

# Interactive Classification for Deep Learning Interpretation

Ángel Cabrera, Fred Hohman, Jason Lin, and Duen Horng Chau

Georgia Institute of Technology, Atlanta GA, USA

{alex.cabrera, fredhohman, jlin401, polo}@gatech.edu

**Abstract.** We present an interactive system enabling users to manipulate images to explore the robustness and sensitivity of deep learning image classifiers. Using modern web technologies to run in-browser inference, users can remove image features using inpainting algorithms and obtain new classifications in real time, which allows them to ask a variety of “what if” questions by experimentally modifying images and seeing how the model reacts. Our system allows users to compare and contrast what image regions humans and machine learning models use for classification, revealing a wide range of surprising results ranging from spectacular failures (e.g., a *water bottle* image becomes a *concert* when removing a person) to impressive resilience (e.g., a *baseball player* image remains correctly classified even without a glove or base). We demonstrate our system at The 2018 Conference on Computer Vision and Pattern Recognition (CVPR) for the audience to try it live. Our system is open-sourced at <https://github.com/poloclub/interactive-classification>. A video demo is available at <https://youtu.be/llub5GcOF6w>.



**Keywords:** Interactive classification · inpainting · interpretable deep learning · image classification

## 1 Introduction

Public trust in artificial intelligence and machine learning is essential to its prevalence and widespread acceptance. To create trust, both researchers and the general public have to understand why models behave the way they do. Existing research has used interactive data visualization as a mechanism for humans to interface with black-box machine learning models [4], revealing how models learn and behave using an interactive dialogue [9] as opposed to static explanations.

### 1.1 Interactive Classification for Interpretation

We have designed and developed an interactive system that allows users to experiment with deep learning image classifiers and explore their robustness and sensitivity. Users are able to remove selected areas of an image in real time with classical computer vision inpainting algorithms, Telea [8] and PatchMatch [2], which allows them to ask a variety of “what if” questions by experimentally

					
Modified Image			Original Image		
Class	↓ Confidence %	Absolute % Change	Class	Confidence %	
liner, ocean liner	51.48	+35.56	dock, dockage, docking facility	52.00	
dock, dockage, docking facility	38.05	-13.95	drilling platform, offshore rig	28.05	
drilling platform, offshore rig	6.53	-21.52	liner, ocean liner	15.92	
container ship, container ship, container vessel	1.46	+1.27	schooner	1.59	
fireboat	0.97	+0.85	pirate, pirate ship	1.40	

**Fig. 1.** The modified image (left), originally classified as *dock* is misclassified as *ocean liner* when the masts of a couple boats are removed from the original image (right). The top five classification scores are tabulated underneath each image.

modifying images and seeing how the model reacts [3]. Allowing a human user to select the regions to be inpainted enables more semantically meaningful regions to be considered for study compared to other techniques where shapes (e.g., gray squares) are placed over random parts of an image to subdue signal to the image classifier [10]. Some existing work uses automated image segmentation techniques, e.g., superpixels [1], and a combinatorial search algorithm to find the most important superpixels; however, there is no guarantee the superpixels are semantically meaningful features to consider [7].

## 1.2 Significance of Our Approach

Through interactive inpainting, our system helps reveal a wide range of surprising results ranging from spectacular failures (e.g., a *water bottle* image becomes a *concert* when removing a person) to impressive resilience (e.g., a *baseball player* image remains correctly classified even without a glove or base). The system also computes class activation maps on demand, which highlight the important semantic regions of an image a model uses for classification [11]. Combining these tools, users can develop qualitative insight into what a model sees and which features impact an image’s classification. Our system is open-sourced at <https://github.com/poloclub/interactive-classification>. A video demo is available at <https://youtu.be/llub5GcOF6w>.

## 2 Usage Examples

### 2.1 Misclassification

In Figure 1 we present an example of our interactive classification system with a deep learning image classifier that exhibits failures when given semantic edits. Consider an image of a couple cars driving next to boats moored at a dock. The original top 5 classification results for this image are shown on the right in Figure 1. The classifier originally correctly produces a top class of *dock*. Using our system, the user highlights the masts of the boats, inpainting them and replacing them with sky. The edits do not alter the meaning of the picture, but unfortunately the classification changes to an inaccurate class. With the masts removed, the modified image (Figure 1 left) is now misclassified as an *ocean liner* instead of the correct classification of a *dock* (Figure 1 right).

Modified Image			Original Image	
Class	↓ Confidence %	Absolute % Change	Class	Confidence %
ballplayer, baseball player	48.34	+36.38	baseball	66.60
baseball	25.47	-41.13	ballplayer, baseball player	11.96
racket, racquet	10.96	+6.20	racket, racquet	4.76
tennis ball	1.77	-0.58	tennis ball	2.35
crutch	1.56	-0.66	crutch	2.22

**Fig. 2.** The modified image (left), originally classified as *baseball* and *ballplayer* remains correctly classified when the ball, glove, and base are removed from the original image (right). The top five classification scores are tabulated underneath each image with decreasing confidence.

### 2.2 Resilient Classification

In Figure 2 we present an example of our interactive classification system with a deep learning image classifier that exhibits resilience to semantic edits. Consider an image instance of a baseball player with a glove throwing a ball at base. The original top 5 classification results for this image are shown on the right in Figure 2. The top 2 class are *baseball* and *ballplayer*. Using our system, the user

highlights the baseball, glove, and base, inpainting and replacing them with the brown dirt. Surprisingly, the modified image remains classified as *baseball player* when these important features are removed from the original image (Figure 2 right), albeit with adjusted probabilities (Figure 2 left).

### 3 System Design

To provide an easily accessible experience, we created an in-browser system by combining multiple novel web technologies. React.js forms the base of the system, providing an interactive and approachable interface that encourages image manipulation. The other core technology is Tensorflow.js, which runs deep learning models in the browser and can classify images in near real-time. We used the SqueezeNet [6] and MobileNet [5] deep learning models for Tensorflow.js since they are designed to be portable and fast. Lastly, we make two different inpainting algorithms available to the user: Telea [8] and PatchMatch [2]. The Telea algorithm uses a fast, heuristic-based method and is implemented fully in JavaScript, allowing the system to be run completely in-browser. For better, more advanced inpainting we utilize the GIMP Resynthesizer plugin, which is an open-source implementation of the PatchMatch algorithm. This is implemented in C and exposed to the system through a Python server.

### 4 Conclusion

We present an interactive system that enables users to explore the robustness and sensitivity of deep learning image classifiers. Using modern web technologies to run in-browser inference and compute class activation maps, users can inpaint images in real time, to ask a variety of “what if” questions by experimentally modifying images and seeing how the model reacts. Our system helps reveal a wide range of surprising classification results ranging from spectacular failures to impressive resilience. Our system is open-sourced. We believe our investigation will help people explore the extent to which humans and machines think alike, and shed light on the advantages and potential pitfalls of deep learning image classifiers.

### Acknowledgments

This work was supported by NSF grants IIS-1563816, CNS-1704701, and TWC-1526254; NASA Space Technology Research Fellowship; and gifts from Google, Symantec, Yahoo, Intel, Microsoft, eBay, Amazon.

### References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence* **34**(11), 2274–2282 (2012)

2. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics-TOG* (2009)
3. Hohman, F., Hodas, N., Chau, D.H.: ShapeShop: Towards Understanding Deep Learning Representations via Interactive Experimentation. In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM (2017)
4. Hohman, F., Kahng, M., Pienta, R., Chau, D.H.: Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2018)
5. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861* (2017)
6. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: SqueezeNet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv:1602.07360* (2016)
7. Ribeiro, M.T., Singh, S., Guestrin, C.: Why should i trust you?: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1135–1144. ACM (2016)
8. Telea, A.: An image inpainting technique based on the fast marching method. *Journal of graphics tools* (2004)
9. Weld, D.S., Bansal, G.: Intelligible Artificial Intelligence. *arXiv:1803.04263* (2018)
10. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: *European conference on computer vision*. pp. 818–833. Springer (2014)
11. Zhou, B., Khosla, A., A., L., Oliva, A., Torralba, A.: Learning Deep Features for Discriminative Localization. *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)