

Supplementary Material

Identifying building locations in the wildland–urban interface before and after fires with convolutional neural networks

Neda K. Kasraee^{A,}, Todd J. Hawbaker^B and Volker C. Radeloff^A*

^ASILVIS Lab, Department of Forest and Wildlife Ecology, University of Wisconsin-Madison, 1630 Linden Drive, Madison, WI 53706, USA

^BUS Geological Survey, Geosciences and Environmental Change Science Center, Denver, CO 80225, USA

*Correspondence to: Email: kasraee@wisc.edu

Disclaimer

Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

Esri CNN Model Arguments

The model has five main arguments: (1) padding, (2) batch size, (3) threshold, (4) return bounding boxes, and (5) tile size.

The padding is the bordering area where the model discards detections. Changing the padding will ultimately change the stride, the amount of movement over each image. We kept the default value of 128 for all three fires.

The batch size indicates the number of image tiles the GPU will process at one time during inferencing. Inferencing refers to the process where a trained model, such as the one used here, makes predictions against new data. The model default is four. The batch size may be limited by GPU memory. The GPU we used ran out of memory with a batch size of four, so we applied a batch size of two for all three fires.

The model outputs a level of confidence for each building prediction, and the threshold assigns the minimum level of confidence the model output must have. For example, if the threshold is set to 0.6, a building is detected only if the model is at least 60% confident the feature is a building. Threshold was the most influential argument, varying for each fire and each year. Once all the other parameters were set, we found that the default of 0.9 was too high for each fire and year. We adjusted the value in increments of .05 at first until we started getting high levels of detection and then adjusted in increments of 0.01 until we found a threshold with the highest accuracy. We found that a building foundation is sometimes left behind after a building is destroyed by a fire, and the foundation sometimes appears like an intact building in

aerial imagery. To prevent false detections of destroyed building foundations, the threshold was made higher post-fire.

The return bounding boxes parameter is either True or False. If set to True, a bounding box is returned around the building. Since we were interested in the buildings only and not the bounding boxes, we kept the default value, False.

For tile size, we used the default value of 512. Adjusting the tile size did not noticeably improve the model results.