A-Faster R-CNN: Generating Hard Positive Examples via Adversary for Traffic Sign Detection

Stephanie Milani and Christoph Mertz

PROBLEM

- Traffic signs must be well-maintained to keep roads safe.
- Road infrastructure inventory and assessment systems that assist with traffic sign maintenance must be able to detect occluded traffic signs.
- **Problem:** How do we robustly model invariances to rare occlusion events?
  - Gather more data? – No, too tedious and time-consuming.
  - Generate all possible occlusions? – No, impossible!

SOLUTION

- We do not have to generate **all** possible occlusions, just difficult ones [1] [2].
- **Proposed Solution:** Generate hard positive examples of occlusions using an adversary.
  - Goal of the detector: Accurately classify the sign in the image.
  - Goal of the adversary: Create examples of occluded signs that are good enough to trick the object detector into misclassifying the sign.
- Integrate with network used by Navlab (Faster R-CNN) for their road infrastructure inventory and assessment system.

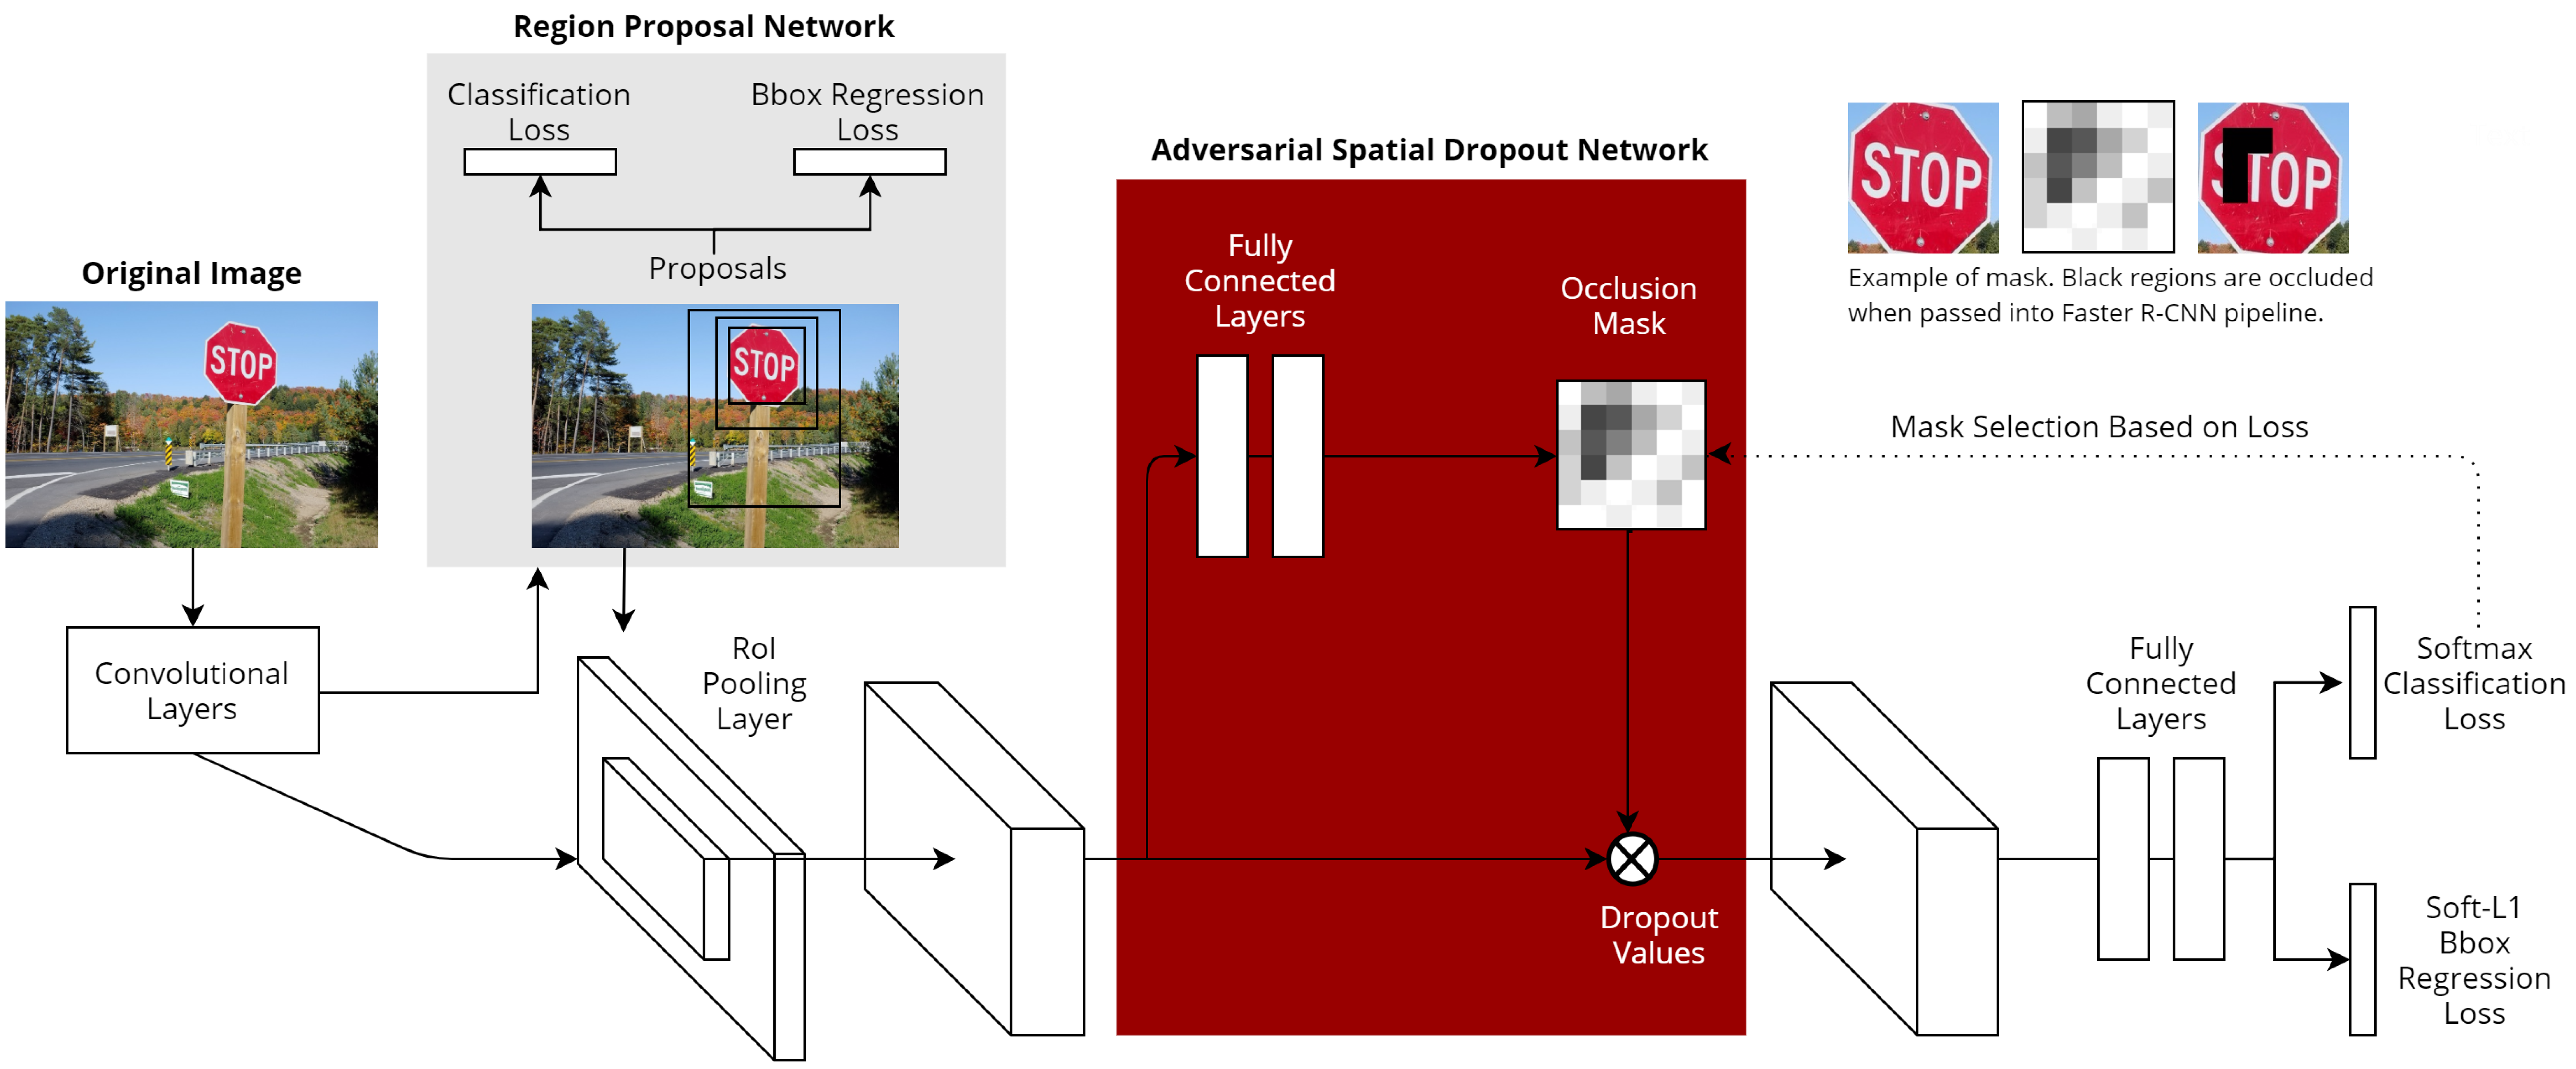

Example of mask. Black regions are occluded when passed into Faster R-CNN pipeline.

Network Design: The architecture of Faster R-CNN integrated with the Adversarial Spatial Dropout Network (ASDN). Given the input of region-based convolutional features proposed by the Region Proposal Network (RPN), the ASDN generates an occlusion mask to indicate which parts of the features to dropout.

METHODS

- Combine the Faster R-CNN [3] architecture with the Adversarial Spatial Dropout Network proposed by Wang et al. [4] to create A-Faster R-CNN.
- Train and test A-Faster R-CNN on the LISA dataset [5].
- Test A-Faster R-CNN on the Navlab dataset – specifically on occluded stop signs that were initially missed by the detector.
- Evaluate the performance of the new network: Which cases is it still unable to classify? Which cases is it now able to classify?

WORKS CITED

[1] A. Shrivastava, A. Gupta, and R. Girshick, "Training region-based object detectors with online hard example mining" in *CVPR*, 2017.

[2] M. Takáč, A. Bijral, P. Richtárik, and N. Srebo, "Mini-batch primal and dual methods for svms" in *Proceedings of the 30th International Conference on Machine Learning*, vol. 28, 2013.

[3] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks" in *CVPR*, 2015.

[4] X. Wang, A. Shrivastava, and A. Gupta, "A-Fast-RCNN: Hard positive generation via adversary for object detection" in *CVPR*, 2017.

[5] A. Møgelmose, M. Trevedi, and T. Moeslund, "Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey" in *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 4, 2012.

FUTURE WORK

- Train and test method on more traffic sign datasets.
- Increase network robustness to sign discoloration and distortion by allowing adversary to further manipulate input.
- Incorporate top-down methods, such as prior knowledge of sign locations.
- Extend to real-time traffic sign detection.

ACKNOWLEDGMENTS

Thank you to Jina Wang, Dat Nguyen, the rest of the Navlab, and Nicholay Topin for their support. Special thanks to Traffic21 for sponsoring this work through the Women in Transportation Fellowship.

traffic21

Carnegie Mellon University