brain
networks
laboratory
Texas A&M University

The Relationship Between Spectral Bias and Saddle Points in Deep Learning

Alex Alvarez
Department of Computer Science and Engineering, Texas A&M University, College Station, Texas
Research Advisor: Yoonsuck Choe

COMPUTER SCIENCE
& ENGINEERING

Introduction

- Neural networks often exhibit spectral bias, prioritizing low-frequency data components over high-frequency ones.
- Saddle points in the loss landscape cause plateaus in the learning curve.
- My research explores whether saddle points are not just obstacles but key regulators of when networks learn higher frequencies.

Methods

- We created a synthetic 1D dataset of sine waves with frequencies 1 Hz, 3 Hz, 7 Hz, and 15 Hz.
- We trained four ReLU-based Multilayer Perceptrons (MLPs): a baseline (2 layers, 50 units), a wider model (2 layers, 100 units), a deeper model (5 layers, 50 units), and the baseline with SGD instead of Adam.
- Models were trained for 10,000 epochs (learning rate 0.01), with loss monitoring, plateau detection via gradient methods, and Fourier analysis of predictions to track frequency learning and jumps.

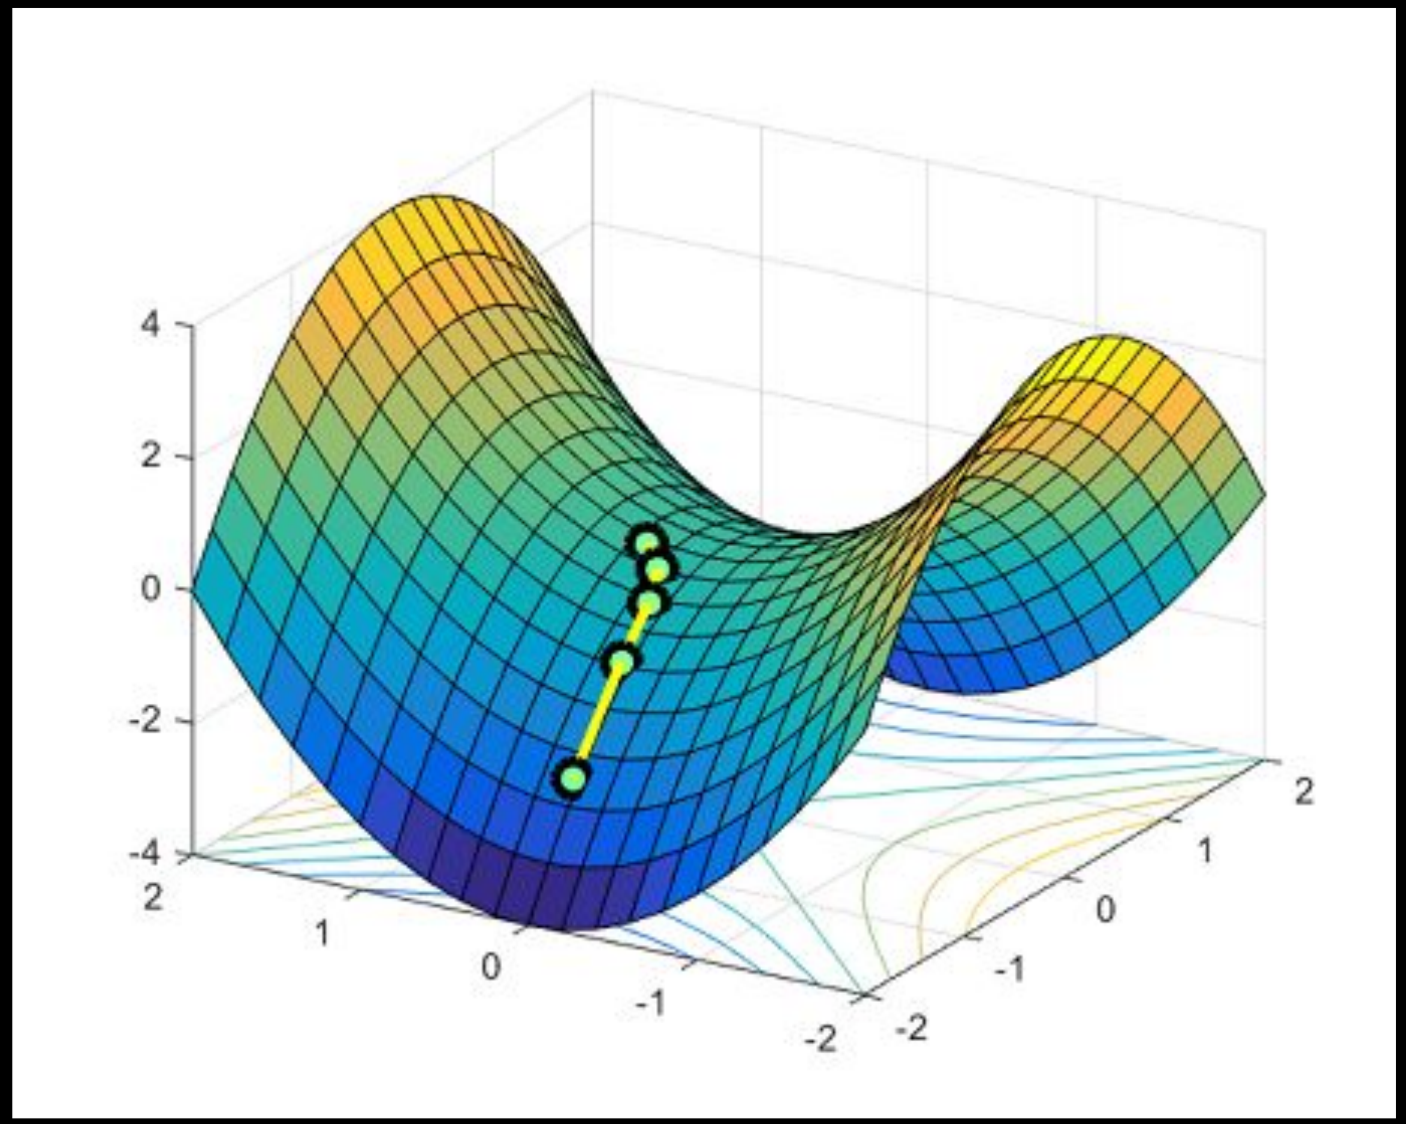

Fig 1. Visualization of a saddle point in the 3d loss landscape.

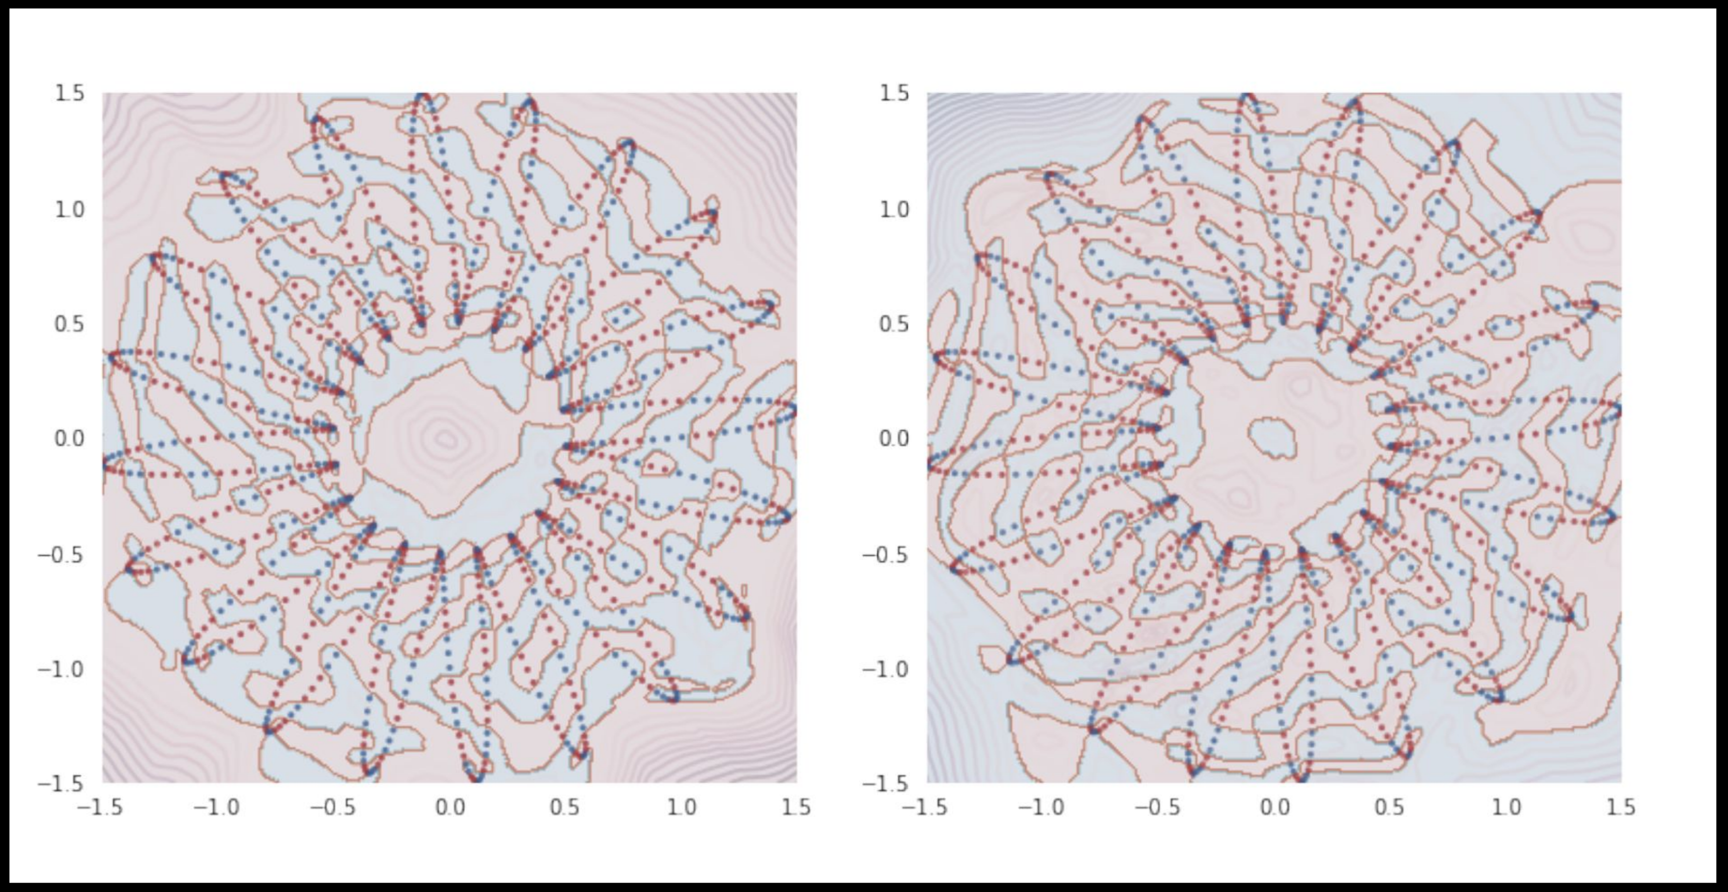

Fig. 2 Two identical networks trained on the same high-frequency sine wave classify training data correctly but generalize differently elsewhere, illustrating spectral bias, aligning with our finding that saddle points may delay high-frequency learning by trapping networks in low-complexity solutions.

Preliminary Results

- The network displayed clear spectral bias, mastering lower frequencies first.
- Plateaus appeared in the loss curve, during which higher-frequency learning stalled.
- Intriguingly, **jumps in frequency learning often occurred near plateau edges**, suggesting that escaping saddle points may trigger progress in capturing more complex, high-frequency features.

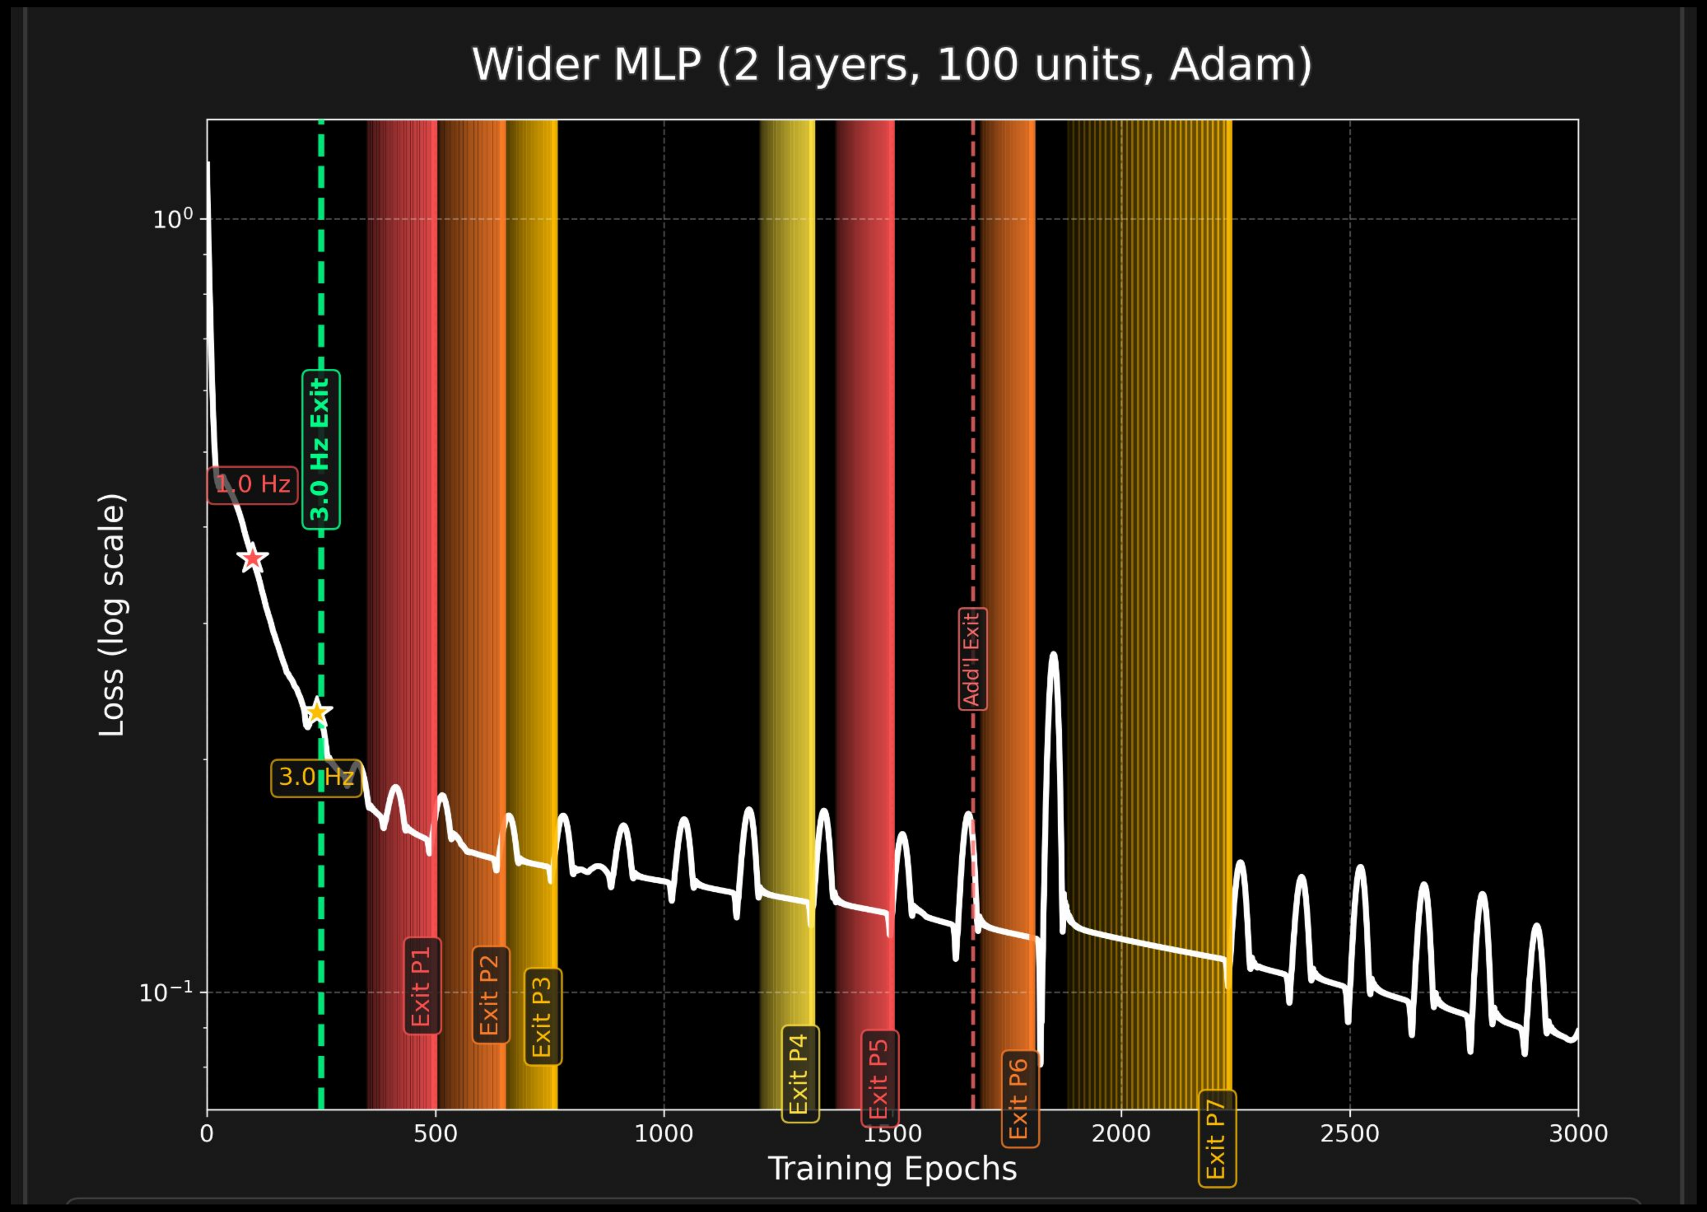

Fig 3. Loss curve with event markers highlighting a saddle point exit coinciding with the model's mastering of 3hz frequency.

Conclusion

- Findings suggest saddle points play a role in spectral bias.
- Plateaus reflect struggles with higher frequencies, and overcoming saddle points allows the network to advance.
- This connection could guide future efforts to enhance training efficiency, perhaps through better optimization methods or network designs.

Future Work

1. Investigating whether similar frequency-dependent learning stages exist in biological neural networks.
2. Developing a mathematical framework that unifies this idea across both artificial and biological learning systems.
3. Exploring how saddle points might represent phase transitions in information compression, connecting to theories of efficient coding in the brain where neural representations become increasingly abstract.

References

A.M. Saxe, J.L. McClelland, & S. Ganguli, A (2019) mathematical theory of semantic development in deep neural networks, Proc. Natl. Acad. Sci. U.S.A. 116 (23) 11537-11546, https://doi.org/10.1073/pnas.1820226116 .

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y. & Courville, A.. (2019). On the Spectral Bias of Neural Networks. *Proceedings of the 36th International Conference on Machine Learning*, in *Proceedings of Machine Learning Research* 97:5301-5310 Available from https://proceedings.mlr.press/v97/rahaman19a.html.