

HYBRID ENERGY-BASED MODELS FOR PHYSICAL AI: PROVABLY STABLE IDENTIFICATION OF PORT-HAMILTONIAN DYNAMICS

✉ **Simone Betteti**

RIAS Lab
The Italian Institute of
Artificial Intelligence for Industry
Turin, 10129, IT
simone.betteti[at]ai4i.it

✉ **Luca Laurenti***

RIAS Lab
The Italian Institute of
Artificial Intelligence for Industry
Turin, 10129, IT
luca.laurenti[at]ai4i.it

April 2, 2026

ABSTRACT

Energy-based models (EBMs) implement inference as gradient descent on a learned Lyapunov function, yielding interpretable, structure-preserving alternatives to black-box neural ODEs and aligning naturally with physical AI. Yet their use in system identification remains limited, and existing architectures lack formal stability guarantees that globally preclude unstable modes. We address this gap by introducing an EBM framework for system identification with stable, dissipative, absorbing invariant dynamics. Unlike classical global Lyapunov stability, absorbing invariance expands the class of stability-preserving architectures, enabling more flexible and expressive EBMs. We extend EBM theory to nonsmooth activations by establishing negative energy dissipation via Clarke derivatives and deriving new conditions for radial unboundedness, exposing a stability-expressivity tradeoff in standard EBMs. To overcome this, we introduce a hybrid architecture with a dynamical visible layer and static hidden layers, prove absorbing invariance under mild assumptions, and show that these guarantees extend to port-Hamiltonian EBMs. Experiments on metric-deformed multi-well and ring systems validate the approach, showcasing how our hybrid EBM architecture combines expressivity with sound and provable safety guarantees by design.

1 INTRODUCTION

Energy-based models (EBMs) are a class of machine-learning architectures in which neural dynamics arise as the gradient flow of a Hopfield-type Lyapunov function. Rooted in the pioneering work of [Cohen and Grossberg \[1983\]](#) and [Hopfield \[1982, 1984\]](#), EBMs encode meaningful data patterns as stable equilibria of a dissipative dynamical system, ensuring convergence towards stored representations through gradient descent on an energy landscape. However, the classical Hopfield construction offered limited expressive capacity for modern learning tasks, with the number of retrievable patterns scaling as $N/\log(N)$ for networks of size N [[McEliece et al., 1987](#)]. Recent generalizations of Hopfield networks [[Krotov and Hopfield, 2020](#), [Hoover et al., 2023](#)] have dramatically expanded their applicability. Modern EBMs preserve the foundational energy-based structure while exhibiting formal parallels with Transformer architectures [[Ramsauer et al., 2021](#)], enabling an exponential scaling of capacity in N [[Demircigil et al., 2017](#)]. Their alignment with contemporary neural architectures has facilitated successful deployments across standard machine-learning domains such as image classification [[Hoover et al., 2024](#)], image generation [[Ambrogioni, 2024](#), [Park et al., 2023](#)], and in-context learning [[Wu et al., 2025](#)]. Canonically, learning in *energy-based models* relies on optimizing the network parameters to map relevant data-class centroids to stable equilibria and their associated basin of attraction in the *energy* landscape.

Despite their dynamical foundation, EBMs have rarely been deployed for system identification, a domain where the underlying objective of capturing the evolution of nonlinear systems from data is inherently dynamical. System

*L.L. is also with the Delft Center for Systems and Control, TU Delft, Delft, 2600 AA, The Netherlands.

identification [Åström and Eykhoff, 1971, Ljung, 2010] sits at the interface of control theory and optimization and has progressively integrated machine-learning components [Chen et al., 2018] to accommodate increasingly complex real-world systems. This emerging direction, now widely referred to as *physical AI*, has driven notable progress in robotics [Perrusquía et al., 2022], power-grid modeling [Chakraborty et al., 2022], and industrial manufacturing [Denno et al., 2018]. Yet, the safety-critical nature of these domains demands not only expressive models, but also structural guarantees of stability and robustness [Dai et al., 2021]. Recent efforts with dissipative neural networks and stable neural ODEs address this need by enforcing parameter constraints [Drgona et al., 2022, Xu and Sivarajan, 2023]. Notably, recent work on learning Port-Hamiltonian dynamics [Massaroli et al., 2020, Roth et al., 2025] has centered on learning optimality and empirical benchmarking, with limited attention to stability guarantees or architectural considerations. In summary, existing neural identifiers either lack formal stability guarantees outside the energy and dissipative frameworks or require strong smoothness and convexity assumptions.

Contributions. This work introduces an energy-based system identification framework that reconciles stability with expressive modeling. Its main contributions are:

- a generalized Lyapunov stability analysis for EBM with nonsmooth activations, establishing energy dissipation via Clarke derivatives and revealing a structural stability-expressivity trade-off;
- a *hybrid EBM architecture* with dissipative, absorbing invariant visible-layer dynamics, guaranteeing bounded trajectories while avoiding the restrictive conditions imposed by fully recurrent EBM, and static hidden layer maps for fast and efficient inference;
- a Port-Hamiltonian extension enabling identification of nonlinear dynamics under state-dependent metrics and rotational perturbations, validated on metric-deformed multi-well and ring systems.

Section II presents the necessary preliminaries. Section III introduces EBM and formulates the problem. Section IV extends classical C^2 dissipation results to locally Lipschitz activations and establishes new conditions for radial unboundedness. Section V develops the hybrid architecture and proves absorbing invariance, including its Port-Hamiltonian generalization. Section VI reports numerical validation, and Section VII concludes with future directions.

2 PRELIMINARIES

Notation. $\mathbf{1}_d$ denotes the d -dimensional vector with all ones, $\mathbf{0}_d$ the d -dimensional vector with all zeros, while $\mathcal{I}_d$ denotes the identity matrix in $\mathbb{R}^{d \times d}$. We denote compact subsets of $\mathbb{R}^d$ as sets $\mathcal{B} \subset \mathbb{R}^d$, and their boundary as $\partial\mathcal{B}$. For two real vectors x, y of the same dimension, $x^\top y$ denotes the standard inner product. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$; for $f \in C^k(\mathbb{R}^d)$, the function is k -times continuously differentiable. For $f \in C^2(\mathbb{R}^d)$, the gradient of f is denoted as ∇f , and the Hessian as $D(\nabla f) = D^2 f$. The partial derivative of f with respect to the variable x_i is denoted as $\partial f / \partial x_i$. Θ bounds the growth of $f \sim \Theta(g)$ through the existence of constants $a, b > 0$ such and a function g such that $a g(x) < f(x) < b g(x)$. Given functions $g : \mathbb{R}^n \rightarrow \mathbb{R}^d$ and $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we denote the composition of the two functions as $f \circ g(y)$, for $y \in \mathbb{R}^n$. The abbreviation a.e. stands for almost everywhere, that is for all $\mathcal{B} \subset \mathbb{R}^d$ except zero measure sets. Given a matrix $A \in \mathbb{R}^{d \times d}$, we denote with $A^\top$ its transpose. In case A is symmetric, $\lambda_{\min}(A), \lambda_{\max}(A) \in \mathbb{R}$ denote its minimum and its maximum eigenvalue. We denote with $B_r(x) \subset \mathbb{R}^d$ the ball of radius $r > 0$ and center $x \in \mathbb{R}^d$.

2.1 Lie and Clarke’s derivatives

Definition 1 (Lie derivative). Let $X : \mathbb{R}^N \rightarrow \mathbb{R}^N$ be a locally Lipschitz continuous vector field and let $\Phi_X : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^N$ denote the associated local flow. For a smooth function $g : \mathbb{R}^N \rightarrow \mathbb{R}$, the Lie derivative of g along X is defined as

$$\mathcal{L}_X g(x) = \frac{d}{dt} \Big|_{t=0} (g \circ \Phi_X^t)(x), \quad (1)$$

and coincides with the directional derivative

$$\mathcal{L}_X g(x) = \nabla_X g(x) = \nabla_x g(x)^\top X(x). \quad (2)$$

Lie derivatives [Lee, 2012, Chapters 3, 9] will be central for the characterization of the negative definiteness of the *energy* function along the trajectories generated by the EBM dynamics. We extend the notion of derivatives to nonsmooth functions by recalling Clarke’s generalized gradient and directional derivative [Clarke, 1975], which apply to any locally Lipschitz function.

Definition 2 (Generalized gradient). *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a locally Lipschitz function. The generalized gradient of f at $x \in \mathbb{R}^N$ is denoted as $\partial f(x)$ and is the convex hull of the set of limits*

$$\lim_{k \rightarrow +\infty} \nabla f(x + h_k), \quad (3)$$

with $h_k \xrightarrow{k \rightarrow +\infty} \mathbb{0}_N$ and f differentiable at $x + h_k \in \mathbb{R}^N$ for all $k \in \mathbb{N}$.

Definition 3 (Generalized directional derivative). *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a locally Lipschitz function and take $v \in \mathbb{R}^N$. Then the Clarke's generalized directional derivative at $x \in \mathbb{R}^N$ is defined as*

$$f^\circ(x; v) = \lim_{h \rightarrow \mathbb{0}_N} \sup_{\delta \rightarrow 0} \frac{f(x + h + \delta v) - f(x + h)}{\delta} \quad (4)$$

$$= \max_{\xi \in \partial f(x)} \xi^\top v. \quad (5)$$

In particular, combining (2) with (5), the Clarke Lie derivative of a locally Lipschitz function f along a vector v is given by

$$\mathcal{L}_v f(x) := f^\circ(x; v) = \max_{\xi \in \partial f(x)} \xi^\top v, \quad (6)$$

extending the classical smooth definition to the nonsmooth setting.

2.2 Conjugate (Legendre) transform

Definition 4 (Conjugate transform). *Let $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}$ be proper and convex. The conjugate transform of $\mathcal{F}$ (also referred to as Legendre transform) is defined as*

$$\mathcal{F}^*(y) = \sup_{x \in \mathbb{R}^N} x^\top y - \mathcal{F}(x). \quad (7)$$

Standard results from convex analysis [Bauschke and Combettes, 2017, Chapter 13] provide information on the convexity and asymptotic growth of the conjugate transform. For polynomially growing convex functions $\mathcal{F}(x) \sim \Theta(\|x\|^q)$, the conjugate $\mathcal{F}^*$ is also convex and $\mathcal{F}^*(y) \sim \Theta(\|y\|^{\frac{q}{q-1}})$.

Remark 1 (Smooth proper convex functions). *If $\mathcal{F} \in C^1(\mathbb{R}^N)$, then any maximizer of (7) satisfies the first-order optimality condition*

$$\begin{aligned} \mathbb{0}_N &= \nabla_x [x^\top y - \mathcal{F}(x)] \\ &= y - \nabla \mathcal{F}(x). \end{aligned} \quad (8)$$

It then follows that $\mathcal{F}^*(y)|_{y=\nabla \mathcal{F}(x)} = x^\top \nabla \mathcal{F}(x) - \mathcal{F}(x)$. Combining this with the growth of the conjugate $\mathcal{F}^*$ yields

$$\begin{aligned} \mathcal{F}^*(y)|_{y=\nabla \mathcal{F}(x)} &\sim \Theta(\|y\|_{|y=\nabla \mathcal{F}(x)}}^{q/(q-1)}) \\ &\sim \Theta(\|x\|^{(q-1)\frac{q}{q-1}}) = \Theta(\|x\|^q), \end{aligned} \quad (9)$$

where the final equivalence uses the fact that the gradient of a q -coercive convex function grows $(q-1)$ -polynomially in $\|x\|$.

3 PROBLEM STATEMENT

EBMs were initially introduced in [Krotov and Hopfield, 2016, 2020] as fully recurrent architectures, where all units interact with one another. Subsequent work specialized EBMs into modular, layered architectures [Hoover et al., 2022, Kozachkov et al., 2025] to improve learning efficiency and compatibility with standard MLP-based frameworks. In this paper, we consider EBMs as defined by the autonomous nonlinear dynamical systems in Definition 5 below, with parametric interactions between units, and that can be trained for machine-learning tasks using standard optimization techniques [Bottou et al., 2018].

Definition 5 (Energy-based models (EBMs)). *Let $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}$ be a proper, convex, and C^1 function, and define $\Psi(x) = \nabla \mathcal{F}(x)$ to be the activation function of the energy-based model. Further, let $W \in \mathbb{R}^{N \times N}$ be the symmetric matrix of parameters and $b \in \mathbb{R}^N$ a bias vector. Then, an energy-based model is defined as a dynamical system*

$$\dot{x} = -x + W\Psi(x) + b, \quad (10)$$

and has associated energy

$$E(x) = -\frac{1}{2}\Psi(x)^\top W\Psi(x) + (x - b)^\top \Psi(x) - \mathcal{F}(x). \quad (11)$$

In classical EBM theory, stability of the dynamics (10) is ensured by proving that the energy function $E(x)$ in (11) is dissipative along the dynamics (10) and radially unbounded. As shown in Section IV, this Lyapunov-based viewpoint imposes restrictive growth conditions on the activation functions, leading to a *stability-expressivity trade-off*. In contrast, this paper relies on a different notion of stability for system 5, which leverages absorbing invariant sets and, as we will show, allows for more expressive *stable-by-design* architectures. An absorbing invariant set, in the sense of Definition 6, is a region that all trajectories are eventually drawn into (absorbing) and from which none can escape once inside (invariant).

Definition 6. (*Absorbing invariant set*) Let $\mathcal{D} \subset \mathbb{R}^N$ be a compact set and $X : \mathbb{R}^N \rightarrow \mathbb{R}^N$ a locally Lipschitz vector field. Consider Φ_X^t the flow induced by X . Then the set $\mathcal{D}$ is

- (i) *absorbing* if $\lim_{t \rightarrow +\infty} \Phi_X^t(x) \in \mathcal{D}$ for all $x \in \mathbb{R}^N$;
- (ii) *invariant* if $X(x)^\top \eta(x) \leq 0$ for all $x \in \partial\mathcal{D}$, with $\eta(x) \in \mathbb{R}^N$ outer normal to $\partial\mathcal{D}$ at x .

An absorbing invariant set as in Definition 6 provides a natural *safe set* for the EBM dynamics (10), confining trajectories to a predictable region of the state space. Dissipativity further enforces non-increasing energy, ruling out divergence and steering trajectories toward energetically stable subsets of $\mathcal{D}$, including any equilibria contained within it. Neural models that guarantee such bounded and predictable behavior are essential for deploying learning-based identifiers in safety-critical settings. Formally, in this paper we consider the following problem:

Problem 1. Given samples $(x, f^*(x)) \in \mathcal{X}$ from an unknown autonomous system $\dot{x} = f^*(x)$, our objective is to identify the underlying dynamics through a EBM (10) $\dot{x} = f(x; W)$ for which there exists an invariant absorbing compact set $\mathcal{D}$ containing the equilibrium points of f^* .

In the following sections we show that Problem 1 can be addressed by relaxing the assumptions typically required to guarantee Lyapunov asymptotic stability for EBMs [Krotov and Hopfield, 2016]. Our approach is non-trivial: it removes the stringent conditions imposed by classical Lyapunov arguments while still yielding a *stable-by-design* neural architecture suitable for reliable system identification. Existing EBM formulations either remain bounded to Lyapunov-based constraints or rely on strongly convex parameterizations, offering far less flexibility than the framework proposed here. In Section IV, we extend standard dissipativity arguments to nonsmooth activations and show that Lyapunov asymptotic stability enforces a sublinearity condition that restricts expressivity. Section V introduces the hybrid EBM architecture, which circumvents this limitation by requiring bounded activation only in the first hidden layer. This boundedness confines the contributions of all deeper hidden layers to compact sets, allowing us to construct a conservative absorbing invariant set $\mathcal{D}$ for the hybrid dynamics.

4 LYAPUNOV STABILITY AND THE STABILITY-EXPRESSIVITY TRADE-OFF

Classical EBM analyses assume $\mathcal{F} \in C^2(\mathbb{R}^N)$ so that the Hessian $D^2\mathcal{F}$ is well defined and positive semidefinite, ensuring energy dissipation. However, practical architectures often use activations that are only locally Lipschitz and thus differentiable almost everywhere. We therefore extend the standard C^2 argument by employing Clarke generalized derivatives [Clarke, 1975], which recover the same negative-definiteness property of the Lie derivative without requiring twice differentiability. For simplicity we take $b = 0_N$, though the results extend to the general case. We begin by establishing dissipativity for nonsmooth activations in Proposition 7.

Proposition 7 (Negative definiteness of the energy derivative). Let $\mathcal{F} \in C^1(\mathbb{R}^N)$ and let $f_E(x) = -x + W\Psi(x)$ be the vector field associated to the EBM dynamics (10). Then for all $x \in \mathbb{R}^N$ and $E(x)$ defined in (11)

$$\mathcal{L}_{f_E} E(x) \leq 0 \quad \text{a.e.} \quad (12)$$

Proof. Since $\Psi = \nabla\mathcal{F}$ is locally Lipschitz, it is differentiable a.e., and at such points

$$\begin{aligned} \nabla E(x) &= -D\Psi(x)W\Psi(x) + D\Psi(x)x \\ &= -D\Psi(x)f_E(x), \end{aligned} \quad (13)$$

where in the first passage we have exploited the symmetry of $D\Psi(x)$ and W . By Rademacher theorem [Evans and Gariepy, 2015, Chapter 6], $D\Psi(x)$ is defined a.e., and by the convexity of $\mathcal{F}$ every element $\Sigma(x) \in \partial(D\Psi)(x)$ of the Clarke generalized Hessian is positive semi-definite, i.e. $\Sigma(x) \succeq 0$ for all $x \in \mathbb{R}^N$. The Clarke generalized gradient of the energy is then

$$\partial E(x) = \{-\Sigma(x)f_E(x) : \Sigma(x) \in \partial(D\Psi)(x)\}. \quad (14)$$

The Clark Lie derivative is then

$$\begin{aligned}
\mathcal{L}_{f_E} E(x) &= \max_{v \in \partial E(x)} v^\top f_E(x) \\
&= \max_{\Sigma(x) \in \partial(D\Psi)(x)} -f_E(x)^\top \Sigma(x) f_E(x) \\
&= \max_{\Sigma(x) \in \partial(D\Psi)(x)} -\|f_E(x)\|_{\Sigma(x)^{1/2}}^2 \leq 0.
\end{aligned} \tag{15}$$

□

The classical EBM literature focuses solely on dissipation and convergence, overlooking the crucial issue of radial unboundedness of the energy. Proposition 8 below provides necessary and sufficient conditions for radial unboundedness in a fully connected EBM architecture, without which trajectories may escape to infinity, producing pathological behaviors such as diverging activity or exploding gradients during training. We consider the realistic case in which $W \in \mathbb{R}^{N \times N}$, learned from data-driven optimization, is not necessarily negative definite and has positive eigenvalues driving the quadratic term of the energy to $-\infty$.

Proposition 8 (Radial unboundedness of the energy derivative). *Let $W \in \mathbb{R}^{N \times N}$ with no zero entries and such that $\lambda_{\max}(W) > 0$. Let $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}$ be a proper, convex and C^1 function. Suppose that $\mathcal{F}(x) \sim \Theta(\|x\|^q)$ with $q \in]1, +\infty[$. Then*

$$\liminf_{\|x\| \rightarrow +\infty} E(x) = +\infty \iff q < 2. \tag{16}$$

Proof. Let $q \in]1, +\infty[$ and consider the decomposition of (11) as $E(x) = E_{Quad}(x) + E_{Ham}(x)$.

- (i) $E_{Quad}(x) = -\frac{1}{2}\Psi(x)^\top W \Psi(x)$: since $\mathcal{F}(x) \sim \Theta(\|x\|^q)$, then we have that $\Psi(x) = \nabla \mathcal{F}(x) \sim \Theta(\|x\|^{q-1})$. Since W has no zero entries, the greatest growing entries of the quadratic product grow as $|z|^{2(q-1)}$ for $z \in \mathbb{R}$, and in fact we have

$$|\nabla \mathcal{F}(x)^\top W \nabla \mathcal{F}(x)| \leq \lambda_{\max}(W) \|\nabla \mathcal{F}(x)\|_2^2.$$

Therefore, $-E_{Quad}(x) \sim \Theta(\|x\|^{2(q-1)})$.

- (ii) $E_{Ham}(x) = x^\top \Psi(x) - \mathcal{F}(x)$: from the convex duality result (Remark 1), $E_{Ham}(x) \sim \Theta(\|x\|^q)$.

Consequently, the energy $E(x) = E_{Quad}(x) + E_{Ham}(x)$ is radially unbounded if

$$\liminf_{\|x\| \rightarrow +\infty} -\|\nabla \mathcal{F}(x)\|_2^2 + x^\top \nabla \mathcal{F}(x) - \mathcal{F}(x) = +\infty. \tag{17}$$

This is equivalent to require that the growth q of E_{Ham} dominates the growth $2(q-1)$ of E_{Quad} , which happens iff

$$q > 2(q-1) \iff q < 2. \tag{18}$$

□

Proposition 8 implies that fully recurrent EBMs are fundamentally limited: stability requires sublinear activations $q \in (1, 2)$, which does not make possible the use of high-degree polynomial activations advocated in [Krotov and Hopfield, 2016, 2020]. Even multilayer feedforward EBMs [Hoover et al., 2022] only soften this restriction, still requiring sublinear activations in every second layer to maintain radial unboundedness. Practical models bypass these constraints with heavy normalization throughout the network. In contrast, the hybrid architecture introduced next achieves stability via invariance with normalization applied solely to the first hidden layer, fully restoring expressivity while guaranteeing bounded trajectories.

5 MULTI-LAYER EBMs AND ABSORBING INVARIANCE

The main contribution of our work is the introduction of a hybrid architecture that preserves both computational expressivity and stability of the EBM by requiring boundedness of the activation only in the first hidden layer, and allowing for arbitrary activations in all the deeper layers. We start by restricting our attention to layered architectures, a standard choice in applications, where each hidden layer interacts only with its immediate neighbors.

Definition 9 (Multi-layer EBMs). Let $L \in \mathbb{N}$ be the number of layers, and denote $N_h \in \mathbb{N}$ the width of the h -layer, with $N = \sum_{h=1}^L N_h$. For each layer, let $x_{(h)} \in \mathbb{R}^{N_h}$ be its state, and define a family $\{\mathcal{F}_h\}_{h=1}^L$ of proper, convex, and C^1 functions such that

$$\mathcal{F}(x) = \sum_{h=1}^L \mathcal{F}_h(x_{(h)}). \quad (19)$$

Define the interaction matrix $W \in \mathbb{R}^{N \times N}$ as the block tridiagonal matrix

$$W = \begin{pmatrix} \mathbb{0}_{N_1 \times N_1} & W_{12} & \mathbb{0}_{N_1 \times N_3} & \cdots & \mathbb{0}_{N_1 \times N_L} \\ W_{12}^\top & \mathbb{0}_{N_2 \times N_2} & W_{23} & \ddots & \mathbb{0}_{N_2 \times N_L} \\ \mathbb{0}_{N_2 \times N_1} & W_{23}^\top & \mathbb{0}_{N_3 \times N_3} & \ddots & \vdots \\ \vdots & \cdots & \ddots & \ddots & W_{(L-1)L} \\ \mathbb{0}_{N_L \times N_1} & \cdots & \cdots & W_{(L-1)L}^\top & \mathbb{0}_{N_L \times N_L} \end{pmatrix}. \quad (20)$$

Then the dynamics of each layer of the EBM architecture are

$$\dot{x}_{(1)} = -x_{(1)} + W_{12}\Psi_2(x_{(2)}), \quad (21)$$

$$\begin{aligned} \dot{x}_{(h)} = & -x_{(h)} + W_{h(h-1)}\Psi_{h-1}(x_{(h-1)}) \\ & + W_{h(h+1)}\Psi_{h+1}(x_{(h+1)}), \quad h = 2, \dots, L-1 \end{aligned} \quad (22)$$

$$\dot{x}_{(L)} = -x_{(L)} + W_{L(L-1)}\Psi_{L-1}(x_{(L-1)}). \quad (23)$$

For system-identification tasks based on 0-step encoding-decoding maps $x \mapsto f(x)$, only the visible layer $x_{(1)}$ from 21 represents measured or reconstructed physical variables. A full multilayer EBM, however, evolves in a higher-dimensional state and converges to an equilibrium determined by all layers, creating two practical difficulties: (i) **temporal alignment**, i.e., determining when the visible state should be matched with data; and (ii) **backpropagation through depth**, since hidden layers introduce recurrent dependencies that complicate gradient computation. To preserve the dynamical structure of the visible layer while removing internal recurrences, we adopt a hybrid model in which hidden-layer activity is computed through feedforward maps, whereas the visible layer retains its full EBM dynamics.

Definition 10 (Hybrid EBM architecture). Let $\{\phi_h^f\}_{h=2}^L$ be a family of continuous functions, $\phi_h^f : \mathbb{R}^{N_{h-1}} \rightarrow \mathbb{R}^{N_h}$, and define the state of each hidden layer of the multi-layer EBM in Definition 9 via the feedforward map

$$x_{(h-1)} \mapsto \phi_h^f(x_{(h-1)}) = x_{(h)}. \quad (24)$$

The projected visible-layer dynamics in (21) are then defined as the projected energy gradient:

$$\begin{aligned} \dot{x}_{(1)} = & -\nabla_{x_{(1)}} E(x_{(1)}, \phi_2^f(x_{(1)}), \dots, \phi_L^f(x_{(L-1)})) \\ = & -\sum_{h=1}^L \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \left[-x_{(h)} + \sum_{k=1}^L W_{hk}\Psi_k(x_{(k)}) \right] \\ = & f_{E^1}(x_{(1)}), \end{aligned} \quad (25)$$

where the interaction matrices $\{W_{hk} \in \mathbb{R}^{N_h \times N_k}\}_{h,k=1}^L$ are defined in (20).

Remark 2 (Implementability of hybrid EBM architecture). The formulation in Definition 10 leads to a highly convenient implementation: once the convex primitives $\{\mathcal{F}_h\}_h$, interaction matrices $\{W_{hk}\}_{h,k}$, and biases $\{b_h\}_h$ are specified, the EBM is fully determined by the energy function $E(x)$, and the neural network action is given by $-\nabla_{x_{(1)}} E(x)$.

Theorem 11 characterizes sufficient conditions under which the hybrid visible-layer dynamics admit a compact absorbing invariant set. The assumptions reflect typical architectural choices: linear visible activation, bounded hidden activations, and feedforward connections consistent with the EBM structure.

Theorem 11 (Absorbing invariance of hybrid EBMs). Let the hybrid EBM from Definition 10 satisfy:

- (i) **linear visible layer**: $\Psi_1(x) = x$;
- (ii) **bounded first hidden layer**: $\Psi_2 : \mathbb{R}^{N_2} \rightarrow K \subset \mathbb{R}^{N_2}$;
- (iii) **feedforward hidden maps**: for $h = 2, \dots, L$

$$\phi_h^f(x_{(h-1)}) = W_{h(h-1)}\Psi_{h-1}(x_{(h-1)}). \quad (26)$$

Then there exists $r > 0$ such that the ball $B_r(0) \subset \mathbb{R}^N$ is f_{E^1} -absorbing invariant.

Proof. Define the **feedback hidden maps**

$$\phi_h^b(x_{(h+1)}) = W_{h(h+1)} \Psi_{(h+1)}(x_{(h+1)}). \quad (27)$$

Using assumption (iii), the dissipation term $-x_{(h)}$ cancels the feedforward term for every hidden layer. Together with (i), the visible dynamics (25) become

$$f_{E^1}(x_{(1)}) = -x_{(1)} + \sum_{h=1}^{L-1} \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \phi_h^b(x_{(h+1)}), \quad (28)$$

where we have used $\partial \Psi_1 / \partial x_{(1)} \equiv \mathcal{I}_{N_1}$ for the dissipation term of the visible layer.

From hypothesis (ii) notice that $\Psi_2(\phi_2^f(x_{(1)})) \in K$ for all $x_{(1)} \in \mathbb{R}^{N_1}$, and consequently

$$\phi_3^f(x_{(2)}) \in K_3 \subset \mathbb{R}^{N_3}. \quad (29)$$

Since $\{\phi_h^f\}_{h=2}^L$ are all continuous functions

- **Boundedness of the feedforward maps ϕ_h^f :**

For all $h = 3, \dots, L$ there exists $K_h \subset \mathbb{R}^{N_h}$ such that the feedforward map

$$\begin{aligned} x_{(h)} &= \phi_h^f(x_{(h-1)}) \\ &= \phi_h^f \circ \dots \circ \phi_2^f(x_{(1)}) \in K_h, \end{aligned} \quad (30)$$

for all $x_{(1)} \in \mathbb{R}^{N_1}$.

- **Boundedness of the feedback maps ϕ_h^b :**

By definition (27), $\{\phi_h^b\}_{h=1}^{L-1}$ are also continuous. Exploiting the boundedness of the feedforward maps, for all $h = 2, \dots, L-1$ there exist $\tilde{K}_h \subset \mathbb{R}^{N_h}$ such that the feedback map

$$\begin{aligned} \phi_h^b(x_{(h+1)}) &= \phi_h^b(\phi_{h+1}^f(x_{(h)})) \\ &= \phi_h^b(\phi_{h+1}^f \circ \dots \circ \phi_2^f(x_{(1)})) \in \tilde{K}_h. \end{aligned} \quad (31)$$

Additionally, observe that $\phi_1^b(x_{(2)}) = W_{12} \Psi_2(x_{(2)})$, and from hypothesis (ii) we recover boundedness of the feedback maps for all $h = 1, \dots, L-1$.

- **Boundedness of the Jacobians $\partial \Psi_h / \partial x_{(1)}$:**

Since $\Psi_2 : \mathbb{R}^{N_2} \rightarrow K$ and $K \subset \mathbb{R}^{N_2}$ is compact, then necessarily $\|\partial \Psi_2 / \partial x_{(1)}\| < M$ uniformly for some $M > 0$. Since Ψ_h are locally Lipschitz, then $\partial \Psi_h / \partial x_{(1)}$ are bounded on every compact in $\mathbb{R}^{N_h}$ for $h \geq 3$.

Express $x_{(h)} = \phi_h^f \circ \dots \circ \phi_2^f(x_{(1)})$ and let

$$\gamma_h(x_{(1)}) = \left\| \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \phi_h^b(\phi_{h+1}^f(x_{(h)})) \right\|_2. \quad (32)$$

For all $h = 1, \dots, L-1$ there exists

$$\begin{aligned} \gamma_h &= \sup_{x_{(1)} \in \mathbb{R}^{N_1}} \gamma_h(x_{(1)}) \\ &= \max_{x_{(h)} \in K_h} \left\| \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \phi_h^b(\phi_{h+1}^f(x_{(h)})) \right\|_2. \end{aligned} \quad (33)$$

Define the radius

$$r = \sum_{h=1}^{L-1} \gamma_h. \quad (34)$$

and the ball $B_r(0) \in \mathbb{R}^{N_1}$, for any $x_{(1)} \in \partial B_r(0)$ ($x_{(1)}/r$ outer normal)

$$\begin{aligned}
x_{(1)}^\top f_{E^1}(x_{(1)}) &= \\
&= -\|x_{(1)}\|_2^2 + x_{(1)}^\top \sum_{h=1}^{L-1} \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \Phi_h^b(\Phi_{h+1}^f(x_{(h)})) \\
&\leq -r^2 + r \left\| \sum_{h=1}^{L-1} \frac{\partial \Psi_h}{\partial x_{(1)}}(x_{(h)}) \Phi_h^b(\Phi_{h+1}^f(x_{(h)})) \right\|_2 \\
&\leq -r^2 + r \left(\sum_{h=1}^{L-1} \gamma_h(x_{(1)}) \right) \\
&\leq -r^2 + r \left(\sum_{h=1}^{L-1} \gamma_h \right) = -r^2 + r^2 = 0.
\end{aligned} \tag{35}$$

Finally, observe that for any $R > r$ we have $x_{(1)}^\top f_{E^1}(x_{(1)}) < 0$ for all $x_{(1)} \in \partial B_R(0)$, and by continuity of f_{E^1} there exists a $\gamma_R > 0$ such that $x_{(1)}^\top f_{E^1}(x_{(1)}) < -\gamma_R$ uniformly in $\partial B_R(0)$. Therefore, for all $\delta t > 0$ there exists a $0 < \varepsilon < R - r$ such that if $x_0 \in \partial B_R(0)$, then $\Phi_{f_{E^1}}^{\delta t}(x_0) \in B_{R-\varepsilon}(0)$. Consequently, $\lim_{t \rightarrow +\infty} \Phi_{f_{E^1}}^{\delta t}(x) \in B_r(0)$ for all $x \in B_R(0)$, and taking $R \rightarrow +\infty$, we conclude. $\square$

Theorem 11 weakens the conditions required in Proposition 8 and enables the use of expressive architectures with arbitrary deep hidden-layer activations while preserving stability, thereby fulfilling the original aim of [Krotov and Hopfield, 2016] to broaden the admissible activation class for EBM. Corollary 12 below establishes that absorbing invariance and dissipativity of the projected visible-layer EBM dynamics (21) are robust under a Port-Hamiltonian transformations, ensuring that the stability guarantees derived for EBM carry over to a larger and practically relevant class of dynamics.

Corollary 12 (Absorbing invariance and dissipativity of Port-Hamiltonian EBMs). *Let $Q : \mathbb{R}^{N_1} \rightarrow \mathbb{R}^{N_1 \times N_1}$ be continuous and uniformly bounded, and assume $Q(x) + Q(x)^\top \succ 0$ for all $x \in \mathbb{R}^{N_1}$. Let $f_{E^1}(x)$ denote the EBM vector field (28), admitting an absorbing invariant ball $B_r(0)$ with radius $r > 0$ defined in (34). Then there exist $q_{\max} > q_{\min} > 0$ and $\rho := \frac{q_{\max}}{q_{\min}} r$ such that the Port-Hamiltonian EBM vector field*

$$\mathcal{H}_E(x) = Q(x)f_{E^1}(x), \tag{36}$$

is absorbing invariant on every ball $B_R(0)$ for all $R \geq \rho$. Furthermore,

$$\mathcal{L}_{\mathcal{H}_E} E(x) \leq 0 \quad \text{a.e.} \tag{37}$$

Proof. Define

$$q_{\min} = \inf_{x \in \mathbb{R}^{N_1} \setminus \{0\}} \frac{x^\top Q(x)x}{x^\top x}; \quad q_{\max} = \max_{x \in \mathbb{R}^{N_1}} \|Q(x)\|_2,$$

with $q_{\min} > 0$ by uniform positive definiteness and $0 < q_{\max} < +\infty$ by boundedness. Then for $x \in \partial B_R(0)$ with $R > r$ we have

$$\begin{aligned}
x^\top \mathcal{H}_E(x) &= x^\top Q(x)f_E(x) \\
&= x^\top Q(x) \left[-x + \sum_{h=1}^{L-1} \frac{\partial \Psi_h}{\partial x}(x_{(h)}) \Phi_h^b(x_{(h+1)}) \right] \\
&\leq -q_{\min} \|x\|_2^2 + q_{\max} r \|x\|_2 \\
&= -R(q_{\min} R - q_{\max} r).
\end{aligned} \tag{38}$$

The inner product is negative for any

$$R > \rho := \frac{q_{\max}}{q_{\min}} r, \tag{39}$$

and following the same approach as in 11, we conclude on the absorbing invariance. Finally, observe that at point of differentiability $\nabla E(x) = f_{E^1}(x)$ (28), and consequently the Clarke Lie derivative of E along $\mathcal{H}_E$ is

$$\begin{aligned}
\mathcal{L}_{\mathcal{H}_E} E(x) &= \max_{v \in \partial E(x)} v^\top \mathcal{H}_E(x) \\
&= \max_{v \in \partial E(x)} v^\top Q(x)v \leq 0,
\end{aligned} \tag{40}$$

by the positive definiteness of $Q(x) + Q(x)^\top$. $\square$

Dynamics	Train Loss	Test Loss
f_m	0.011	0.002
f_e	0.047	0.042

Table 1: Summary of train and test loss at the end of $n = 5000$ learning epochs for system identification of f_m and f_e using a Port-Hamiltonian EBM $\mathcal{H}_E$.

We have proven how absorbing invariance and dissipativity results for hybrid EBMs extends seamlessly to a broader class of Port-Hamiltonian [Van der Schaft, 2007] energy-based systems, where the dissipative gradient flow is composed with a state-dependent metric. In this setting, the EBM vector field is premultiplied by a continuous matrix field $Q(x)$ whose symmetric part remains uniformly positive definite, thereby preserving dissipativity while introducing geometry-dependent rotations. Port-Hamiltonian models capture dynamics evolving in non-Euclidean geometries, as well as systems subjected to rotational or gyroscopic perturbations, and play a central role in geometric control, robotics, and structure-preserving numerical schemes. Therefore, our proposed Port-Hamiltonian hybrid EBM architecture naturally provides a safe and *ready-to-use* data-driven paradigm for system identification in many relevant and safety-critical real-world settings.

6 MODEL VALIDATION

Setting. The hybrid EBM validation focuses on identifying nonlinear dynamics generated by the gradient of nontrivial potentials pre-multiplied by a state-dependent metric. To showcase the ability of EBMs to identify rich, strongly nonlinear dynamics we benchmark the proposed framework on two characteristic systems on $\mathbb{R}^2$:

1. a **multi-well potential** $V_m : \mathbb{R}^2 \rightarrow \mathbb{R}$ exhibiting multiple convex-concave regions

$$V_m(x, y) = \alpha(x^2 + \beta y^2) + \omega(\sin(\gamma x) \cos(y)), \quad (41)$$

2. an **exotic potential** $V_e : \mathbb{R}^2 \rightarrow \mathbb{R}$ characterized by multiple local minima and a ring-shaped invariant set where the vector field vanishes

$$V_e(x, y) = \alpha(x^2 + y^2 - 1)^2 + \omega(\sin(\gamma x) \cos(\zeta y)) + \xi \sin(x^2 - y^2) + \theta \cos(\rho xy) + \delta(x^2 y - y^3). \quad (42)$$

where the parameters $\alpha, \beta, \omega, \gamma, \zeta, \xi, \theta, \delta > 0$ are arbitrarily chosen to emphasize specific qualities of the dynamics, i.e., ruggedness of the potential vs induced rotation. Both potentials are evaluated under a non-Euclidean metric deformation, implemented through a smooth sin-cos dependent matrix

$$Q(x) = \begin{pmatrix} 1 + \sin(x)/2 & 0.3 \cos(y) \\ 0.3 \cos(y) & 1 + \cos(x + y)/2 \end{pmatrix}, \quad (43)$$

leading to ground-truth dynamics $f_m(x, y) = Q(x, y) \nabla V_m(x, y)$ and $f_e(x, y) = Q(x, y) \nabla V_e(x, y)$.

Methods. *Data generation:* We sample 2000 initial conditions in $[-2, 2]^2$ and generate trajectories of length $T = 10$ using Euler integration with $\delta t = 0.01$. States and vector-field evaluations $((x_t, y_t), f(x_t, y_t))$ form the dataset, which is split into 1600 training and 400 test trajectories. *Architectures:* For the multi-well system we use a hybrid EBM with one softmax hidden layer ($M = 128$). For the exotic system we adopt a two-layer architecture (softmax layer with $M = 512$ and polynomial layer with $M = 128$). In all cases, the visible layer is dynamical, hidden layers are feedforward, and the metric is learned via an MLP. *Training:* The model is trained with a composite loss combining weighted field error, short-horizon rollout error, and ℓ_2 regularization. Optimization uses AdamW, and hyperparameters are selected through Weights & Biases (wandb) sweeps².

Numerical results. The hybrid EBM learns a Port-Hamiltonian representation $\mathcal{H}_E(x, y)$ of these fields. Table 1 reports the final train and test losses after 5000 learning epochs for the identification of the two systems. In both cases, the hybrid EBM achieves low reconstruction error, with test losses remaining in the same range as the training error despite the presence of nonlinear metric distortions. For the multi-well system f_m , the learned energy E reproduces the geometry of the true potential V_m , as shown in Figures 1(a)–(b). Despite a different magnitude scaling, the convex–concave structure is accurately captured. The reconstructed vector field and trajectories (Figures 1(c)–(d)) closely match the true dynamics: quiver fields are visually indistinguishable, and trajectories initialized at the same state converge to the same minima, confirming the model’s ability to recover both local flow geometry and global attractor structure.

²Code available at: https://github.com/sim1bet/EBM_StableSystemIdentification




Figure 1: **EBM identification of multi-well potential dynamics.** (a) Ground-truth potential $V_m(x, y)$. (b) EBM energy $E(x, y)$ accurately recovering the landscape geometry. (c) True vector field. (d) Reconstructed field with true and EBM trajectories, showing consistent flow patterns and convergence to the correct attractors.




Figure 2: **EBM identification of exotic potential dynamics.** (a) Ground-truth potential $V_e(x, y)$. (b) EBM energy $E(x, y)$ accurately recovering the landscape ring-geometry. (c) True vector field. (d) Reconstructed field with true and EBM trajectories. EBM dynamics are consistent with the true vector field and reliably recover non-linear, rotational components.

For the exotic system f_e , which features a ring-shaped zero-field region and strong metric-induced distortions, the hybrid EBM accurately recovers both the global basin structure and the central ring (Figures 2(a)–(b)). The reconstructed vector field and trajectories (Figures 2(c)–(d)) remain consistent with the true dynamics, capturing rotational components and preserving the qualitative behavior predicted by the theoretical stability analysis. These results demonstrate that the proposed Port-Hamiltonian hybrid EBM reliably identifies nonlinear flows on non-Euclidean geometries while maintaining structural stability and interpretability through its energy-based formulation.

Direct computation of the invariance radius r in (39) is generally intractable for practical EBM architectures: the width of the first hidden layer determines the dimension of the activation function image, and the corresponding search space grows exponentially. In our setting, the smallest first hidden layer has $M = 128$ units, and its activation co-domain is $[0, 1]^M$, making an exhaustive evaluation of the quantities entering (39) computationally prohibitive. To obtain a meaningful and computationally feasible estimate, we instead compute an *expansion radius* $r_{\text{ex}} < r$ using only the data-driven range explored during training. Specifically, since the dataset $\mathcal{X}$ is generated synthetically over the mesh $\mathcal{A} = [-2, 2]^2$, we forward-propagate all points in $\mathcal{A}$ through the first hidden layer and use the restricted activation image $\phi_2^f(\mathcal{A})$ to evaluate the maximal value of the terms entering (39). This leads to a conservative but sound and computationally tractable estimate r_{ex} that suffices to certify absorbing invariance over the domain explored by the learned dynamics. As illustrated in Figure 3, the expansion radius r_{ex} is largely insensitive to network depth (one hidden layer for V_m versus two for V_e) prior to metric deformation. Following multiplication by the factor $q_{\text{max}}/q_{\text{min}}$, the radius increases by approximately one order of magnitude.

7 CONCLUSIONS

Energy-based models bridge dynamical systems theory and neural networks, offering expressive models that are *stable-by-design*. This work advances their foundations along three dimensions. First, we generalize classical EBM stability analysis to locally Lipschitz activations, establishing energy dissipation via Clarke derivatives and deriving conditions for radial unboundedness beyond the C^2 setting. Second, we introduce a hybrid EBM architecture with dynamic visible layer and static feedforward-feedback hidden layer maps. We prove absorbing invariance of the hybrid EBM dynamics under mild assumptions and extend this guarantee to port-Hamiltonian flows, enabling the modeling of systems on non-Euclidean geometries. Finally, we validate the framework on metric-deformed multi-well and ring-




Figure 3: **Expansiveness radius.** Maximal radius of expansiveness given the data $\mathcal{X}$. The radii have similar magnitude for both the (a) multi-well potential $V_m(x, y)$ and the (b) exotic potential $V_e(x, y)$.

shaped potentials, demonstrating accurate vector-field reconstruction, faithful recovery of energy-landscape geometry, and reliable trajectory tracking. These results highlight hybrid EBM as physically grounded, stable-by-design models suited for data-driven identification in nonlinear, safety critical domains. Future work will explore integrating control inputs to enable safe and certifiable identification, further advancing energy-based modeling as a core methodology for physical AI.

8 ACKNOWLEDGMENTS

The authors gratefully acknowledge the contribution of IT4LIA AI Factory project EHPC-AIF-2026PG01-103.

References

- M. A. Cohen and S. Grossberg. Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5):815–826, September 1983. ISSN 2168-2909. doi:[10.1109/tsmc.1983.6313075](https://doi.org/10.1109/tsmc.1983.6313075).
- J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, April 1982. ISSN 1091-6490. doi:[10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554).
- J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81(10):3088–3092, May 1984. ISSN 1091-6490. doi:[10.1073/pnas.81.10.3088](https://doi.org/10.1073/pnas.81.10.3088).
- R. McEliece, E. Posner, E. Rodemich, and S. Venkatesh. The capacity of the hopfield associative memory. *IEEE Transactions on Information Theory*, 33(4):461–482, July 1987. ISSN 1557-9654. doi:[10.1109/tit.1987.1057328](https://doi.org/10.1109/tit.1987.1057328).
- D. Krotov and J. J. Hopfield. Large associative memory problem in neurobiology and machine learning. In *International Conference on Learning Representations*, 2020. doi:[10.48550/arXiv.2008.06996](https://doi.org/10.48550/arXiv.2008.06996).
- B. Hoover, Y. Liang, B. Pham, R. Panda, H. Strobelt, D. H. Chau, M. Zaki, and D. Krotov. Energy transformer. *Advances in neural information processing systems*, 36:27532–27559, 2023. doi:[10.48550/arXiv.2302.07253](https://doi.org/10.48550/arXiv.2302.07253).
- H. Ramsauer, B. Schäfl, J. Lehner, P. Seidl, M. Widrich, L. Gruber, M. Holzleitner, T. Adler, D. Kreil, M. K. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2021. doi:[10.48550/arXiv.2008.02217](https://doi.org/10.48550/arXiv.2008.02217).
- M. Demircigil, J. Heusel, M. Lowe, S. Uppang, , and F. Vermet. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2):288–299, May 2017. ISSN 1572-9613. doi:[10.1007/s10955-017-1806-y](https://doi.org/10.1007/s10955-017-1806-y).
- B. Hoover, D. H. Chau, H. Strobelt, P. Ram, and D. Krotov. Dense associative memory through the lens of random features. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 23549–23576. Curran Associates, Inc., 2024. doi:[10.52202/079017-0742](https://doi.org/10.52202/079017-0742).
- L. Ambrogioni. In search of dispersed memories: Generative diffusion models are associative memory networks. *Entropy*, 26(5):381, April 2024. ISSN 1099-4300. doi:[10.3390/e26050381](https://doi.org/10.3390/e26050381).
- G. Y. Park, J. Kim, B. Kim, S. W. Lee, and J. C. Ye. Energy-based cross attention for bayesian context update in text-to-image diffusion models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 76382–76408. Curran Associates, Inc., 2023. doi:doi.org/10.48550/arXiv.2306.09869.
- W. Wu, T. Y. Hsiao, J. Y. C. Hu, W. Zhang, and H. Liu. In-context learning as conditioned associative memory retrieval. In *Forty-second International Conference on Machine Learning*, 2025.
- K. J. Åström and P. Eykhoff. System identification—a survey. *Automatica*, 7(2):123–162, March 1971. ISSN 0005-1098. doi:[10.1016/0005-1098\(71\)90059-8](https://doi.org/10.1016/0005-1098(71)90059-8).
- L. Ljung. Perspectives on system identification. *Annual Reviews in Control*, 34(1):1–12, April 2010. ISSN 1367-5788. doi:[10.1016/j.arcontrol.2009.12.001](https://doi.org/10.1016/j.arcontrol.2009.12.001).
- R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. doi:[10.48550/arXiv.1806.07366](https://doi.org/10.48550/arXiv.1806.07366).
- A. Perrusquía, R. Garrido, and W. Yu. Stable robot manipulator parameter identification: A closed-loop input error approach. *Automatica*, 141:110294, July 2022. ISSN 0005-1098. doi:[10.1016/j.automatica.2022.110294](https://doi.org/10.1016/j.automatica.2022.110294).
- R. Chakraborty, H. Jain, and G. S. Seo. A review of active probing-based system identification techniques with applications in power systems. *International Journal of Electrical Power & Energy Systems*, 140:108008, September 2022. ISSN 0142-0615. doi:[10.1016/j.ijepes.2022.108008](https://doi.org/10.1016/j.ijepes.2022.108008).
- P. Denno, C. Dickerson, and J. A. Harding. Dynamic production system identification for smart manufacturing systems. *Journal of Manufacturing Systems*, 48:192–203, July 2018. ISSN 0278-6125. doi:[10.1016/j.jmsy.2018.04.006](https://doi.org/10.1016/j.jmsy.2018.04.006).
- H. Dai, B. Landry, L. Yang, M. Pavone, and R. Tedrake. Lyapunov-stable neural-network control. *arXiv*, 2021. doi:[10.48550/ARXIV.2109.14152](https://doi.org/10.48550/ARXIV.2109.14152).
- J. Drgona, A. Tuor, S. Vasishth, and D. Vrabie. Dissipative deep neural dynamical systems. *IEEE Open Journal of Control Systems*, 1:100–112, 2022. ISSN 2694-085X. doi:[10.1109/ojcsys.2022.3186838](https://doi.org/10.1109/ojcsys.2022.3186838).

- Y. Xu and S. Sivaranjani. Learning dissipative neural dynamical systems. *IEEE Control Systems Letters*, 7:3531–3536, 2023. ISSN 2475-1456. doi:[10.1109/lcsys.2023.3337851](https://doi.org/10.1109/lcsys.2023.3337851).
- S. Massaroli, M. Poli, M. Bin, J. Park, A. Yamashita, and H. Asama. Stable neural flows. *arXiv*, 2020. doi:[10.48550/ARXIV.2003.08063](https://doi.org/10.48550/ARXIV.2003.08063).
- F. J. Roth, D. K. Klein, M. Kannapinn, J. Peters, and O. Weeger. Stable port-hamiltonian neural networks. *arXiv*, 2025. doi:[10.48550/ARXIV.2502.02480](https://doi.org/10.48550/ARXIV.2502.02480).
- J. M. Lee. *Introduction to Smooth Manifolds*. Springer New York, 2012. ISBN 9781441999825. doi:[10.1007/978-1-4419-9982-5](https://doi.org/10.1007/978-1-4419-9982-5).
- F. H. Clarke. Generalized gradients and applications. *Transactions of the American Mathematical Society*, 205:247–247, 1975. ISSN 0002-9947. doi:[10.1090/s0002-9947-1975-0367131-6](https://doi.org/10.1090/s0002-9947-1975-0367131-6).
- H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer International Publishing, 2017. ISBN 9783319483115. doi:[10.1007/978-3-319-48311-5](https://doi.org/10.1007/978-3-319-48311-5).
- D. Krotov and J. J. Hopfield. Dense associative memory for pattern recognition. In *Advances in neural information processing systems*, volume 29, 2016. doi:[10.48550/arXiv.1606.01164](https://doi.org/10.48550/arXiv.1606.01164).
- B. Hoover, D. H. Chau, H. Strobelt, and D. Krotov. A universal abstraction for hierarchical hopfield networks. In *The Symbiosis of Deep Learning and Differential Equations II*, 2022.
- L. Kozachkov, J. J. Slotine, and D. Krotov. Neuron–astrocyte associative memory. *Proceedings of the National Academy of Sciences*, 122(21), May 2025. ISSN 1091-6490. doi:[10.1073/pnas.2417788122](https://doi.org/10.1073/pnas.2417788122).
- L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2): 223–311, January 2018. ISSN 1095-7200. doi:[10.1137/16m1080173](https://doi.org/10.1137/16m1080173).
- L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions, Revised Edition*. Chapman and Hall/CRC, April 2015. ISBN 9781482242393. doi:[10.1201/b18333](https://doi.org/10.1201/b18333).
- A. Van der Schaft. *Port-Hamiltonian systems: an introductory survey*. EMS Press, May 2007. ISBN 9783985475384. doi:[10.4171/022-3/65](https://doi.org/10.4171/022-3/65).