MAX PLANCK GESELLSCHAFT

SECAI School of Embedded Composite Artificial Intelligence

Relative Geometry of Neural Forecasters: Linking Accuracy and Alignment in Learned Dynamics

Deniz Kucukahmetler*[,1,2], Maximilian Jean Hemmann[3], Julian Mosig von Aehrenfeld[3], Maximilian Amthor[3], Christian Deubel[3], Nico Scherf[1,5], Diaaeldin Taha[6]

[1] Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany
[2] School of Embedded Composite Artificial Intelligence (SECAI) Dresden/Leipzig, Germany
[3] Leipzig University, Germany
[5] Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig , Germany
[6] Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

We study neural forecasters for dynamical systems through the lens of representational alignment. We introduce anchor-based, geometry-agnostic relative embeddings that remove rotational and scaling ambiguities, enabling robust cross-seed and cross-architecture comparison. Across diverse periodic, quasi-periodic, and chaotic systems, we observe consistent family-level patterns: MLPs align with MLPs, RNNs with RNNs, and ESNs show reduced alignment on chaotic dynamics, while Transformers often align weakly but still perform well. Alignment generally correlates with forecasting accuracy, yet high accuracy can coexist with low alignment. Relative embeddings thus offer a simple, reproducible basis for comparing learned dynamics.

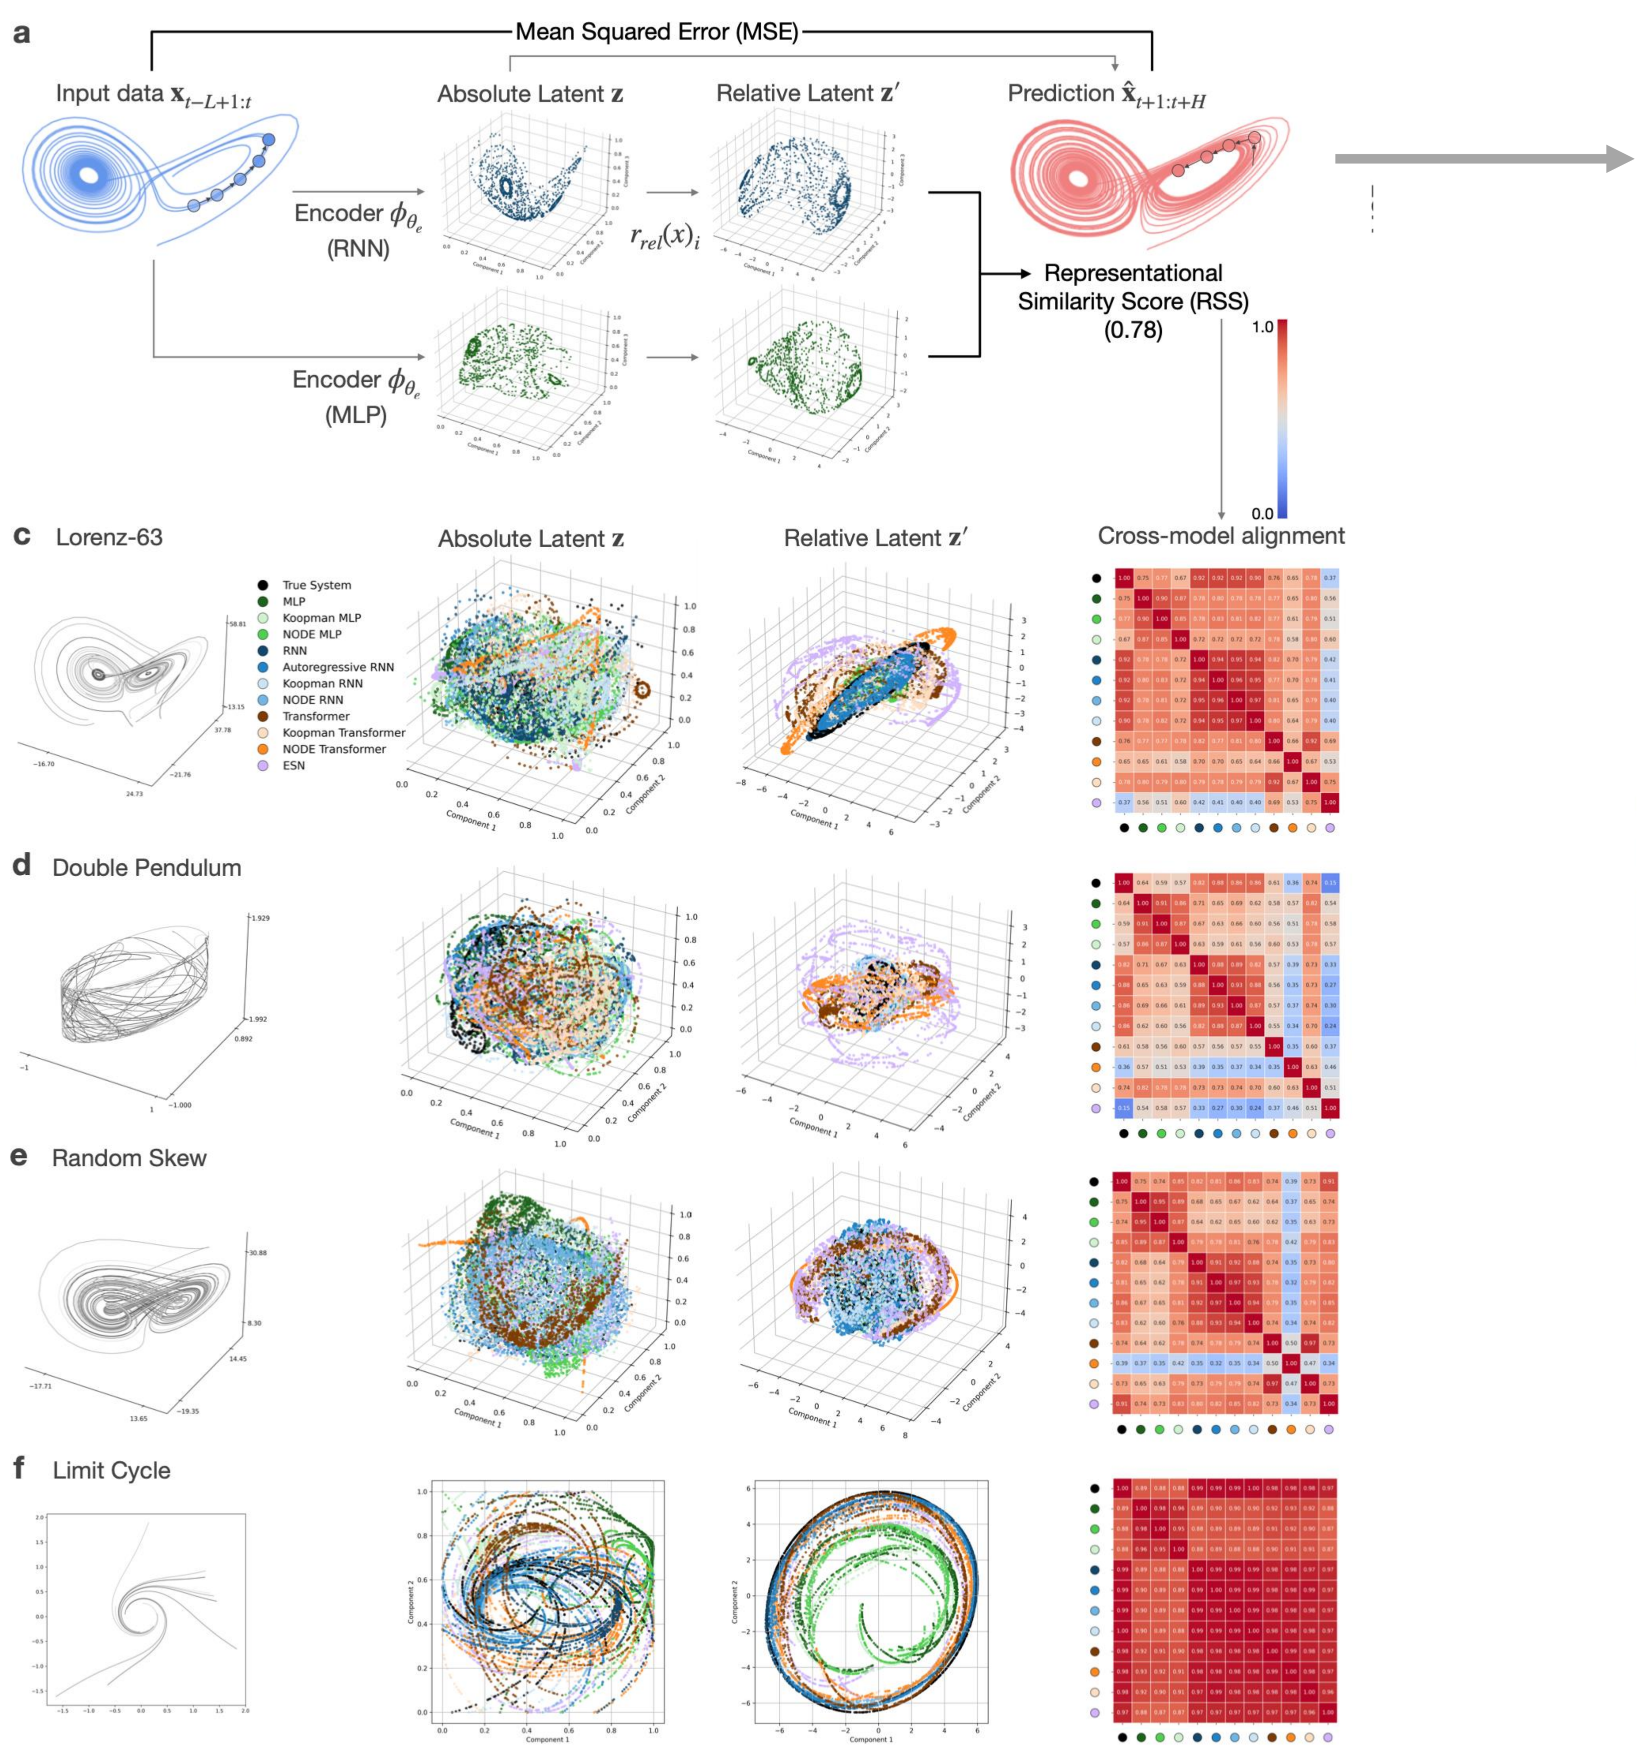

Figure 1: Relative embeddings reveal consistent geometric structure across model families while removing rotational and scaling ambiguities. Six systems (Lorenz, double pendulum, skew-product; trajectories can differ markedly across seeds due to sensitivity to initial conditions; POD wake, limit cycle, Hopf). Rows show trajectories, absolute embeddings, relative embeddings (PCA), and cross-model similarity heatmaps (avg. over five seeds). Relative embeddings reduce geometric variability, enabling direct comparison across forecasters.

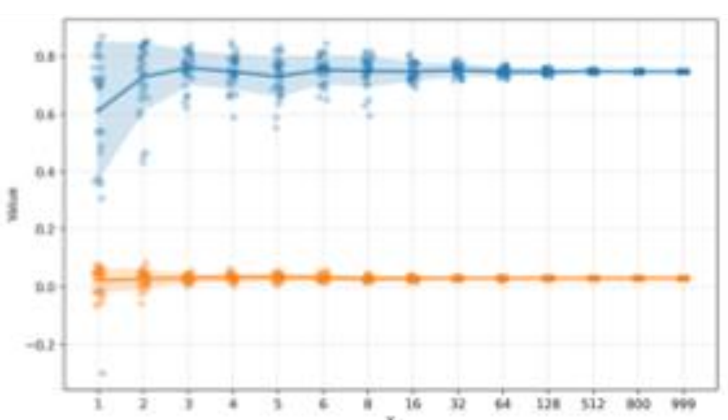

Figure 3: Anchor ablation and baseline. (Blue) Alignment vs. number of anchors K; lines show mean over 30 repeats. Stabilization occurs for K ≥16; we choose K = 80 (vertical marker) for the main experiments. (Orange) Random baseline with disjoint anchor sets across spaces, yielding near-zero alignment.

Latent Space Alignment

Neural forecasters can predict complex dynamics well, but their latent spaces are hard to compare due to arbitrary geometric differences. To address this, we adapt the relative-embedding alignment framework of Moschella et al. (2023), which represents **each latent state by its similarities to a fixed set of anchor points** rather than by absolute coordinates. This yields geometry-agnostic, anchor-based latent representations that remove rotational and scaling ambiguities, enabling direct and stable comparison of forecasters across models, seeds, and architectures.

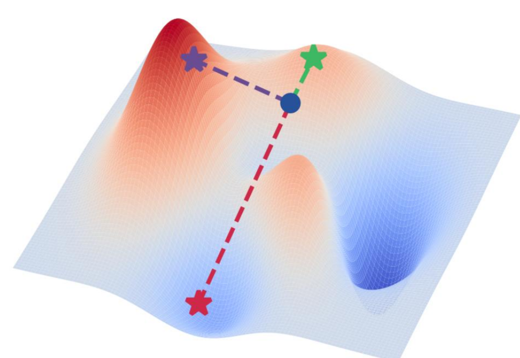

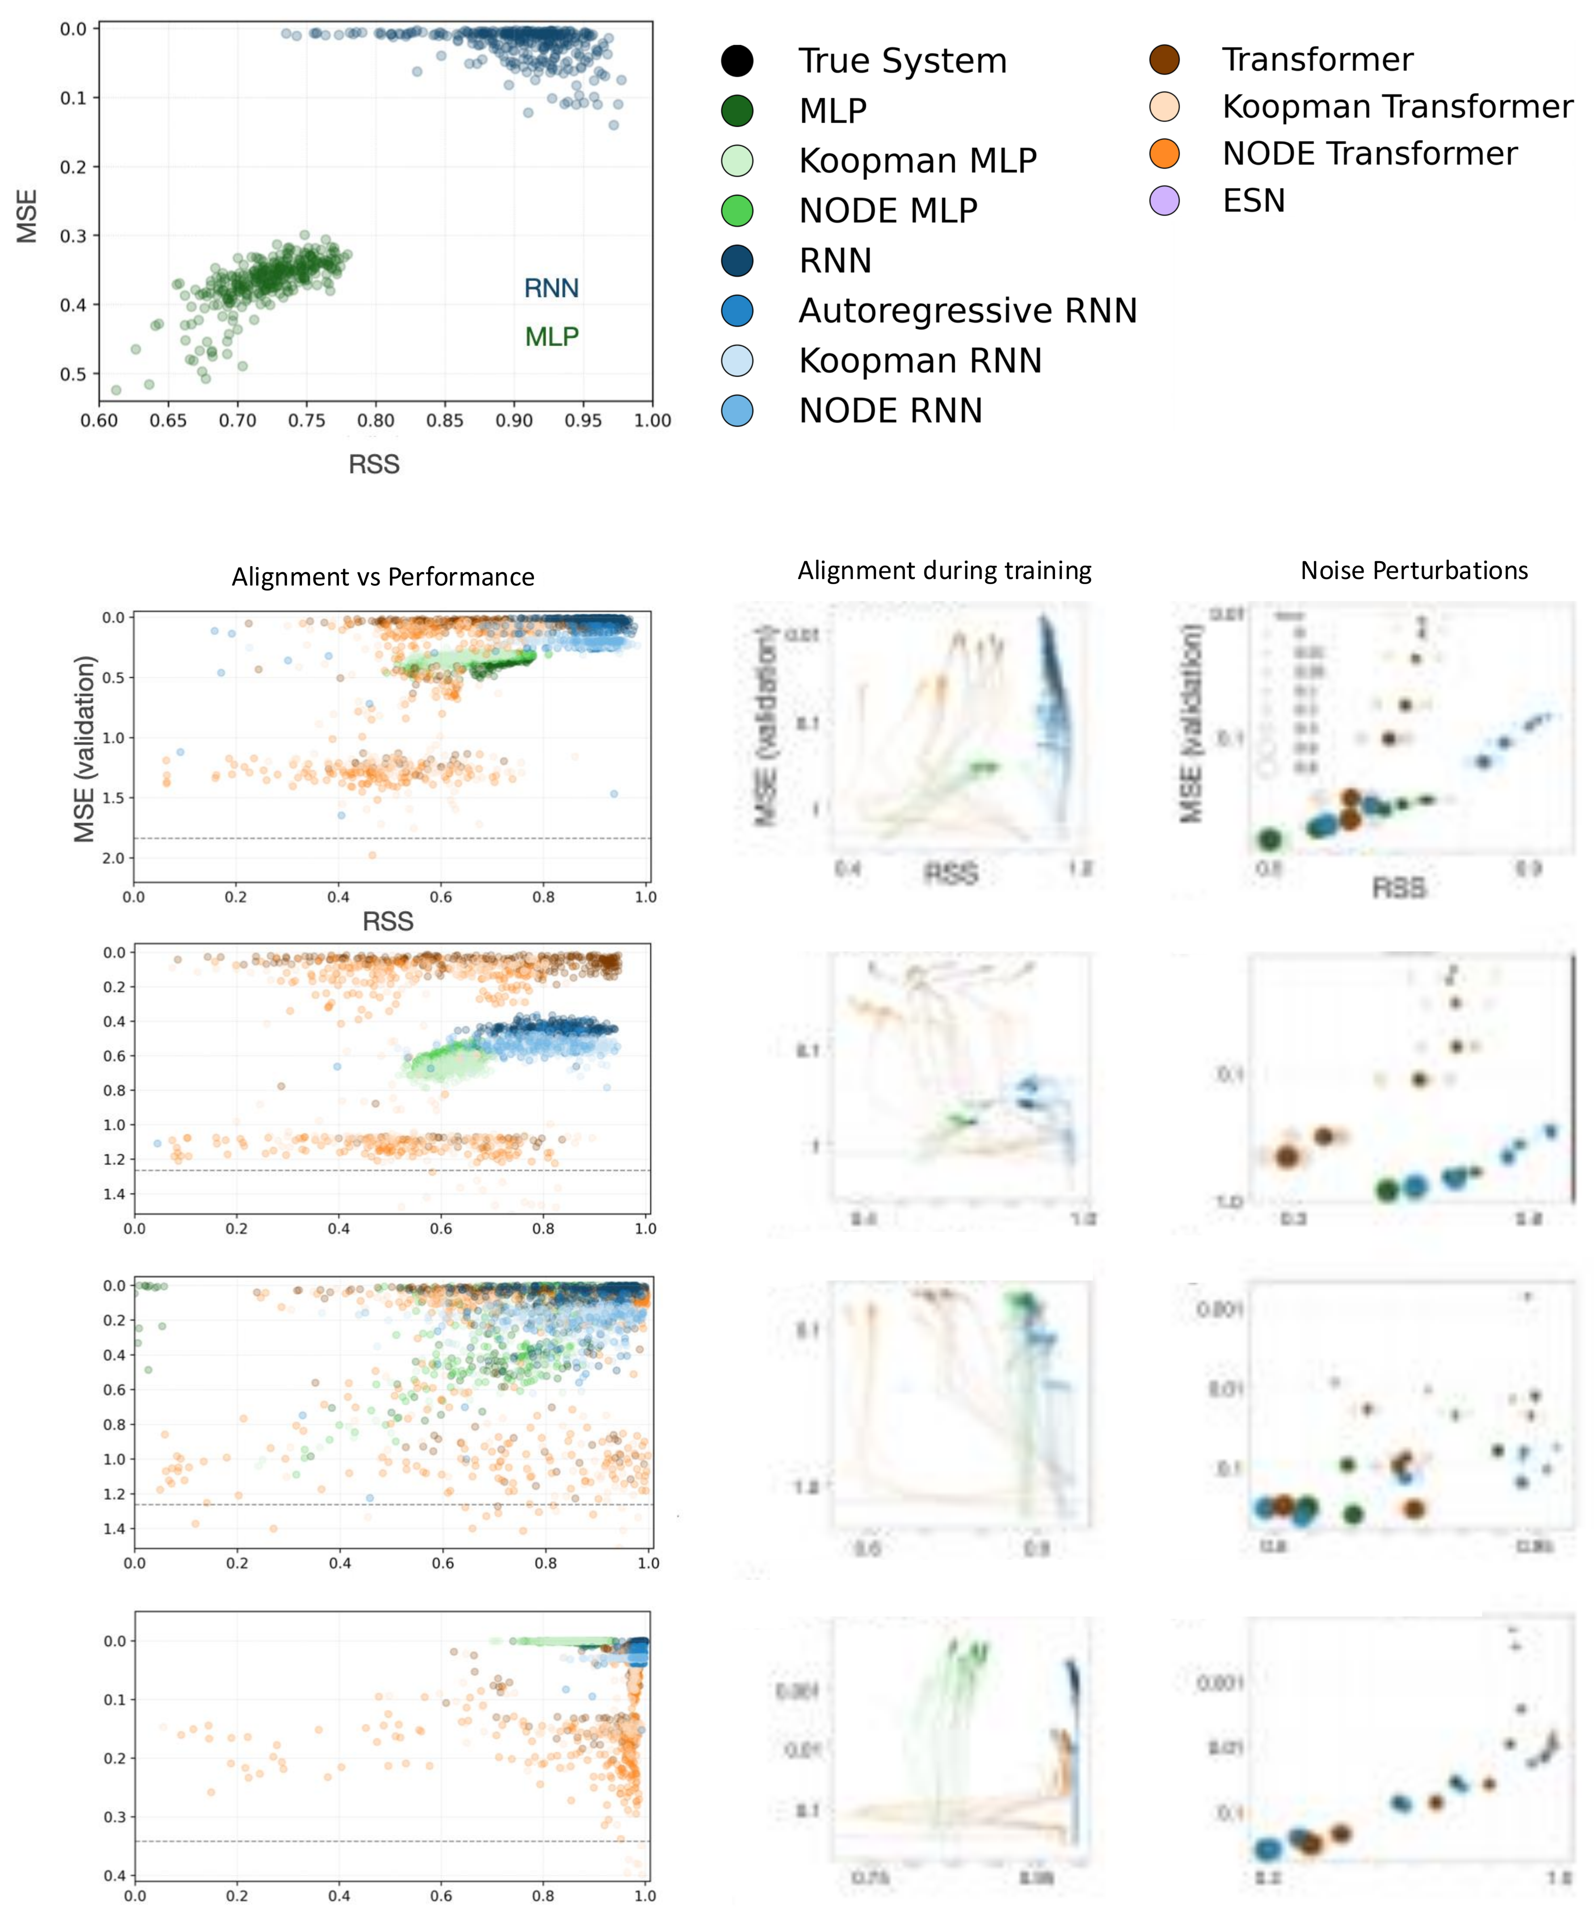

Figure 2: Alignment–performance trade-offs across model families and dynamical systems. At the end of training, RNNs consistently combined high forecasting accuracy with strong alignment, MLPs showed a clear positive relationship between the two, while transformers achieved superior accuracy despite weaker alignment. Under increasing noise perturbations, all models exhibited reduced alignment and degraded forecasting, revealing distinct family-specific balances between predictive performance and representational stability.

References

Kucukahmetler, D., Hemmann, M. J., Mosig von Aehrenfeld, J., Amthor, M., Deubel, C., Scherf, N., & Taha, D. (submitted). Relative Geometry of Neural Forecasters: Linking Accuracy and Alignment in Learned Dynamics. Submitted to Transactions on Machine Learning Research (TMLR).

Moschella, L., Maiorca, V., Fumero, M., Norelli, A., Locatello, F., & Rodolà, E. (2023). *Relative representations enable zero-shot latent space communication.* In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR).*

[*] Corresponding author: kucukahm@cbs.mpg.de

MAX PLANCK INSTITUTE FOR HUMAN COGNITIVE AND BRAIN SCIENCES
UNIVERSITÄT LEIPZIG
MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES
SPONSORED BY THE Federal Ministry of Education and Research
DAAD Zuse Schools Konrad Zuse Schools of Excellence in Artificial Intelligence