ShapeShop

Open-sourced on Github
github.com/fredhohman/shapeshop

Towards Understanding Deep Learning Representations via Interactive Experimentation

Fred Hohman
fredhohman@gatech.edu

Nathan Hodas
nathan.hodas@pnnl.gov

Polo Chau
polo@gatech.edu

Pacific Northwest NATIONAL LABORATORY

Georgia Tech

Summary

ShapeShop is an interactive system for visualizing and understanding what representations a neural network model has learned.

Built using standard web technologies, ShapeShop allows users to experiment with and compare deep learning models to help explore and understand the relationship between data and its learned representations.

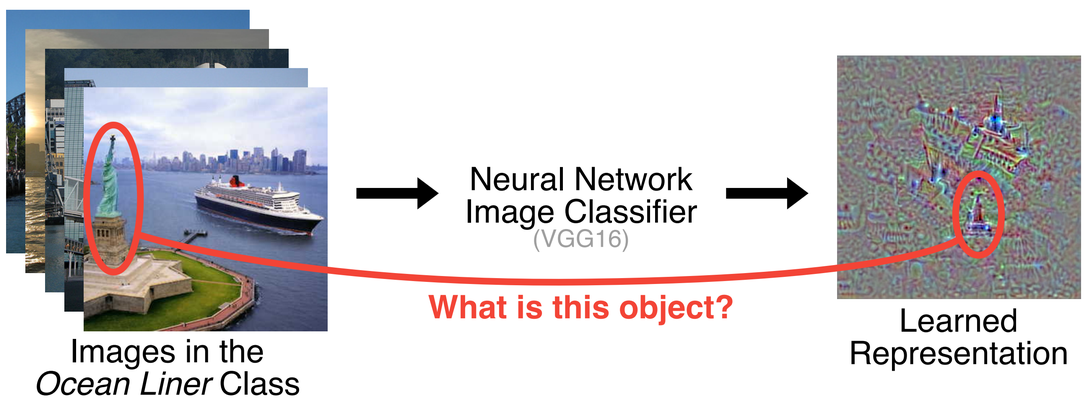

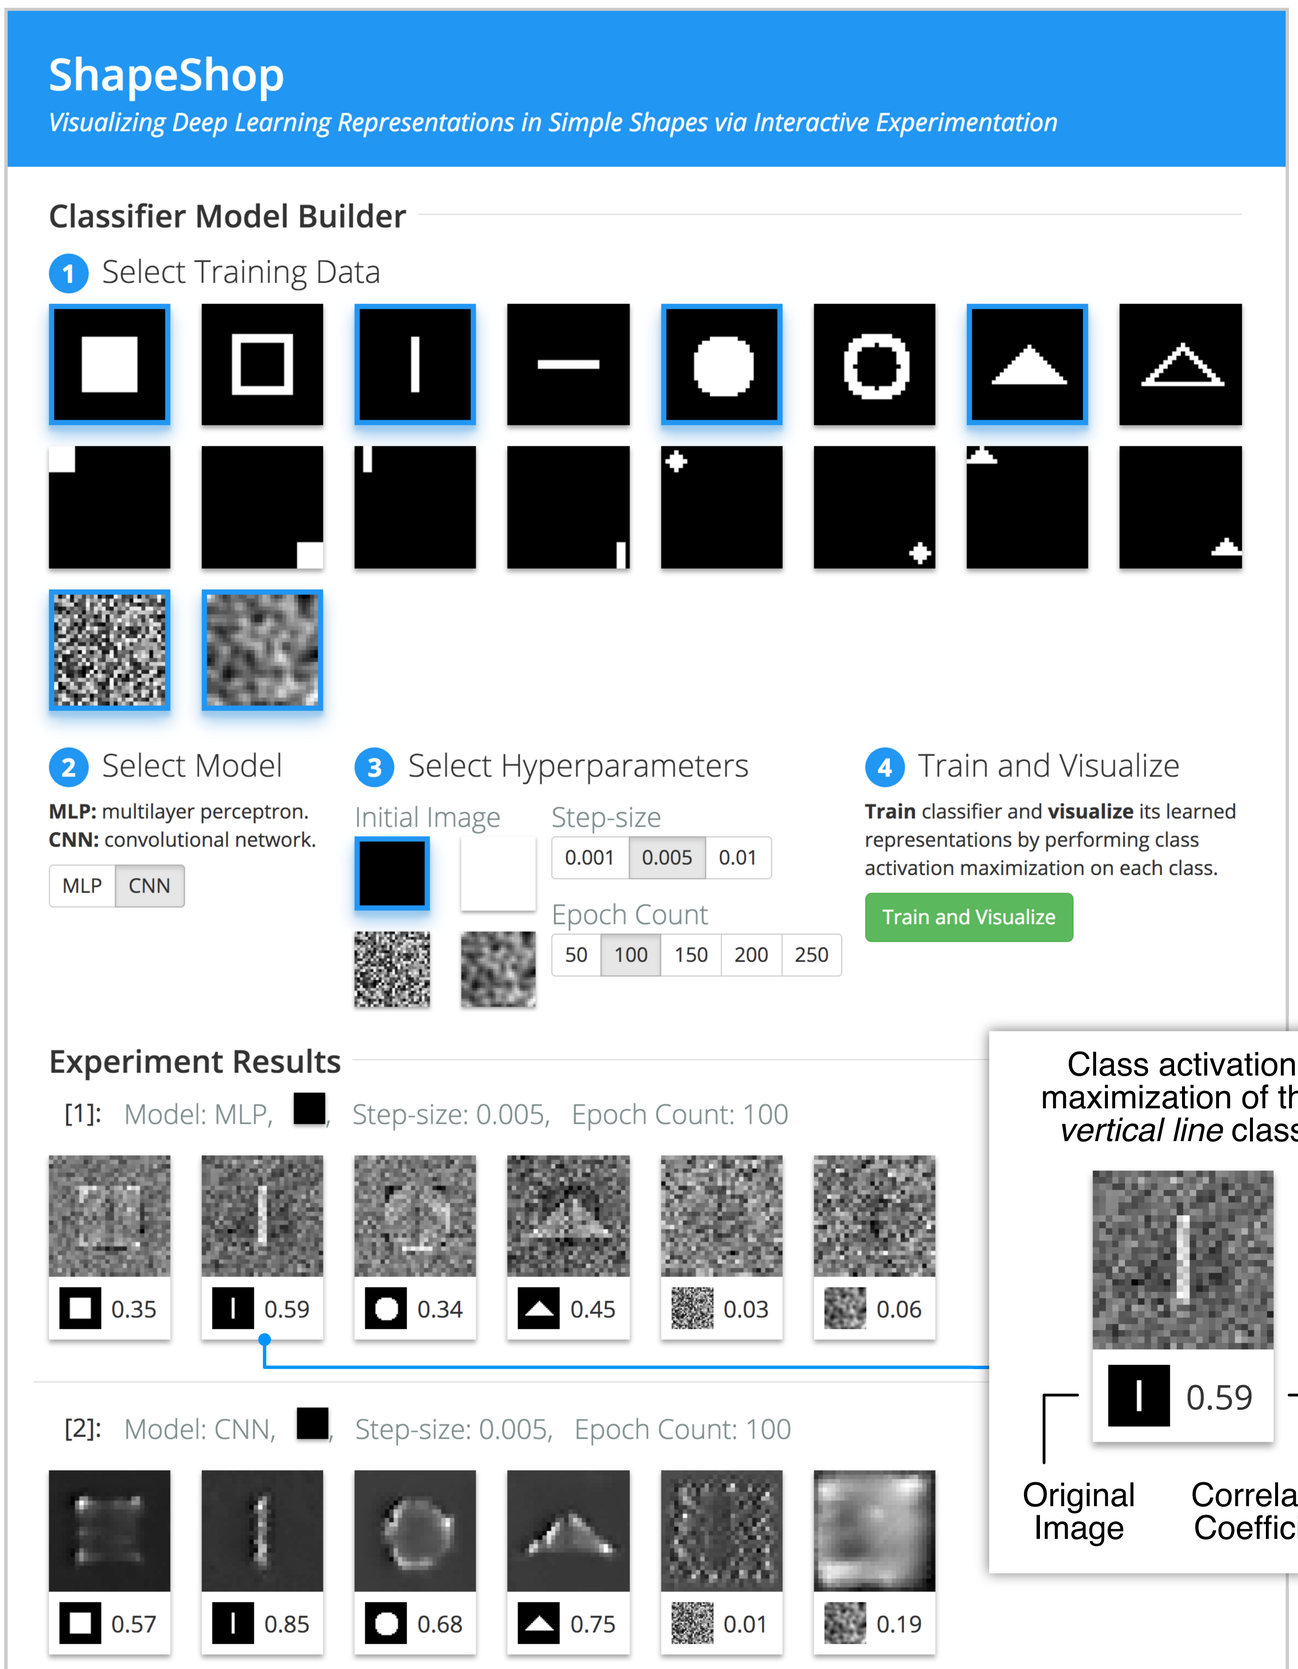

Example

Addition of *diverse data* produces more human recognizable representations.

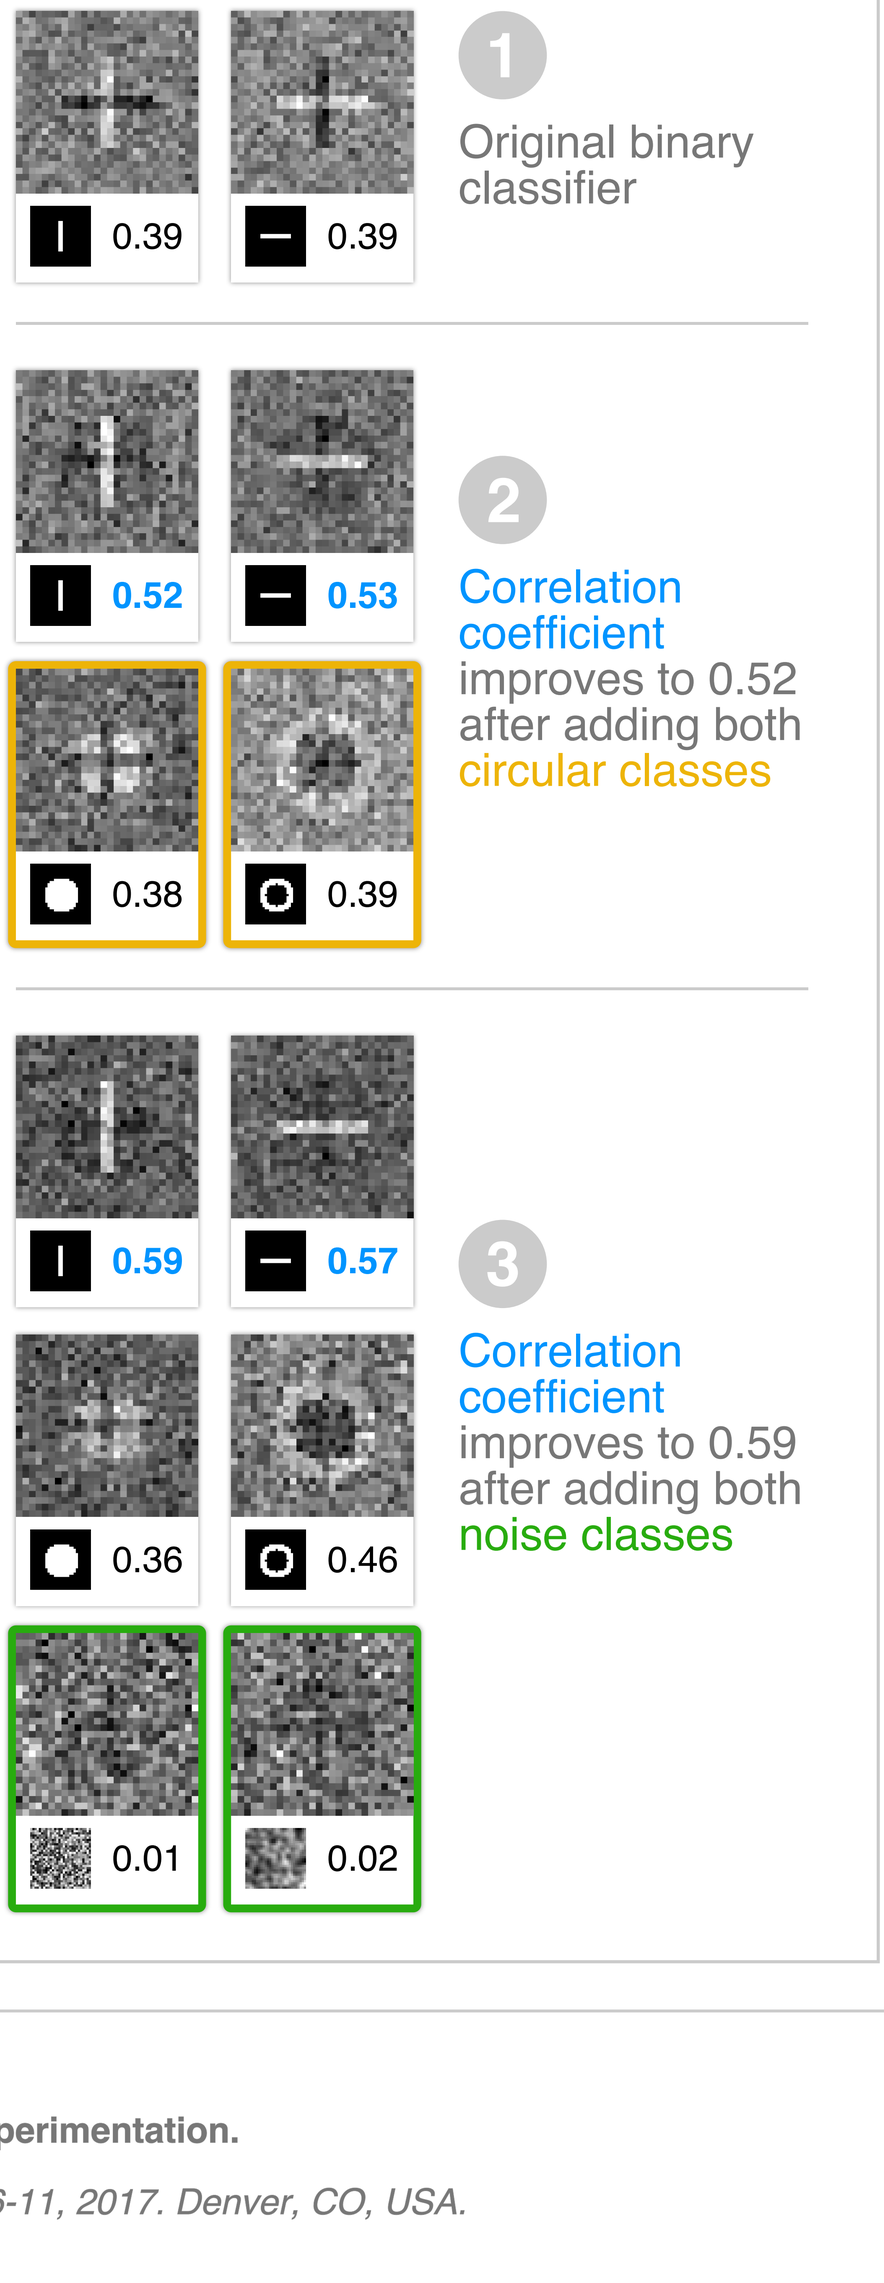

Ongoing Work

Incorporating human feedback to improve interpretability via semantic highlighting.

Yellow = ✓
Red = ✗


Ocean Liner Class

This work supported in part by grants NSF IIS-1563816 and NIH BD2K U54EB020404.

ShapeShop: Towards Understanding Deep Learning Representations via Interactive Experimentation.
Fred Hohman, Nathan Hodas, Duen Horng (Polo) Chau.
Extended Abstracts, ACM Conference on Human Factors in Computing Systems (CHI). May 6-11, 2017. Denver, CO, USA.