Metriplex: A New Class of Physics Neural Networks with Exact Conservation Laws

Sehaj Randhir Singh
Department of Electrical and Computer Engineering,
NYU Tandon School of Engineering (independent researcher)
Brooklyn, NY, USA
sehajrsinghs@gmail.com

***Abstract*—We introduce the metriplectic layer, a neural layer class for physical AI whose invariants hold *exactly, per layer, to machine precision* – the discrete, learnable form of GENERIC dynamics. A balanced-weight flux channel conserves $\sum h$ for any nonlinearity; a Cayley-transformed positive-semidefinite metric contracts $\|h\|$ exactly, with a learnable gate the network must open only if its world is open; an optional Cayley rotation conserves norm for reversible subsystems. The same layer instantiates as an MLP, an RNN, and a GNN. Across eight benchmark domains vs. ResMLP, LSTM, NeuralODE, HNN, and plain-GNN baselines, trained as one-step maps and evaluated by autoregressive rollout: latent drift at the machine-precision floor (10^{-13} structural, 10^{-7} float32) for every metriplectic model – baselines have no conserved quantity by construction; the field GNN conserves the predicted mass exactly on advection–diffusion; and the recurrent instantiation stays exactly conserving over $2\times$ the training horizon. Exactness is not free and is reported honestly: on smooth small-state ODEs NeuralODE fits one-step transitions more tightly, the metriplectic models sit in the MLP-family band, and structure shows where it mismatches the physics – circulation (which conserves the wrong quantity for Kepler) measurably hurts the rollout there, while on the spring the circulation channel, which conserves the energy-like $\|h\|$, fits measurably better than flux, while on Kepler it conserves the wrong quantity and hurts. On the higher-dimensional conservative domains the structure pays: the metriplectic model is the tightest of the MLP-family baselines on a 6-D coupled-spring chain (HNN included), and on a chaotic double pendulum its error compounds at the lowest growth of the family—below NeuralODE and well below HNN.**

***Index Terms*—physical AI; structure-preserving neural networks; metriplectic systems; exact conservation laws; neural integrators**

I. Introduction

Physical AI – learned world models and simulators that roll out dynamics for long horizons – has a defining failure mode: networks fit one-step transitions well, then drift off the manifold of physically possible states as errors compound. Soft physics penalties make the violation smaller but never zero. Structure-preserving alternatives exist for the conservative case (Hamiltonian and symplectic networks), but the physical world is mostly open: friction, diffusion, and control dissipate. The metriplectic (GENERIC) formalism generalizes Hamiltonian dynamics to open systems, and neural metriplectic networks exist – but they learn continuous-time vector fields, whose invariants hold only in the limit of perfect integration. This paper's contribution is the missing discrete construction: a layer whose invariants are algebraic, holding for every parameter value, at every step, in the map that actually runs.

II. The Layer

Let $h \in \mathbb{R}^n$ be the latent state. One metriplectic step is

$$h' = h + \Delta t\, W\, \phi(h) - \Delta t\, \gamma\, (I - M/2)(I + M/2)^{-1} h, \qquad M = BB^{\mathsf T} \succeq 0$$

with W column-balanced ($\mathbb{1}^{\mathsf T} W = 0$) and γ learned through a sigmoid, initialized so the layer is born conservative.

- **Balanced flux:** $\mathbb{1}^{\mathsf T} h' = \mathbb{1}^{\mathsf T} h$ exactly, for *any* nonlinearity ϕ – the discrete-divergence form of a conservation law (the continuous analog is the divergence-free condition of Hamiltonian flows).
- **Metric friction:** $Q = (I - M/2)(I + M/2)^{-1}$ has spectrum in $(-1, 1]$, so $\|Qh\| \le \|h\|$ exactly – the second law as structure; the learned subspace B decides *where* energy is dissipated, the sign is not learned. Removing this channel is an exact ablation.
- **Optional rotation:** the Cayley transform of a skew matrix is orthogonal, conserving $\|h\|$ for reversible subsystems (off by default, since it conserves norm, not mass).

The class instantiates as an MLP (depth is integration time), an RNN (the layer is the cell, applied at every rollout step), and a GNN (flux message passing on a mesh, conserving the *predicted field*'s mass). The unconstrained baseline removes the balanced projection and friction channel while keeping every other parameter, so comparisons are controlled tests of structure.

III. Results

Protocol: one-step MSE training on transitions (128 trajectories, subsampled to 2048 per epoch, AdamW, 150 epochs, 5 seeds) and autoregressive rollout evaluation on 32 held-out trajectories; ground truth from independent solvers (closed forms, RK4, spectral). Conservation is measured structurally (float64 core rollouts) and physically (drift of each domain's invariants on the predicted trajectories; ground truth drifts by 10^{-5}–10^{-8}).

Three findings. (1) *Exactness survives deployment*: the metriplectic models' latent drift is at the machine-precision floor (Table I) and the GNN conserves the predicted field's mass exactly while the plain GNN leaks a substantial fraction;

TABLE I
CONSERVATION IN DEPLOYMENT: MAXIMAL PER-STEP STRUCTURAL DRIFT OF THE LAYER'S CONSERVED QUANTITY OVER THE FULL ROLLOUT (FLOAT64). "N/A": THE MODEL HAS NO CONSERVED QUANTITY BY CONSTRUCTION. FOR THE FIELD MODELS THE QUANTITY IS THE PREDICTED MASS ITSELF (PHYSICAL); FOR THE MECHANICAL MODELS IT IS LATENT $\sum h$ (STRUCTURAL).

Domain	Model	Structural drift
spring	Metriplectic (latent)	2.17e-14
	ResMLP / LSTM / NeuralODE / HNN	n/a
damped spring	Metriplectic (latent)	2.12e-14
	ResMLP / LSTM / NeuralODE	n/a
pendulum	Metriplectic (latent)	1.29e-14
	ResMLP / LSTM / NeuralODE / HNN	n/a
kepler	Metriplectic (latent)	2.90e-14
	ResMLP / LSTM / NeuralODE / HNN	n/a
coupled springs	Metriplectic (latent)	5.04e-14
	ResMLP / LSTM / NeuralODE / HNN	n/a
double pendulum	Metriplectic (latent)	4.77e-14
	ResMLP / LSTM / NeuralODE / HNN	n/a
advection-diff.	Metriplectic GNN (field mass)	2.22e-16 (structural), 0.0148 (deployment)
advection-diff.	Plain GNN	measured rollout mass drift 0.5782

TABLE II
FINAL ROLLOUT RMSE (NORMALIZED STATES, 5 SEEDS). BOLD = BEST.

Domain	Metriplectic	ResMLP	LSTM	NeuralODE	HNN
spring	0.5948	0.4862	0.9906	0.1752	0.9133
damped spring	0.2320	0.2281	1.0057	0.0695	—
pendulum	0.5046	0.3290	0.7514	0.1184	0.5222
damped pendulum	0.6554	0.5655	0.8593	0.1180	—
kepler	0.9574	0.7857	0.5706	0.0703	0.5079
coupled springs	1.0551	1.0810	1.3607	0.5403	1.6511
double pendulum	1.7095	1.7651	1.6172	0.2382	0.7425

baselines have no conserved quantity by construction. (2) *Structure shows in the shape of the error curve*:NeuralODE, a continuous-time map in the raw state space, is the tightest one-step fitter; the metriplectic model sits in the MLP-family band and on the spring its horizon growth ratio (3.290) is far below NeuralODE (7.690), LSTM (4.701), and HNN (10.014), and of the same order as ResMLP (2.770). The higher-dimensional conservative domains sharpen the payoff: on the coupled-spring chain (6-D, three modes) the metriplectic model is the tightest of the MLP-family baselines (1.0551 vs. 1.0810, 1.3607, and 1.6511), and on the chaotic double pendulum its error compounds at 3.872 — the lowest growth in the family, versus 5.764 and 8.273 for the continuous-time baselines that fit one-step transitions more tightly (0.2382, 0.7425). The structural evidence is the channel ablation: on the spring the two conservative channels are neutral (flux 0.5854 vs. circulation 0.4573), while on Kepler switching to the channel that conserves the wrong quantity costs 0.3962 $\to$ 1.3727. (3) *The second-law channel is born closed and measured as such*: the learned gate ends at $\sigma(\gamma) = 0.0029$ on the damped spring – essentially closed, because the flux channel suffices to fit the transitions; the guarantee is the sign of the channel, not its learned value. The metriplectic RNN conserves its latent norm over $2\times$ the training horizon (per-step drift 3.83e-16, final error 0.9486 vs. 0.4395 for the LSTM, which carries no conserved quantity).

IV. HONEST LIMITATIONS

The friction channel contracts the state norm – right for these benchmarks, but not the full GENERIC degeneracy. The conserved quantity is a *learned latent invariant*: the layer guarantees some quantity is conserved exactly, and the network must discover the physically meaningful one (the field GNN is the case where the invariant lives in the predicted state and is therefore physically exact). HNN is a serious competitor on the conservative domains where it is defined but is undefined on every open system. Extending exactness from latent structure to named physical invariants is the central open problem.

REFERENCES

[1] A. Hernández, S. Massaroli, K. Yuan, M. G. Códoba, J. Ott, and P. Stinis, "Metriplectic operator networks for learning dissipative dynamics," NeurIPS, 2023.

[2] A. Gruber, K. Lee, and N. Trask, "Reversible and irreversible bracket-based dynamics for deep graph learning," ICLR, 2023.

[3] M. Grmela and H. C. ́'Ottinger, "Dynamics and thermodynamics of complex fluids. I. Development of a general formalism," Phys. Rev. E, vol. 56, pp. 6620–6632, 1997.

[4] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," NeurIPS, 2019.

[5] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, "Neural ordinary differential equations," NeurIPS, 2018.

[6] B. Chang, M. Chen, E. Haber, and E. H. Chi, "AntisymmetricRNN: a dynamical system view on recurrent neural networks," ICLR, 2019.

[7] E. Hairer, C. Lubich, and G. Wanner, Geometric Numerical Integration, 2nd ed. Springer, 2006.