



# Scene Classification with Residual CNN's


Josh King, Vayu Kishore, Filippo Ranalli

## Problem Statement

Given a picture of a scene, determine which of 365 classes it belongs to.

## Dataset


- 365 classes
- 5000 training examples per class
- 100 validation examples per class
- 900 test examples per class
- Resolution of 256 x 256
  - We scale down to 64 x 64 due to memory limitations
- Extra "Challenge" data exists
  - 25,000 training examples per class
- Higher Resolution images exist
  - 512+ x 512+



Home Office



Archaeological Excavation



Beach

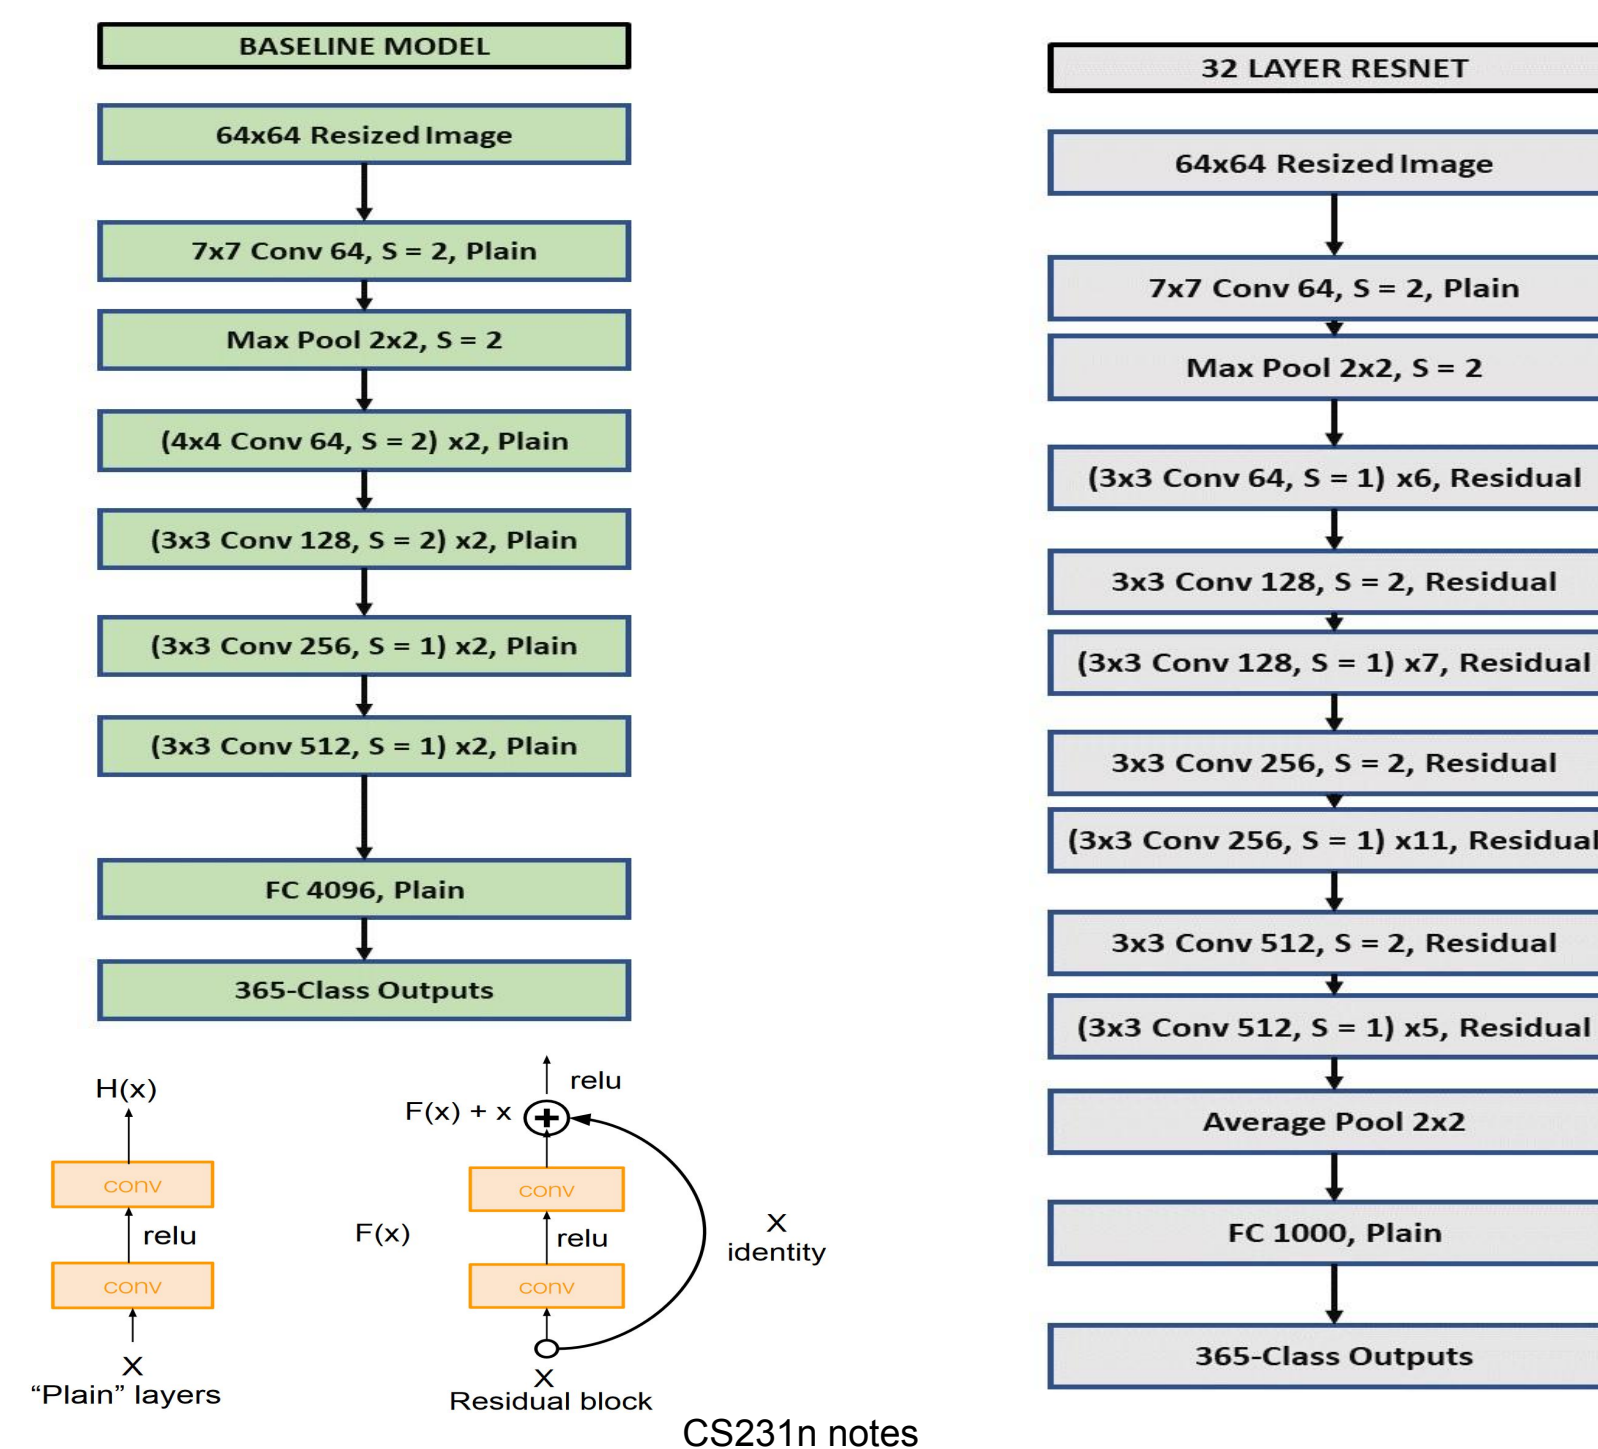


Sandbar

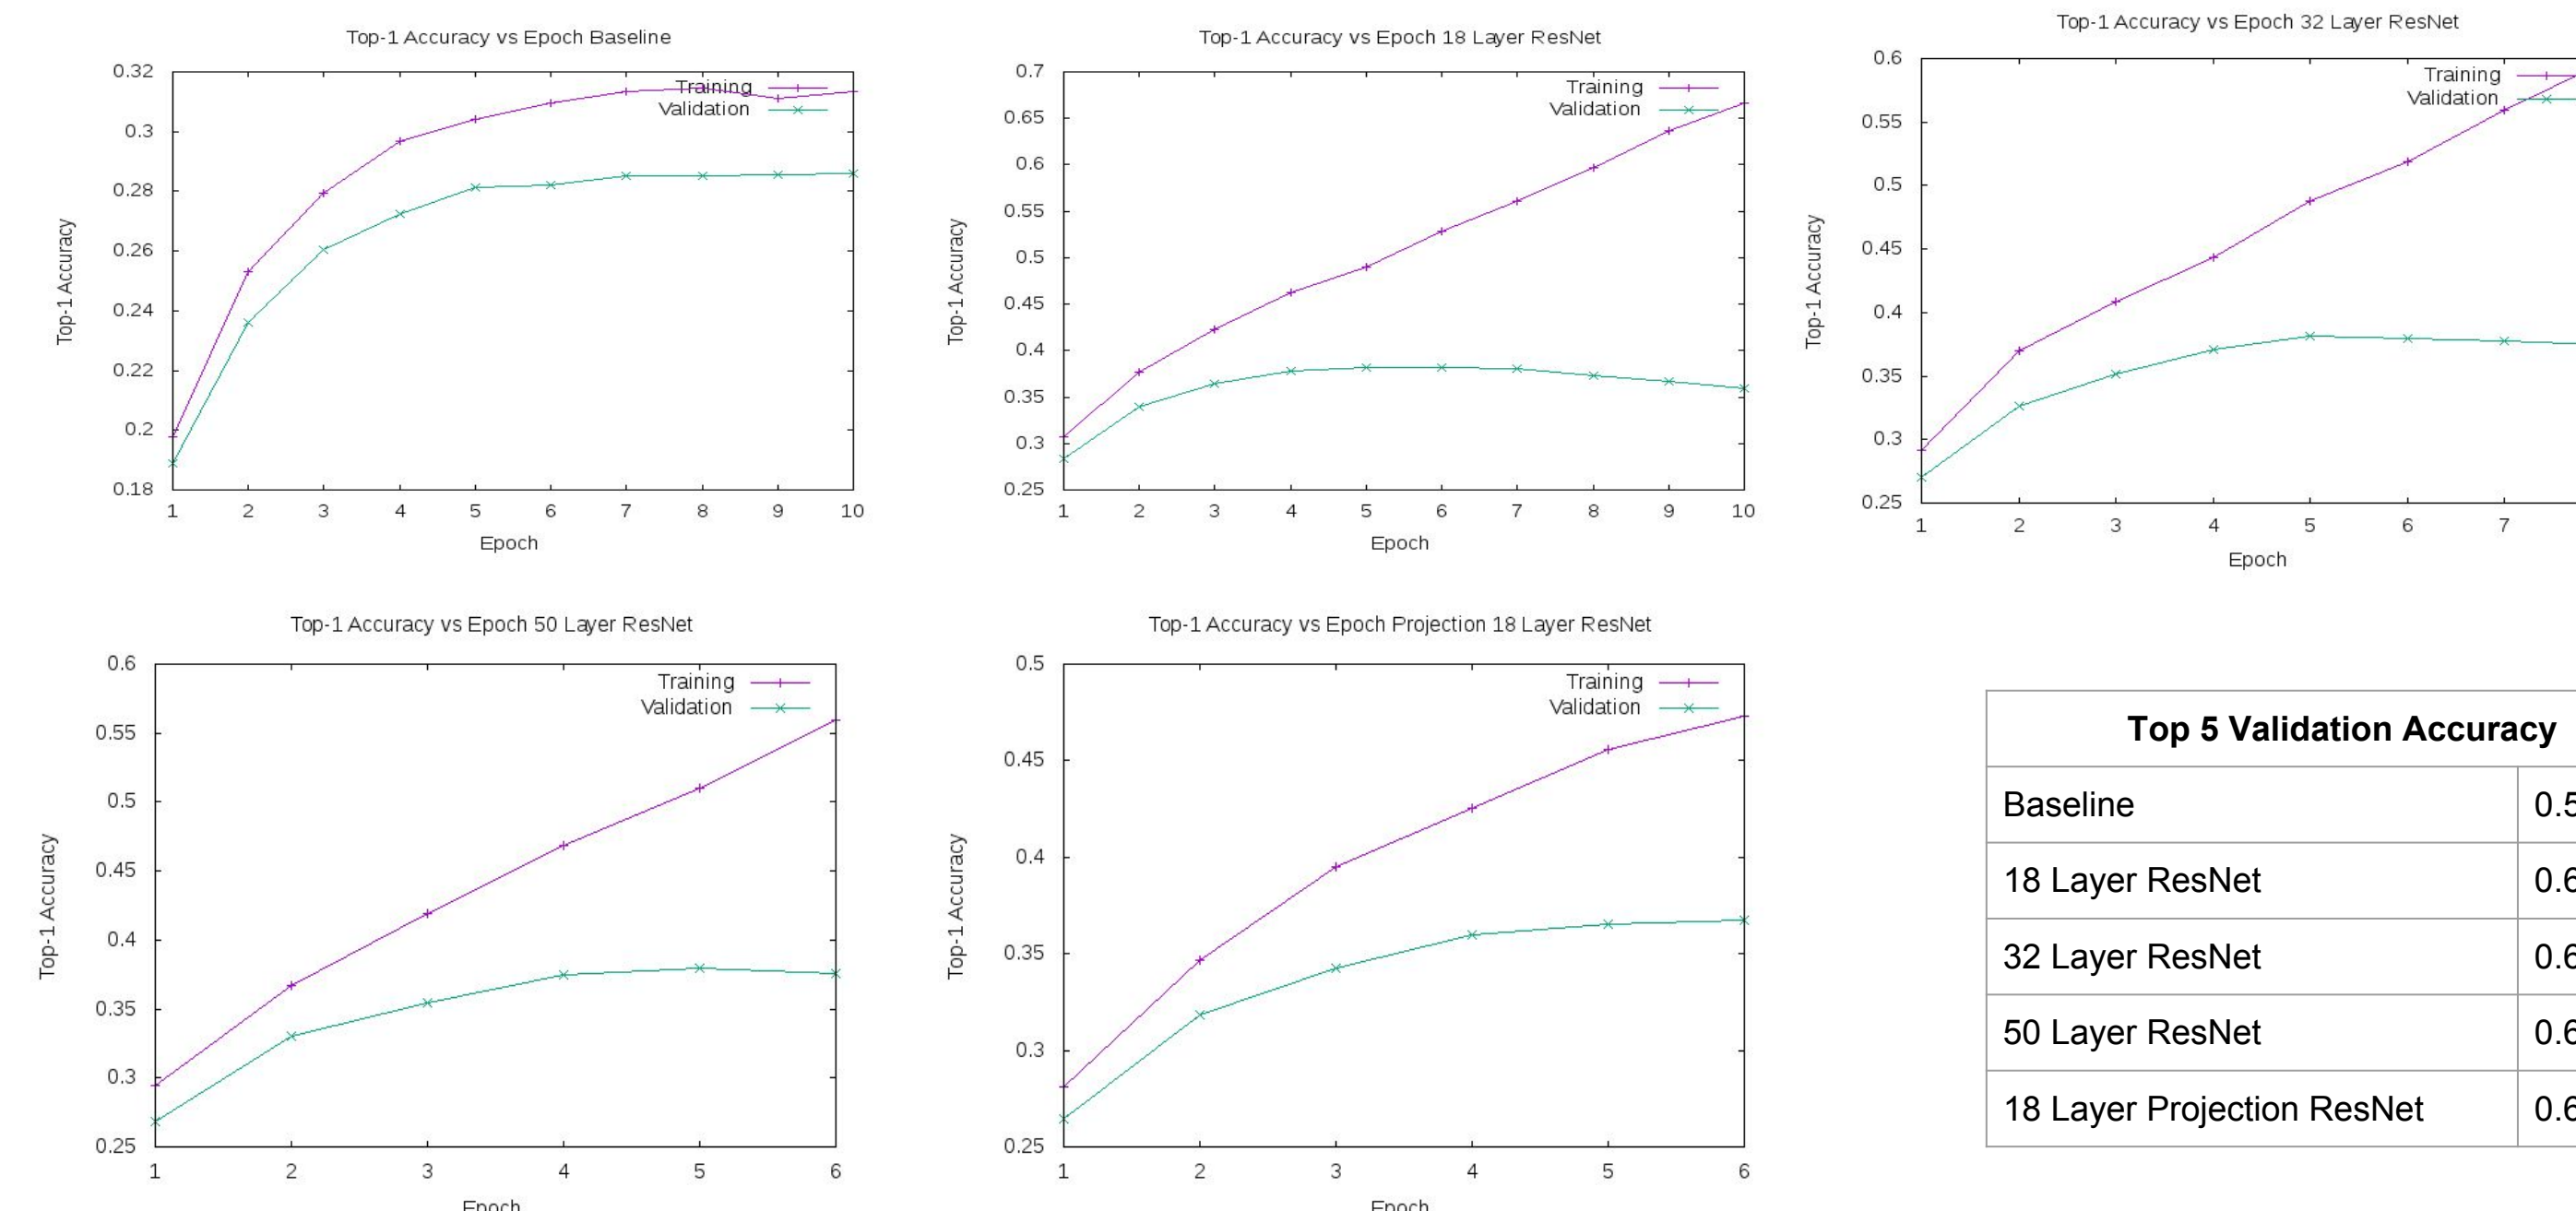
## References

- CS231n 2017 class notes.
- He, Zhang, Ren, Sun, "Deep Residual Learning for Image Recognition".
- Zhou, Khosla, Lapedriza, Torralba, Oliva, "Places: An Image Database for Deep Scene Understanding"
- <http://places2.csail.mit.edu/>

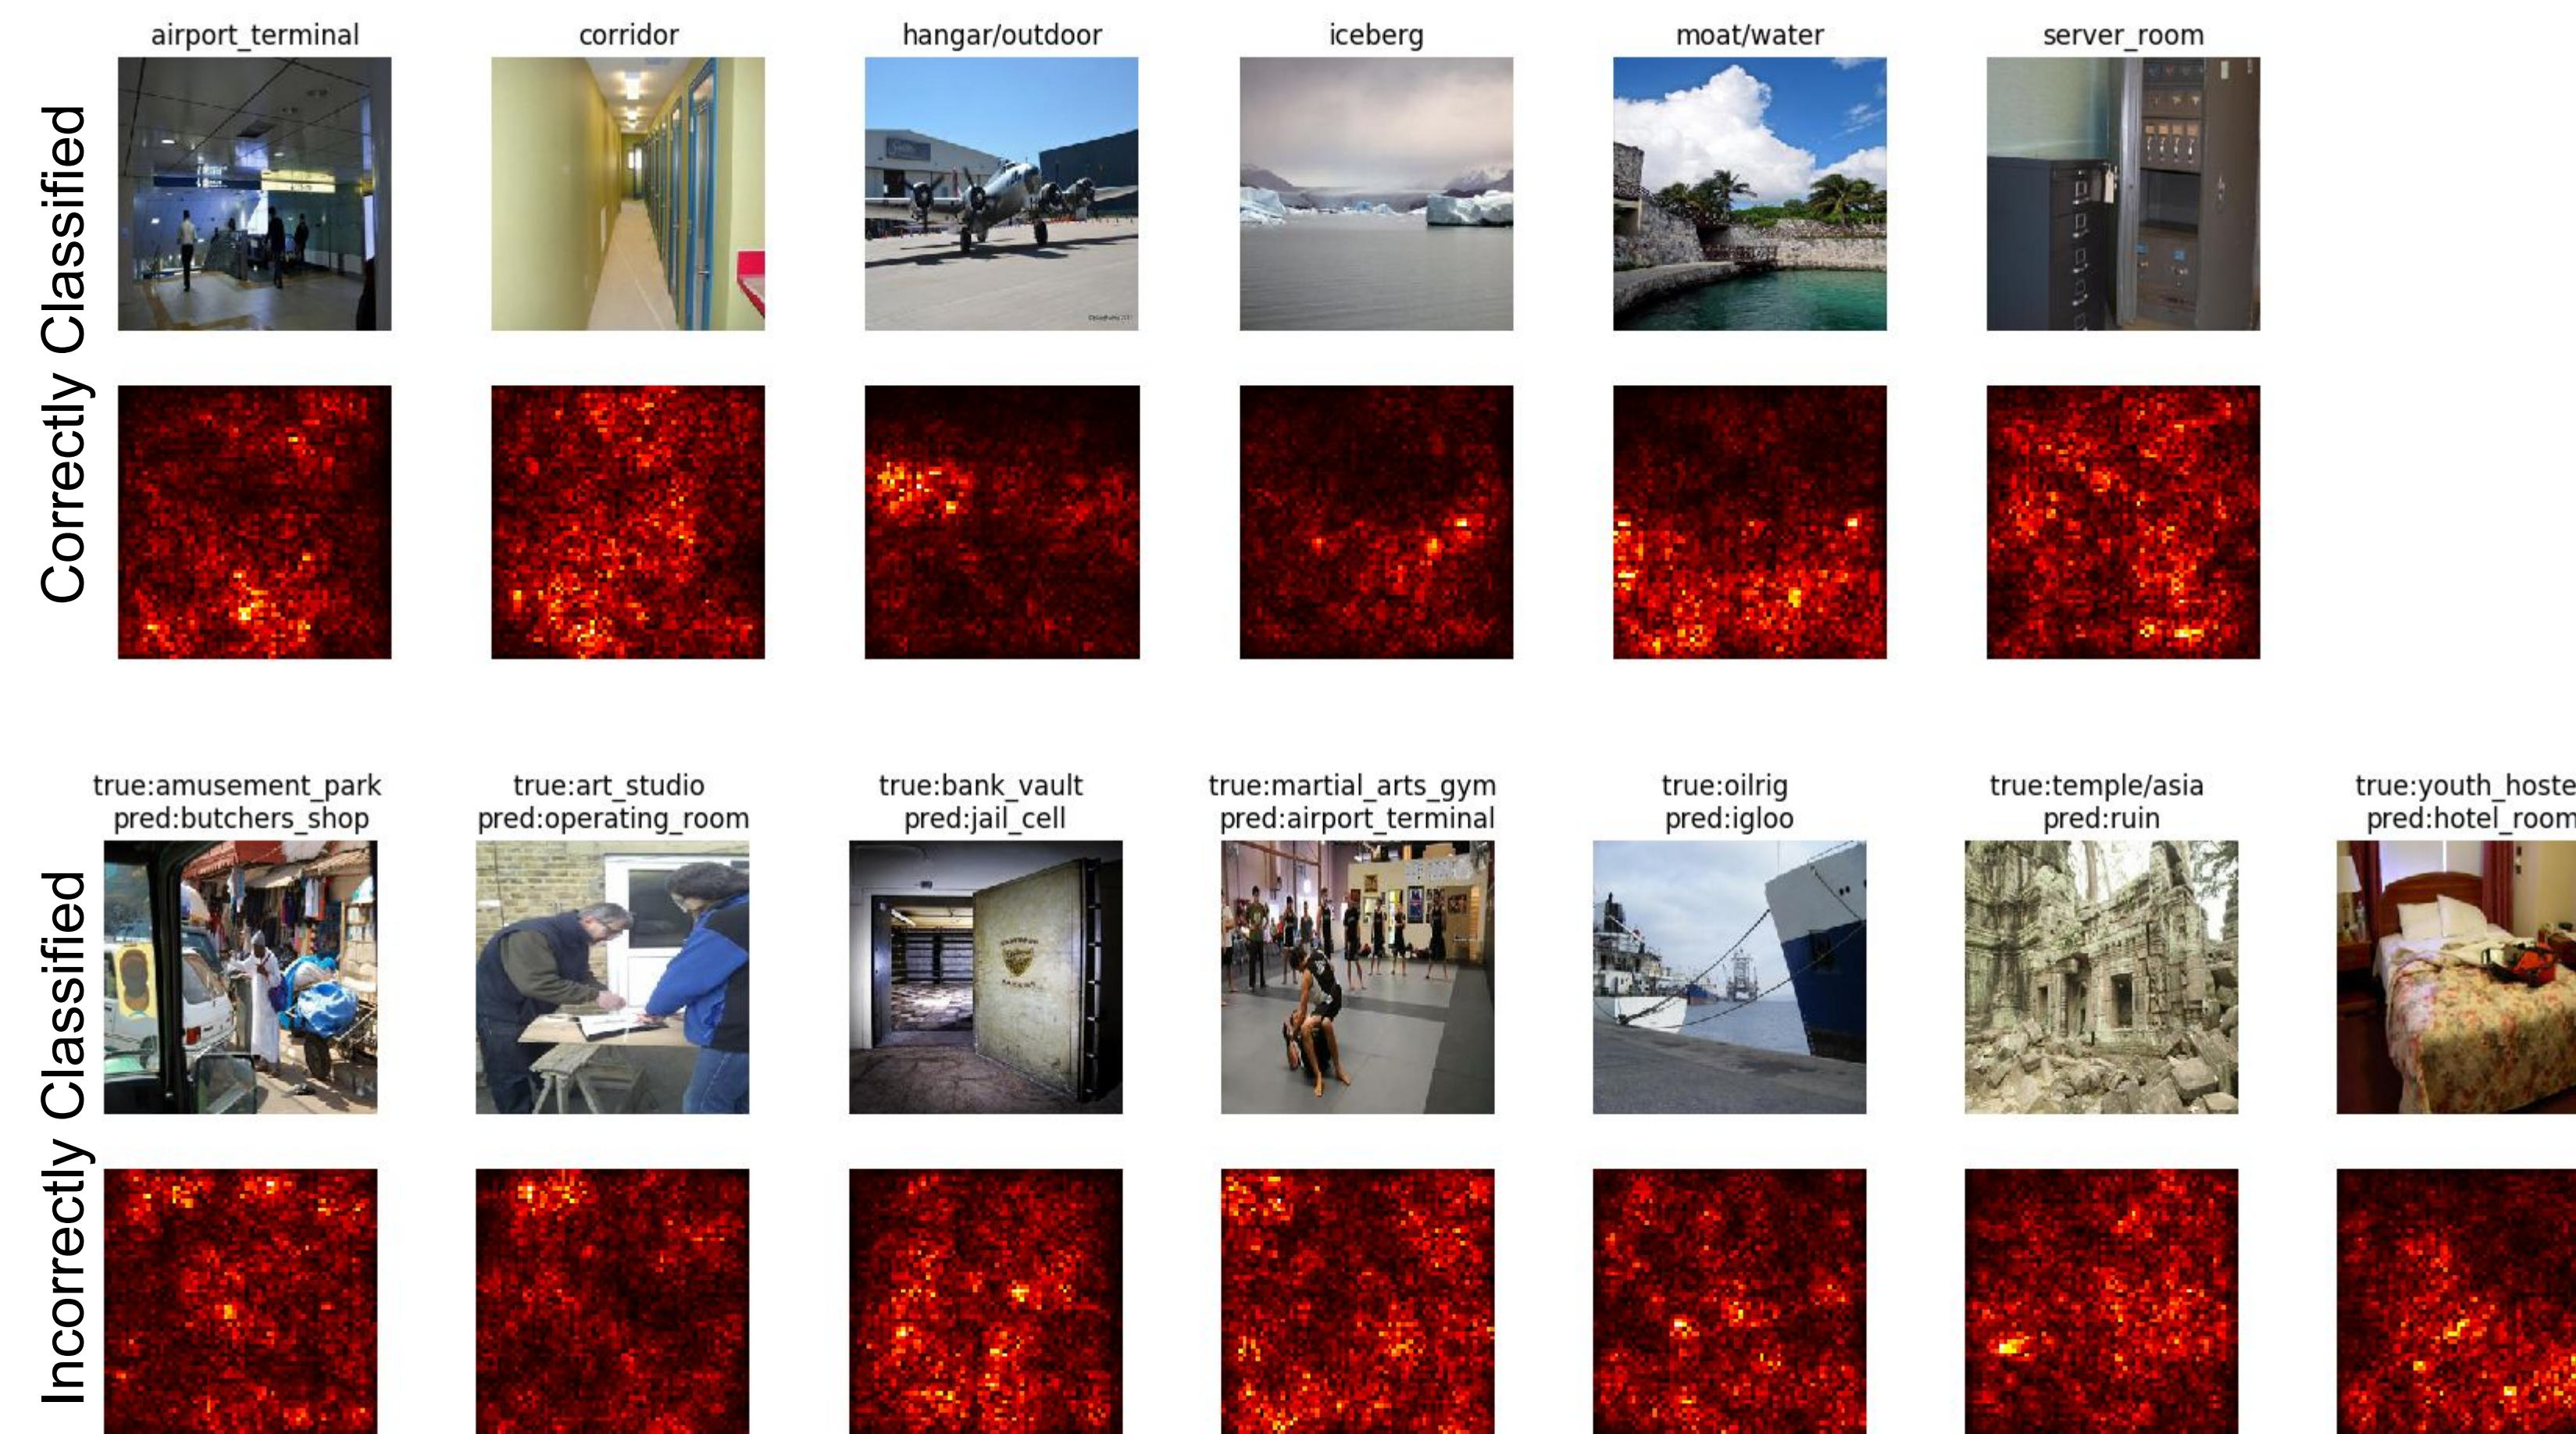
## Models



## Training and Validation Accuracy



## Example Results and Saliency Maps



## Discussion

- State of the Art classifiers get 90% top 5 accuracy, much higher than our models.
- We attribute this our inability to train on entire higher resolution dataset.
- Learning curves indicate overfitting, even in the presence of strong L2 regularization and dropout.
- Found top-5 val. accuracy to be more informative than top-1 as classes may be ambiguous (hotel vs hostel) and are not mutually exclusive (ex: temple and ruin).
- Based on visual inspection, significant portion of misclassification seem reasonable.
- Some labelling issues in validation dataset.
- Saliency maps seem generally more diffuse than saliency maps for object detection, indicating that the classification generated by model is sensitive to the entire image. Some exceptions where presence of a particular object defines a scene.