Accelerating Predictive TRANSP

with Space-Time Profile Regularization and Staged Solver Controls

for the NSTX-U digital twin project

Surrogate development for fast temperature evolution

A. Y. Pankin, J. Breslau, M. Gorelenkova,
M. Scotto D'Abusco, M. Churchill

Princeton Plasma Physics Laboratory

APS meeting

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, as part of the Genesis Mission National Science and Technology Challenges under Award Number DE-AC02-09CH11466

Core message: Improved full-physics turnaround while developing a timestep-consistent surrogate for faster profile evolution.

Motivation and main message

- **Digital-twin goal:** evaluate actuator and scenario changes before the next discharge.
- **Predictive bottleneck:** stiff transport and imperfect input profiles increase nonlinear solver cost.
- **Filtering task:** regularize density and temperature inputs, then stage the transition into prediction.
- **Surrogate task:** learn electron- and ion-temperature rates from PT_SOLVER evolution.
- **Physics reference:** retain TRANSP for validation, physical constraints and full-solver fallback.

Core message: Improved full-physics turnaround while developing a timestep-consistent surrogate for faster profile evolution.

Why profile artifacts increase predictive cost

$$R/L_T = -R\,\partial_r \ln T$$

- **Starting point:** between-shot predictions use recent fitted density and temperature profiles.
- **Artifacts:** spikes, temporal jumps and conversion errors can distort local gradients.
- **Stiff response:** small profile perturbations can produce large changes in gradient-dependent fluxes.
- **Numerical cost:** backtracking and rejected steps repeat expensive transport-model evaluations.
- **Additional sensitivity:** weak gradients can make the diffusion/convection split poorly conditioned.
- **Objective:** reduce input-induced numerical stiffness while retaining physically relevant gradients.

Causal space-time profile regularization

$$x_k = \arg\min_{x \in \mathbb{R}^N} \left[\frac{1}{2}(x - y_k)^T W_k (x - y_k) + \frac{\lambda_t}{2} \|x - p_{k-1}^{\text{ref}}\|_2^2 + \frac{\lambda_\rho}{2} \|D_2 x\|_2^2 + \mathcal{S}_{\text{mon}}(x) \right]$$

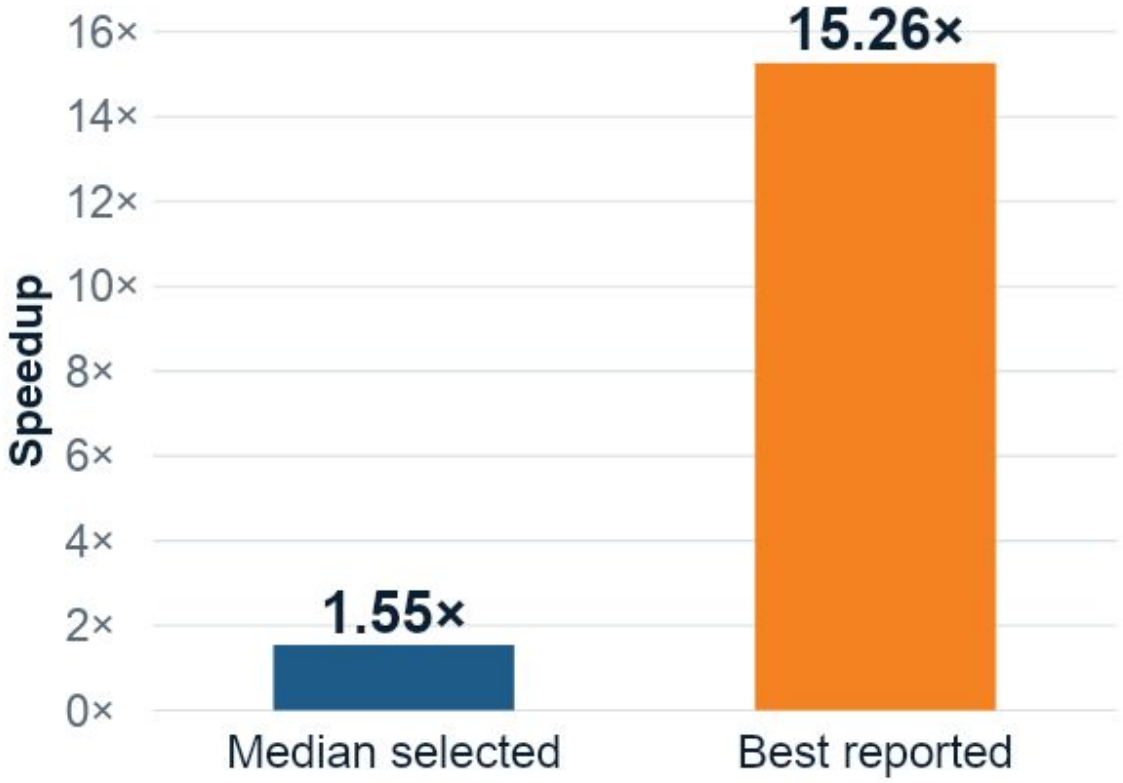
- **Causal input:** combine the current fit with the previous accepted profile, using no future frames.
- **Data fidelity:** Huber robust weights reduce outlier influence while retaining the measured profile.
- **Temporal anchor:** suppress artificial jumps between successive profile frames.
- **Radial curvature:** penalize second differences to suppress short-scale radial wiggles.
- **Soft shape control:** reduce spurious outward increases while retaining pedestal and off-axis structure.
- **Physical properties:** condition solver inputs with positivity and relevant gradients retained.

Staged controls and consistent convergence checks

Predictive phase	Solver action
Regularized initial state	Start from conditioned density and temperature profiles.
Transition to prediction	Stage time-step and convergence controls to reduce retries.
Established prediction	Restore standard solver controls after profile adjustment.

- **Residual consistency:** include stabilization, moving geometry, implicit coupling and source terms.
- **Verification:** coupled diffusion, manufactured pedestal profiles and stiffness scans test the advance.
- **Iteration budget:** the cap decreased from 100 to 30 in target workflows.
- **Retained physics:** the governing transport equations and physical closures continue to define the prediction.

Filtering and staged controls reduce predictive runtime



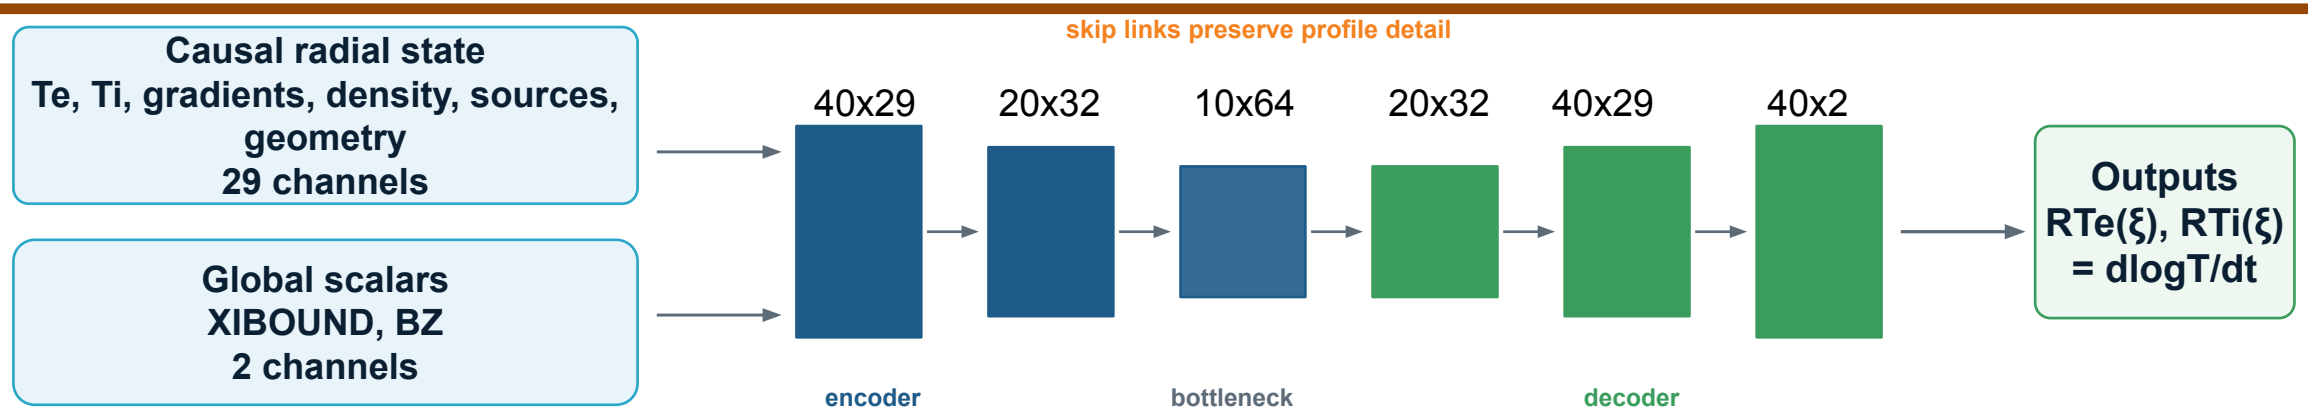
- **Combined method:** regularized profiles plus staged convergence.
- **Best reported result:** 15.26× speedup, exceeding an order of magnitude.
- **Median selected result:** 1.55× speedup.
- **Physical meaning:** conditioning reduces numerical difficulty without changing the transport equations.
- **Comparison requirement:** evaluate runtime over the same physical simulation endpoint.
- **Median combined temperature error:** 2.57%

Surrogate evolution uses logarithmic temperature rates

$$T_s^{n+1} = T_s^n \exp[R_{T_s}, (t_b - t_a)], \qquad s \in \{e, i\}$$

- **Task:** predict the electron- and ion-temperature evolution inside TRANSP.
- **Earlier network:** finite log-temperature increments could persist when TRANSP used very small timesteps.
- **Revised target:** learn logarithmic temperature rates from the PT_SOLVER profile advances.
- **Integration:** TRANSP applies the actual timestep to the predicted rates.
- **Small-step consistency:** the predicted increment vanishes as the timestep tends to zero.
- **Positive temperatures:** the exponential update retains positivity for positive inputs and finite rates.

Geometry-aware spatial and temporal architecture



U-Net stage	Grid × channels
Input	40 × 29
Encoder	20 × 32
Bottleneck	10 × 64
Decoder	20 × 32
Restored grid	40 × 29
Output	40 × 2

- **Causal radial state:** temperatures, gradients, density, sources and geometry.
- **Spatial backbone:** a 1D U-Net captures local and multiscale radial response.
- **Global conditioning:** FiLM (Feature-wise Linear Modulation) uses boundary location and toroidal field.
- **Skip connections:** retain fine profile structure through encoding and decoding.
- **Temporal branch:** a bounded causal TCN correction combines with the U-Net through a gate.
- **Outputs:** two radial rate fields feed the actual TRANSP timestep integration.

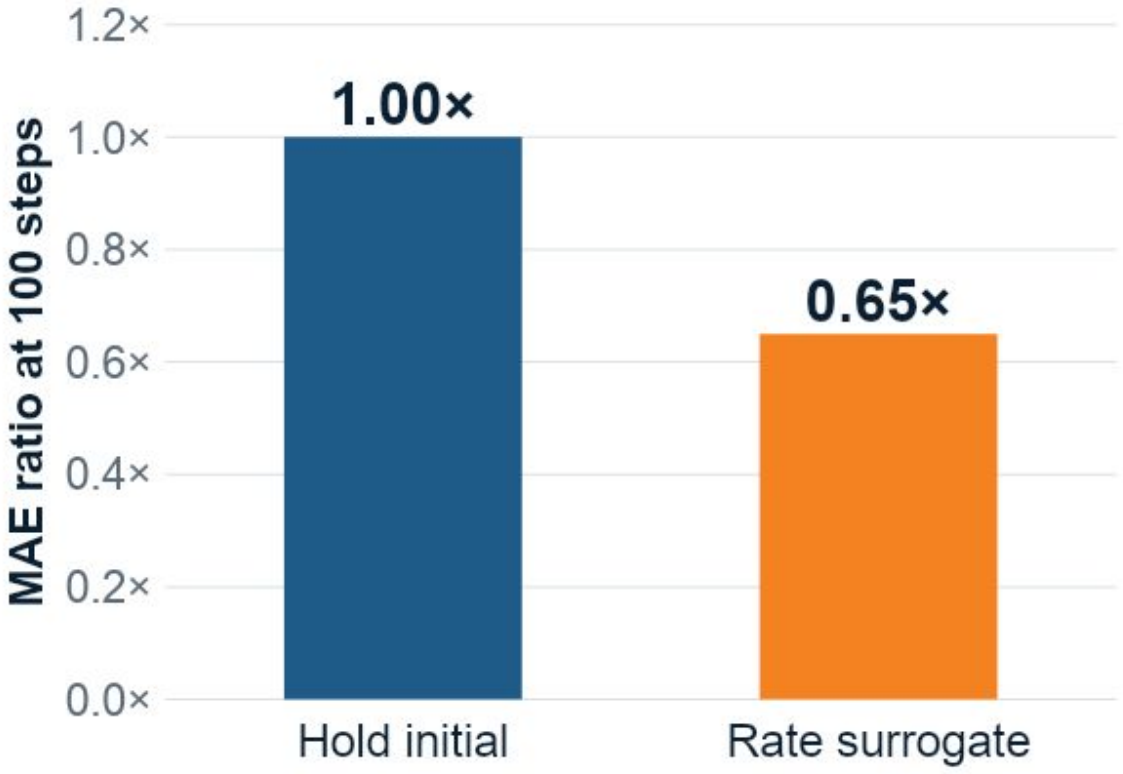
Physics-informed training constrains profile evolution

$$\mathcal{L} = \mathcal{L}_{R_T} + 0.50\mathcal{L}_T + 0.25\mathcal{L}_{\nabla T} + 0.25\mathcal{L}_{L_T^{-1}} + 0.10\mathcal{L}_E + 0.02\mathcal{L}_{\text{trust}} + 0.05\mathcal{L}_{\text{bnd}}$$

Loss term	Weight	Physical or numerical role
Log-temperature rate	1.00	Huber loss on the rate targets
Temperature profiles	0.50	Volume-weighted Te/Ti agreement
Spatial gradients	0.25	Radial temperature-gradient agreement
Inverse scale length	0.25	Geometry-aware temperature scale lengths
Energy balance	0.10	Q_{IE}-free thermal-energy balance
Trust penalty	0.02	Discourage large log-temperature updates
Boundary loss	0.05	Outermost active radial point

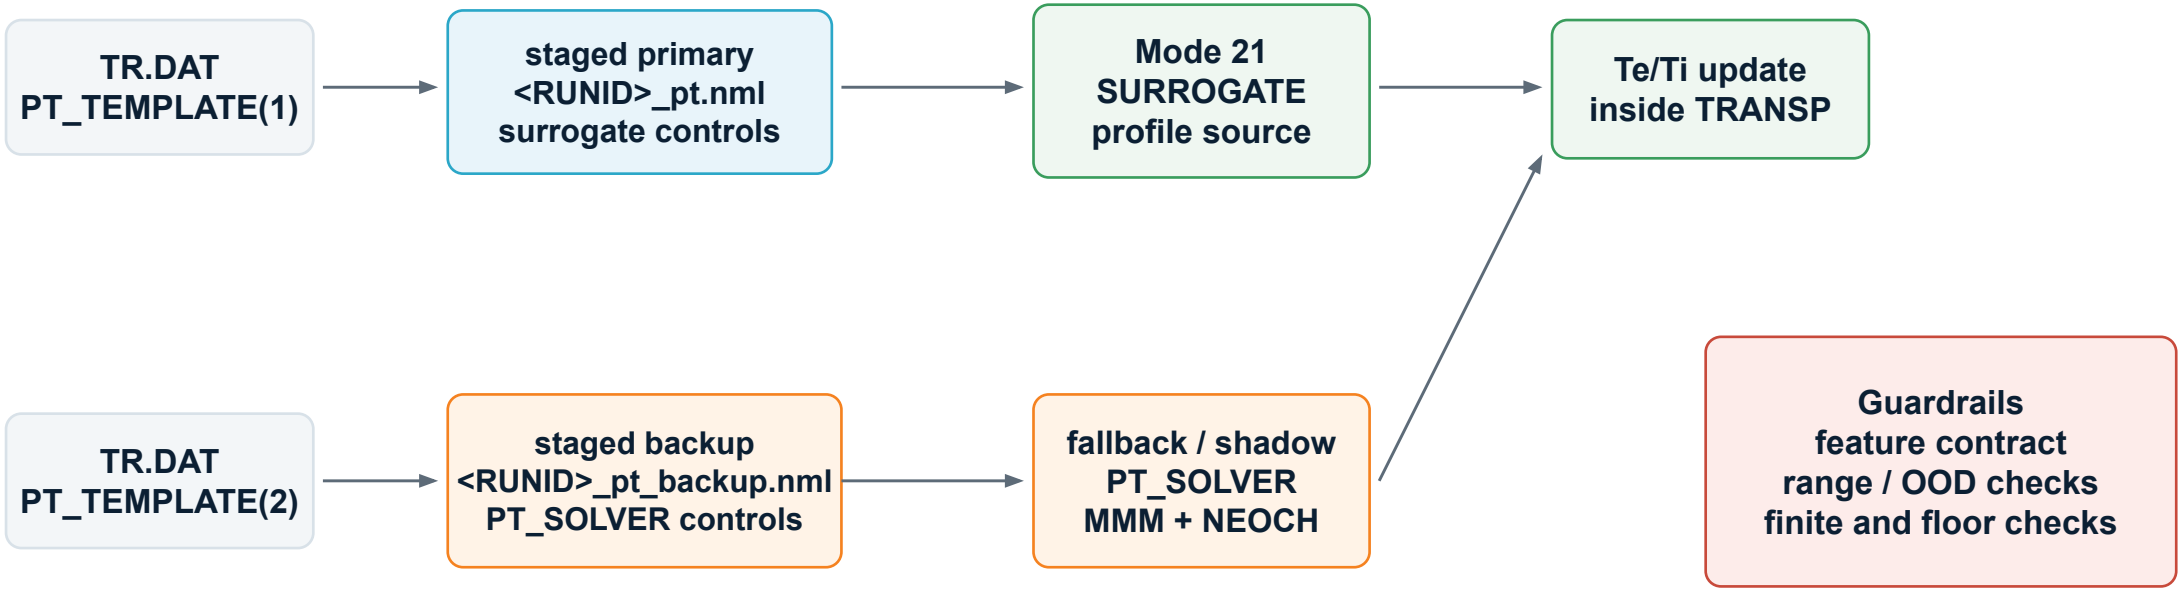
- **Physical-space losses:** use the actual timestep to reconstruct the temperature profiles.

Preliminary surrogate validation



- **Validation population:** 45,894 samples.
- **One-step accuracy:** 2.93 eV combined Te/Ti MAE and 0.426% relative profile error.
- **Rollouts:** about 435k forecasts to 100 steps, with zero reported failed rollouts.
- **Long-horizon comparison:** MAE is about 35% below hold-initial at 100 steps.
- **Remaining drift:** recursive forecasts still differ from teacher-forced/oracle references.
- **Status:** preliminary results require consistent invalid-tail cutoffs and revalidation.

TRANSP integration, fallback and validation steps



$$\text{Guard Metric } (S_{\text{guard}}) : \text{Dimensionless score measuring departure from training } \log T \text{ and } |\Delta \log T|$$
$$\text{Operational Impact} : \text{Intercepts } 100\% \text{ of hard failures}$$

- **Rate outputs:** dlog(T_e)/dt and dlog(T_i)/dt. TRANSP applies its actual timestep.
- **Auditable selection:** logs identify the active profile source and selected model.
- **Next steps:** revalidate cutoffs, export the model, test live TRANSP, then evaluate the locked test set.

Summary: filtering and surrogate development

- **Filtering:** causal input regularization and staged controls reduce predictive startup cost.
- **Demonstration:** 15.26× best speedup, 1.55× median selected speedup, and 27/35 variants below 5% combined T_e/T_i error.
- **Surrogate:** logarithmic-rate targets support timestep-consistent temperature evolution.
- **Physics-informed model:** geometry-aware U-Net, temporal correction and profile-space loss penalties.
- **Initial validation:** 2.93 eV one-step combined MAE and about 0.65× hold-initial MAE at 100 steps.
- **Integration and next steps:** guarded full-solver fallback, consistent cutoffs, live TRANSP tests and locked-test evaluation.

Together, profile regularization and physics-informed surrogate move TRANSP toward the goal of the NSTX-U digital twin on reliable predictions, delivered in time to guide the next discharge