG on all comparable rows, with speedups from 1.13 to 1.37.
- Against GPBiCG(m, ℓ), the result is mixed: 3 wins, 1 near tie, 3 losses.

Interpretation

The k-phase is useful, but not universally. The supported conclusion is not domination; it is parameter sensitivity plus recoverable performance.

II

Live adaptivity

A poor initial triplet is easier to rescue than a completely parameter-free interface.

One complete cycle, then one possible decision

Observe

Keep $\theta_q = (m_q, \ell_q, k_q)$ fixed during cycle q. For each phase $p \in \{B, G, L\}$ with $n_{q,p} > 0$ updates,

$$\rho_{q,p} = \left(\frac{\|r_{q,p}^{\text{out}}\|_2}{\|r_{q,p}^{\text{in}}\|_2} \right)^{1/n_{q,p}}.$$

For the whole cycle,

$$g_q = \frac{\|r_q^{\text{out}}\|_2}{\|r_q^{\text{in}}\|_2}.$$

Propose and check

Move one update from the phase with largest $\rho_{q,p}$ to the one with smallest $\rho_{q,p}$.

- Accept only if $g_q < 1$, all phase scores are < 1, and the triplet is feasible.
- Distinct best/worst phases; no mid-cycle change.
- Otherwise retain the triplet.

Feasible set: $\Theta = \{1, \ldots, 7\}^3$. A move preserves $m + \ell + k$. Zero residual terminates; invalid scores do not trigger a switch.

Experiment I: two productive cycles locked after a move. Experiment II adds fixed confidence/cost gates. These controller constants are not matrix-wise tuned; the interface, not every internal rule, is parameter-free. Details: appendix.

An accepted transfer at a cycle boundary

Illustration for the prescribed-start rule, at an unlocked decision. The full random-start rule has an additional confidence/cost gate.

Observed cycle

$$\theta_q = (2, 4, 3), \qquad g_q = 0.84,$$
$$(\rho_{q,B}, \rho_{q,G}, \rho_{q,L}) = (0.91, 0.97, 0.88).$$

- Smallest score: L (0.88).
- Largest score: G (0.97).
- Cycle and all phases contract.

Next complete cycle

Transfer one update from G to L:

$$\theta' = (2, 4-1, 3+1) = (2, 3, 4).$$

It lies in $\Theta = \{1, \ldots, 7\}^3$ and keeps $m + \ell + k = 9$.

$$\boxed{\theta_{q+1} = (2, 3, 4)}$$

Run the next complete cycle.

Accepting a transfer is an empirical decision, not a guarantee of future speedup.

Worked example retained from the original ICMASE controller explanation, page 17. These numbers illustrate the rule, not a new measured trajectory. In experiment I, a move then locks adaptation for two productive cycles.

Experiment I Shows Selective Rescue, Not Free Speedup

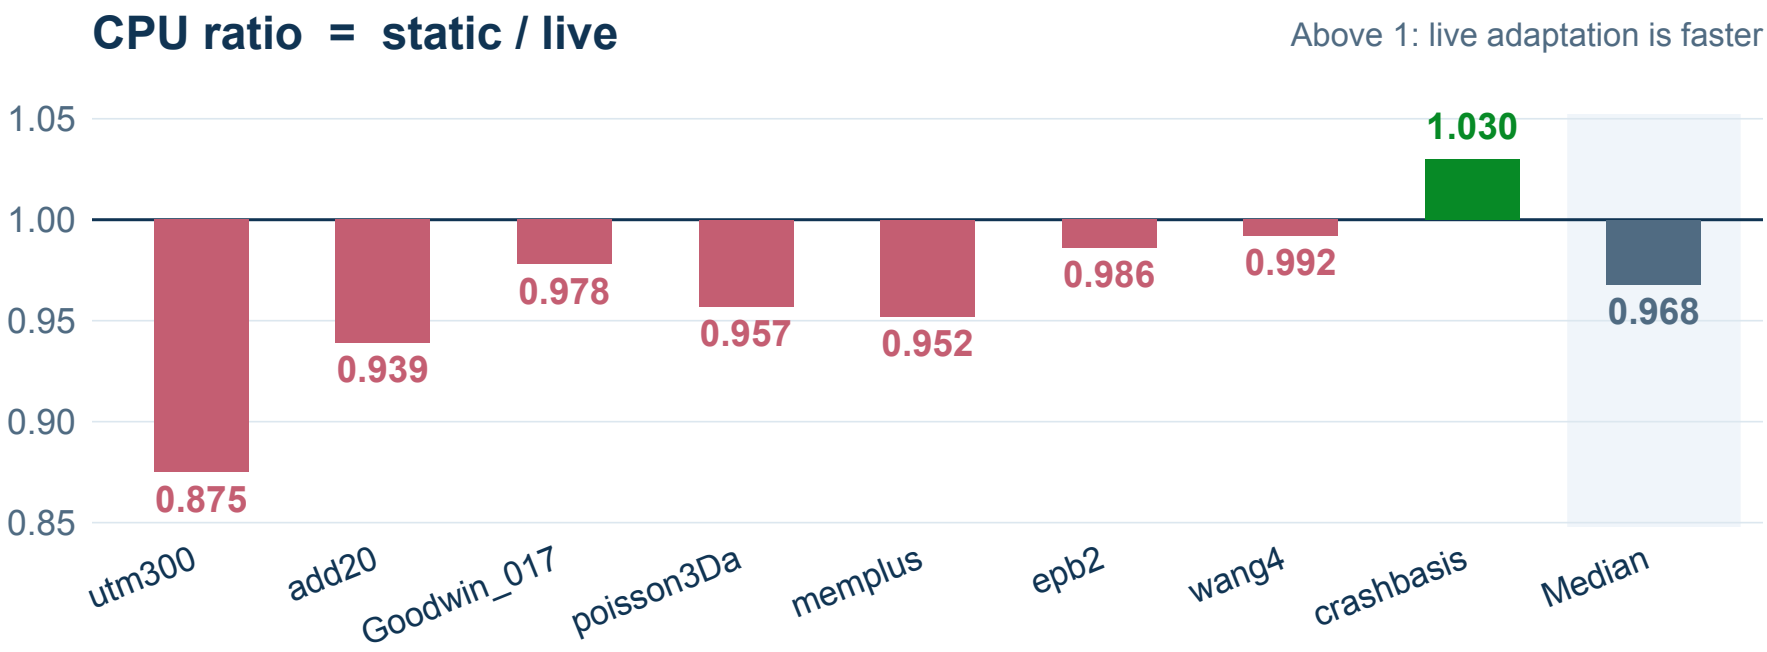

Figure 3: Prescribed-start rescue experiment. Ratios above one favour the live run over the static run from the same triplet.

- 2744 paired starts; transitions occur in about 25% of runs.
- Median static/live $= 0.968$: overhead is visible when the initial triplet is already acceptable.
- Clear rescue on `crashbasis`: static/live $= 1.030$, with live runs faster in 69.1% of prescribed starts.
- Interpretation: the controller is a safeguard. It is not expected to win when the prescribed triplet is already good.

Experiment II Quantifies the Price of Blind Starts

Matrix	Conv.	PF mean (s)	Best MLK/PF	Best classical/PF	Classical winner
utm300	66%	0.0051	0.63	0.95	BiCGSTAB
add20	100%	0.0315	0.76	0.87	GPBiCG
Goodwin_017	49%	0.1605	0.53	0.85	BiCGSTAB
poisson3Da	100%	0.0885	0.93	**1.01**	BiCGSTAB
memplus	97%	0.6646	0.83	0.96	GPBiCG
epb2	99%	0.2277	0.88	0.90	BiCGSTAB
wang4	93%	0.3000	0.67	0.67	BiCGSTAB
crashbasis	100%	1.1870	0.80	0.84	BiCGSTAB

Table 3: Fully parameter-free (PF) experiment. Best MLK is the offline oracle from the completed 7^3 grid; the classical reference is the fastest of GPBiCG, BiCGSTAB and BiCGSTAB2. Each ratio is comparator CPU divided by PF mean CPU, so values above one favour PF.

- Best MLK/PF measures the price of removing triplet tuning: its median ratio is 0.78, so the tuned offline oracle remains faster on all matrices.
- Best classical/PF is the strict no-tuning comparison, with the classical baseline selected independently on every matrix.
- PF wins that test on `poisson3Da`; its median ratio is $\simeq 0.885$, corresponding to a median CPU penalty of about 13–14%.
- Thus the result is reduced triplet sensitivity at a measurable cost, not superiority over the best possible tuned triplet.

III

Learning

Learn from other matrices, then choose one complete triplet before solving.

Training matrices and evaluation matrices

Partition	Matrices	Groups
Training	659	95
Article cases, excluded from fitting	8	–

- Group related matrices before splitting, rather than splitting runs.
- Keep the eight article matrices out of fitting and model selection.
- A training row describes a matrix, not an individual solver run.

The eight article cases illustrate behaviour; they are already exposed, not a new blind test.

Frozen audited corpus, October 2026. Groups are collection/provenance proxies; the training corpus is not certified nonsymmetric-only. Complete 826-matrix identity/group list: bundled training_matrices.csv.

What is learned from each training matrix?

$$\Theta = \{1, \ldots, 7\}^3, \quad \theta = (m, \ell, k) \in \Theta : \qquad 343 \text{ candidate triplets.}$$

1. Run each candidate and independently check convergence.
2. Use three runs to obtain a median native CPU time.
3. Divide by the best verified median for that training matrix.
4. Assign a fixed penalty to failed candidates.

Training supplies 343 relative-cost targets per matrix. Prediction does not run those 343 solves again.

Targets are capped at 10; failure penalty 10 is dimensionless. Native training labels were collected under concurrency. Physical tolerance 10^{-12}. Exact formula and failure handling: appendix.

How is a new matrix represented?

Features	Information supplied to the model
13	Size, sparsity, diagonal dominance, entry scale/signs, $\|A\mathbf{1}\|_2$.
26	Add structural and row/column scaling statistics.
43	Add deterministic probes: 11 matrix/transpose-vector products.

- No matrix name and no new solver result is an input.
- These frozen descriptors depend on A and the $b = A\mathbf{1}$ contract.

Richer information is not free: descriptor extraction belongs to selection time.

Frozen base/enriched descriptor implementations. Complete definitions and zero-denominator conventions: appendix. No arbitrary new-RHS feature is claimed. Related work: Sun, Jing & Xu, "A New Matrix Feature Selection Strategy in Machine Learning Models for Certain Krylov Solver Prediction," Journal of Classification 42, 72–89 (2025), DOI 10.1007/s00357-024-09484-0.

How is the triplet selected?

$$z(A) \xrightarrow{\text{frozen forest}} (\widehat{y}_\theta)_{\theta\in\Theta} \xrightarrow{\text{smallest score}} \widehat{\theta} = (m, \ell, k).$$

- ExtraTrees predicts the cost of all 343 triplets; choose the lowest score.
- A practical tabular baseline with simple native inference; no neural-network comparison was run.
- Exact ties use a ranking learned from fitting TRAIN only.
- Two different matrices may legitimately receive the same triplet.

The model chooses a combination; it does not predict three unrelated integers.

192 trees, maximum depth 12, minimum leaf size 2, seed 20261018; no new hyperparameter search. Geurts, Ernst and Wehenkel, Machine Learning 63, 3–42 (2006), DOI 10.1007/s10994-006-6226-1.

Prediction and solve: the complete algorithm

1. **Describe:** compute the frozen features of A.
2. **Predict:** evaluate the 343 forest scores.
3. **Select:** retain the minimum-score triplet.
4. **Solve:** start from $x_0 = 0$ with this fixed triplet.
5. **Verify:** check $\|b - Ax\|_2 / \|b\|_2 \leq 10^{-12}$ independently.
6. **Account:** retain selection, solve, verification and total CPU.

No test-matrix sweep, no cached solution, and no convergence inferred from a prediction.

Model and matrix already resident. In the new matched experiment every request recomputes features and predicts again; no cached decision. Offline training and first model/matrix loading are separate costs.

IV

Measured results

Keep the article’s five families. Measure selection and resolution separately.

Matched numerical comparison

- Same eight article matrices; $b = A\mathbf{1}$, $x_0 = 0$; no preconditioner.
- Physical tolerance 10^{-12}, iteration cap 2000, one thread.
- **360 complete requests**: nine arms, eight matrices, five repeats.

Arms	Selection rule
Five article families	GPBiCG, BiCGSTAB, BiCGSTAB2, historical ML/MLK parameter references.
Fixed control	GPBiCG$(1, 1, 1)$.
Three learners	Frozen 13-, 26-, 43-feature predictions.

Historical ML/MLK configurations are fixed before timing and rerun here, not new October oracles. No CGS arm. Native recurrence objects unchanged; shuffled arm order per repetition. Models/matrices resident; all requests reset the solution.

Residual histories: the 13-feature learner

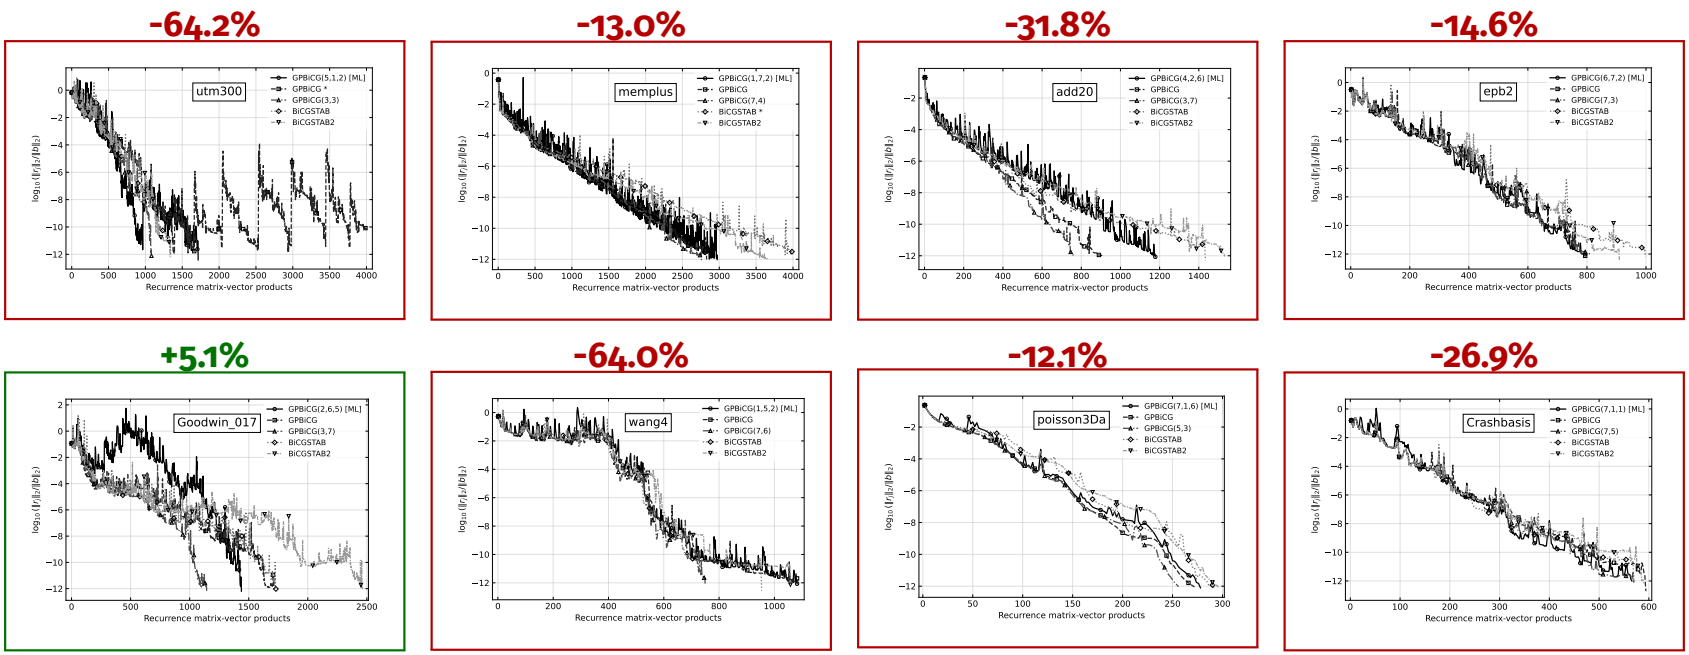

- The MLK curve uses the **predicted** triplet, not the best test triplet.
- Same generator, method styles and matrix order as the article.
- The next slides enlarge the informative cases and show their paid costs.

Badges: complete-request CPU gain vs fastest verified classical (green = faster, red = slower). Five-repeat residual histories; physical convergence checked independently. Horizontal axis: two recurrence matrix actions per update; not CPU time.

Goodwin: a useful predicted triplet

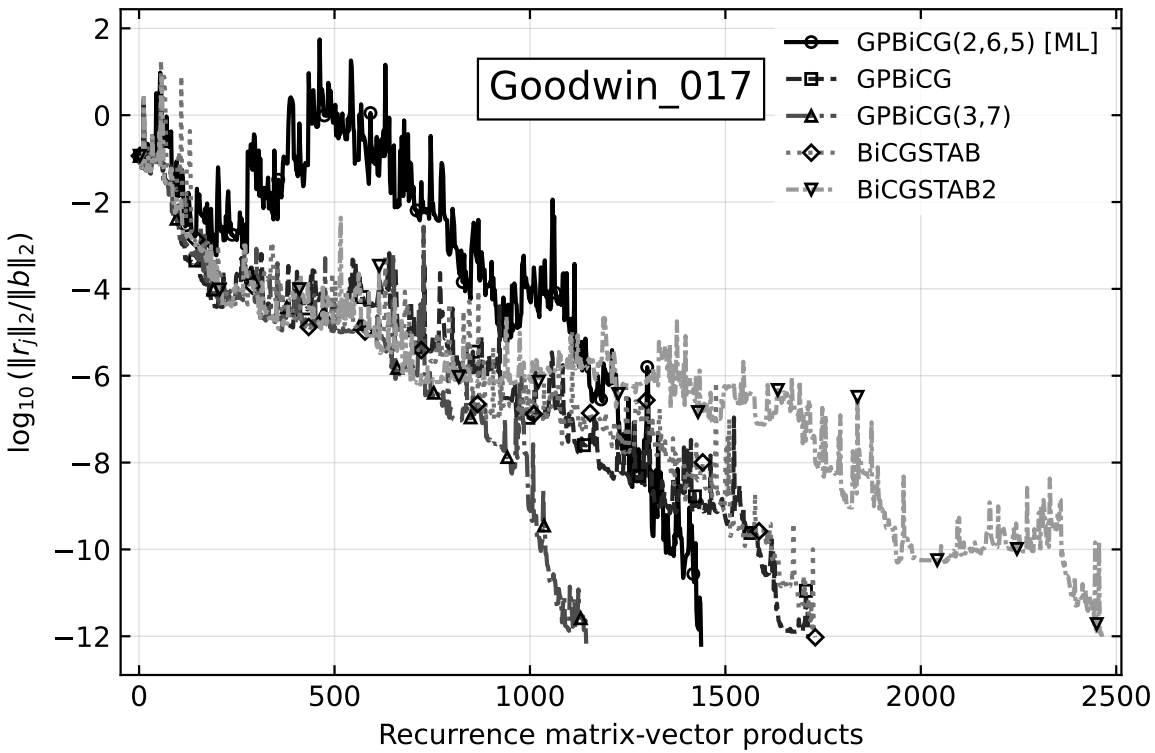

Component	CPU
Selection	0.402 ms
Native solve	0.11990 s
Check + overhead	0.167 ms
Total	0.12047 s
BiCGSTAB	0.12696 s
CPU reduction	+5.11%

The 13-feature model chooses $(2, 6, 5)$ and retains a local gain after paying for selection.

New matched-session means, five repeats; native features and native forest. Best verified classical: BiCGSTAB. No cached decision. The historical offline ML reference is shown, not treated as no-tuning.

memplus: richer features, a small paid advantage

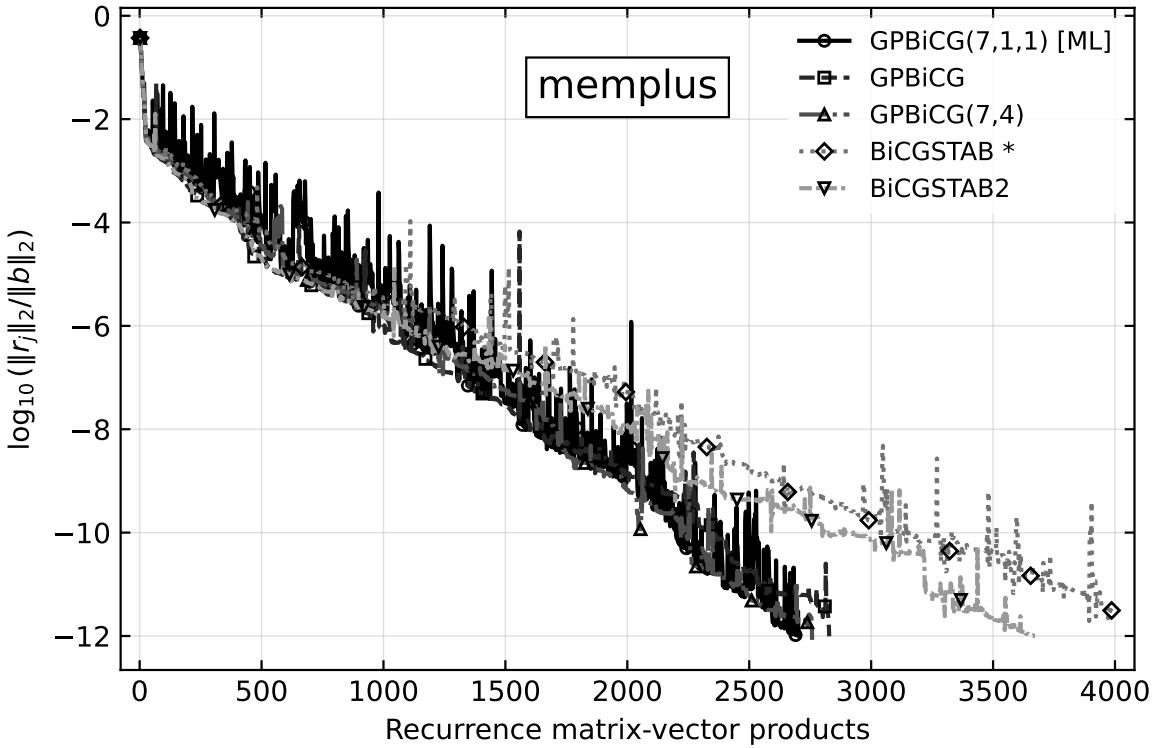

Component	CPU
Selection	8.730 ms
Native solve	0.56266 s
Check + overhead	0.234 ms
Total	0.57162 s
GPBiCG	0.57786 s
CPU reduction	+1.08%

The 43-feature model chooses $(7, 1, 1)$; its roughly 1% paid margin remains small.

One fixed 43-feature learner, distinct from the Goodwin learner; no per-matrix model oracle. Best verified classical: GPBiCG. Features remain Python; forest inference is native. Five-repeat mean, not a significance claim.

utm300: selection can erase a solver-only gain

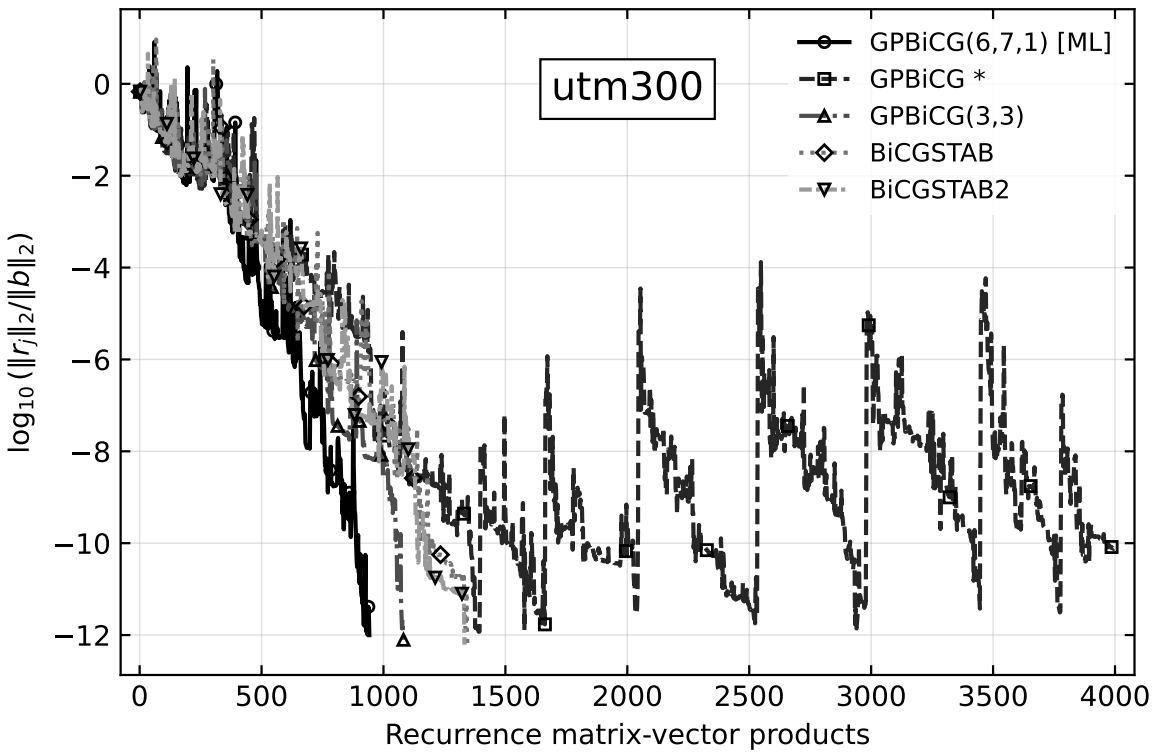

Component	CPU
Selection	1.029 ms
Native solve	0.00406 s
Check + overhead	0.060 ms
Total	0.00515 s
BiCGSTAB	0.00441 s
CPU reduction	-16.95%

The solve is cheaper, but the complete predicted route is slower than BiCGSTAB.

The same frozen 43-feature learner chooses $(6, 7, 1)$. Fresh descriptors and prediction are paid on every request. No cherry-picked cache regime. All methods use the same resident accounting.

Selection time and solve time are different costs

$$T_{\text{total}} = T_{\text{sel}} + T_{\text{solve}} + T_{\text{check}} + T_{\text{dispatch}}.$$

Matrix	Triplet	T_{sel} ms	T_{solve} s	Other ms	T_{total} s	Gain %
utm300	5,1,2	0.178	0.00700	0.056	0.00724	-64.2
add20	4,2,6	0.264	0.03465	0.129	0.03504	-31.8
Goodwin_017	2,6,5	0.402	0.11990	0.167	0.12047	+5.1
poisson3Da	7,1,6	1.472	0.08696	0.368	0.08880	-12.1
memplus	1,7,2	0.483	0.65217	0.241	0.65289	-13.0
epb2	6,7,2	0.626	0.22769	0.265	0.22858	-14.6
wang4	1,5,2	0.634	0.32125	0.274	0.32216	-64.0
crashbasis	7,1,1	3.671	1.07743	1.287	1.08239	-26.9

- **Selection**: features, forest and minimum-score choice.
- **Solve**: complete native invocation, including its internal checks.
- **Other**: independent check plus dispatch; not counted twice.

13-feature model; selection/Other in ms, solve/total in s. Gain vs fastest verified GPBiCG/BiCGSTAB/BiCGSTAB2. Means sum before rounding. Initial loading and offline training are outside this resident request scope.

All eight matrices: local gains, not general superiority

Features	Median selection (ms)	Wins vs classical
13	0.554	1/8
26	8.052	0/8
43	9.294	1/8

- A single learner is evaluated on all eight matrices, including losses.
- No learner passed the original grouped qualification procedure.
- All 24 paid predictions remain slower than the offline parameter references.

The positive evidence is local, paid and reproducible; the general ML claim remains unproved.

360-request matched session, five repetitions per arm. Selection statistic: median of eight matrix-specific mean CPU times. Counts are descriptive, not statistical tests. Full three-model tables and curves: appendix.

V

Conclusions

Separate the static contribution from preliminary parameter-selection strategies.

Conclusions

- **Static extension:** a third phase creates useful matrix-dependent choices.
- **Live adaptation:** some poor initial triplets can be rescued.
- **Learned initialization:** some recommendations are useful, but selection has a measurable cost.
- **Limits:** no universal no-tuning advantage and no qualified learned policy.

Parameter-sensitivity reduction is the research objective; universal solver dominance is not the claim.

Static and live evidence: original article campaigns. Learned initialization: separate frozen-model 360-request comparison. No tuning on article outcomes and no pooled timing sessions.

Complete CPU comparison: original article families

Matrix	MLK ref.	ML ref.	GPBiCG	BiCGSTAB	BiCGSTAB2	ML13	ML26	ML43
utm300	0.00362	0.00416	0.01570†	0.00441	0.00535	0.00724	0.00762	0.00515
add20	0.02722	0.02098	0.02658	0.03146	0.04213	0.03504	0.03811	0.03254
Goodwin_017	0.07979	0.09188	0.13921	0.12696	0.19743	0.12047	0.22268	0.19781
poisson3Da	0.08217	0.07632	0.08371	0.07921	0.08903	0.08880	0.09732	0.09670
memplus	0.55718	0.52120	0.57786	0.62820†	0.72720	0.65289	0.67351	0.57162
epb2	0.21724	0.19525	0.21096	0.19953	0.23323	0.22858	0.23366	0.24114
wang4	0.20310	0.19014	0.29664	0.19648	0.27558	0.32216	0.33431	0.32499
crashbasis	1.01518	0.99965	1.13964	0.85264	1.08050	1.08239	1.17176	1.27706

- Seconds per complete resident request, mean of five repetitions.
- ML/MLK references use the historical parameters, rerun in this session.
- A dagger retains an unverified arm; it cannot count as a speedup.

No July timings mixed into the October table. Offline parameter search and model training are not charged to the resident request; parameterized references are not no-tuning controls.

All eight matrices: the 26-feature learner

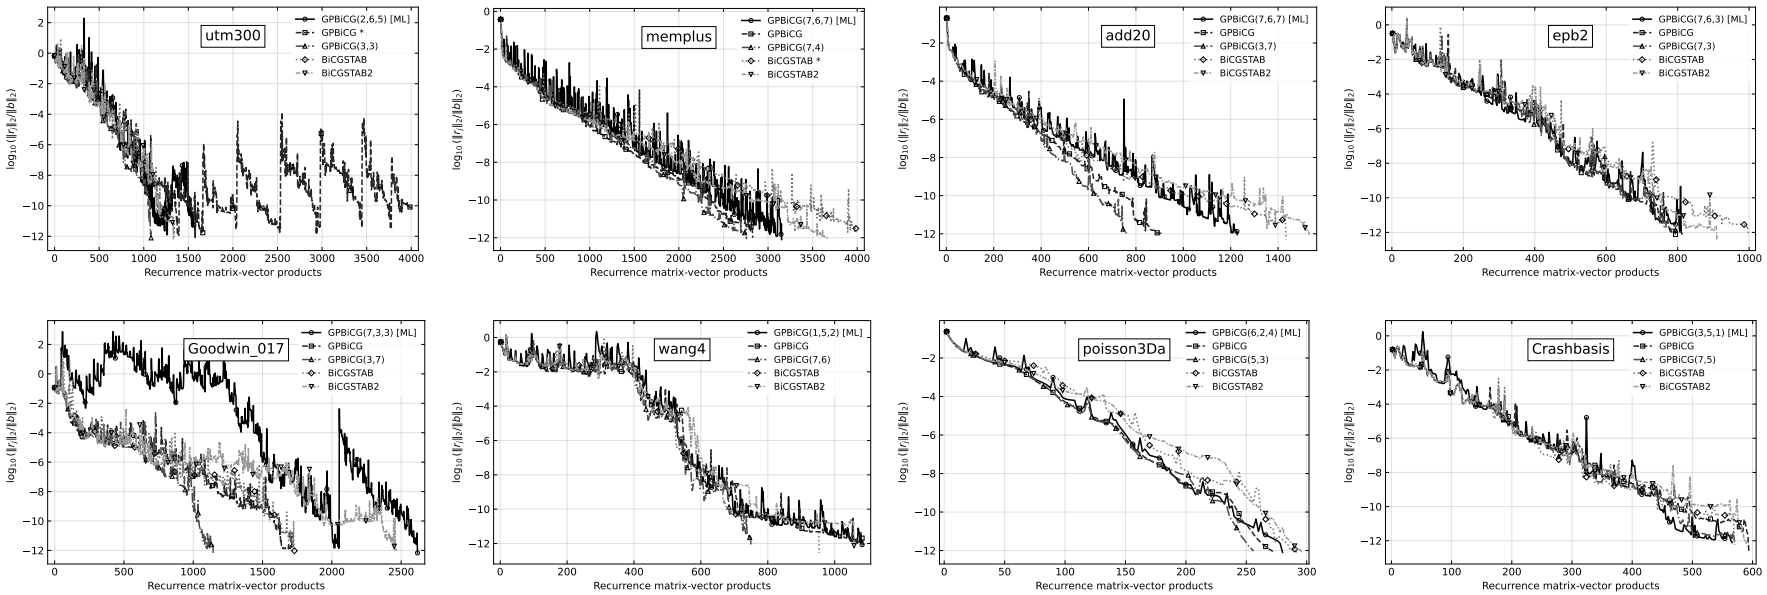

Same matrix order, five families and plotter; predictions frozen before timing. Recursive residual histories, physical candidate/final checks, two recurrence matrix actions per logged update. All failures remain visible.

All eight matrices: the 43-feature learner

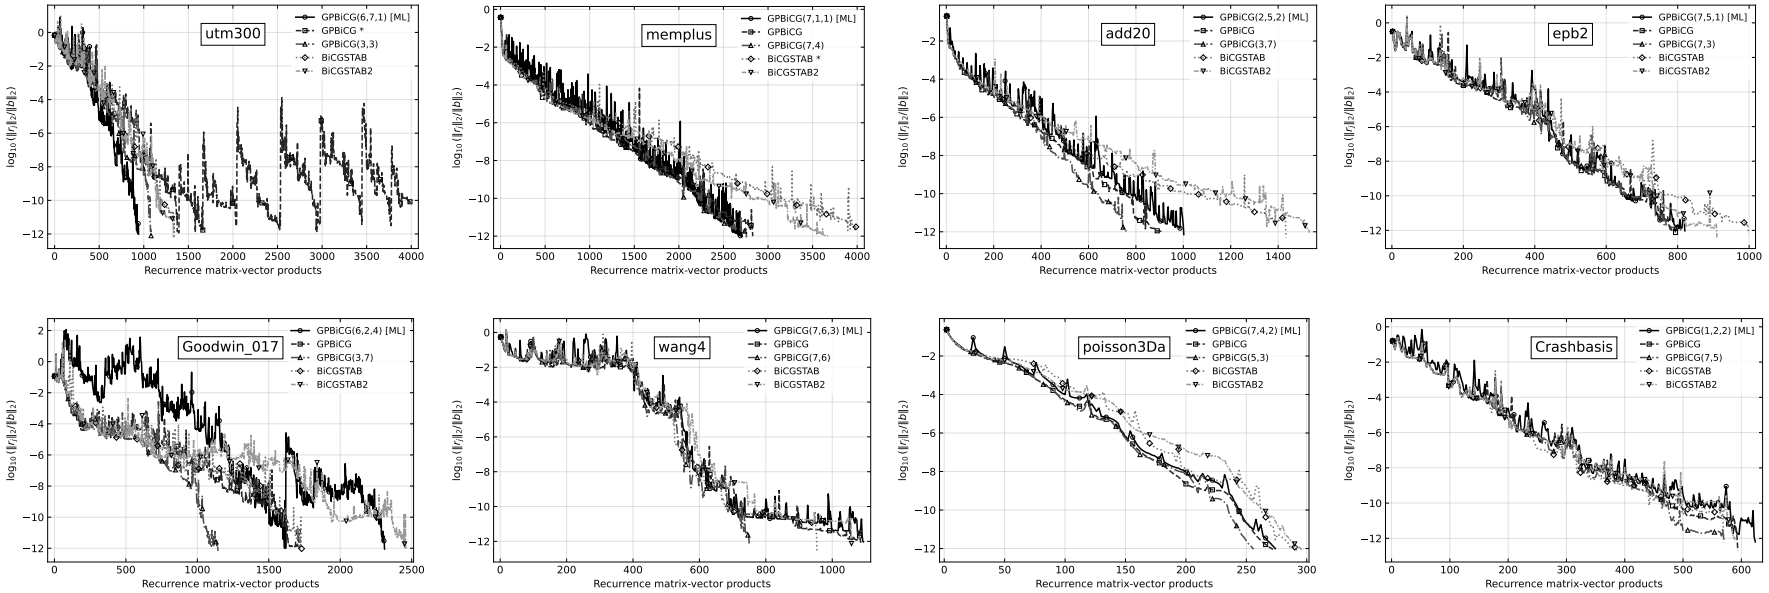

Same 43-feature learner across every case. The 11 descriptor matrix actions belong to recommendation cost, not the solver-trajectory axis. No case-wise minimum between models.

Complete 43-feature CPU decomposition

Matrix	Triplet	T_{sel} ms	T_{solve} s	Other ms	T_{total} s	Gain %
utm300	6,7,1	1.029	0.00406	0.060	0.00515	-16.9
add20	2,5,2	1.845	0.03054	0.155	0.03254	-22.4
Goodwin_017	6,2,4	6.499	0.19114	0.168	0.19781	-55.8
poisson3Da	7,4,2	11.363	0.08501	0.334	0.09670	-22.1
memplus	7,1,1	8.730	0.56266	0.234	0.57162	+1.1
epb2	7,5,1	9.881	0.23098	0.284	0.24114	-20.9
wang4	7,6,3	9.858	0.31484	0.289	0.32499	-65.4
crashbasis	1,2,2	31.485	1.24437	1.204	1.27706	-49.8

Selection/Other in milliseconds; solve/total in seconds. Other = independent physical check plus residual dispatch interval. Every request recomputes descriptors; no decision/solution cache. Positive gain favours ML.

Complete 26-feature CPU decomposition

Matrix	Triplet	T_{sel} ms	T_{solve} s	Other ms	T_{total} s	Gain %
utm300	2,6,5	0.831	0.00673	0.063	0.00762	-73.0
add20	7,6,7	1.629	0.03636	0.117	0.03811	-43.4
Goodwin_017	7,3,3	5.794	0.21672	0.169	0.22268	-75.4
poisson3Da	6,2,4	9.228	0.08773	0.365	0.09732	-22.9
memplus	7,6,7	7.617	0.66564	0.254	0.67351	-16.6
epb2	7,6,3	8.571	0.22482	0.268	0.23366	-17.1
wang4	1,5,2	8.488	0.32551	0.313	0.33431	-70.1
crashbasis	3,5,1	24.038	1.14650	1.221	1.17176	-37.4

Same matched session and units. All eight matrices retained, including losses. This model does not inherit the result of the 13- or 43-feature model.

ML criteria 1/9: scope, admissibility and integrity

Frozen experimental scope

Real, finite, nonempty square matrices; no preconditioner; $b = A\mathbf{1} \neq 0$, $x_0 = 0$; $\Theta = \{1, \ldots, 7\}^3$. The eight article cases are excluded from model fitting and selection.

- Canonical CSR: combine duplicates, remove explicit zeros, sort indices.
- Match raw-file and canonical sparse-content SHA-256 hashes.
- Require finite features, aligned identities, candidate channels and labels.
- TRAIN: 659 matrices / 95 groups; validation: 167 / 26 disjoint groups.
- No matrix names, hashes, family labels or solver outcomes as features.

Frozen enriched protocol and code. Groups are provenance proxies, not certified independent physical families or a nonsymmetric-only guarantee. Input hashes protect identity, not generalization.

ML criteria 2/9: all 13 base descriptors

Let $N = \mathrm{nnz}(A)$, r_i, c_i be row/column nonzero counts, $d_i = |A_{ii}|$, $s_i = \sum_j |A_{ij}|$. μ, σ are population mean/std.

No.	Descriptor
1	$\log(1+n)$
2	$\log(1+N)$
3	N/n^2
4–6	$\mu(r),\ \sigma(r),\ \max_i r_i$
7–8	$\sigma(c),\ \max_i c_i$
9	Fraction of rows with $d_i > 0$
10	Fraction of rows with $d_i \geq s_i - d_i$
11	$\log(1+\max_{ij} \lvert A_{ij}\rvert)$
12	Fraction of stored nonzeros that are negative
13	$\log(1+\lVert A\mathbf{1}\rVert_2)$

Exact base-feature implementation archived with this version. Logarithms here are natural logs. These descriptors are measurements, not convergence certificates. Zero rows contribute to descriptor 10 under the stated weak inequality.

ML criteria 3/9: all 13 structural additions

Set $C = A/\|A\|_\infty$, $s_i = \sum_j |C_{ij}|$, $h_j = \sum_i |C_{ij}|$ and $a_i = |C_{ii}|/s_i$ (zero if $s_i = 0$).

No.	Descriptor
14–18	Mean, std, 10th/50th/90th percentiles of a_i
19	Fraction of negative diagonal entries
20–21	$\sigma(s)/\mu(s)$ and $\sigma(h)/\mu(h)$
22–24	Mean, 90th percentile, maximum of sampled $\lvert i-j\rvert/\max(n-1,1)$
25	Mean of $\lvert(C\mathbf{1})_i\rvert/s_i$ (zero if $s_i = 0$)
26	$\log_{10}\|A\|_\infty$

- Use up to 100,000 evenly spaced integer CSR positions, including endpoints.
- The bandwidth maximum is sampled, not certified.
- Scale safely in two steps: largest entry, then resulting infinity norm.

Population std; zero denominators return zero. The original matrix is not changed. Descriptor 26 is evaluated as the sum of the two scaling logarithms to avoid overflow.

ML criteria 4/9: all 17 probe additions

Use $C = A/\|A\|_\infty$ and $v = (\varepsilon_1, \ldots, \varepsilon_n)^T/\sqrt{n}$, $\varepsilon_i \in \{-1, 1\}$, generated with NumPy default_rng seed 20261018. Write $a = Cv$, $t = C^T v$.

$$27: \frac{\|a - t\|_2}{\|a\|_2 + \|t\|_2}, \qquad 28: \frac{\|C^T a - Ct\|_2}{\|C^T a\|_2 + \|Ct\|_2},$$

$$29: \frac{v^T a}{\|a\|_2}, \qquad 30: \log_{10} \max(\|a\|_2, 10^{-300}).$$

Start $u_0 = C\mathbf{1}/\|C\mathbf{1}\|_2$. For $j = 1, \ldots, 6$, set $w_j = Cu_{j-1}$ and record the following 12 values (31–42):

$$\log_{10} \max(\|w_j\|_2, 10^{-300}), \qquad \frac{u_{j-1}^T w_j}{\|w_j\|_2}; \quad u_j = \frac{w_j}{\|w_j\|_2}.$$

Descriptor 43: first exact-zero action index divided by 6; use 1 if none.

- Zero denominators give zero; if $w_j = 0$, set $u_j = 0$ for subsequent actions.
- An exact zero first occurring at step 6 shares the encoding 1.
- Total: 11 matrix/transpose actions, not MLK iterations or free warm-up progress.

Directional indicators only: no certified symmetry, normality, eigenvalue or convergence estimate. Zero/nonfinite campaign RHS and nonfinite feature vectors are rejected.

ML criteria 5/9: labels, forest and triplet ranking

- A candidate succeeds only if **all three** recorded independent physical checks pass at 10^{-12}.
- Successful target: median native CPU divided by the best robust median for that matrix, clipped to $[1, 10]$.
- Failure target: 10 dimensionless; a matrix with no robust choice has 343 targets equal to 10.
- One input row per matrix; 343 outputs with fixed triplet identities.
- Same ExtraTrees: 192 trees, depth 12, leaf size 2, all features eligible, seed 20261018, one thread; squared-error criterion.
- Every fitting group gets equal total sample weight.
- Select the lowest predicted score. Exact ties use fitting-TRAIN ranking: grouped success first, then grouped relative cost, then channel index.

No hyperparameter search, imposed diversity, inference oracle or article-based tie-breaking. The per-matrix oracle is used only for labels/evaluation, not prediction.

ML criteria 6/9: the exact qualification inequalities

Let I_g be matrices in group g, $\mathcal{G}$ the groups, $N = \sum_g |I_g|$. Policy p chooses θ_i^p. Set $s_i^p = 1$ for three verified runs, otherwise 0; $z_i^p = y_{i,\theta_i^p}$ is its capped relative-cost label.

$$S_{\mathrm{mat}}(p) = \frac{1}{N} \sum_i s_i^p, \qquad S_{\mathrm{grp}}(p) = \frac{1}{|\mathcal{G}|} \sum_g \frac{1}{|I_g|} \sum_{i \in I_g} s_i^p,$$

$$R_{\mathrm{grp}}(p) = \frac{1}{|\mathcal{G}|} \sum_g \frac{1}{|I_g|} \sum_{i \in I_g} z_i^p.$$

Controls $\mathcal{C}$: fixed111 and the fitting-TRAIN global triplet. All three inequalities must hold:

$$S_{\mathrm{mat}}(p) \geq \max_{c \in \mathcal{C}} S_{\mathrm{mat}}(c),$$
$$S_{\mathrm{grp}}(p) \geq \max_{c \in \mathcal{C}} S_{\mathrm{grp}}(c),$$
$$R_{\mathrm{grp}}(p) \leq 0.99 \min_{c \in \mathcal{C}} R_{\mathrm{grp}}(c).$$

Apply to held-out predictions, not training-fit predictions. The chosen variant's outer-OOF gate AND the nested-policy gate must pass. This is a fixed empirical rule, not a confidence bound or a theorem.

ML criteria 7/9: nesting, controls and stopping rules

1. Four outer TRAIN group folds; three inner group folds within each fitting part.
2. Inner folds select 13/26/43 features or a fixed control, using only fitting information.
3. Outer held-out groups assess the entire selection procedure.
4. Among eligible variants: highest grouped success, then lowest grouped cost, then variant name.
5. If none qualifies, choose between fixed111 and TRAIN-global by grouped success, cost, then name.
6. Freeze the final decision before external validation. No refit, article tuning or post-hoc gate relaxation.

Exactly three predeclared feature sets. Validation feature-permutation is a seeded negative control, not a selectable model. Group-bootstrap intervals remain exploratory. Require a complete panel; STOP/deadline saves checkpoints and forbids partial-panel fitting.

ML criteria 8/9: serial verification and paid costs

The later serial study is exploratory: no retraining or policy promotion.

- Freeze all 24 recommendations before native results.
- Five serial repetitions per distinct configuration; one worker/thread; same qualified binary, no PC, $b = A\mathbf{1}$, $x_0 = 0$.
- Tolerance 10^{-12}, cap 2000, native timeout 30 s. Only 5/5 verified repetitions define a successful comparison time.
- Identical triplets share native observations; each policy pays its own features/inference.
- Warm CPU: input/RHS handling, features, prediction, native total, solution read/hash and independent physical verification.
- Cold CPU adds model deserialization. Never add nested native setup/check subtotals twice.

CPU and wall are separate; interpreter startup and orchestration are not fully charged OS wall time. Preserve failed-run wall costs and unknown CPU as missing. Positive generalization still needs untouched external cases.

ML criteria 9/9: do not mix the experimental contracts

Earlier HGB model

Log median process/read/RHS/verification wall target; three verified runs; otherwise 300 s penalty. Grouped cost gain $\geq 1\%$, no lower matrix or group success versus TRAIN-global. Not the enriched-forest target.

Short-observation study

32 MLK111 steps then a cold restart, at most 1968 further steps. Failure penalty 10; grouped deployment-wall proxy, not isolated CPU. Two launches can cost 60 s versus 30 s for direct runs.

Live adaptivity is a separate rule

Cycle/phase contraction, admissible move, distinct best/worst phase and two-productive-cycle lock; confidence/cost gate for random starts. Exact live formulas: A7–A9. Do not substitute ML qualification for these tests.

Prescribed-start live test: 10^{-10}; random-start and serial ML: 10^{-12}. No cross-protocol CPU speedup, hidden tuning on exposed cases or universal-solver claim.