7 Learned folds

R. Heller · Date: 2026-07-29 · Last changed: 2026-11-09

Can a network learn the fold and the unfold at once?

Chapter status: planned. The items below describe work still to be implemented; this chapter does not provide a trained autoencoder or benchmark results.

- A `torch` autoencoder with a two-unit bottleneck. (`keras` would pull TensorFlow and break clean-clone reproducibility.)
- The only method here that learns an explicit inverse: encoder = low-dimensional representation, decoder = reconstruction in the input space.
- Reconstruction loss is the fold–unfold identity made computable.
- Compare the learned decoder against the true unfolding, not just against the input.