

EFFICIENTLY PARAMETERIZED NEURAL METRIPLECTIC SYSTEMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Metriplectic systems are learned from data in a way that scales quadratically in both the size of the state and the rank of the metriplectic data. Besides being provably energy conserving and entropy stable, the proposed approach comes with approximation results demonstrating its ability to accurately learn metriplectic dynamics from data as well as an error estimate indicating its potential for generalization to unseen timescales when approximation error is low. Examples are provided which illustrate performance in the presence of both full state information as well as when entropic variables are unknown, confirming that the proposed approach exhibits superior accuracy and scalability without compromising on model expressivity.

1 INTRODUCTION

The theory of metriplectic, also called GENERIC, systems Morrison (1986); Grmela and Öttinger (1997) provides a principled formalism for encoding dissipative dynamics in terms of complete thermodynamical systems that conserve energy and produce entropy. By formally expressing the reversible and irreversible parts of state evolution with separate algebraic brackets, the metriplectic formalism provides a general mechanism for maintaining essential conservation laws while simultaneously respecting dissipative effects. Thermodynamic completeness implies that any dissipation is caught within a metriplectic system through the generation of entropy, allowing for a holistic treatment which has already found use in modeling fluids Morrison (1984; 2009), plasmas Kaufman and Morrison (1982); Materassi (2012), and kinetic theories Kaufman (1984); Holm et al. (2008).

From a machine learning point of view, the formal separation of conservative and dissipative effects makes metriplectic systems highly appealing for the development of phenomenological models. Given data which is physics-constrained or exhibits some believed properties, a metriplectic system can be learned to exhibit the same properties with clearly identifiable conservative and dissipative parts. This allows for a more nuanced understanding of the governing dynamics via an evolution equation which reduces to an idealized Hamiltonian system as the dissipation is taken to zero. Moreover, elements in the kernel of the learned conservative part are immediately understood as Casimir invariants, which are special conservation laws inherent to the phase space of solutions, and are often useful for understanding and exerting control on low-dimensional structure in the system. On the other hand, the same benefit of metriplectic structure as a “direct sum” of reversible and irreversible parts makes it challenging to parameterize in an efficient way, since delicate degeneracy conditions must be enforced in the system for all time. In fact, there are no methods at present for learning general metriplectic systems which scale optimally with both dimension and the rank of metriplectic data—an issue which this work directly addresses.

Precisely, metriplectic dynamics on the finite or infinite dimensional phase space \mathcal{P} are generated by a free energy function(al) $F : \mathcal{P} \rightarrow \mathbb{R}$, $F = E + S$ defined in terms of a pair $E, S : \mathcal{P} \rightarrow \mathbb{R}$ representing energy and entropy, respectively, along with two algebraic brackets $\{\cdot, \cdot\}, [\cdot, \cdot] : C^\infty(\mathcal{P}) \times C^\infty(\mathcal{P}) \rightarrow C^\infty(\mathcal{P})$ which are bilinear derivations on $C^\infty(\mathcal{P})$ with prescribed symmetries and degeneracies $\{S, \cdot\} = [E, \cdot] = 0$. Here $\{\cdot, \cdot\}$ is an antisymmetric Poisson bracket, which is a Lie algebra realization on functions, and $[\cdot, \cdot]$ is a degenerate metric bracket which is symmetric and positive semi-definite. When $\mathcal{P} \subset \mathbb{R}^n$ for some $n > 0$, these brackets can be identified with symmetric matrix fields $\mathbf{L} : \mathcal{P} \rightarrow \text{Skew}_n(\mathbb{R})$, $\mathbf{M} : \mathcal{P} \rightarrow \text{Sym}_n(\mathbb{R})$ satisfying $\{F, G\} = \nabla F \cdot \mathbf{L} \nabla G$ and $[F, G] = \nabla F \cdot \mathbf{M} \nabla G$ for all functions $F, G \in C^\infty(\mathcal{P})$ and all states $\mathbf{x} \in \mathcal{P}$. Using the degeneracy conditions along with $\nabla \mathbf{x} = \mathbf{I}$ and abusing notation slightly then leads the standard equations for

metriplectic dynamics,

$$\dot{\mathbf{x}} = \{\mathbf{x}, F\} + [\mathbf{x}, F] = \{\mathbf{x}, E\} + [\mathbf{x}, S] = \mathbf{L}\nabla E + \mathbf{M}\nabla S,$$

which are provably energy conserving and entropy producing. To see why this is the case, recall that $\mathbf{L}^\top = -\mathbf{L}$. It follows that the infinitesimal change in energy satisfies

$$\dot{E} = \dot{\mathbf{x}} \cdot \nabla E = \mathbf{L}\nabla E \cdot \nabla E + \mathbf{M}\nabla S \cdot \nabla E = -\mathbf{L}\nabla E \cdot \nabla E + \nabla S \cdot \mathbf{M}\nabla E = 0,$$

and therefore energy is conserved along the trajectory of \mathbf{x} . Similarly, the fact that $\mathbf{M}^\top = \mathbf{M}$ is positive semi-definite implies that

$$\dot{S} = \dot{\mathbf{x}} \cdot \nabla S = \mathbf{L}\nabla E \cdot \nabla S + \mathbf{M}\nabla S \cdot \nabla S = -\nabla E \cdot \mathbf{L}\nabla S + \mathbf{M}\nabla S \cdot \nabla S = |\nabla S|_{\mathbf{M}}^2 \geq 0,$$

so that entropy is nondecreasing along \mathbf{x} as well. Geometrically, this means that the motion of a trajectory \mathbf{x} is everywhere tangent to the level sets of energy and transverse to those of entropy, reflecting the fact that metriplectic dynamics are a combination of noncanonical Hamiltonian ($\mathbf{M} = \mathbf{0}$) and generalized gradient ($\mathbf{L} = \mathbf{0}$) motions. Note that these considerations also imply the Lyapunov stability of metriplectic trajectories, as can be seen by taking E as a Lyapunov function. Importantly, this also implies that metriplectic trajectories which start in the (often compact) set $K = \{\mathbf{x} \mid E(\mathbf{x}) \leq E(\mathbf{x}_0)\}$ remain there for all time.

In phenomenological modeling, the entropy S is typically chosen from Casimirs of the Poisson bracket generated by \mathbf{L} , i.e. those quantities $C \in C^\infty(\mathcal{P})$ such that $\mathbf{L}\nabla C = \mathbf{0}$. On the other hand, the method which will be presented here, termed neural metriplectic systems (NMS), allows for all of the metriplectic data $\mathbf{L}, \mathbf{M}, E, S$ to be approximated simultaneously, removing the need for Casimir invariants to be known or assumed ahead of time. The only restriction inherent to NMS is that the metriplectic system being approximated is nondegenerate (c.f. Definition 3.1), a mild condition meaning that the gradients of energy and entropy cannot vanish at any point $\mathbf{x} \in \mathcal{P}$ in the phase space. It will be shown that NMS alleviates known issues with previous methods for metriplectic learning, leading to easier training, superior parametric efficiency, and better generalization performance.

Contributions. The proposed NMS method for learning metriplectic models offers the following advantages over previous state-of-the-art: (1) It approximates arbitrary nondegenerate metriplectic dynamics with optimal quadratic scaling in both the problem dimension n and the rank r of the irreversible dynamics. (2) It produces realistic, thermodynamically consistent entropic dynamics from unobserved entropy data. (3) It admits universal approximation and error accumulation results given in Proposition 3.7 and Theorem 3.9. (4) It yields exact energy conservation and entropy stability by construction, allowing for effective generalization to unseen timescales.

2 PREVIOUS AND RELATED WORK

Previous attempts to learn metriplectic systems from data separate into “hard” and “soft” constrained methods. Hard constrained methods enforce metriplectic structure by construction, so that the defining properties of metriplecticity cannot be violated. Conversely, methods with soft constraints relax some aspects of metriplectic structure in order to produce a wider model class which is easier to parameterize. While hard constraints are the only way to truly guarantee appropriate generalization in the learned surrogate, the hope of soft constrained methods is that the resulting model is “close enough” to metriplectic that it will exhibit some of the favorable characteristics of metriplectic systems, such as energy and entropy stability. Some properties of the methods compared in this work are summarized in Table 1.

Soft constrained methods. Attempts to learn metriplectic systems using soft constraints rely on relaxing the degeneracy conditions $\mathbf{L}\nabla S = \mathbf{M}\nabla E = \mathbf{0}$. This is the approach taken in Hernández et al. (2021b), termed SPNN, which learns an almost-metriplectic model parameterized with generic neural networks through a simple L^2 penalty term in the training loss, $\mathcal{L}_{\text{pen}} = |\mathbf{L}\nabla E|^2 + |\mathbf{M}\nabla S|^2$. This widens the space of allowable network parameterizations for the approximate operators \mathbf{L}, \mathbf{M} . While the resulting model violates the first and second laws of thermodynamics, the authors show that reasonable trajectories are still obtained, at least when applied within the range of timescales used for training. A similar approach is taken in Hernández et al. (2021a), which targets larger problems and develops an almost-metriplectic model reduction strategy based on the same core idea.

Hard constrained methods. Perhaps the first example of learning metriplectic systems from data was given in González et al. (2019) in the context of system identification. Here, training data is assumed to come from a finite element simulation, so that the discrete gradients of energy and entropy can be approximated as $\nabla E(\mathbf{x}) = \mathbf{A}\mathbf{x}$, $\nabla S(\mathbf{x}) = \mathbf{B}\mathbf{x}$. Assuming a fixed form for \mathbf{L} produces a constrained learning problem for the constant matrices \mathbf{M} , \mathbf{A} , \mathbf{B} which is solved to yield a provably metriplectic surrogate model. Similarly, the work Ruiz et al. (2021) learns \mathbf{M} , \mathbf{E} given \mathbf{L} , \mathbf{S} by considering a fixed block-wise decoupled form which trivializes the degeneracy conditions, i.e. $\mathbf{L} = [\star \mathbf{0}; \mathbf{0} \mathbf{0}]$ and $\mathbf{M} = [\mathbf{0} \mathbf{0}; \mathbf{0} \star]$. This line of thought is continued in Xu et al. (2022) and Xu et al. (2023), both of which learn metriplectic systems with neural network parameterizations by assuming this decoupled block structure. A somewhat broader class of metriplectic systems are considered in Gruber et al. (2023b) using tools from exterior calculus, with the goal of learning metriplectic dynamics on graph data. This leads to a structure-preserving network surrogate which scales linearly in the size of the graph domain, but also cannot express arbitrary metriplectic dynamics due to the specific choices of model form for \mathbf{L} , \mathbf{M} .

A particularly inspirational method for learning general metriplectic systems was given in Lee et al. (2021) and termed GNODE, building on parameterizations of metriplectic operators developed in Öttinger (2014). GNODE parameterizes learnable reversible and irreversible bracket generating matrices \mathbf{L} , \mathbf{M} in terms of state-independent tensors $\xi \in (\mathbb{R}^n)^{\otimes 3}$ and $\zeta \in (\mathbb{R}^n)^{\otimes 4}$: for $1 \leq \alpha, \beta, \gamma, \mu, \nu \leq n$, the authors choose $L_{\alpha\beta}(\mathbf{x}) = \sum_{\gamma} \xi_{\alpha\beta\gamma} \partial^{\gamma} S$ and $M_{\alpha\beta}(\mathbf{x}) = \sum_{\mu, \nu} \zeta_{\alpha\mu, \beta\nu} \partial^{\mu} E \partial^{\nu} E$, where $\partial^{\alpha} F = \partial F / \partial x_{\alpha}$, ξ is totally antisymmetric, and ζ is symmetric between the pairs (α, μ) and (β, ν) but antisymmetric within each of these pairs. The key idea here is to exchange the problem of enforcing degeneracy conditions $\mathbf{L}\nabla E = \mathbf{M}\nabla S = \mathbf{0}$ in matrix fields \mathbf{L} , \mathbf{M} with the problem of enforcing symmetry conditions in tensor fields ξ , ζ , which is comparatively easier but comes at the expense of underdetermining the problem. In GNODE, enforcement of these symmetries is handled by a generic learnable 3-tensor $\tilde{\xi} \in (\mathbb{R}^n)^{\otimes 3}$ along with learnable matrices $\mathbf{D} \in \text{Sym}_r(\mathbb{R})$ and $\mathbf{\Lambda}^s \in \text{Skew}_n(\mathbb{R})$ for $1 \leq s \leq r \leq n$, leading to the final parameterizations $\xi_{\alpha\beta\gamma} = \frac{1}{3!} (\tilde{\xi}_{\alpha\beta\gamma} - \tilde{\xi}_{\alpha\gamma\beta} + \tilde{\xi}_{\beta\gamma\alpha} - \tilde{\xi}_{\beta\alpha\gamma} + \tilde{\xi}_{\gamma\alpha\beta} - \tilde{\xi}_{\gamma\beta\alpha})$ and $\zeta_{\alpha\mu, \beta\nu} = \sum_{s,t} \Lambda_{\alpha\mu}^s D_{st} \Lambda_{\beta\nu}^t$. Along with learnable energy and entropy functions E , S parameterized by multi-layer perceptrons (MLPs), the data \mathbf{L} , \mathbf{M} learned by GNODE guarantees metriplectic structure in the surrogate model and leads to successful learning of metriplectic systems in some simple cases of interest. However, note that this is a highly redundant parameterization requiring $\binom{n}{3} + r\binom{n}{2} + \binom{r+1}{2} + 2$ learnable scalar functions, which exhibits $\mathcal{O}(n^3 + rn^2)$ scaling in the problem size because of the necessity to compute and store $\binom{n}{3}$ entries of ξ and $r\binom{n}{2}$ entries of $\mathbf{\Lambda}$. Additionally, the assumption of state-independence in the bracket generating tensors ξ , ζ is somewhat restrictive, limiting the class of problems to which GNODE can be applied.

A related approach to learning metriplectic dynamics with hard constraints was seen in Zhang et al. (2022), which proposed a series of GFINN architectures depending on how much of the information \mathbf{L} , \mathbf{M} , \mathbf{E} , \mathbf{S} is assumed to be known. In the case that \mathbf{L} , \mathbf{M} are known, the GFINN energy and entropy are parameterized with scalar-valued functions $f \circ \mathbf{P}_{\ker \mathbf{A}}$ where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ (E or S) is learnable and $\mathbf{P}_{\ker \mathbf{A}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is orthogonal projection onto the kernel of the (known) operator \mathbf{A} (\mathbf{L} or \mathbf{M}). It follows that the gradient $\nabla(f \circ \mathbf{P}_{\ker \mathbf{A}}) = \mathbf{P}_{\ker \mathbf{A}} \nabla f(\mathbf{P}_{\ker \mathbf{A}})$ lies in the kernel of \mathbf{A} , so that the degeneracy conditions are guaranteed at the expense of constraining the model class of potential energies/entropies. Alternatively, in the case that all of \mathbf{L} , \mathbf{M} , \mathbf{E} , \mathbf{S} are unknown, GFINNs use learnable scalar functions f for E , S parameterized by MLPs as well as two matrix fields \mathbf{Q}^E , $\mathbf{Q}^S \in \mathbb{R}^{r \times n}$ with rows given by $\mathbf{q}_s^f = (\mathbf{S}_s^f \nabla f)^{\top}$ for learnable skew-symmetric matrices \mathbf{S}_s^f , $1 \leq s \leq r$, $f = E, S$. Along with two triangular $(r \times r)$ matrix fields \mathbf{T}_L , \mathbf{T}_M , this yields the parameterizations $\mathbf{L}(\mathbf{x}) = \mathbf{Q}^S(\mathbf{x})^{\top} (\mathbf{T}_L(\mathbf{x})^{\top} - \mathbf{T}_L(\mathbf{x})) \mathbf{Q}^S(\mathbf{x})$ and $\mathbf{M}(\mathbf{x}) = \mathbf{Q}^E(\mathbf{x})^{\top} (\mathbf{T}_M(\mathbf{x})^{\top} \mathbf{T}_M(\mathbf{x})) \mathbf{Q}^E(\mathbf{x})$, which necessarily satisfy the symmetries and degeneracy conditions required for metriplectic structure. GFINNs are shown to both increase expressivity over the GNODE method as well as decrease redundancy, since the need for an explicit order-3 tensor field is removed and the reversible and irreversible brackets are allowed to depend explicitly on the state \mathbf{x} . However, GFINNs still exhibit cubic scaling through the requirement of $rn(n-1) + r^2 + 2 = \mathcal{O}(rn^2)$ learnable functions, which is well above the theoretical minimum required to express a general metriplectic system and leads to difficulties in training the resulting models.

Model reduction. Finally, it is worth mentioning the closely related line of work involving model reduction for metriplectic systems, which began with Öttinger (2015). As remarked there, preserving metriplecticity in reduced-order models (ROMs) exhibits many challenges due to its delicate requirements on the kernels of the involved operators. There are also hard and soft constrained approaches: the already mentioned Hernández et al. (2021a) aims to learn a close-to-metriplectic data-driven ROM by enforcing degeneracies by penalty, while Gruber et al. (2023a) directly enforces metriplectic structure in projection-based ROMs using exterior algebraic factorizations. The parameterizations of metriplectic data presented here are related to those presented in Gruber et al. (2023a), although NMS does not require access to nonzero components of $\nabla E, \nabla S$.

3 FORMULATION AND ANALYSIS

The proposed formulation of the metriplectic bracket-generating operators \mathbf{L}, \mathbf{M} used by NMS is based on the idea of exploiting structure in the tensor fields ξ, ζ to reduce the necessary number of degrees of freedom. In particular, it will be shown that the degeneracy conditions $\mathbf{L}\nabla S = \mathbf{M}\nabla E = \mathbf{0}$ imply more than just symmetry constraints on these fields, and that taking these additional constraints into account allows for a more compact representation of metriplectic data. Following this, results are presented which show that the proposed formulation is universally approximating on nondegenerate systems (c.f. Definition 3.1) and admits a generalization error bound in time.

3.1 EXTERIOR ALGEBRA

Developing these metriplectic expressions will require some basic facts from exterior algebra, of which more details can be found in, e.g., Tu (2017, Chapter 19). The basic objects in the exterior algebra $\bigwedge V$ over the vector space V are multivectors, which are formal linear combinations of totally antisymmetric tensors on V . More precisely, if $I(V)$ denotes the two-sided ideal of the free tensor algebra $T(V)$ generated by elements of the form $v \otimes v$ ($v \in V$), then the exterior algebra is the quotient space $\bigwedge V \simeq T(V)/I(V)$ equipped with the antisymmetric wedge product operation \wedge . This graded algebra is equipped with natural projection operators $P^k : \bigwedge V \rightarrow \bigwedge^k V$ which map between the full exterior algebra and the k^{th} exterior power of V , denoted $\bigwedge^k V$, whose elements are homogeneous k -vectors. More generally, given an n -manifold M with tangent bundle TM , the exterior algebra $\bigwedge(TM)$ is the algebra of multivector fields whose fiber over $x \in M$ is given by $\bigwedge T_x M$.

For the present purposes, it will be useful to develop a correspondence between bivectors $B \in \bigwedge^2(\mathbb{R}^n)$ and skew-symmetric matrices $\mathbf{B} \in \text{Skew}_n(\mathbb{R})$, which follows directly from Riesz representation in terms of the usual Euclidean dot product. More precisely, supposing that e_1, \dots, e_n are the standard basis vectors for \mathbb{R}^n , any bivector $B \in \bigwedge^2 T\mathbb{R}^n$ can be represented as $B = \sum_{i < j} B^{ij} e_i \wedge e_j$ where $B^{ij} \in \mathbb{R}$ denote the components of B . Define a grade-lowering action of bivectors on vectors through right contraction (see e.g. Section 3.4 of Dorst et al. (2007)), expressed for any vector v and basis bivector $e_i \wedge e_j$ as $(e_i \wedge e_j) \cdot v = (e_j \cdot v)e_i - (e_i \cdot v)e_j$. It follows that the action of B is equivalent to

$$B \cdot v = \sum_{i < j} B^{ij} ((e_j \cdot v)e_i - (e_i \cdot v)e_j) = \sum_{i < j} B^{ij} v_j e_i - \sum_{j < i} B^{ji} v_j e_i = \sum_{i,j} B^{ij} v_j e_i = \mathbf{B}v,$$

where $\mathbf{B}^\top = -\mathbf{B} \in \mathbb{R}^{n \times n}$ is a skew-symmetric matrix representing B , and we have re-indexed under the second sum and applied that $B^{ij} = -B^{ji}$ for all i, j . Since the kernel of this action is the zero bivector, it is straightforward to check that this string of equalities defines an isomorphism $\mathcal{M} : \bigwedge^2 \mathbb{R}^n \rightarrow \text{Skew}_n(\mathbb{R})$ from the 2^{nd} exterior power of \mathbb{R}^n to the space of skew-symmetric $(n \times n)$ -matrices over \mathbb{R} : in what follows, we will write $\mathbf{B} \simeq B$ rather than $\mathbf{B} = \mathcal{M}(B)$ for notational convenience. Note that a correspondence in the more general case of bivector/matrix fields follows in the usual way via the fiber-wise extension of \mathcal{M} .

3.2 LEARNABLE METRIPECTIC OPERATORS

It is now possible to explain the proposed NMS formulation. First, note the following key definition which prevents the consideration of unphysical examples.

Definition 3.1. A metriplectic system on $K \subset \mathbb{R}^n$ generated by the data $\mathbf{L}, \mathbf{M}, E, S$ will be called *nondegenerate* provided $\nabla E, \nabla S \neq \mathbf{0}$ for all $\mathbf{x} \in K$.

With this, the NMS parameterizations for metriplectic operators are predicated on an algebraic result proven in Appendix A.

Lemma 3.2. Let $K \subset \mathbb{R}^n$. For all $\mathbf{x} \in K$, the operator $\mathbf{L} : K \rightarrow \mathbb{R}^{n \times n}$ satisfies $\mathbf{L}^\top = -\mathbf{L}$ and $\mathbf{L}\nabla S = \mathbf{0}$ for some $S : K \rightarrow \mathbb{R}, \nabla S \neq \mathbf{0}$, provided there exists a non-unique bivector field $\mathbf{A} : U \rightarrow \bigwedge^2 \mathbb{R}^n$ and equivalent matrix field $\mathbf{A} \simeq \mathbf{A}$ such that

$$\mathbf{L} \simeq \left(\mathbf{A} \wedge \frac{\nabla S}{|\nabla S|^2} \right) \cdot \nabla S = \mathbf{A} - \frac{1}{|\nabla S|^2} \mathbf{A} \nabla S \wedge \nabla S.$$

Similarly, for all $\mathbf{x} \in K$ a positive semi-definite operator $\mathbf{M} : K \rightarrow \mathbb{R}^{n \times n}$ satisfies $\mathbf{M}^\top = \mathbf{M}$ and $\mathbf{M}\nabla E = \mathbf{0}$ for some $E : K \rightarrow \mathbb{R}, \nabla E \neq \mathbf{0}$, provided there exists a non-unique matrix-valued $\mathbf{B} : K \rightarrow \mathbb{R}^{n \times r}$ and symmetric matrix-valued $\mathbf{D} : K \rightarrow \mathbb{R}^{r \times r}$ such that $r \leq n$ and

$$\begin{aligned} \mathbf{M} &= \sum_{s,t} D_{st} \left(\mathbf{b}^s \wedge \frac{\nabla E}{|\nabla E|^2} \right) \cdot \nabla E \otimes \left(\mathbf{b}^t \wedge \frac{\nabla E}{|\nabla E|^2} \right) \cdot \nabla E \\ &= \sum_{s,t} D_{st} \left(\mathbf{b}^s - \frac{\mathbf{b}^s \cdot \nabla E}{|\nabla E|^2} \nabla E \right) \left(\mathbf{b}^t - \frac{\mathbf{b}^t \cdot \nabla E}{|\nabla E|^2} \nabla E \right)^\top, \end{aligned}$$

where \mathbf{b}^s denotes the s^{th} column of \mathbf{B} . Moreover, using $\mathbf{P}_f^\perp = \left(\mathbf{I} - \frac{\nabla f \nabla f^\top}{|\nabla f|^2} \right)$ to denote the orthogonal projector onto $\text{Span}(\nabla f)^\perp$, these parameterizations of \mathbf{L}, \mathbf{M} are equivalent to the matricized expressions $\mathbf{L} = \mathbf{P}_S^\perp \mathbf{A} \mathbf{P}_S^\perp$ and $\mathbf{M} = \mathbf{P}_E^\perp \mathbf{B} \mathbf{D} \mathbf{B}^\top \mathbf{P}_E^\perp$.

Remark 3.3. Observe that the projections appearing in these expressions are the minimum necessary for guaranteeing the symmetries and degeneracy conditions necessary for a general metriplectic structure. In particular, conjugation by \mathbf{P}_f^\perp respects symmetry and ensures that both the input and output to the conjugated matrix field lie in $\text{Span}(\nabla f)^\perp$. Note that this implies that the class of nondegenerate metriplectic systems cannot be parameterized in a way that scales sub-quadratically in the state dimension n .

Lemma 3.2 demonstrates specific parameterizations for \mathbf{L}, \mathbf{M} that hold for any nondegenerate metriplectic data and are core to the NMS method for learning metriplectic dynamics. While generally underdetermined, these expressions are in a sense maximally specific given no additional information, since there is nothing available in the general metriplectic formalism to determine the matrix fields $\mathbf{A}, \mathbf{B} \mathbf{D} \mathbf{B}^\top$ on $\text{Span}(\nabla S), \text{Span}(\nabla E)$, respectively. The following Theorem, also proven in Appendix A, provides a rigorous correspondence between metriplectic systems and these particular parameterizations.

Theorem 3.4. The data $\mathbf{L}, \mathbf{M}, E, S$ form a nondegenerate metriplectic system in the state variable $\mathbf{x} \in K$ if and only if there exist a skew-symmetric $\mathbf{A} : K \rightarrow \text{Skew}_n(\mathbb{R})$, symmetric positive semidefinite $\mathbf{D} : K \rightarrow \text{Sym}_r(\mathbb{R})$, and generic $\mathbf{B} : K \rightarrow \mathbb{R}^{n \times r}$ such that

$$\dot{\mathbf{x}} = \mathbf{L}\nabla E + \mathbf{M}\nabla S = \mathbf{P}_S^\perp \mathbf{A} \mathbf{P}_S^\perp \nabla E + \mathbf{P}_E^\perp \mathbf{B} \mathbf{D} \mathbf{B}^\top \mathbf{P}_E^\perp \nabla S.$$

Remark 3.5. As mentioned, the proposed parameterizations for \mathbf{L}, \mathbf{M} are not one-to-one but properly contain the set of valid nondegenerate metriplectic systems in E, S . This is due to both the orthogonal projections employed and the lack of a guaranteed Jacobi identity for \mathbf{L} . Recall that the Jacobi identity is necessary for a true Poisson manifold structure and hence a truly metriplectic system. For $1 \leq i, j, k \leq n$, the Jacobi identity is given in components as $\sum_\ell L_{i\ell} \partial^\ell L_{jk} + L_{j\ell} \partial^\ell L_{ki} + L_{k\ell} \partial^\ell L_{ij} = 0$. However, this requirement is not often enforced in algorithms for learning general metriplectic (or even symplectic) systems, since it is considered subordinate to energy conservation and it is well known that both qualities cannot hold simultaneously in general Zhong and Marsden (1988).

3.3 SPECIFIC PARAMETERIZATIONS

Now that Theorem 3.4 has provided a model class which is rich enough to express any desired metriplectic system, it remains to discuss what NMS actually learns. First, note that it is unlikely to

be the case that E, S are known *a priori*, so it is beneficial to allow these functions to be learnable alongside the governing operators L, M . For simplicity, energy and entropy E, S are parameterized as scalar-valued MLPs with \tanh activation, although any desired architecture could be chosen for this task. The skew-symmetric matrix field $A = A_{\text{tri}} - A_{\text{tri}}^\top$ used to build L is parameterized through its strictly lower-triangular part A_{tri} using a vector-valued MLP with output dimension $\binom{n}{2}$, which guarantees that the mapping $A_{\text{tri}} \mapsto A$ above is bijective. Similarly, the symmetric matrix field $D = K_{\text{chol}} K_{\text{chol}}^\top$ is parameterized through its lower-triangular Cholesky factor K_{chol} , which is a vector-valued MLP with output dimension $\binom{r+1}{2}$. While this choice does not yield a bijective mapping $K_{\text{chol}} \mapsto D$ unless, e.g., D is assumed to be positive definite with diagonal entries of fixed sign, this does not hinder the method in practice. In fact, D should not be positive definite in general, as this is equivalent to claiming that M is positive definite on vectors tangent to the level sets of E . Finally, the generic matrix field B is parameterized as a vector-valued MLP with output dimension nr . Remarkably, the exterior algebraic expressions in Lemma 3.2 require less redundant operations than the corresponding matricized expressions in Theorem 3.4, and therefore the expressions from Lemma 3.2 are used when implementing NMS. [This can be seen from the proof of Lemma 3.2 in Appendix A, where matricization of these expressions requires computing a term which is identically zero due to symmetry considerations.](#) Figure 1 summarizes the proposed NMS architecture.

Remark 3.6. With these choices, the NMS parameterization of metriplectic systems requires only $(1/2)((n+r)^2 - (n-r)) + 2$ learnable scalar functions, in contrast to $\binom{n}{3} + r\binom{n}{2} + \binom{r+1}{2} + 2$ for the GNODE approach in Lee et al. (2021) and $rn(n-1) + r^2 + 2$ for the GFINN approach in Zhang et al. (2022). In particular, NMS is quadratic in both n, r with no decrease in model expressivity, in contrast to the cubic scaling of previous methods.

Table 1: Properties of the metriplectic architectures compared.

Name	Physics Bias	Restrictive	Scale
NODE	None	No	Linear
SPNN	Soft	No	Quadratic
GNODE	Hard	Yes	Cubic
GFINN	Hard	No	Cubic
NMS	Hard	No	Quadratic

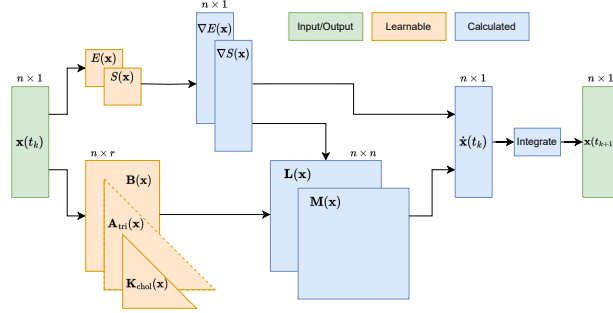


Figure 1: A visual depiction of the NMS architecture.

3.4 APPROXIMATION AND ERROR

Besides offering a compact parameterization of metriplectic dynamics, the expressions used in NMS also exhibit desirable approximation properties which guarantee a reasonable bound on state error over time. To state this precisely, first note the following universal approximation result proven in Appendix A.

Proposition 3.7. *Let $K \subset \mathbb{R}^n$ be compact and $E, S : K \rightarrow \mathbb{R}$ be continuous such that $L\nabla S = M\nabla E = \mathbf{0}$ and $\nabla E, \nabla S \neq \mathbf{0}$ for all $\mathbf{x} \in K$. For any $\varepsilon > 0$, there exist two-layer neural network functions $\tilde{E}, \tilde{S} : K \rightarrow \mathbb{R}$, $\tilde{L} : K \rightarrow \text{Skew}_n(\mathbb{R})$ and $\tilde{M} : K \rightarrow \text{Sym}_n(\mathbb{R})$ such that $\nabla \tilde{E}, \nabla \tilde{S} \neq \mathbf{0}$ on K , \tilde{M} is positive semi-definite, $\tilde{L}\nabla \tilde{S} = \tilde{M}\nabla \tilde{E} = \mathbf{0}$ for all $\mathbf{x} \in K$, and each approximate function is ε -close to its given counterpart on K . Moreover, if L, M have $k \geq 0$ continuous derivatives on K then so do \tilde{L}, \tilde{M} .*

Remark 3.8. The assumption $\mathbf{x} \in K$ of the state remaining in a compact set V is not restrictive when at least one of $E, -S : \mathbb{R}^n \rightarrow \mathbb{R}$, say E , has bounded sublevel sets. In this case, letting $\mathbf{x}_0 = \mathbf{x}(0)$ it follows from $\dot{E} \leq 0$ that $E(\mathbf{x}(t)) = E(\mathbf{x}_0) + \int_0^t \dot{E}(\mathbf{x}(\tau)) d\tau \leq E(\mathbf{x}_0)$, so that the entire trajectory $\mathbf{x}(t)$ lies in the (closed and bounded) compact set $K = \{\mathbf{x} \mid E(\mathbf{x}) \leq E(\mathbf{x}_0)\} \subset \mathbb{R}^n$.

Leaning on Proposition 3.7 and classical universal approximation results in Li (1996), it is further possible to establish the following error estimate also proven in Appendix A which gives an idea of the error accumulation rate that can be expected from this method.

Theorem 3.9. Suppose $\mathbf{L}, \mathbf{M}, E, S$ are nondegenerate metriplectic data such that \mathbf{L}, \mathbf{M} have at least one continuous derivative, E, S have Lipschitz continuous gradients, and at least one of $E, -S$ have bounded sublevel sets. For any $\varepsilon > 0$, there exist nondegenerate metriplectic data $\tilde{\mathbf{L}}, \tilde{\mathbf{M}}, \tilde{E}, \tilde{S}$ defined by two-layer neural networks such that, for all $T > 0$,

$$\left(\int_0^T |\mathbf{x} - \tilde{\mathbf{x}}|^2 dt \right)^{\frac{1}{2}} \leq \varepsilon \left| \frac{b}{a} \right| \sqrt{e^{2aT} - 2e^{aT} + T + 1},$$

where $a, b \in \mathbb{R}$ are constants depending on both sets of metriplectic data and $\dot{\tilde{\mathbf{x}}} = \tilde{\mathbf{L}}(\tilde{\mathbf{x}})\nabla\tilde{E}(\tilde{\mathbf{x}}) + \tilde{\mathbf{M}}(\tilde{\mathbf{x}})\nabla\tilde{S}(\tilde{\mathbf{x}})$.

Remark 3.10. Theorem 3.9 provides a bound on state error over time under the assumption that the approximation error in the metriplectic networks can be controlled. On the other hand, notice that Theorem 3.9 can also be understood as a generic error bound on any trained metriplectic networks $\tilde{\mathbf{L}}, \tilde{\mathbf{M}}, \tilde{E}, \tilde{S}$ provided universal approximation results are not invoked in the estimation leading to εb .

This result confirms that the error in the state \mathbf{x} for a fixed final time T tends to zero with the approximation error in the networks $\tilde{\mathbf{L}}, \tilde{\mathbf{M}}, \tilde{E}, \tilde{S}$, as one would hope based on the approximation capabilities of neural networks. More importantly, Theorem 3.9 also bounds the generalization error of any trained metriplectic network for an appropriate (and possibly large) ε equal to the maximum approximation error on K , where the learned metriplectic trajectories are confined for all time. With this theoretical guidance, the remaining goal of this work is to demonstrate that NMS is also practically effective at learning metriplectic systems from data and exhibits reasonable generalization to unseen timescales.

4 ALGORITHM

Similar to previous approaches in Lee et al. (2021) and Zhang et al. (2022), the learnable parameters in NMS are calibrated using data along solution trajectories to a given dynamical system. This brings up an important question regarding how much information about the system in question is realistically present in the training data. While many systems have a known metriplectic form, it is not always the case that one will know metriplectic governing equations for a given set of training data. Therefore, two approaches are considered in the experiments below corresponding to whether full or partial state information is assumed available during NMS training. More precisely, the state $\mathbf{x} = (\mathbf{x}^o, \mathbf{x}^u)$ will be partitioned into “observable” and “unobservable” variables, where \mathbf{x}^u may be empty in the case that full state information is available. In a partially observable system \mathbf{x}^o typically contains positions and momenta while \mathbf{x}^u contains entropy or other configuration variables which are more difficult to observe during physical experiments. In both cases, NMS will learn a metriplectic system in \mathbf{x} according to the procedure described in Algorithm 1.

Note that the batch-wise training strategy in Algorithm 1 requires initial conditions for \mathbf{x}^u in the partially observed case. There are several options for this, and two specific strategies will be considered here. Suppose the data are drawn from the training interval $[0, T]$ with initial state \mathbf{x}_0 and final state \mathbf{x}_T . The first strategy sets $\mathbf{x}_0^u = \mathbf{0}$, $\mathbf{x}_T^u = \mathbf{1}$ (where $\mathbf{1}$ is the all ones vector), and $\mathbf{x}_s^u = \mathbf{1}/s$, $0 < s < T$, so that the unobserved states are initially assumed to lie on a straight line. **In the absence of other information, this ensures that the entropy will be satisfy the monotonicity constraint necessary for metriplectic structure.** The second strategy is more sophisticated, and involves training a diffusion model to predict the distribution of \mathbf{x}^u given \mathbf{x}^o . Specific details of this procedure are given in Appendix H. **Note that other strategies are also possible, such as allowing the unobserved initial states to be trainable as in Chen et al. (2020).**

5 EXAMPLES

The goal of the following experiments is to show that NMS is effective even when entropic information cannot be observed during training, yielding superior performance when compared to previous methods including GNODE, GFINN, and SPNN discussed in Section 2. The metrics considered for this purpose will be mean absolute error (MAE) and mean squared error (MSE) defined in the usual way as the average ℓ^1 (resp. squared ℓ^2) error between the discrete states $\mathbf{x}, \tilde{\mathbf{x}} \in \mathbb{R}^{n \times n_s}$. For

Algorithm 1 Training neural metriplectic systems

```

1: Input: snapshot data  $\mathbf{X} \in \mathbb{R}^{n \times n_s}$ , each column  $\mathbf{x}_s = \mathbf{x}(t_s, \boldsymbol{\mu}_s)$ , target rank  $r \geq 1$ , batch size
    $n_b \geq 1$ .
2: Initialize networks  $\mathbf{A}_{\text{tri}}, \mathbf{B}, \mathbf{K}_{\text{chol}}, E, S$ , and loss  $L = 0$ 
3: for step in  $N_{\text{steps}}$  do
4:   Randomly draw batch  $P = \{(t_{s_k}, \mathbf{x}_{s_k})\}_{k=1}^{n_b}$ 
5:   for  $(t, \mathbf{x})$  in  $P$  do
6:     Evaluate  $\mathbf{A}_{\text{tri}}(\mathbf{x}), \mathbf{B}(\mathbf{x}), \mathbf{K}_{\text{chol}}(\mathbf{x}), E(\mathbf{x}), S(\mathbf{x})$ 
7:     Automatically differentiate  $E, S$  to obtain  $\nabla E(\mathbf{x}), \nabla S(\mathbf{x})$ 
8:     Form  $\mathbf{A}(\mathbf{x}) = \mathbf{A}_{\text{tri}}(\mathbf{x}) - \mathbf{A}_{\text{tri}}(\mathbf{x})^\top$  and  $\mathbf{D}(\mathbf{x}) = \mathbf{K}_{\text{chol}}(\mathbf{x})\mathbf{K}_{\text{chol}}(\mathbf{x})^\top$ 
9:     Build  $\mathbf{L}(\mathbf{x}), \mathbf{M}(\mathbf{x})$  according to Lemma 3.2
10:    Evaluate  $\dot{\mathbf{x}} = \mathbf{L}(\mathbf{x})\nabla E(\mathbf{x}) + \mathbf{M}(\mathbf{x})\nabla S(\mathbf{x})$ 
11:    Randomly draw  $n_1, \dots, n_l$  and form  $t_j = t + n_j \Delta t$  for all  $j$ 
12:     $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_l = \text{ODEsolve}(\dot{\mathbf{x}}, t_1, \dots, t_l)$ 
13:     $L += l^{-1} \sum_j \text{Loss}(\mathbf{x}_j, \tilde{\mathbf{x}}_j)$ 
14:   end for
15:   Rescale  $L = |P|^{-1} L$ 
16:   Update  $\mathbf{A}_{\text{tri}}, \mathbf{B}, \mathbf{K}_{\text{chol}}, E, S$  through gradient descent on  $L$ .
17: end for

```

brevity, many implementation details have been omitted here and can be found in Appendix B. An additional experiment showing the effectiveness of NMS in the presence of both full and partial state information can be found in Appendix D.

Remark 5.1. To facilitate a more equal parameter count between the compared metriplectic methods, the results of the experiments below were generated using the alternative parameterization $\mathbf{D} = \mathbf{K}\mathbf{K}^\top$ where $\mathbf{K} : K \rightarrow \mathbb{R}^{r \times r'}$ is full and $r' \geq r$. Of course, this change does not affect metriplecticity since \mathbf{D} is still positive semi-definite for each $\mathbf{x} \in K$.

5.1 TWO GAS CONTAINERS

The first test of NMS involves two ideal gas containers separated by movable wall which is removed at time t_0 , allowing for the exchange of heat and volume. In this example, the motion of the state $\mathbf{x} = (q \ p \ S_1 \ S_2)^\top$ is governed by the metriplectic equations:

$$\begin{aligned} \dot{q} &= \frac{p}{m}, & \dot{p} &= \frac{2}{3} \left(\frac{E_1(\mathbf{x})}{q} - \frac{E_2(\mathbf{x})}{2L - q} \right), \\ \dot{S}_1 &= \frac{9N^2 k_B^2 \alpha}{4E_1(\mathbf{x})} \left(\frac{1}{E_1(\mathbf{x})} - \frac{1}{E_2(\mathbf{x})} \right), & \dot{S}_2 &= -\frac{9N^2 k_B^2 \alpha}{4E_1(\mathbf{x})} \left(\frac{1}{E_1(\mathbf{x})} - \frac{1}{E_2(\mathbf{x})} \right), \end{aligned}$$

where (q, p) are the position and momentum of the separating wall, S_1, S_2 are the entropies of the two subsystems, and the internal energies E_1, E_2 are determined from the Sackur-Tetrode equation for ideal gases, $S_i/Nk_B = \ln(\hat{c}V_i E_i^{3/2})$, $1 \leq i \leq 2$. Here, m denotes the mass of the wall, $2L$ is the total length of the system, and V_i is the volume of the i^{th} container. As in Lee et al. (2021; 2022) $Nk_B = 1$ and $\alpha = 0.5$ fix the characteristic macroscopic unit of entropy while $\hat{c} = 102.25$ ensures the argument of the logarithm defining E_i is dimensionless. This leads to the total entropy $S(\mathbf{x}) = S_1 + S_2$ and the total energy $E(\mathbf{x}) = (1/2m)p^2 + E_1(\mathbf{x}) + E_2(\mathbf{x})$, which are guaranteed to be nondecreasing and constant, respectively.

The primary goal here is to verify that NMS can accurately and stably predict gas container dynamics without the need to observe the entropic variables S_1, S_2 . To that end, NMS has been compared to GNODE, SPNN, and GFINN on the task of predicting the trajectories of this metriplectic system over time, with results displayed in Table 2. More precisely, given an initial condition \mathbf{x}_0 and an interval $0 < t_{\text{train}} < t_{\text{valid}} < t_{\text{test}}$, each method is trained on partial state information (in the case of NMS) or full state information (in the case of the others) from the interval $[0, t_{\text{train}}]$ and validated on $(t_{\text{train}}, t_{\text{valid}}]$ before state errors in q, p only are calculated on the whole interval $[0, t_{\text{test}}]$. As can be seen from Table 2 and Figure 2, NMS is remarkably accurate over unseen timescales even in this unfair comparison, avoiding the unphysical behavior which often hinders soft-constrained methods

Table 2: Prediction errors for x^o measured in MSE and MAE on the interval $[0, t_{\text{test}}]$ in the two gas containers example (left) and on the test set in the thermoelastic double pendulum example (right).

	NODE	SPNN	GNODE	GFINN	NMS		NODE	SPNN	GNODE	GFINN	NMS
MSE	.12 \pm .04	.13 \pm .10	.16 \pm .10	.07 \pm .03	.01 \pm .02	MSE	.41 \pm .01	.42 \pm .01	.42 \pm .01	.40 \pm .03	.38 \pm .03
MAE	.25 \pm .10	.26 \pm .14	.25 \pm .13	.13 \pm .03	.08 \pm .06	MAE	.48 \pm .04	.47 \pm .03	.46 \pm .04	.43 \pm .07	.42 \pm .07

like SPNN. The energy and instantaneous entropy plots in Figure 2 further confirm that the strong enforcement of metriplectic structure guaranteed by NMS leads to correct energetic and entropic dynamics for all time.

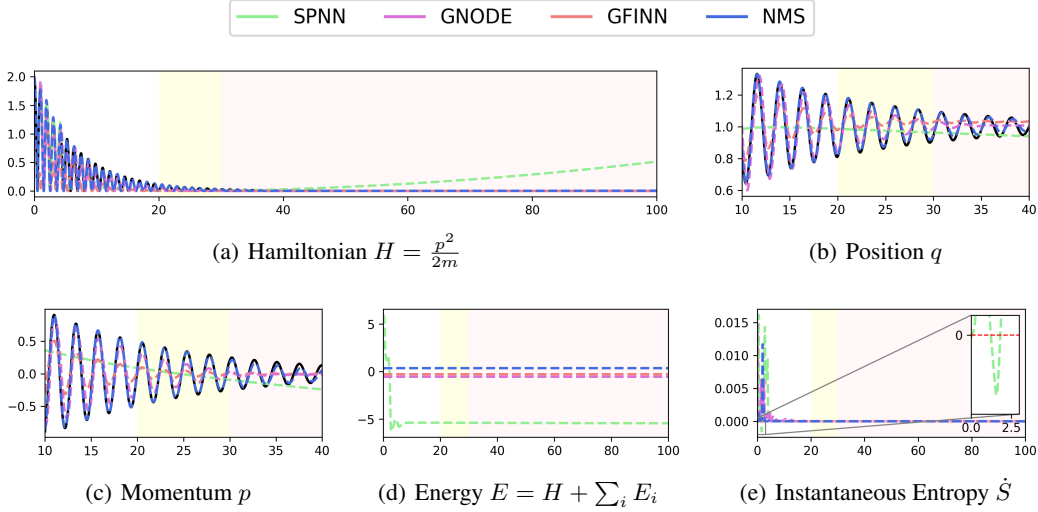


Figure 2: The ground-truth and predicted position, momentum, instantaneous entropy, and energies for the two gas containers example in the training (white), validation (yellow), and testing (red) regimes.

5.2 THERMOELASTIC DOUBLE PENDULUM

Next, consider the thermoelastic double pendulum from Romero (2009) with 10-dimensional state variable $\mathbf{x} = (q_1 \ q_2 \ p_1 \ p_2 \ S_1 \ S_2)^\top$, which represents a highly challenging benchmark for metriplectic methods. The equations of motion in this case are given for $1 \leq i \leq 2$ as

$$\dot{\mathbf{q}}_i = \frac{\mathbf{p}_i}{m_i}, \quad \dot{\mathbf{p}}_i = -\partial_{\mathbf{q}_i}(E_1(\mathbf{x}) + E_2(\mathbf{x})), \quad \dot{S}_1 = \kappa(T_1^{-1}T_2 - 1), \quad \dot{S}_2 = \kappa(T_1T_2^{-1} - 1),$$

where $\kappa > 0$ is a thermal conductivity constant (set to 1), m_i is the mass of the i^{th} spring (also set to 1) and $T_i = \partial_{S_i} E_i$ is its absolute temperature. In this case, $\mathbf{q}_i, \mathbf{p}_i \in \mathbb{R}^2$ represent the position and momentum of the i^{th} mass, while S_i represents the entropy of the i^{th} pendulum. As before, the total entropy $S(\mathbf{x}) = S_1 + S_2$ is the sum of the entropies of the two springs, while defining the internal energies $E_i(\mathbf{x}) = (1/2)(\ln \lambda_i)^2 + \ln \lambda_i + e^{S_i - \ln \lambda_i} - 1$, $\lambda_1 = |\mathbf{q}_i|$, $\lambda_2 = |\mathbf{q}_2 - \mathbf{q}_1|$, leads to the total energy $E(\mathbf{x}) = (1/2m_1)|\mathbf{p}_1|^2 + (1/2m_2)|\mathbf{p}_2|^2 + E_1(\mathbf{x}) + E_2(\mathbf{x})$.

The task in this case is prediction across initial conditions. As in Zhang et al. (2022), 100 trajectories are drawn from the ranges in Appendix B and integrated over the interval $[0, 40]$ with $\Delta t = 0.1$, with an 80/10/10 split for training/validation/testing. Here all compared models are trained using full state information. As seen in Table 2, NMS is again the most performant, although all models struggle to approximate the dynamics over the entire training interval. It is also notable that the training time of NMS is greatly decreased relative to GNODE and GFINN due to its improved quadratic scaling; a representative study to this effect is given in Appendix G.

6 CONCLUSION

Neural metriplectic systems (NMS) have been considered for learning finite-dimensional metriplectic dynamics from data. Making use of novel non-redundant parameterizations for metriplectic operators, NMS provably approximates arbitrary nondegenerate metriplectic systems with generalization error bounded in terms of the operator approximation quality. Benchmark examples have shown that NMS is both more scalable and more accurate than previous methods, including when only partial state information is observed. Future work will consider extensions of NMS to infinite-dimensional metriplectic systems with the aim of addressing its main limitation: the difficulty of scaling NMS (among all present methods for metriplectic learning) to realistic, 3-D problems of the size that would be considered in practice. A promising direction is to consider the use of NMS in model reduction, where sparse, large-scale systems are converted to small, dense systems through a clever choice of encoding/decoding.

REFERENCES

- Z. Chen, J. Zhang, M. Arjovsky, and L. Bottou. Symplectic recurrent neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BkgYPREtPr>.
- J. R. Dormand and P. J. Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- L. Dorst, D. Fontijne, and S. Mann. *Geometric Algebra for Computer Science: An Object-oriented Approach to Geometry*. Morgan Kaufmann, Amsterdam, 2007.
- D. González, F. Chinesta, and E. Cueto. Thermodynamically consistent data-driven computational mechanics. *Continuum Mechanics and Thermodynamics*, 31(1):239–253, 2019.
- M. Grmela and H. C. Öttinger. Dynamics and thermodynamics of complex fluids. i. development of a general formalism. *Phys. Rev. E*, 56:6620–6632, Dec 1997. doi: 10.1103/PhysRevE.56.6620. URL <https://link.aps.org/doi/10.1103/PhysRevE.56.6620>.
- A. Gruber, M. Gunzburger, L. Ju, and Z. Wang. Energetically consistent model reduction for metriplectic systems. *Computer Methods in Applied Mechanics and Engineering*, 404:115709, 2023a. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2022.115709>. URL <https://www.sciencedirect.com/science/article/pii/S0045782522006648>.
- A. Gruber, K. Lee, and N. Trask. Reversible and irreversible bracket-based dynamics for deep graph neural networks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b.
- Q. Hernández, A. Badías, D. González, F. Chinesta, and E. Cueto. Deep learning of thermodynamics-aware reduced-order models from data. *Computer Methods in Applied Mechanics and Engineering*, 379:113763, 2021a.
- Q. Hernández, A. Badías, D. González, F. Chinesta, and E. Cueto. Structure-preserving neural networks. *Journal of Computational Physics*, 426:109950, 2021b.
- D. D. Holm, V. Putkaradze, and C. Tronci. Kinetic models of oriented self-assembly. *Journal of Physics A: Mathematical and Theoretical*, 41(34):344010, aug 2008. doi: 10.1088/1751-8113/41/34/344010. URL <https://dx.doi.org/10.1088/1751-8113/41/34/344010>.
- A. N. Kaufman. Dissipative hamiltonian systems: A unifying principle. *Physics Letters A*, 100(8):419–422, 1984. ISSN 0375-9601. doi: [https://doi.org/10.1016/0375-9601\(84\)90634-0](https://doi.org/10.1016/0375-9601(84)90634-0). URL <https://www.sciencedirect.com/science/article/pii/0375960184906340>.
- A. N. Kaufman and P. J. Morrison. Algebraic structure of the plasma quasilinear equations. *Physics Letters A*, 88(8):405–406, 1982. ISSN 0375-9601. doi: [https://doi.org/10.1016/0375-9601\(82\)90664-8](https://doi.org/10.1016/0375-9601(82)90664-8). URL <https://www.sciencedirect.com/science/article/pii/0375960182906648>.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- K. Lee, N. Trask, and P. Stinis. Machine learning structure preserving brackets for forecasting irreversible processes. *Advances in Neural Information Processing Systems*, 34:5696–5707, 2021.
- K. Lee, N. Trask, and P. Stinis. Structure-preserving sparse identification of nonlinear dynamics for data-driven modeling. In *Mathematical and Scientific Machine Learning*, pages 65–80. PMLR, 2022.

- X. Li. Simultaneous approximations of multivariate functions and their derivatives by neural networks with one hidden layer. *Neurocomputing*, 12(4):327–343, 1996.
- H. Lim, M. Kim, S. Park, and N. Park. Regular time-series generation using sgm. *arXiv preprint arXiv:2301.08518*, 2023.
- E. Materassi, M.; Tassi. Metriplectic framework for dissipative magneto-hydrodynamics. *Physica D: Nonlinear Phenomena*, 2012.
- P. Morrison. Thoughts on brackets and dissipation: old and new. In *Journal of Physics: Conference Series*, volume 169, page 012006. IOP Publishing, 2009.
- P. J. Morrison. Some observations regarding brackets and dissipation. Technical Report PAM-228, Center for Pure and Applied Mathematics, University of California, Berkeley, 1984.
- P. J. Morrison. A paradigm for joined hamiltonian and dissipative systems. *Physica D: Nonlinear Phenomena*, 18(1):410–419, 1986. ISSN 0167-2789. doi: [https://doi.org/10.1016/0167-2789\(86\)90209-5](https://doi.org/10.1016/0167-2789(86)90209-5). URL <https://www.sciencedirect.com/science/article/pii/0167278986902095>.
- H. C. Öttinger. Irreversible dynamics, onsager-casimir symmetry, and an application to turbulence. *Phys. Rev. E*, 90:042121, Oct 2014.
- H. C. Öttinger. Preservation of thermodynamic structure in model reduction. *Phys. Rev. E*, 91:032147, Mar 2015.
- I. Romero. Thermodynamically consistent time-stepping algorithms for non-linear thermomechanical systems. *International Journal for Numerical Methods in Engineering*, 79(6):706–732, 2023/05/14 2009.
- D. Ruiz, D. Portillo, and I. Romero. A data-driven method for dissipative thermomechanics. *IFAC-PapersOnLine*, 54(19):315–320, 2021.
- X. Shang and H. C. Öttinger. Structure-preserving integrators for dissipative systems based on reversible–irreversible splitting. *Proceedings of the Royal Society A*, 476(2234):20190446, 2020.
- Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. *CoRR*, abs/2011.13456, 2020. URL <https://arxiv.org/abs/2011.13456>.
- S. Särkkä and A. Solin, editors. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.
- L. W. Tu. *Differential Geometry: Connections, Curvature, and Characteristic Classes*. Springer International Publishing, 2017.
- P. Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- B. Xu, Y. Chen, T. Matsubara, and T. Yaguchi. Learning generic systems using neural symplectic forms. In *International Symposium on Nonlinear Theory and Its Applications*, number A2L-D-03 in IEICE Proceeding Series, pages 29–32. The Institute of Electronics, Information, and Communication Engineers (IEICE), 2022.
- B. Xu, Y. Chen, T. Matsubara, and T. Yaguchi. Equivalence class learning for GENERIC systems. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023.
- Z. Zhang, Y. Shin, and G. Em Karniadakis. Gfinns: Generic formalism informed neural networks for deterministic and stochastic dynamical systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 380(2229):20210207, 2022.
- G. Zhong and J. E. Marsden. Lie-poisson hamilton-jacobi theory and lie-poisson integrators. *Physics Letters A*, 133(3):134–139, 1988.

A PROOF OF THEORETICAL RESULTS

This Appendix provides proof of the analytical results in Section 3 of the body. First, the parameterizations of L, M in terms of exterior algebra are established.

Proof of Lemma 3.2. First, it is necessary to check that the operators L, M parameterized this way satisfy the symmetries and degeneracy conditions claimed in the statement. To that end, recall that $\mathbf{a} \wedge \mathbf{b} \simeq \mathbf{a}\mathbf{b}^\top - \mathbf{b}\mathbf{a}^\top$, meaning that $(\mathbf{a}\mathbf{b}^\top - \mathbf{b}\mathbf{a}^\top)^\top \simeq \mathbf{b} \wedge \mathbf{a} = -\mathbf{a} \wedge \mathbf{b}$. It follows that $\mathbf{A}^\top \simeq \tilde{\mathbf{A}} = -\mathbf{A}$ where $\tilde{\mathbf{A}}$ denotes the reversion of \mathbf{A} , i.e., $\tilde{\mathbf{A}} = \sum_{i < j} A^{ij} \mathbf{e}_j \wedge \mathbf{e}_i$. Therefore, we may write

$$\mathbf{L}^\top \simeq \tilde{\mathbf{A}} - \frac{1}{|\nabla S|^2} \widetilde{\mathbf{A} \nabla S \wedge \nabla S} = -\mathbf{A} + \frac{1}{|\nabla S|^2} \mathbf{A} \nabla S \wedge \nabla S \simeq -\mathbf{L},$$

showing that $\mathbf{L}^\top = -\mathbf{L}$. Moreover, using that

$$(\mathbf{b} \wedge \mathbf{c}) \cdot \mathbf{a} = -\mathbf{a} \cdot (\mathbf{b} \wedge \mathbf{c}) = (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c},$$

it follows that

$$\mathbf{L} \nabla S = \mathbf{A} \cdot \nabla S - \frac{1}{|\nabla S|^2} (\mathbf{A} \nabla S \wedge \nabla S) \cdot \nabla S = \mathbf{A} \nabla S - \mathbf{A} \nabla S = \mathbf{0},$$

since $\nabla S \cdot \mathbf{A} \nabla S = -\nabla S \cdot \mathbf{A} \nabla S = 0$. Moving to the case of M , notice that $M = D_{st} \mathbf{v}^s \otimes \mathbf{v}^t$ for a particular choice of \mathbf{v} , meaning that

$$\mathbf{M}^\top = \sum_{s,t} D_{st} (\mathbf{v}^s \otimes \mathbf{v}^t)^\top = \sum_{s,t} D_{st} \mathbf{v}^t \otimes \mathbf{v}^s = \sum_{t,s} D_{ts} \mathbf{v}^s \otimes \mathbf{v}^t = \sum_{s,t} D_{st} \mathbf{v}^s \otimes \mathbf{v}^t = \mathbf{M},$$

since D is a symmetric matrix. Additionally, it is straightforward to check that, for any $1 \leq s \leq r$,

$$\mathbf{v}^s \cdot \nabla E = \left(\mathbf{b}^s - \frac{\mathbf{b}^s \cdot \nabla E}{|\nabla E|^2} \nabla E \right) \cdot \nabla E = \mathbf{b}^s \cdot \nabla E - \mathbf{b}^s \cdot \nabla E = 0.$$

So, it follows immediately that

$$\mathbf{M} \nabla E = \sum_{s,t} D_{st} (\mathbf{v}^s \otimes \mathbf{v}^t) \cdot \nabla E = \sum_{s,t} D_{st} (\mathbf{v}^t \cdot \nabla E) \mathbf{v}^s = \mathbf{0}.$$

Now, observe that

$$\begin{aligned} \mathbf{L} &= \mathbf{A} - \frac{1}{|\nabla S|^2} (\mathbf{A} \nabla S (\nabla S)^\top - \nabla S (\mathbf{A} \nabla S)^\top) \\ &= \mathbf{A} - \frac{1}{|\nabla S|^2} (\mathbf{A} \nabla S (\nabla S)^\top + \nabla S (\nabla S)^\top \mathbf{A}) \\ &= \left(\mathbf{I} - \frac{\nabla S (\nabla S)^\top}{|\nabla S|^2} \right) \mathbf{A} \left(\mathbf{I} - \frac{\nabla S (\nabla S)^\top}{|\nabla S|^2} \right) = \mathbf{P}_S^\perp \mathbf{A} \mathbf{P}_S^\perp, \end{aligned}$$

since $\mathbf{A}^\top = -\mathbf{A}$ and hence $\mathbf{v}^\top \mathbf{A} \mathbf{v} = 0$ for all $\mathbf{v} \in \mathbb{R}^n$. Similarly, it follows that for every $1 \leq s \leq r$,

$$\mathbf{P}_E^\perp \mathbf{b}^s = \mathbf{b}^s - \frac{\mathbf{b}^s \cdot \nabla E}{|\nabla E|^2} \nabla E,$$

and therefore M is expressible as

$$\mathbf{M} = \sum_{s,t} D_{st} (\mathbf{P}_E^\perp \mathbf{b}^s) (\mathbf{P}_E^\perp \mathbf{b}^t)^\top = \mathbf{P}_E^\perp \mathbf{B} \mathbf{D} \mathbf{B}^\top \mathbf{P}_E^\perp. \quad \square$$

With Lemma 3.2 established, the proof of Theorem 3.4 is straightforward.

Proof of Theorem 3.4. The “if” direction follows immediately from Lemma 3.2. Now, suppose that \mathbf{L} and \mathbf{M} define a metriplectic system, meaning that the mentioned symmetries and degeneracy conditions hold. Then, it follows from $\mathbf{L}\nabla S = \mathbf{0}$ that the projection $\mathbf{P}_S^\perp \mathbf{L} \mathbf{P}_S^\perp = \mathbf{L}$ leaves \mathbf{L} invariant, so that choosing $\mathbf{A} = \mathbf{L}$ yields $\mathbf{P}_S^\perp \mathbf{A} \mathbf{P}_S^\perp = \mathbf{L}$. Similarly, from positive semi-definiteness and $\mathbf{M}\nabla E = \mathbf{0}$ it follows that $\mathbf{M} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top = \mathbf{P}_E^\perp \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top \mathbf{P}_E^\perp$ for some column-orthonormal $\mathbf{U} \in \mathbb{R}^{N \times r}$ and positive diagonal $\mathbf{\Lambda} \in \mathbb{R}^{r \times r}$. Therefore, choosing $\mathbf{B} = \mathbf{U}$ and $\mathbf{D} = \mathbf{\Lambda}$ yields $\mathbf{M} = \mathbf{P}_E^\perp \mathbf{B} \mathbf{D} \mathbf{B}^\top \mathbf{P}_E^\perp$, as desired. \square

Looking toward the proof of Proposition 3.7, we also need to establish the following Lemmata which give control over the orthogonal projectors $\mathbf{P}_E^\perp, \mathbf{P}_S^\perp$. First, we recall how control over the L^∞ norm $|\cdot|_\infty$ of a matrix field gives control over its spectral norm $|\cdot|$.

Lemma A.1. *Let $\mathbf{A} : K \rightarrow \mathbb{R}^{n \times n}$ be a matrix field defined on the compact set $K \subset \mathbb{R}^n$ with m continuous derivatives. Then, for any $\varepsilon > 0$ there exists a two-layer neural network $\tilde{\mathbf{A}} : K \rightarrow \mathbb{R}^{n \times n}$ such that $\sup_{\mathbf{x} \in K} |\mathbf{A} - \tilde{\mathbf{A}}| < \varepsilon$ and $\sup_{\mathbf{x} \in K} |\nabla^k \mathbf{A} - \nabla^k \tilde{\mathbf{A}}|_\infty < \varepsilon$ for $1 \leq k \leq m$ where ∇^k is the (total) derivative operator of order k .*

Proof. This will be a direct consequence of Corollary 2.2 in Li (1996) provided we show that $|\mathbf{A}| \leq c|\mathbf{A}|_\infty$ for some $c > 0$. To that end, if $\sigma_1 \geq \dots \geq \sigma_r > 0$ ($r \leq n$) denote the nonzero singular values of $\mathbf{A} - \tilde{\mathbf{A}}$, it follows that for each $\mathbf{x} \in K$,

$$|\mathbf{A} - \tilde{\mathbf{A}}| = \sigma_1 \leq \sqrt{\sigma_1^2 + \dots + \sigma_r^2} = \sqrt{\sum_{i,j} |A_{ij} - \tilde{A}_{ij}|^2} = |\mathbf{A} - \tilde{\mathbf{A}}|_F.$$

On the other hand, it also follows that

$$|\mathbf{A} - \tilde{\mathbf{A}}|_F = \sqrt{\sum_{i,j} |A_{ij} - \tilde{A}_{ij}|^2} \leq \sqrt{\sum_{i,j} \max_{i,j} |A_{ij} - \tilde{A}_{ij}|} = n \sqrt{\max_{i,j} |A_{ij} - \tilde{A}_{ij}|} = n |\mathbf{A} - \tilde{\mathbf{A}}|_\infty,$$

and therefore the desired inequality holds with $c = n$. Now, for any $\varepsilon > 0$ it follows from Li (1996) that there exists a two layer network $\tilde{\mathbf{A}}$ with m continuous derivatives such that $\sup_{\mathbf{x} \in K} |\mathbf{A} - \tilde{\mathbf{A}}|_\infty < \varepsilon/n$ and $\sup_{\mathbf{x} \in K} |\nabla^k \mathbf{A} - \nabla^k \tilde{\mathbf{A}}|_\infty < \varepsilon/n < \varepsilon$ for all $1 \leq k \leq m$. Therefore, it follows that

$$\sup_{\mathbf{x} \in K} |\mathbf{A} - \tilde{\mathbf{A}}| \leq n \sup_{\mathbf{x} \in K} |\mathbf{A} - \tilde{\mathbf{A}}|_\infty < n \frac{\varepsilon}{n} = \varepsilon,$$

completing the argument. \square

Next, we bound the deviation in the orthogonal projectors $\mathbf{P}_E^\perp, \mathbf{P}_S^\perp$.

Lemma A.2. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be such that $\nabla f \neq \mathbf{0}$ on the compact set $K \subset \mathbb{R}^n$. For any $\varepsilon > 0$, there exists a two-layer neural network $\tilde{f} : K \rightarrow \mathbb{R}$ such that $\nabla \tilde{f} \neq \mathbf{0}$ on K , $\sup_{\mathbf{x} \in K} |f - \tilde{f}| < \varepsilon$, $\sup_{\mathbf{x} \in K} |\nabla f - \nabla \tilde{f}| < \varepsilon$, and $\sup_{\mathbf{x} \in K} |\mathbf{P}_f^\perp - \mathbf{P}_{\tilde{f}}^\perp| < \varepsilon$.*

Proof. Denote $\nabla f = \mathbf{v}$ and consider any $\tilde{\mathbf{v}} : K \rightarrow \mathbb{R}$. Since $|\mathbf{v}| \leq |\tilde{\mathbf{v}}| + |\mathbf{v} - \tilde{\mathbf{v}}|$, it follows for all $\mathbf{x} \in K$ that whenever $|\mathbf{v} - \tilde{\mathbf{v}}| < (1/2) \inf_{\mathbf{x} \in K} |\mathbf{v}|$,

$$|\tilde{\mathbf{v}}| \geq |\mathbf{v}| - |\mathbf{v} - \tilde{\mathbf{v}}| > |\mathbf{v}| - \frac{1}{2} \inf_{\mathbf{x} \in K} |\mathbf{v}| > 0,$$

so that $\tilde{\mathbf{v}} \neq \mathbf{0}$ in K , and since the square function is monotonic,

$$\inf_{\mathbf{x} \in K} |\tilde{\mathbf{v}}|^2 \geq \inf_{\mathbf{x} \in K} \left(|\mathbf{v}| - \frac{1}{2} \inf_{\mathbf{x} \in K} |\mathbf{v}| \right)^2 = \frac{1}{4} \inf_{\mathbf{x} \in K} |\mathbf{v}|^2.$$

On the other hand, we also have $|\tilde{\mathbf{v}}| \leq |\mathbf{v}| + |\mathbf{v} - \tilde{\mathbf{v}}| < |\mathbf{v}| + (1/2) \inf_{\mathbf{x} \in K} |\mathbf{v}|$, so that, adding and subtracting $\tilde{\mathbf{v}}\mathbf{v}^\top$ and applying Cauchy-Schwarz, it follows that for all $\mathbf{x} \in K$,

$$|\mathbf{v}\mathbf{v}^\top - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top| \leq |\mathbf{v} - \tilde{\mathbf{v}}||\mathbf{v}| + |\tilde{\mathbf{v}}||\mathbf{v} - \tilde{\mathbf{v}}| \leq 2 \max\{|\mathbf{v}|, |\tilde{\mathbf{v}}|\} |\mathbf{v} - \tilde{\mathbf{v}}| < \left(2|\mathbf{v}| + \inf_{\mathbf{x} \in K} |\mathbf{v}| \right) |\mathbf{v} - \tilde{\mathbf{v}}|.$$

Now, by Corollary 2.2 in Li (1996), for any $\varepsilon > 0$ there exists a two-layer neural network $\tilde{f} : K \rightarrow \mathbb{R}$ such that

$$\sup_{\mathbf{x} \in K} |\mathbf{v} - \nabla \tilde{f}| < \min \left\{ \frac{1}{2} \inf_{\mathbf{x} \in K} |\mathbf{v}|, \frac{\inf_{\mathbf{x} \in K} |\mathbf{v}|^2}{2 \sup_{\mathbf{x} \in K} |\mathbf{v}| + \inf_{\mathbf{x} \in K} |\mathbf{v}|} \frac{\varepsilon}{4}, \varepsilon \right\} \leq \varepsilon,$$

and also $\sup_{\mathbf{x} \in K} |f - \tilde{f}| < \varepsilon$. Letting $\tilde{\mathbf{v}} = \nabla \tilde{f}$, it follows that for all $\mathbf{x} \in K$,

$$\left| \mathbf{P}_f^\perp - \mathbf{P}_{\tilde{f}}^\perp \right| = \left| \frac{\mathbf{v}\mathbf{v}^\top}{|\mathbf{v}|^2} - \frac{\tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top}{|\tilde{\mathbf{v}}|^2} \right| \leq \frac{|\mathbf{v}\mathbf{v}^\top - \tilde{\mathbf{v}}\tilde{\mathbf{v}}^\top|}{\min\{|\mathbf{v}|^2, |\tilde{\mathbf{v}}|^2\}} \leq \frac{2|\mathbf{v}| + \inf_{\mathbf{x} \in K} |\mathbf{v}|}{\min\{|\mathbf{v}|^2, |\tilde{\mathbf{v}}|^2\}} |\mathbf{v} - \tilde{\mathbf{v}}|,$$

and therefore, taking the supremum of both sides and applying the previous work yields the desired estimate,

$$\sup_{\mathbf{x} \in K} \left| \mathbf{P}_f^\perp - \mathbf{P}_{\tilde{f}}^\perp \right| \leq 4 \frac{2 \sup_{\mathbf{x} \in K} |\mathbf{v}| + \inf_{\mathbf{x} \in K} |\mathbf{v}|}{\inf_{\mathbf{x} \in K} |\mathbf{v}|^2} \sup_{\mathbf{x} \in K} |\mathbf{v} - \tilde{\mathbf{v}}| < \varepsilon. \quad \square$$

With these intermediate results established, the proof of the approximation result Proposition 3.7 proceeds as follows.

Proof of Proposition 3.7. Recall from Theorem 3.4 that we can write $\mathbf{L} = \mathbf{P}_S^\perp (\mathbf{A}_{\text{tri}} - \mathbf{A}_{\text{tri}}^\top) \mathbf{P}_S^\perp$ and similarly for $\tilde{\mathbf{L}}$. Notice that, by adding and subtracting $\mathbf{P}_S^\perp \mathbf{A}_{\text{tri}} \mathbf{P}_S^\perp$ and $\mathbf{P}_{\tilde{S}}^\perp \tilde{\mathbf{A}}_{\text{tri}} \mathbf{P}_{\tilde{S}}^\perp$, it follows that for all $\mathbf{x} \in K$,

$$\begin{aligned} & \left| \mathbf{P}_S^\perp \mathbf{A}_{\text{tri}} \mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp \tilde{\mathbf{A}}_{\text{tri}} \mathbf{P}_{\tilde{S}}^\perp \right| \\ &= \left| (\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp) \mathbf{A}_{\text{tri}} \mathbf{P}_S^\perp + \mathbf{P}_{\tilde{S}}^\perp (\mathbf{A}_{\text{tri}} - \tilde{\mathbf{A}}_{\text{tri}}) \mathbf{P}_S^\perp + \mathbf{P}_{\tilde{S}}^\perp \tilde{\mathbf{A}}_{\text{tri}} (\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp) \right| \\ &\leq |\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp| |\mathbf{A}_{\text{tri}}| + |\mathbf{A}_{\text{tri}} - \tilde{\mathbf{A}}_{\text{tri}}| + |\tilde{\mathbf{A}}_{\text{tri}}| |\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp| \\ &\leq 2 \max\{|\mathbf{A}_{\text{tri}}|, |\tilde{\mathbf{A}}_{\text{tri}}|\} |\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp| + |\mathbf{A}_{\text{tri}} - \tilde{\mathbf{A}}_{\text{tri}}| \end{aligned}$$

where we have used that $\mathbf{P}_S^\perp, \mathbf{P}_{\tilde{S}}^\perp$ have unit spectral norm. By Lemma A.1, for any $\varepsilon > 0$ there exists a two layer neural network $\tilde{\mathbf{A}}_{\text{tri}}$ such that $\sup_{\mathbf{x} \in K} |\mathbf{A}_{\text{tri}} - \tilde{\mathbf{A}}_{\text{tri}}| < \frac{\varepsilon}{4}$, and by Lemma A.2 there exists a two-layer network \tilde{S} with $\nabla \tilde{S} \neq \mathbf{0}$ on K such that

$$\sup_{\mathbf{x} \in K} |\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp| < \min \left\{ \varepsilon, \max \left\{ \sup_{\mathbf{x} \in K} |\mathbf{A}_{\text{tri}}|, \sup_{\mathbf{x} \in K} |\tilde{\mathbf{A}}_{\text{tri}}| \right\}^{-1} \frac{\varepsilon}{8} \right\}.$$

It follows that $\tilde{S}, \nabla \tilde{S}$ are ε -close to $S, \nabla S$ on K and

$$\sup_{\mathbf{x} \in K} \left(2 \max\{|\mathbf{A}_{\text{tri}}|, |\tilde{\mathbf{A}}_{\text{tri}}|\} |\mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp| \right) < \frac{\varepsilon}{4}.$$

Therefore, the estimate

$$\sup_{\mathbf{x} \in K} |\mathbf{L} - \tilde{\mathbf{L}}| \leq 2 \sup_{\mathbf{x} \in K} \left| \mathbf{P}_S^\perp \mathbf{A}_{\text{tri}} \mathbf{P}_S^\perp - \mathbf{P}_{\tilde{S}}^\perp \tilde{\mathbf{A}}_{\text{tri}} \mathbf{P}_{\tilde{S}}^\perp \right| < 2 \left(\frac{\varepsilon}{4} + \frac{\varepsilon}{4} \right) = \varepsilon,$$

implies that $\tilde{\mathbf{L}}$ is ε -close to \mathbf{L} on K as well.

Moving to the case of \mathbf{M} , we see that for all $\mathbf{x} \in K$, by writing $\mathbf{M} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top = \mathbf{K}_{\text{chol}} \mathbf{K}_{\text{chol}}^\top$ for $\mathbf{K}_{\text{chol}} = \mathbf{U} \mathbf{\Lambda}^{1/2}$ and repeating the first calculation with \mathbf{K}_{chol} in place of \mathbf{A}_{tri} and \mathbf{P}_E^\perp in place of \mathbf{P}_S^\perp ,

$$\begin{aligned} & \left| \mathbf{P}_E^\perp \mathbf{K}_{\text{chol}} \mathbf{K}_{\text{chol}}^\top \mathbf{P}_E^\perp - \mathbf{P}_{\tilde{E}}^\perp \tilde{\mathbf{K}}_{\text{chol}} \tilde{\mathbf{K}}_{\text{chol}}^\top \mathbf{P}_{\tilde{E}}^\perp \right| \\ &\leq 2 \max\{|\mathbf{K}_{\text{chol}}|, |\tilde{\mathbf{K}}_{\text{chol}}|\} |\mathbf{P}_E^\perp - \mathbf{P}_{\tilde{E}}^\perp| + |\mathbf{K}_{\text{chol}} \mathbf{K}_{\text{chol}}^\top - \tilde{\mathbf{K}}_{\text{chol}} \tilde{\mathbf{K}}_{\text{chol}}^\top|. \end{aligned}$$

Moreover, if $\|\mathbf{K}_{\text{chol}} - \tilde{\mathbf{K}}_{\text{chol}}\| < (1/2) \inf_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\|$ for all $\mathbf{x} \in K$ then similar arguments as used in the proof of Lemma A.2 yield the following estimate for all $\mathbf{x} \in K$,

$$\begin{aligned} \|\mathbf{K}_{\text{chol}} \mathbf{K}_{\text{chol}}^\top - \tilde{\mathbf{K}}_{\text{chol}} \tilde{\mathbf{K}}_{\text{chol}}^\top\| &\leq 2 \max\{\|\mathbf{K}_{\text{chol}}\|, \|\tilde{\mathbf{K}}_{\text{chol}}\|\} \|\mathbf{K}_{\text{chol}} - \tilde{\mathbf{K}}_{\text{chol}}\| \\ &\leq \left(2\|\mathbf{K}_{\text{chol}}\| + \inf_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\|\right) \|\mathbf{K}_{\text{chol}} - \tilde{\mathbf{K}}_{\text{chol}}\|. \end{aligned}$$

As before, we now invoke Lemma A.1 to construct a two-layer lower-triangular network $\tilde{\mathbf{K}}_{\text{chol}}$ such that

$$\sup_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}} - \tilde{\mathbf{K}}_{\text{chol}}\| < \min \left\{ \frac{1}{2} \inf_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\|, \left(2 \sup_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\| + \inf_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\|\right)^{-1} \frac{\varepsilon}{2} \right\},$$

as well as (using Lemma A.2) a network \tilde{E} satisfying $\nabla \tilde{E} \neq \mathbf{0}$ on K and

$$\sup_{\mathbf{x} \in K} \|\mathbf{P}_E^\perp - \mathbf{P}_{\tilde{E}}^\perp\| < \min \left\{ \varepsilon, \max \left\{ \sup_{\mathbf{x} \in K} \|\mathbf{K}_{\text{chol}}\|, \sup_{\mathbf{x} \in K} \|\tilde{\mathbf{K}}_{\text{chol}}\| \right\}^{-1} \frac{\varepsilon}{4} \right\}.$$

Again, it follows that $\tilde{E}, \nabla \tilde{E}$ are ε -close to $E, \nabla E$ on K , and by the work above we conclude

$$\sup_{\mathbf{x} \in K} \|\mathbf{M} - \tilde{\mathbf{M}}\| = \sup_{\mathbf{x} \in K} \|\mathbf{P}_E^\perp \mathbf{K}_{\text{chol}} \mathbf{K}_{\text{chol}}^\top \mathbf{P}_E^\perp - \mathbf{P}_{\tilde{E}}^\perp \tilde{\mathbf{K}}_{\text{chol}} \tilde{\mathbf{K}}_{\text{chol}}^\top \mathbf{P}_{\tilde{E}}^\perp\| < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon,$$

as desired. \square

It is now possible to give a proof of the error bound in Theorem 3.9. Recall the $L^2([0, T])$ error metric $\|\mathbf{x}\|$ and Lipschitz constant L_f , defined for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and Lipschitz continuous functions f as

$$\|\mathbf{x}\|^2 = \int_0^T \|\dot{\mathbf{x}}\|^2 dt, \quad |f(\mathbf{x}) - f(\mathbf{y})| \leq L_f \|\mathbf{x} - \mathbf{y}\|.$$

Proof of Theorem 3.9. First, note that the assumption that one of $E, -S$ (without loss of generality, say E) has bounded sublevel sets implies bounded trajectories for the state \mathbf{x} as in Remark 3.8, so we may assume $\mathbf{x} \in K$ for some compact $K \subset \mathbb{R}^n$. Moreover, for any $\varepsilon > 0$ it follows from Proposition 3.7 that there are approximate networks \tilde{E}, \tilde{S} which are ε -close to E, S on K . Additionally, it follows that \tilde{E}, \tilde{S} have nonzero gradients $\nabla \tilde{E}, \nabla \tilde{S}$ which are also ε -close to the true gradients $\nabla E, \nabla S$ on K . This implies that for each $\mathbf{x} \in K$, $E = \tilde{E} + (E - \tilde{E}) \leq \tilde{E} + \varepsilon$, so it follows that the sublevel sets $\{\mathbf{x} \mid \tilde{E}(\mathbf{x}) \leq m\} \subseteq \{\mathbf{x} \mid E(\mathbf{x}) \leq m + \varepsilon\}$ are also bounded. Therefore, we may assume (by potentially enlarging K) that both $\mathbf{x}, \tilde{\mathbf{x}} \in K$ lie in the compact set K for all time.

Now, let $\mathbf{y} = \mathbf{x} - \tilde{\mathbf{x}}$. The next goal is to bound the following quantity:

$$\begin{aligned} |\dot{\mathbf{y}}| &= \left| \mathbf{L}(\mathbf{x}) \nabla E(\mathbf{x}) + \mathbf{M}(\mathbf{x}) \nabla S(\mathbf{x}) - \tilde{\mathbf{L}}(\tilde{\mathbf{x}}) \nabla \tilde{E}(\tilde{\mathbf{x}}) - \tilde{\mathbf{M}}(\tilde{\mathbf{x}}) \nabla \tilde{S}(\tilde{\mathbf{x}}) \right| \\ &= \left| \left(\mathbf{L}(\mathbf{x}) \nabla E(\mathbf{x}) - \tilde{\mathbf{L}}(\tilde{\mathbf{x}}) \nabla E(\tilde{\mathbf{x}}) \right) + \left(\mathbf{M}(\mathbf{x}) \nabla S(\mathbf{x}) - \tilde{\mathbf{M}}(\tilde{\mathbf{x}}) \nabla S(\tilde{\mathbf{x}}) \right) \right| =: |\dot{\mathbf{y}}_E + \dot{\mathbf{y}}_S|. \end{aligned}$$

To that end, notice that by adding and subtracting $\mathbf{L}(\mathbf{x}) \nabla E(\tilde{\mathbf{x}}), \tilde{\mathbf{L}}(\mathbf{x}) \nabla E(\tilde{\mathbf{x}}), \tilde{\mathbf{L}}(\tilde{\mathbf{x}}) \nabla E(\tilde{\mathbf{x}})$, it follows that

$$\begin{aligned} \dot{\mathbf{y}}_E &= \mathbf{L}(\mathbf{x}) (\nabla E(\mathbf{x}) - \nabla E(\tilde{\mathbf{x}})) + \left(\mathbf{L}(\mathbf{x}) - \tilde{\mathbf{L}}(\mathbf{x}) \right) \nabla E(\tilde{\mathbf{x}}) \\ &\quad + \left(\tilde{\mathbf{L}}(\mathbf{x}) - \tilde{\mathbf{L}}(\tilde{\mathbf{x}}) \right) \nabla E(\tilde{\mathbf{x}}) + \tilde{\mathbf{L}}(\tilde{\mathbf{x}}) (\nabla E(\tilde{\mathbf{x}}) - \nabla \tilde{E}(\tilde{\mathbf{x}})). \end{aligned}$$

By Proposition 3.7 there exists a two-layer neural network $\tilde{\mathbf{L}}$ with one continuous derivative such that $\sup_{\mathbf{x} \in K} \|\mathbf{L} - \tilde{\mathbf{L}}\| < \varepsilon$, which implies that $\tilde{\mathbf{L}}$ is Lipschitz continuous with (uniformly well-approximated) Lipschitz constant. Using this fact along with the assumed Lipschitz continuity of ∇E and the approximation properties of the network \tilde{E} already constructed then yields

$$|\dot{\mathbf{y}}_E| \leq \left(L_{\nabla E} \sup_{\mathbf{x} \in K} \|\mathbf{L}\| + L_{\tilde{\mathbf{L}}} \sup_{\mathbf{x} \in K} \|\nabla E\| \right) |\mathbf{y}| + \varepsilon \left(\sup_{\mathbf{x} \in K} \|\tilde{\mathbf{L}}\| + \sup_{\mathbf{x} \in K} \|\nabla E\| \right) =: a_E |\mathbf{y}| + \varepsilon b_E.$$

Similarly, by adding and subtracting $\mathbf{M}(\mathbf{x})\nabla S(\tilde{\mathbf{x}})$, $\tilde{\mathbf{M}}(\mathbf{x})\nabla S(\tilde{\mathbf{x}})$, $\tilde{\mathbf{M}}(\tilde{\mathbf{x}})\nabla S(\tilde{\mathbf{x}})$, it follows that

$$\begin{aligned}\dot{\mathbf{y}}_S &= \mathbf{M}(\mathbf{x})(\nabla S(\mathbf{x}) - \nabla S(\tilde{\mathbf{x}})) + (\mathbf{M}(\mathbf{x}) - \tilde{\mathbf{M}}(\mathbf{x}))\nabla S(\tilde{\mathbf{x}}) \\ &\quad + (\tilde{\mathbf{M}}(\mathbf{x}) - \tilde{\mathbf{M}}(\tilde{\mathbf{x}}))\nabla S(\tilde{\mathbf{x}}) + \tilde{\mathbf{M}}(\tilde{\mathbf{x}})(\nabla S(\tilde{\mathbf{x}}) - \nabla \tilde{S}(\tilde{\mathbf{x}})).\end{aligned}$$

By Proposition 3.7, there exists a two-layer network $\tilde{\mathbf{M}}$ with one continuous derivative such that $\sup_{\mathbf{x} \in K} \|\mathbf{M} - \tilde{\mathbf{M}}\| < \varepsilon$, with $\tilde{\mathbf{M}}$ Lipschitz continuous for the same reason as before. It follows from this and $\sup_{\mathbf{x} \in K} \|\nabla S - \nabla \tilde{S}\| < \varepsilon$ that

$$|\dot{\mathbf{y}}_S| \leq \left(L_{\nabla S} \sup_{\mathbf{x} \in K} \|\mathbf{M}\| + L_{\tilde{\mathbf{M}}} \sup_{\mathbf{x} \in K} \|\nabla S\| \right) |\mathbf{y}| + \varepsilon \left(\sup_{\mathbf{x} \in K} \|\tilde{\mathbf{M}}\| + \sup_{\mathbf{x} \in K} \|\nabla S\| \right) =: a_S |\mathbf{y}| + \varepsilon b_S.$$

Now, recall that $\partial_t |\mathbf{y}| = |\mathbf{y}|^{-1} (\dot{\mathbf{y}} \cdot \mathbf{y}) \leq |\dot{\mathbf{y}}|$ by Cauchy-Schwarz, and therefore the time derivative of $|\mathbf{y}|$ is bounded by

$$\partial_t |\mathbf{y}| \leq |\dot{\mathbf{y}}_E| + |\dot{\mathbf{y}}_S| = (a_E + a_S) |\mathbf{y}| + \varepsilon (b_E + b_S) =: a |\mathbf{y}| + b.$$

This implies that $\partial_t |\mathbf{y}| - a |\mathbf{y}| \leq b$, so multiplying by the integrating factor e^{-at} and integrating in time yields

$$|\mathbf{y}(t)| \leq \varepsilon b \int_0^t e^{a(t-\tau)} d\tau = \varepsilon \frac{b}{a} (e^{at} - 1),$$

where we used that $\mathbf{y}(0) = \mathbf{0}$ since the initial condition of the trajectories is shared. Therefore, the L^2 error in time can be approximated by

$$\|\mathbf{y}\|^2 = \int_0^T |\mathbf{y}|^2 dt \leq \varepsilon^2 \frac{b^2}{a^2} (e^{2aT} - 2e^{aT} + T + 1),$$

establishing the conclusion. \square

B EXPERIMENTAL AND IMPLEMENTATION DETAILS

This Appendix records additional details related to the numerical experiments in Section 5. For each benchmark problem, a set of trajectories is manufactured given initial conditions by simulating ODEs with known metriplectic structure. [The metriplectic forms of these systems are given in the subsections below.](#) For the experiments in Table 2, only the observable variables are used to construct datasets, since entropic information is assumed to be unknown. Algorithm 2 summarizes the training of the dynamics models used for comparison with NMS.

Algorithm 2 Training dynamics models

- 1: **Input:** snapshot data $\mathbf{X} \in \mathbb{R}^{n \times n_s}$, each column $\mathbf{x}_s = \mathbf{x}(t_s, \boldsymbol{\mu}_s)$, target rank $r \geq 1$
 - 2: Initialize loss $L = 0$ and networks with parameters Θ
 - 3: **for** step in N_{steps} **do**
 - 4: Randomly draw an initial condition $(t_{0_k}, \mathbf{x}_{0_k})$ where $k \in n_s$
 - 5: $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_l = \text{ODEsolve}(\mathbf{x}_{0_k}, \dot{\mathbf{x}}, t_1, \dots, t_l)$
 - 6: Compute the loss $L((\mathbf{x}_1^0, \dots, \mathbf{x}_l^0), (\tilde{\mathbf{x}}_0^0, \dots, \tilde{\mathbf{x}}_l^0))$
 - 7: Update the model parameters Θ via SGD
 - 8: **end for**
-

For each compared method, integrating the ODEs is done via the Dormand–Prince method (dopri5) Dormand and Prince (1980) with relative tolerance 10^{-7} and absolute tolerance 10^{-9} . The loss is evaluated by measuring the discrepancy between the ground truth observable states \mathbf{x}^o and the approximate observable states $\tilde{\mathbf{x}}^o$ in the mean absolute error (MAE) and mean squared error (MSE) metrics, [which are defined in the usual way via the \$\ell^1, \ell^2\$ norms respectively for true data \$\mathbf{x}\(t_i\)\$ and approximate data \$\tilde{\mathbf{x}}\(t_i\)\$:](#)

$$MAE(\tilde{\mathbf{x}}, \mathbf{x}) = \frac{1}{n_t} \sum_{i=1}^{n_t} |\tilde{\mathbf{x}}(t_i) - \mathbf{x}(t_i)|_1, \quad MSE(\tilde{\mathbf{x}}, \mathbf{x}) = \frac{1}{n_t} \sum_{i=1}^{n_t} |\tilde{\mathbf{x}}(t_i) - \mathbf{x}(t_i)|_2.$$

The model parameters Θ (i.e., the weights and biases) are updated by using Adamax Kingma and Ba (2014) with an initial learning rate of 0.01. The number of training steps is set as 30,000, and the model parameters resulting in the best performance for the validation set are chosen for testing. Specific information related to the experiments in Section 5 is given in the subsections below.

For generating the results reported in Table 2, we implemented the proposed algorithm in Python 3.9.12 and PyTorch 2.0.0. Other required information is provided with the accompanying code. All experiments are conducted on Apple M2 Max chips with 96 GB memory. To provide the mean and the standard deviation, experiments are repeated three times with varying random seeds for all considered methods.

B.1 TWO GAS CONTAINERS

The metriplectic equations of motion in the two gas containers example are given in terms of the state $\mathbf{x} = (q \ p \ S_1 \ S_2)^\top$. With the energy $E(\mathbf{x}) = (1/2m)p^2 + E_1(\mathbf{x}) + E_2(\mathbf{x})$ and entropy $S(\mathbf{x}) = S_1 + S_2$ defined in Section 5, it follows that $\dot{\mathbf{x}} = \mathbf{L}\nabla E(\mathbf{x}) + \mathbf{M}(\mathbf{x})\nabla S$ where $\nabla S = (0 \ 0 \ 1 \ 1)^\top$ and

$$\mathbf{L} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{M}(\mathbf{x}) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & T_1^{-2} & -(T_1 T_2)^{-1} \\ 0 & 0 & -(T_1 T_2)^{-1} & T_2^{-2} \end{pmatrix},$$

$$\nabla E(\mathbf{x}) = \begin{pmatrix} -\left(\frac{e^{2S_1}}{q^5}\right)^{\frac{1}{3}} + \left(\frac{e^{2S_2}}{(2-q)^5}\right)^{\frac{1}{3}} \\ \frac{3}{2}p \\ E_1(\mathbf{x}) \\ E_2(\mathbf{x}) \end{pmatrix}, \quad T_i(\mathbf{x}) = \partial_{S_i} E_i(\mathbf{x}).$$

As mentioned in the body, the two gas container (TGC) problem tests models' predictive capability (i.e., extrapolation in time). To this end, one simulated trajectory is obtained by solving an IVP with a known TGC system and an initial condition, and the trajectory of the observable variables is split into three subsequences, $[0, t_{\text{train}}]$, $(t_{\text{train}}, t_{\text{val}}]$, and $(t_{\text{val}}, t_{\text{test}}]$ for training, validation, and test with $0 < t_{\text{train}} < t_{\text{val}} < t_{\text{test}}$.

In the experiment, a sequence of 100,000 timesteps is generated using the Runge–Kutta 4th-order (RK4) time integrator with a step size 0.001. The initial condition is given as $\mathbf{x} = (1, 2, 103.2874, 103.2874)$ following Shang and Öttinger (2020). The training/validation/test split is defined by $t_{\text{train}} = 20$, $t_{\text{val}} = 30$, and $t_{\text{test}} = 100$. For a fair comparison, all considered models are set to have a similar number of model parameters, $\sim 2,000$. The specifications of the network architectures are:

- NMS: The total number of model parameters is 1959. The functions \mathbf{A}_{tri} , \mathbf{B} , \mathbf{K}_{chol} , E , S are parameterized as MLPs with the Tanh nonlinear activation function. The MLPs parameterizing \mathbf{A}_{tri} , \mathbf{B} , \mathbf{K}_{chol} , E are specified as 1 hidden layer with 10 neurons, and the one parameterizing S is specified as 3 hidden layers with 25 neurons.
- NODE: The total number of model parameters is 2179. The black-box NODE is parameterized as an MLP with the Tanh nonlinear activation function, 4 hidden layers and 25 neurons.
- SPNN: The total number of model parameters is 1954. The functions E and S are parameterized as MLPs with the Tanh nonlinear activation function; each MLP is specified as 3 hidden layers and 20 neurons. The two 2-tensors defining \mathbf{L} and \mathbf{M} are defined as learnable 3×3 matrices.
- GNODE: The total number of model parameters is 2343. The functions E and S are parameterized as MLPs with the Tanh nonlinear activation function; each MLP is specified as 2 hidden layers and 30 neurons. The matrices and 3-tensors required to learn \mathbf{L} and \mathbf{M} are defined as learnable 3×3 matrices and $3 \times 3 \times 3$ tensor.
- GFINN: The total number of model parameters is 2065. The functions E and S are parameterized as MLPs with Tanh nonlinear activation function; each MLP is specified as 2

hidden layers and 20 neurons. The matrices to required to learn \mathbf{L} and \mathbf{M} are defined as K learnable 3×3 matrices, where K is set to 2.

B.2 THERMOELASTIC DOUBLE PENDULUM

The equations of motion in this case are given for $1 \leq i \leq 2$ as

$$\dot{\mathbf{q}}_i = \frac{\mathbf{p}_i}{m_i}, \quad \dot{\mathbf{p}}_i = -\partial_{\mathbf{q}_i}(E_1(\mathbf{x}) + E_2(\mathbf{x})), \quad \dot{S}_1 = \kappa(T_1^{-1}T_2 - 1), \quad \dot{S}_2 = \kappa(T_1T_2^{-1} - 1),$$

where $\kappa > 0$ is a thermal conductivity constant (set to 1), m_i is the mass of the i^{th} spring (also set to 1) and $T_i = \partial_{S_i} E_i$ is its absolute temperature. In this case, $\mathbf{q}_i, \mathbf{p}_i \in \mathbb{R}^2$ represent the position and momentum of the i^{th} mass, while S_i represents the entropy of the i^{th} pendulum. As before, the total entropy $S(\mathbf{x}) = S_1 + S_2$ is the sum of the entropies of the two springs, while defining the internal energies

$$E_i(\mathbf{x}) = \frac{1}{2}(\ln \lambda_i)^2 + \ln \lambda_i + e^{S_i - \ln \lambda_i} - 1, \quad \lambda_1 = |\mathbf{q}_i|, \quad \lambda_2 = |\mathbf{q}_2 - \mathbf{q}_1|,$$

leads to the total energy $E(\mathbf{x}) = (1/2m_1)|\mathbf{p}_1|^2 + (1/2m_2)|\mathbf{p}_2|^2 + E_1(\mathbf{x}) + E_2(\mathbf{x})$. This system is metriplectic with matrix fields \mathbf{L}, \mathbf{M} , given by

$$\mathbf{L}(\mathbf{x}) = \begin{pmatrix} 0 & 0 & \mathbf{I} & 0 & 0 \\ 0 & 0 & 0 & \mathbf{I} & 0 \\ -\mathbf{I} & 0 & 0 & 0 & 0 \\ 0 & -\mathbf{I} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{M}(\mathbf{x}) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & T_1^{-1}T_2(\mathbf{x}) & -1 \\ 0 & 0 & 0 & 0 & -1 & T_1T_2^{-1}(\mathbf{x}) \end{pmatrix},$$

where $T_i = \partial_{S_i} E_i$.

The thermoelastic double pendulum experiment tests model prediction across initial conditions. In this case, 100 trajectories are generated by varying initial conditions that are randomly sampled from $[0.1, 1.1] \times [-0.1, 0.1] \times [2.1, 2.3] \times [-0.1, 0.1] \times [-1.9, 2.1] \times [0.9, 1.1] \times [-0.1, 0.1] \times [0.9, 1.1] \times [0.1, 0.3] \subset \mathbb{R}^{10}$. Each trajectory is obtained from the numerical integration of the ODEs using an RK4 time integrator with step size 0.02 and the final time $T = 40$, resulting in the trajectories of length 2,000. The resulting 100 trajectories are split into 80/10/10 for training/validation/test sets. For a fair comparison, all considered models are again set to have similar number of model parameters, $\sim 2,000$. The specifications of the network architectures are:

- NMS: The total number of model parameters is 2201. The functions $\mathbf{A}, \mathbf{B}, \mathbf{K}, E, S$ are parameterized as MLPs with the Tanh nonlinear activation function. The MLPs parameterizing are specified as 1 hidden layer with 15 neurons.
- NODE: The total number of model parameters is 2005. The black-box NODE is parameterized as an MLP with the Tanh nonlinear activation function, 2 hidden layers and 35 neurons.
- SPNN: The total number of model parameters is 2362. The functions E and S are parameterized as MLPs with the Tanh nonlinear activation function; each MLP is specified as 3 hidden layers and 20 neurons. The two 2-tensors defining \mathbf{L} and \mathbf{M} are defined as learnable 3×3 matrices.
- GNODE: The total number of model parameters is 2151. The functions E and S are parameterized as MLPs with the Tanh nonlinear activation function; each MLP is specified as 2 hidden layers and 15 neurons. The matrices and 3-tensors required to learn \mathbf{L} and \mathbf{M} are defined as learnable 3×3 matrices and $3 \times 3 \times 3$ tensor.
- GFINN: The total number of model parameters is 2180. The functions E and S are parameterized as MLPs with Tanh nonlinear activation function; each MLP is specified as 2 hidden layers and 15 neurons. The matrices to required to learn \mathbf{L} and \mathbf{M} are defined as K learnable 3×3 matrices, where K is set to 2.

C ADDITIONAL FIGURES: THERMOELASTIC DOUBLE PENDULUM

Figure 3 depicts trajectories of the ground-truth and the predicted position (\mathbf{q}), momentum (\mathbf{p}), entropy (S), and energy (E). Figures show that the trajectories produced by NMS follow the ground-truth trajectories closer than those produced by the baseline methods. Figures 3(g) and 3(h) demonstrate that the trajectories of S and E produced by GNODE, GFINN, and NMS are physically-consistent (i.e., non-decreasing entropy and conserved energy).

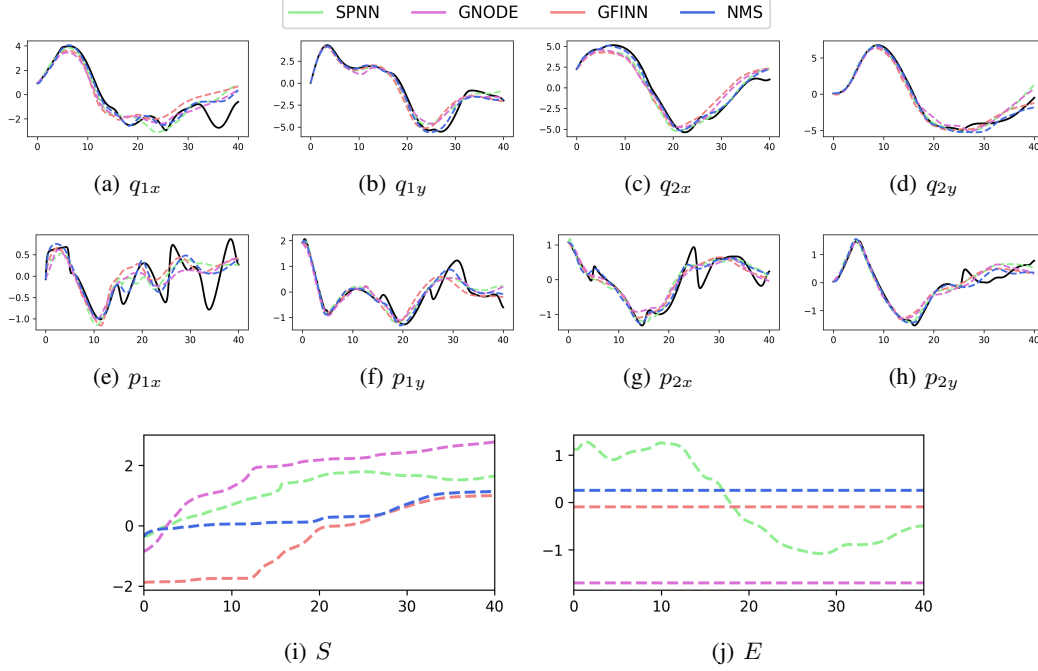


Figure 3: The ground-truth and predicted position, momentum, entropy, and energy for the thermoelastic double pendulum. Figures illustrate that the trajectories generated by NMS follow the ground-truth ones closer than those of other methods do.

D ADDITIONAL EXPERIMENT: DAMPED NONLINEAR OSCILLATOR

Consider a damped nonlinear oscillator of variable dimension with state $\mathbf{x} = (\mathbf{q} \ \mathbf{p} \ S)^\top$, whose motion is governed by the metriplectic system

$$\dot{\mathbf{q}} = \frac{\mathbf{p}}{m}, \quad \dot{\mathbf{p}} = k \sin \mathbf{q} - \gamma \mathbf{p}, \quad \dot{S} = \frac{\gamma |\mathbf{q}|^2}{mT}.$$

Here $\mathbf{q}, \mathbf{p} \in \mathbb{R}^n$ denote the position and momentum of the oscillator, S is the entropy of a surrounding thermal bath, and the constant parameters m, γ, T are the mass, damping rate, and (constant) temperature. This leads to the total energy $E(\mathbf{x}) = (1/2m)|\mathbf{p}|^2 - k \cos \mathbf{q} + TS$, which is readily seen to be constant along solutions $\mathbf{x}(t)$. This system has a metriplectic expression in terms of E, S and the matrix fields

$$\mathbf{L}(\mathbf{x}) = \begin{pmatrix} \mathbf{0} & -\mathbf{I} & 0 \\ -\mathbf{I} & \mathbf{0} & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{M}(\mathbf{x}) = \begin{pmatrix} \mathbf{0} & \mathbf{0} & 0 \\ \mathbf{0} & \mathbf{0} & 0 \\ 0 & 0 & \frac{\gamma |\mathbf{q}|^2}{mT} \end{pmatrix}$$

It is now verified that NMS can accurately and stably predict the dynamics of a nonlinear oscillator $\mathbf{x} = (\mathbf{q} \ \mathbf{p} \ S)^\top$ in the case that $n = 1, 2$, both when the entropy S is observable as well as when it is not. As before, the task considered is prediction in time, although all compared methods NODE,

GNODE, and $\text{NMS}_{\text{known}}$ are now trained on full state information from the training interval, and test errors are computed over the full state \mathbf{x} on the extrapolation interval $(t_{\text{valid}}, t_{\text{test}}]$, which is 150% longer than the training interval. In addition, another NMS model, NMS_{diff} , was trained using only the partial state information $\mathbf{x}^o = (\mathbf{q}, \mathbf{p})^\top$ and tested under the same conditions, with the initial guess for \mathbf{x}^u generated as in Appendix H. As can be seen in Table 3, NMS is more accurate than GNODE or NODE in both the 1-D and 2-D nonlinear oscillator experiments, improving on previous results by up to two orders of magnitude. Remarkably, NMS produces more accurate entropic dynamics even in the case where the entropic variable S is unobserved during NMS training and observed during the training of other methods. This illustrates another advantage of the NMS approach: because of the reasonable initial data for S produced by the diffusion model, the learned metriplectic system produced by NMS remains performant even when metriplectic governing equations are unknown and only partial state information is observed.

Table 3: Experimental results for the benchmark problems with respect to MSE and MAE. The best scores are in boldface.

	1-D D.N.O.		T.G.C.		2-D D.N.O.	
	MSE	MAE	MSE	MAE	MSE	MAE
NMS_{diff}	.0170	.1132	.0045	.0548	.0275	.1456
NMS_{known}	.0239	.1011	.0012	.0276	.0018	.0357
NODE	.0631	.2236	.0860	.2551	.0661	.2096
GNODE	.0607	.1976	.0071	.0732	.2272	.4267

To describe the experimental setup precisely, data is collected from a single trajectory with initial condition as $\mathbf{x} = (\mathbf{z}, \mathbf{0}, 0)$ following Lee et al. (2021). The path is calculated at 180,000 steps with a time interval of 0.001, and is then split into training/validation/test sets as before using $t_{\text{train}} = 60$, $t_{\text{val}} = 90$ and $t_{\text{test}} = 180$. Specifications of the networks used for the experiments in Table 3 are:

- NMS: The total number of parameters is 154. The number of layers for $\mathbf{A}_{\text{tri}}, \mathbf{B}, \mathbf{K}_{\text{chol}}, E, S$ is selected from $\{1, 2, 3\}$ and the number of neurons per layer from $\{5, 10, 15\}$. The best hyperparameters are 1 hidden layer with 5 neurons for each network function.
- GNODE: The total number of model parameters is 203. The number of layers and number of neurons for each network is chosen from the same ranges as for NMS. The best hyperparameters are 1 layer with 10 neurons for each network function.
- NODE: The total number of model parameters is 3003. The NODE architecture is formed by stacking MLPs with Tanh activation functions. The number of blocks is chosen from $\{3, 4, 5\}$ and the number of neurons of each MLP from $\{30, 40, 50\}$. The best hyperparameters are 4 and 30 for the number of blocks and number of neurons, respectively.

E ADDITIONAL EXPERIMENT: HAMILTONIAN NEURAL NETWORKS FOR TWO GAS CONTAINER

As briefly mentioned in the main body, from the definition of metriplectic dynamics, restricting $M = 0$ leads to Hamiltonian systems. In this section, we present the results of training Hamiltonian systems only on the observable variables (\mathbf{q}, \mathbf{p}) from the two gas container example problems. We use the same dataset used in Section 5.1 and the same experimental setup described in Section B.1.

For parameterization of the Hamiltonian dynamics, we explicitly consider the canonical Poisson matrix $\mathbf{L} = \mathbf{J} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$ as is standard in Hamiltonian neural networks (HNNs), along with an MLP to parameterize the Hamiltonian function $H = H_{\Theta}(\mathbf{q}, \mathbf{p})$, resulting in the dynamics

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} \partial_{\mathbf{q}} H \\ \partial_{\mathbf{p}} H \end{bmatrix} = \mathbf{L} \nabla H. \quad (1)$$

Figure 4 presents the trajectories of position and momentum predicted by the trained HNN. It essentially shows that imposing an inductive bias (i.e., energy conservation), which does not match with the underlying physical system (i.e., dissipation), results in catastrophic failure in prediction. Figure 4(c) confirms that the learned Hamiltonian is kept constant over the numerical simulation.




Figure 4: [Hamiltonian neural networks] The ground-truth and predicted trajectories of position and momentum for the two gas containers example, along with the learned Hamiltonian. Notice that the incorrect inductive bias of conservation leads to catastrophic failure in network prediction.

F ADDITIONAL EXPERIMENT: A DAMPED THERMOELASTIC ROD

As an additional benchmark problem, we consider a damped thermoelastic rod, which discretizes an infinite-dimensional metriplectic system. Consider a 1-dimensional elastic rod, with coordinate $s \in [0, l]$, which evolves following the dynamics defined as:

$$\begin{bmatrix} \dot{q} \\ \dot{p} \\ \dot{S} \end{bmatrix} = \begin{bmatrix} \frac{p}{m} \\ V'(q) - \gamma \frac{p}{m} \\ \gamma \frac{p^2}{m^2} \end{bmatrix} = L \nabla E(q, p, e) + M(q, p, e) \nabla S. \quad (2)$$

Here, V denotes a given potential function, γ denotes a constant controlling the rate of dissipation, $m = m(s)$ denotes the mass density, and $e = e(s)$ denotes the internal energy representing the conversion of mechanical energy into heat. The energy function is given by

$$E(q, p, e) = H(q, p) + S(e) = \int_0^\ell \left(\frac{p(s)^2}{2m(s)} + V(q(s)) \right) + \int_0^\ell e(s), \quad (3)$$

where $H(q, p)$ denotes the Hamiltonian.

In numerical computation, this continuous system is discretized, leading to s being represented by a vector $s \in \mathbb{R}^N$, where N is the number of discretization points. This discretization results in a vectorized representation of q, p , denoted as $q, p \in \mathbb{R}^N$, and a discretized version of Eq. (2), where the state variables are $[q, p, S]$. In particular, the resulting metriplectic system is given in terms of $\nabla S = (0 \ 0 \ 1)$ and

$$L(x) = \begin{pmatrix} 0 & I & 0 \\ -I & 0 & 0 \\ 0^\top & 0^\top & 0 \end{pmatrix}, \quad M(x) = \gamma \begin{pmatrix} 0 & 0 & 0 \\ 0 & I & -\frac{p}{m} \\ 0 & -\frac{p^\top}{m} & \left(\frac{|p|}{m}\right)^2 \end{pmatrix}, \quad \nabla E(x) = \begin{pmatrix} V'(q) \\ \frac{p}{m} \\ 1 \end{pmatrix}.$$

In the experiments, we largely follow the setup used in the two gas container experiments, i.e., testing models' predictive capability through temporal extrapolation. One simulated trajectory is obtained by solving an IVP with a known thermoelastic rod system, and the trajectory of observed variables is then split into three subsequences, $[0, t_{\text{train}}]$, $(t_{\text{train}}, t_{\text{val}}]$, and $(t_{\text{val}}, t_{\text{test}}]$ for training, validation, and testing with $0 < t_{\text{train}} < t_{\text{val}} < t_{\text{test}}$. In the experiment, we set $t_{\text{train}} = 30$, $t_{\text{val}} = 40$, and $t_{\text{test}} = 50$. We train the proposed NMS on spatial discretizations of varying size $N = \{50, 75, 100\}$, which results in systems of dimension $\{101, 151, 201\}$, respectively. The performance of NMS is then compared to the performance of GFNNs, which are the current state-of-the-art in metriplectic learning.

Figure 5 shows the ground truth and predicted trajectories of the first and the last components of q and p for the case of $N = 50$, where it can be seen that the trajectories predicted by NMS are significantly

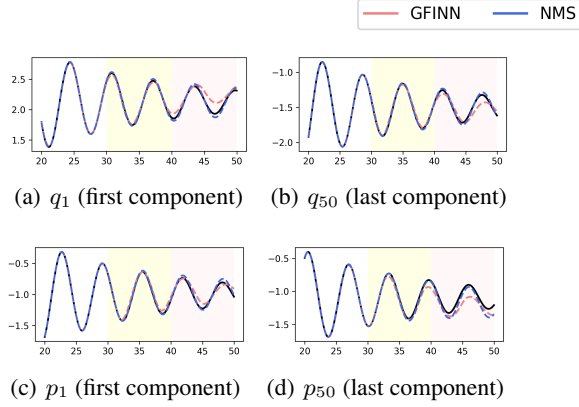


Figure 5: The ground-truth and predicted trajectories of positions and momenta for the thermoelastic rod are depicted with $N = 50$.




Figure 6: MSEs (in the \log_{10} scale) of predictions obtained by NMS and GFINN for varying N .

better aligned with the ground truth trajectories compared to those of GFINN, particularly in the extrapolation regime ($40 < t < 50$). Figure 6 shows the results of training GFINN and NMS on thermoelastic rod systems with varying N , where GFINN and NMS instances with comparable model sizes are considered. The results suggest that the proposed NMS is parameter efficient, supporting the theoretical conclusions in Section 3.

G SCALING STUDY

To compare the scalability of the proposed NMS architecture design with existing architectures, different realizations of GNODE, GFINN, and NMS are generated by varying the dimension of the state variables, $n = \{1, 5, 10, 15, 20, 30, 50\}$. The specifications of these models (i.e., hyperparameters) are set so that the number of model parameters is kept similar between each method for smaller values of n . For example, for $n = 1, 5$ the number of model parameters is $\sim 20,000$ for each architecture. The results in Figure 7(a) confirm that GNODE scales cubically in n while both GFINN and NMS scale quadratically. Note that only a constant scaling advantage of NMS over GFINN can be seen from this plot, since r is fixed during this study.

It is also worthwhile to investigate the computational timings of these three models. Considering the same realizations of the models listed above, i.e., the model instances for varying $n = \{1, 5, 10, 15, 20, 30, 50\}$, 1,000 random samples of states $\{\mathbf{x}^{(i)}\}_{i=1}^{1,000}$ are generated. These samples are then fed to the dynamics function $\mathbf{L}(\mathbf{x}^{(i)})\nabla E(\mathbf{x}^{(i)}) + \mathbf{M}(\mathbf{x}^{(i)})\nabla S(\mathbf{x}^{(i)})$ for $i = 1, \dots, 1000$, and the computational wall time of the function evaluation via PyTorch’s profiler API is measured. The results of this procedure are displayed in Figure 7(b). Again, it is seen that the proposed NMSs require less computational resources than GNODEs and GFINNs.

H DIFFUSION MODEL FOR UNOBSERVED VARIABLES

Recent work in Lim et al. (2023) suggests the benefits of performing time-series generation using a diffusion model. This Appendix describes how this technology is used to generate initial conditions for the unobserved NMS variables in the experiments corresponding to Table 3. More precisely, we describe how to train a conditional diffusion model which generates values for unobserved variables \mathbf{x}^u given values for the observed variables \mathbf{x}^o .

Training and sampling: Recall that diffusion models add noise with the following stochastic differential equation (SDE):

$$d\mathbf{x}(t) = \mathbf{f}(t, \mathbf{x}(t))dt + g(t)d\mathbf{w}, \quad t \in [0, 1],$$

where $\mathbf{w} \in \mathbb{R}^{\dim(\mathbf{x})}$ is a multi-dimensional Brownian motion, $\mathbf{f}(t, \cdot) : \mathbb{R}^{\dim(\mathbf{x})} \rightarrow \mathbb{R}^{\dim(\mathbf{x})}$ is a vector-valued drift term, and $g : [0, 1] \rightarrow \mathbb{R}$ is a scalar-valued diffusion function.

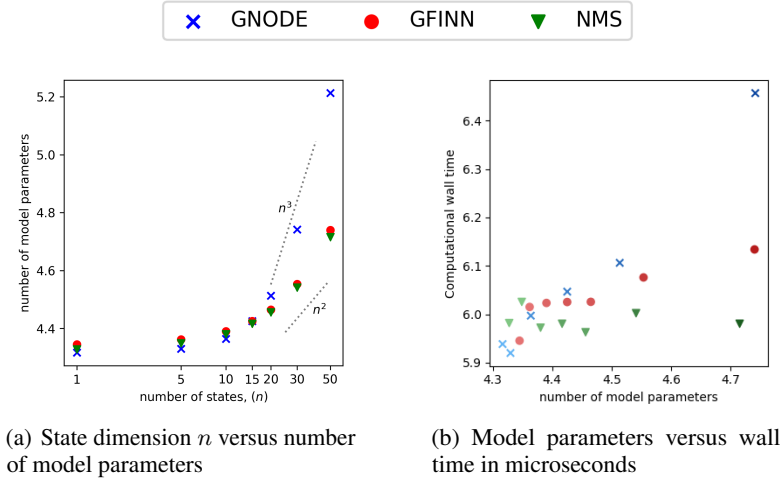


Figure 7: A study of the scaling behavior of GNODE, GFINN, and NMS.

For the forward SDE, there exists a corresponding reverse SDE:

$$d\mathbf{x}(t) = [\mathbf{f}(t, \mathbf{x}(t)) - g^2(t)\nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t))]dt + g(t)d\bar{\mathbf{w}},$$

which produces samples from the initial distribution at $t = 0$. This formula suggests that if the score function, $\nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t))$, is known, then real samples from the prior distribution $p(\mathbf{x}) \sim \mathcal{N}(\mu, \sigma^2)$ can be recovered, where μ, σ vary depending on the forward SDE type.

In order for a model M_θ to learn the score function, it has to optimize the following loss:

$$L(\theta) = \mathbb{E}_t\{\lambda(t)\mathbb{E}_{\mathbf{x}(t)}[\|M_\theta(t, \mathbf{x}(t)) - \nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t))\|_2^2]\},$$

where t is uniformly sampled over $[0, 1]$ with an appropriate weight function $\lambda(t) : [0, 1] \rightarrow \mathbb{R}$. However, using the above formula is computationally prohibitive. Thanks to Vincent (2011), this loss can be substituted with the following denoising score matching loss:

$$L^*(\theta) = \mathbb{E}_t\{\lambda(t)\mathbb{E}_{\mathbf{x}(0)}\mathbb{E}_{\mathbf{x}(t)|\mathbf{x}(0)}[\|M_\theta(t, \mathbf{x}(t)) - \nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t)|\mathbf{x}(0))\|_2^2]\}.$$

Since score-based generative models use an affine drift term, the transition kernel $p(\mathbf{x}(t)|\mathbf{x}(0))$ follows a certain Gaussian distribution Särkkä and Solin (2019), and therefore the gradient term $\nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t)|\mathbf{x}(0))$ can be analytically calculated.

Experimental details On the other hand, the present goal is to generate unobserved variables \mathbf{x}^u given values for the observed variables $\mathbf{x}^o = (\mathbf{q}, \mathbf{p})$, i.e., conditional generation. Therefore, our model has to learn the conditional score function, $\nabla_{\mathbf{x}^u(t)}\log p(\mathbf{x}^u(t)|\mathbf{x}^o)$. For example, in the damped nonlinear oscillator case, $S(t)$ is initialized as a perturbed $t \in [0, 1]$, from which the model takes the concatenation of $\mathbf{q}, \mathbf{p}, S(t)$ as inputs and learns conditional the score function $\nabla_{S(t)}\log(S(t)|\mathbf{q}, \mathbf{p})$.

For the experiments in Table 3, diffusion models are trained to generate \mathbf{x}^u variables on three benchmark problems: the damped nonlinear oscillator, two gas containers, and thermolastic double pendulum. On each problem, representative parameters such as mass or thermal conductivity are varied, with the total number of cases denoted by N . Full trajectory data of length T is then generated using a standard numerical integrator (e.g., dopri5), before it is evenly cut into $\lfloor T/L \rfloor$ pieces of length L . Let V, U denote the total number of variables and the number of unobserved variables, respectively. It follows that the goal is to generate U unobserved variables given $V - U$ observed ones, i.e., the objective is to generate data of shape $(NT/L, L, U)$ conditioned on data of shape $(NT/L, L, V - U)$. After the diffusion model has been trained for this task, the output data is reshaped into size (N, T, U) , which is used to initialize the NMS model. Note that the NODE and GNODE methods compared to NMS in Table 3 use full state information for their training, i.e., $\mathbf{x}^u = \emptyset$ in these cases, making it comparatively easier for these methods to learn system dynamics.

As in other diffusion models e.g. Song et al. (2020), a U-net architecture is used, modifying 2-D convolutions to 1-D ones and following the detailed hyperparameters described in Song et al. (2020). Note the following *probability flow* ODE seen in Song et al. (2020):

$$d\mathbf{x}(t) = \left[\mathbf{f}(t, \mathbf{x}(t)) - \frac{1}{2}g^2(t)\nabla_{\mathbf{x}(t)}\log p(\mathbf{x}(t)) \right] dt,$$

Although models trained to mimic the probability flow ODE do not match the performance of the forward SDE's result in the image domain, the authors of Lim et al. (2023) observe that the probability flow ODE outperforms the forward SDE in the time-series domain. Therefore, the probability flow ODE is used with the default hyperparameters of Song et al. (2020).