Neuro-Visualizer: A Novel Auto-Encoder-Based Loss Landscape Visualization Method With an Application in Knowledge-Guided Machine Learning

Mohannad Elhamod and Anuj Karpatne

VIRGINIA TECH

KGML Knowledge-guided Machine Learning

ICML International Conference On Machine Learning

Loss Landscape Visualization

- As the complexity of neural networks has increased exponentially in recent years, researchers are showing great interest in ***qualitatively*** studying the properties of their loss landscape.
- Loss landscape visualization has recently been used to assess neural networks in several ways, including their generalization performance and training convergence.

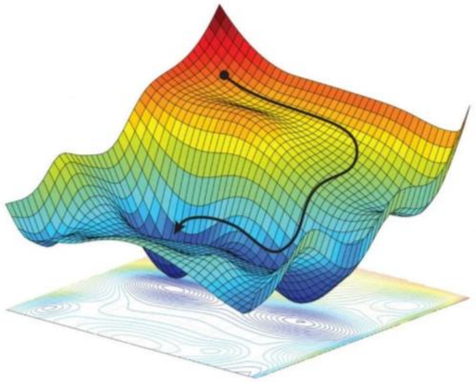

(Amini et. al. 2019)

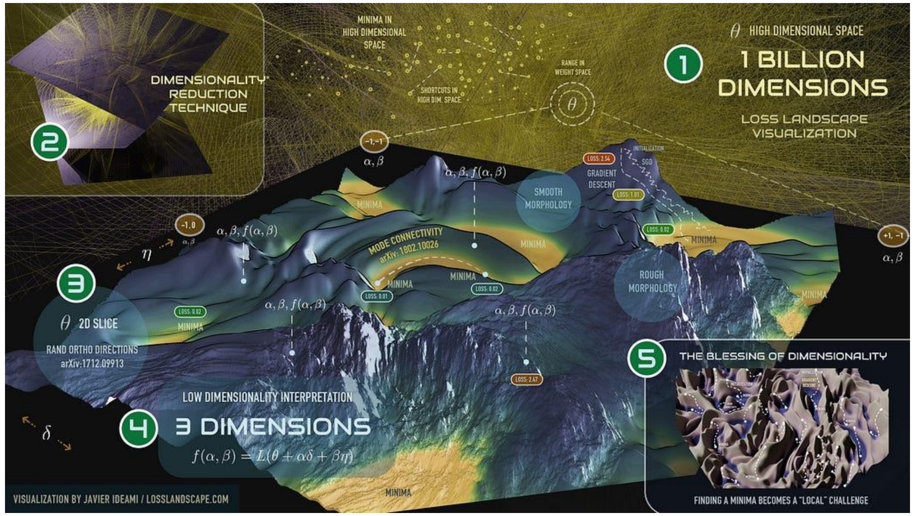

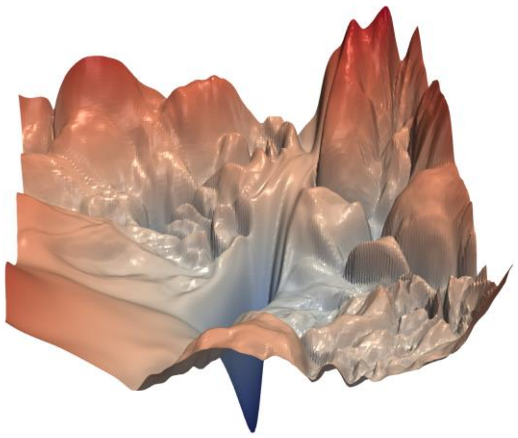

(Li et. al. 2018)

Loss Landscape Visualization

- However, current approaches use linear techniques to visualize 2D planar slices, leading to severe limitations:
    - **<u>High dimensional trajectories are not planar</u>**, leading to inaccurate visualizations.
    - Additional complications arise regarding the choice of the location, orientation, and the scale of the visualization plane and its grid.

- ***<u>Problem #1:</u>*** *Current loss landscape visualization methods suffer from limitations and inaccuracies.*

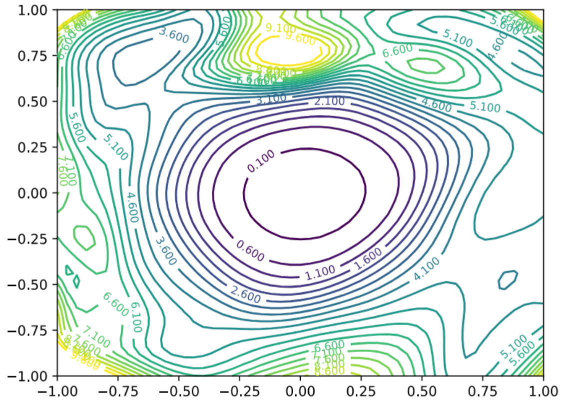

(Li et. al. 2018)

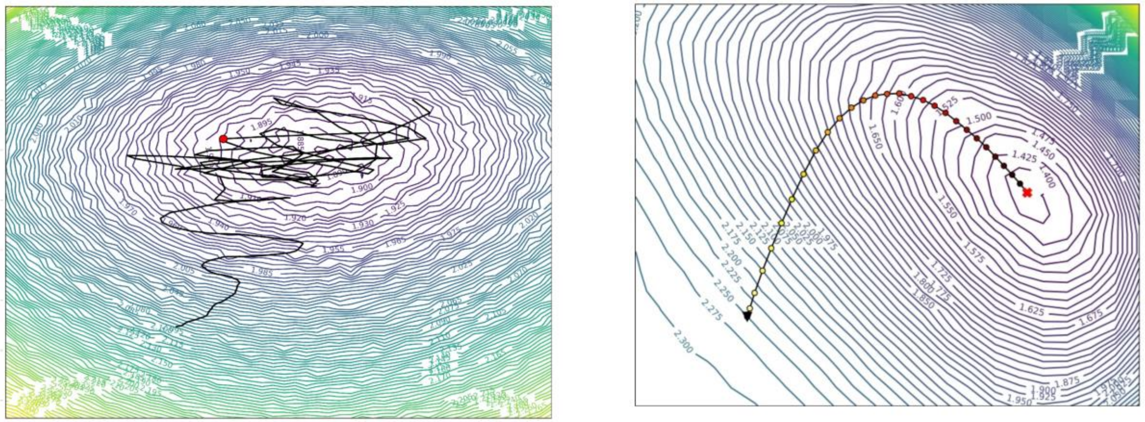

(Chatzimichailidis et. al. 2019)

Knowledge-Guided ML (KGML)

- The KGML framework aims at achieving better model generalizability by imbuing neural networks with domain knowledge (e.g., through regularization terms).
- This framework has been applied in many fields, including biology and physics.

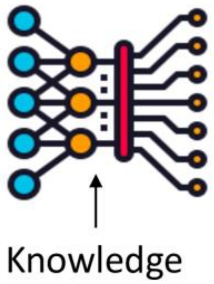

Knowledge-Guided ML
(Karpatne et. al. 2024)

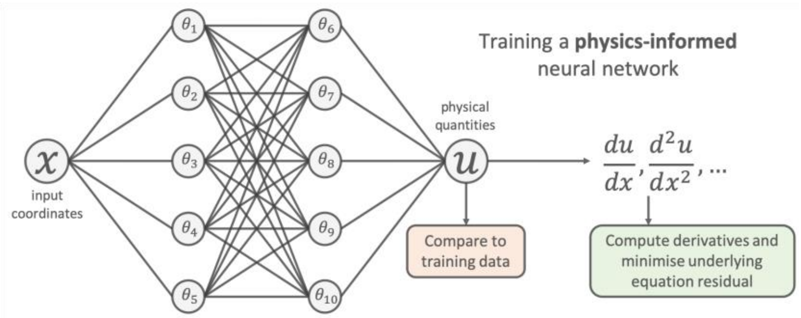

PINN
(Courtesy of Ben Moseley)

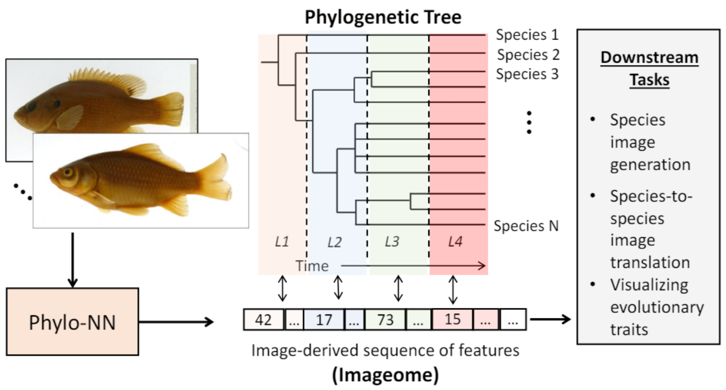

Phylo-NN
(Elhamod et. al. 2023)

Knowledge-Guided ML (KGML)

- In this work, we focus on 2 applications:
  - ***PINNs*** for solving differential equations.
  - ***CoPhy-PGNN*** for eigen-decomposition.

- Despite its resounding success, there is still a lack of a comprehensive understanding of how the KGML framework affects the model optimization process.

- ***Problem #2:*** *The mechanisms underlying the KGML framework are still not fully understood.*

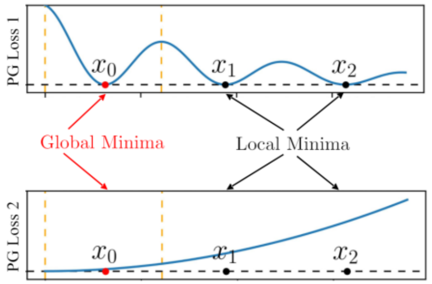

CoPhy-PGNN
(Elhamod and Bu et. al. 2022)

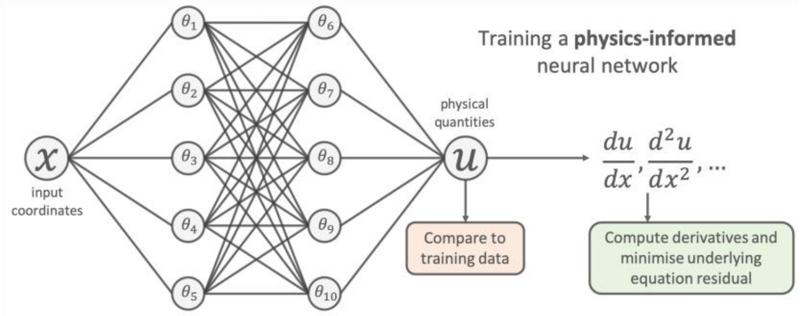

PINN
(Courtesy of Ben Moseley)

Research Goals

- ***<u>Problem #1:</u>*** *Current loss landscaper visualization methods suffer from <u>limitations and inaccuracies</u>.*

- ***<u>Problem #2:</u>*** *The mechanisms underlying the KGML framework are still not fully understood.*

- ***<u>Research Goals:</u>***
  1. *To d*evise a loss landscape visualization approach that addresses the limitations of current methods.
  2. To use the new approach to better understand the mechanisms underlying the KGML framework.

Main Idea

- ***Current approach:***

Projecting the high-dimensional training trajectory on a **2D plane**.

- ***Proposed approach:***

Finding a **2D manifold** that faithfully captures the high-dimensional training trajectory.

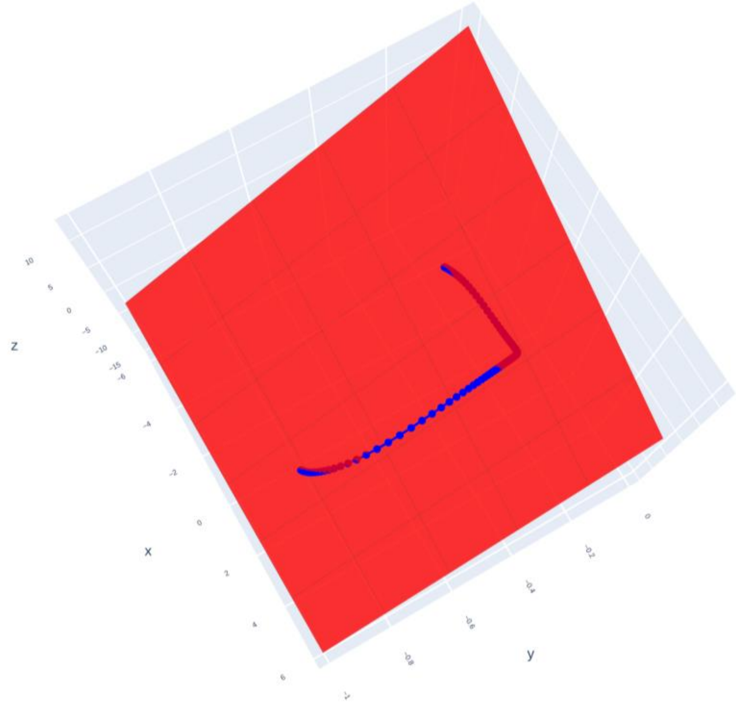

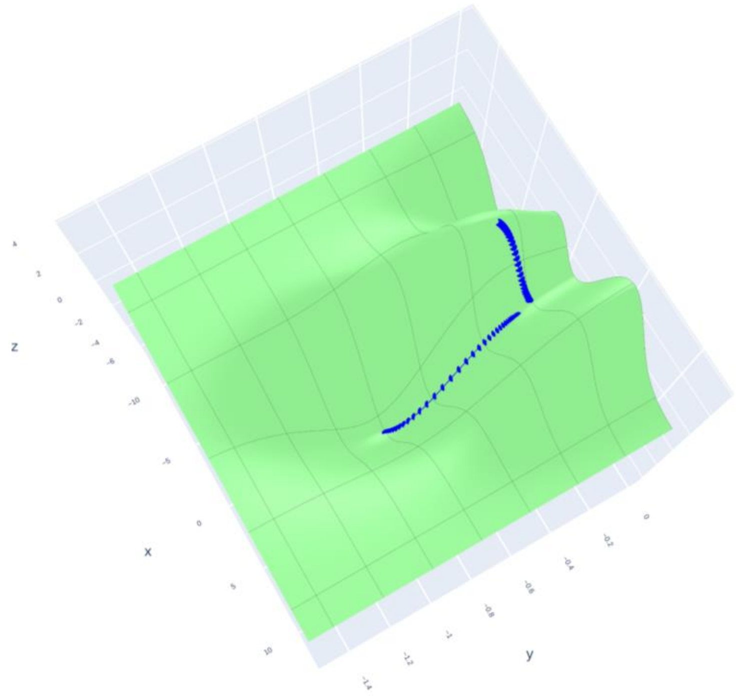

Main Idea

- ***Current approach:***

Using PCA.

- ***Proposed approach:***

Using an **auto-encoder.**

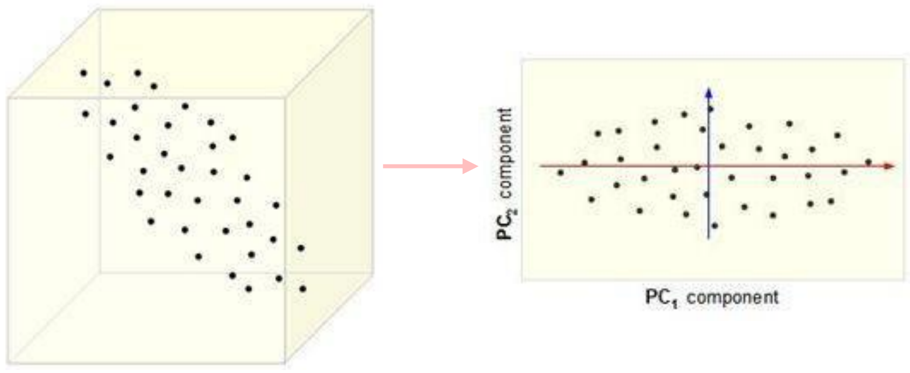

PCA
(Courtesy of David Zeleney)

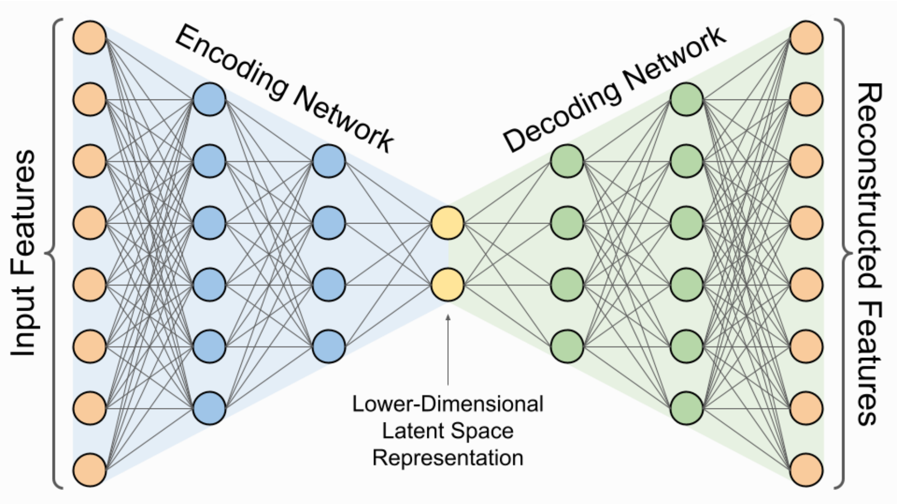

Auto-encoder
(Courtesy of Etorezone)

Neuro-Visualizer

- Guided by the reconstruction loss, the auto-encoder is trained and learns a 2D embedding that is sufficient to faithfully reconstruct the training trajectory in the original high-dimensional parameter space.

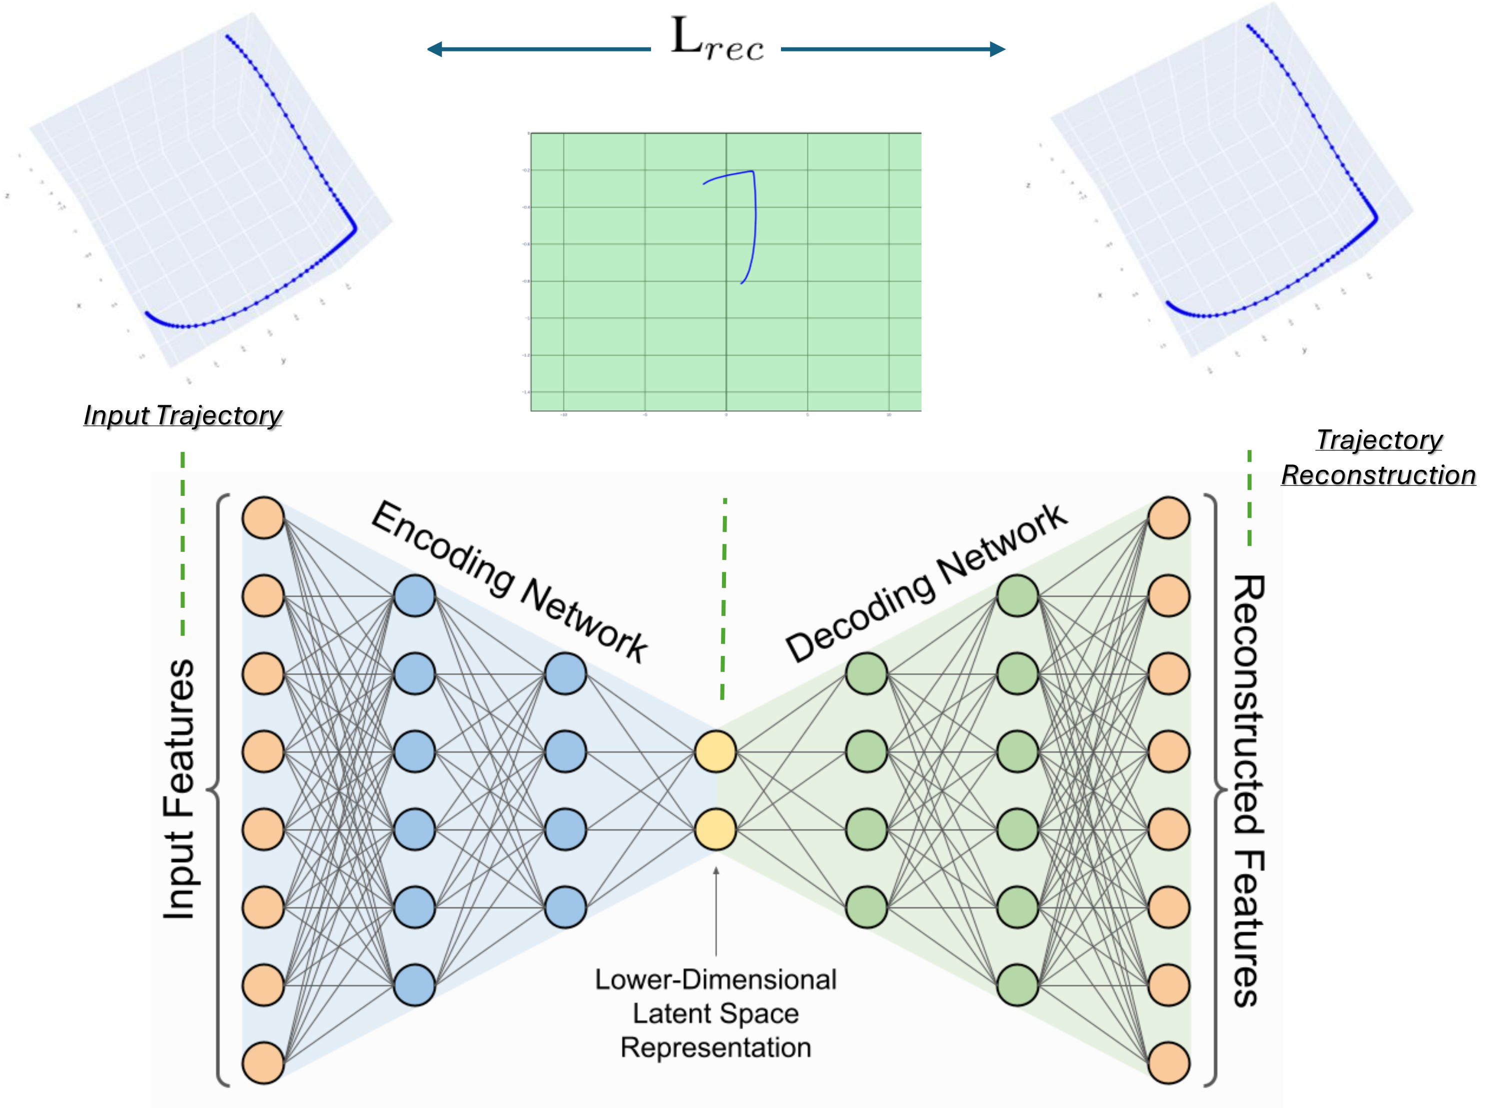

Auto-encoder
(Courtesy of Etorezone)

Neuro-Visualizer

- Traversing and sampling ***the Neuro-Visualizer***'s latent space provides the grid layout.
- Using the decoder, the grid is projected back into high dimensional space as a 2D manifold.
- Using the reconstructed manifold, the grid's loss heatmap is now calculated and visualized.

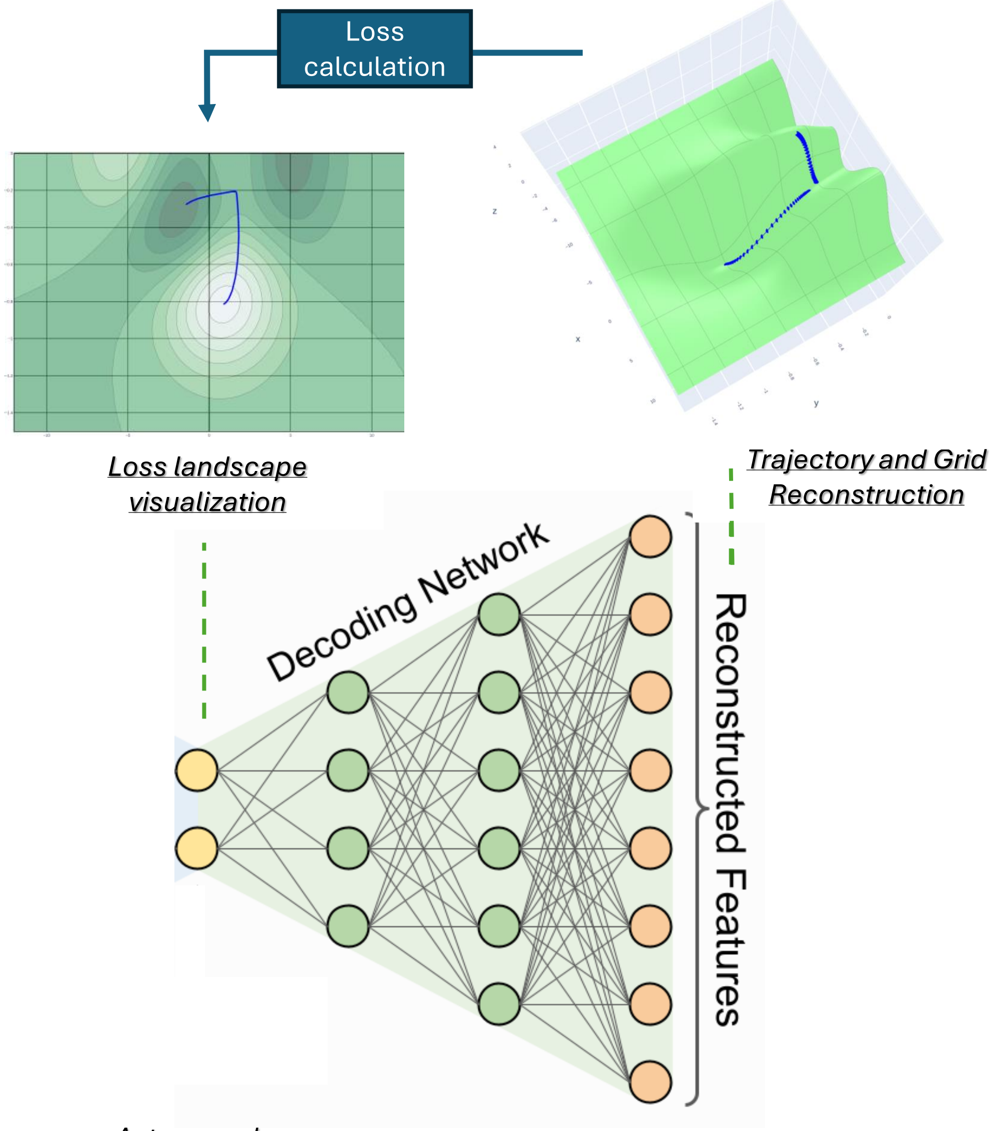

Loss landscape visualization

Trajectory and Grid Reconstruction

Auto-encoder
(Courtesy of Etorezone)

Neuro-Visualizer

- Imposing additional constraints as loss functions allows the learned manifold to be molded to have certain desired properties.

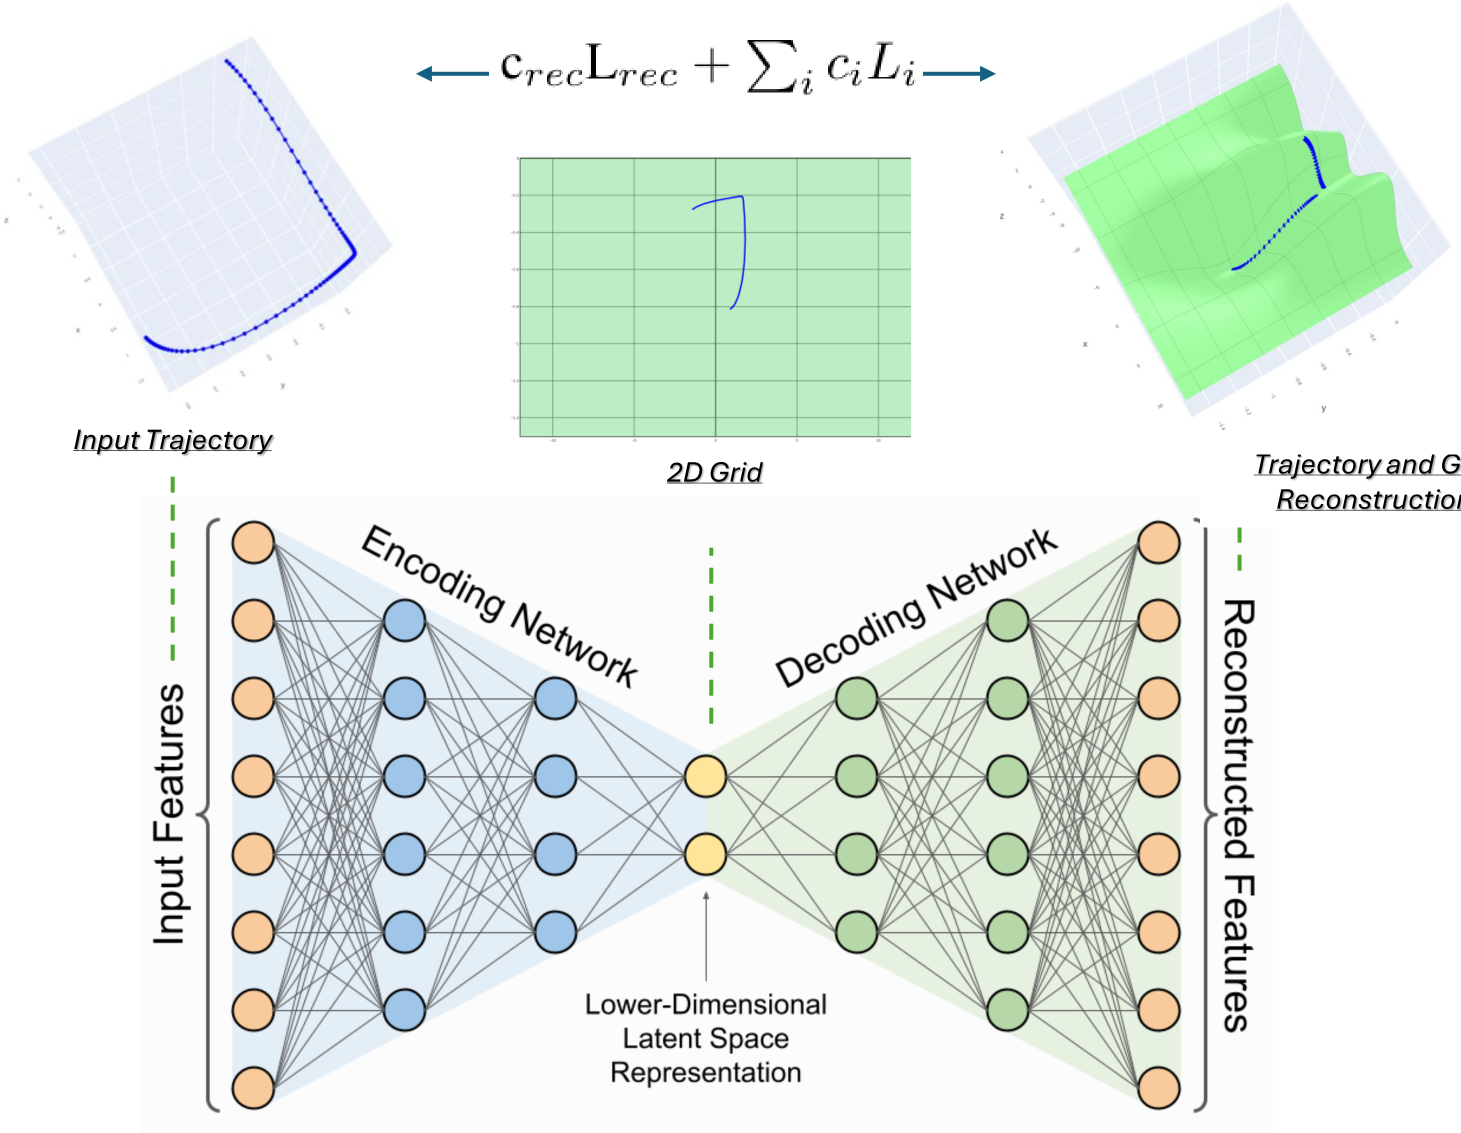

Auto-encoder
(Courtesy of Etorezone)

Neuro-Visualizer

- Imposing additional constraints as loss functions allows the learned manifold to be molded to have certain desired properties.

- **<u>Examples:</u>**

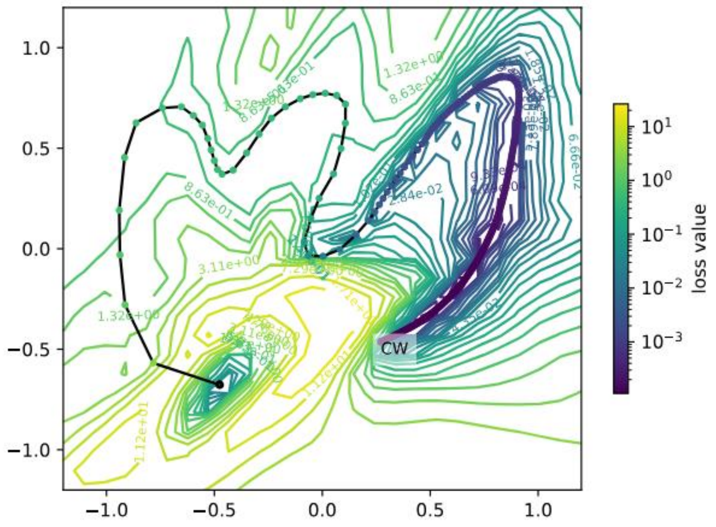

<u>No constraints</u>

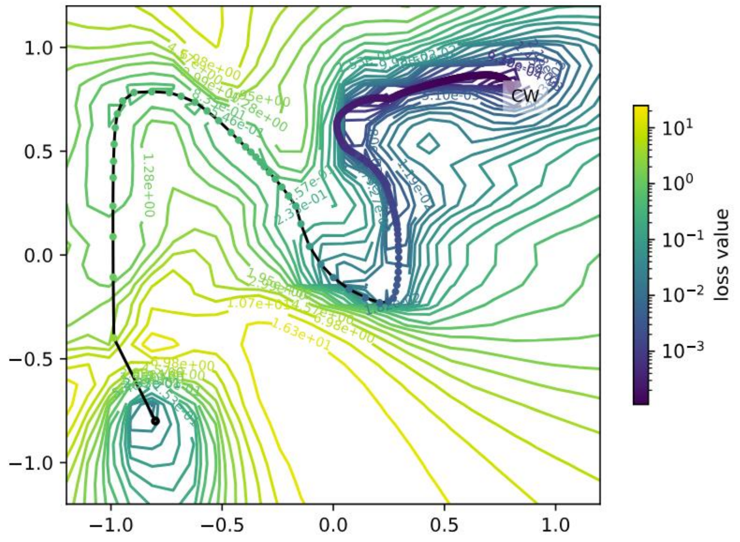

<u>Location Anchoring/Pinning</u>

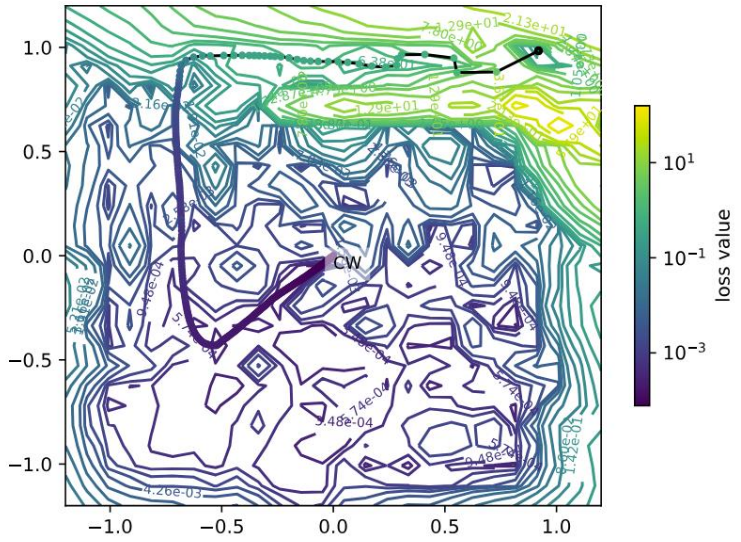

<u>Grid scaling</u>

Results
CoPhy-PGNN (Physics-Guided Neural Nets)

- *Neuro-Visualizer* offers much richer details.
- *Neuro-Visualizer* fits the trajectory more accurately.

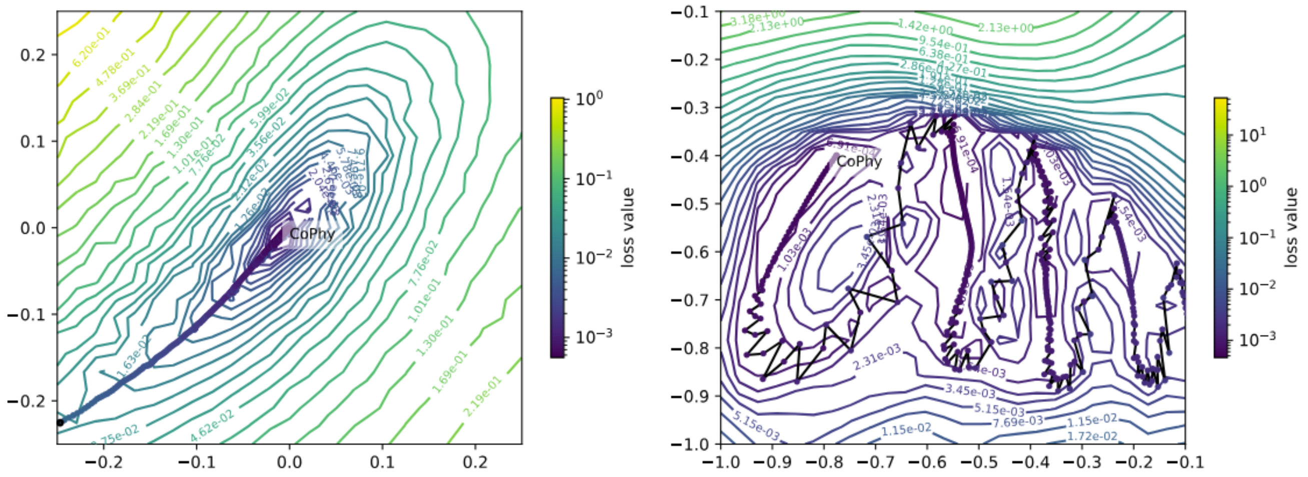

(a) PCA (zoomed-in)

(b) *Neuro-Visualizer* (zoomed-in)

Metric	*Neuro-Visualizer*	PCA	Kernel-PCA	UMAP
$e_{\mathbf{relative}}$	0.0095	1.6782	4.7250	0.4295
$e_{\mathbf{proj}}$	0.0005	0.2832	0.0865	0.2307

Results:
PINNs (Physics-Informed Neural Nets)

- How does β impact the convergence of a PINN when solving a Convection PDE?

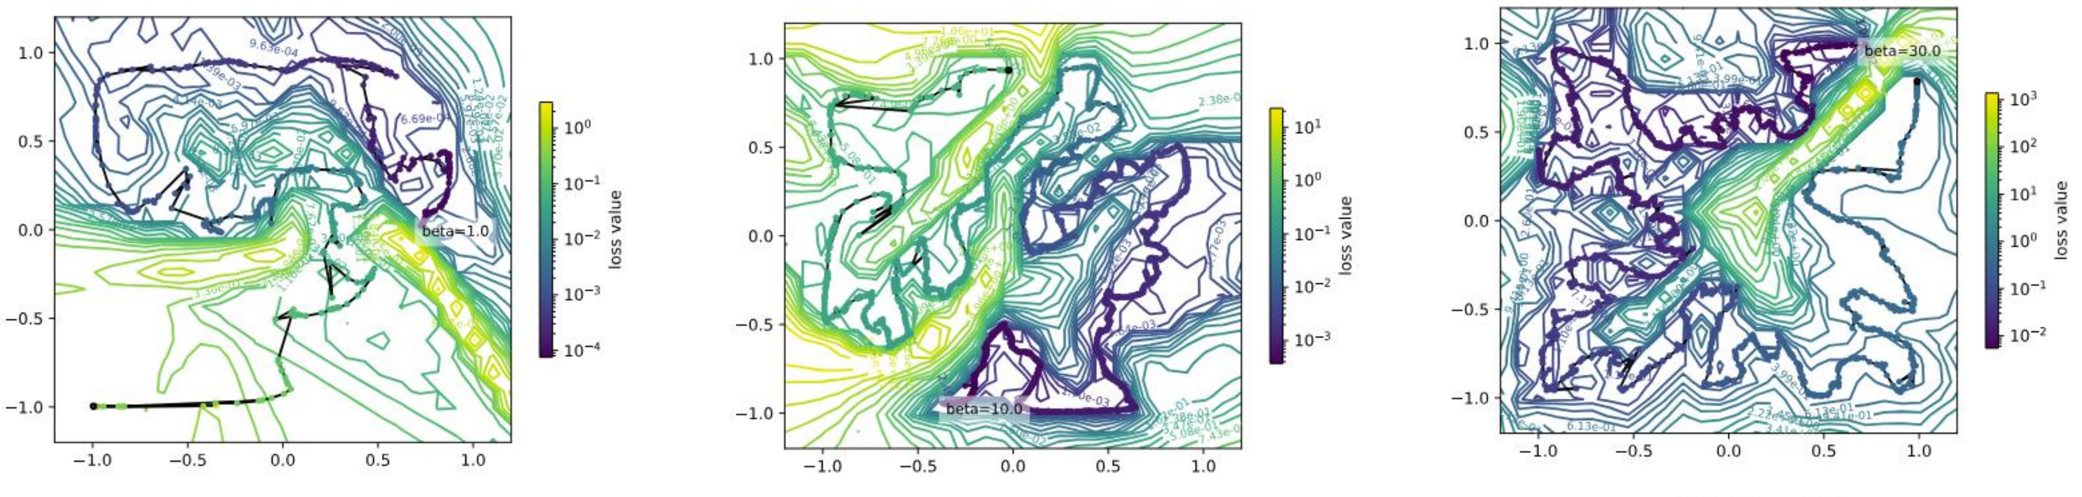

(a) $\beta = 1$ (b) $\beta = 10$ (c) $\beta = 30$

Conclusion: Optimization becomes more difficult as β increases. This is due to a corresponding increase in loss landscape complexity.

Results:
PINNs (Physics-Informed Neural Nets)

- How does the choice of weight balancing technique impact a PINN's convergence?

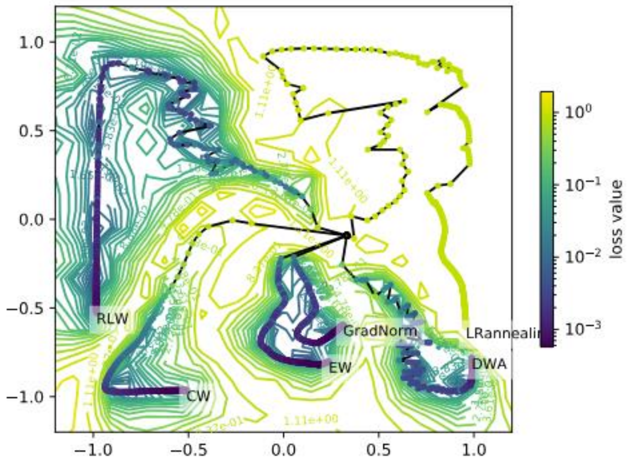

(b) L_{test}

Conclusion: Different techniques arrive at different minima. There is no single technique that universally outperforms the rest. Feasibility highly depends on the differential equation in question.

Thank you!

https://github.com/elhamod/NeuroVisualizer

https://proceedings.mlr.press/v235/elhamod24a.html

elhamod@vt.edu

SCAN ME

VIRGINIA TECH

KGML
Knowledge-guided Machine Learning

ICML
International Conference On Machine Learning