



Disaster Response: Leveraging Wide Area Earth Observation with Machine Learning

Olga Isupova (1), Danil Kuzin (2), Brooke Simmons (3), and Steven Reece (1)

(1) Department of Engineering Science, Oxford University, Oxford, UK (reece@robots.ox.ac.uk, olga.isupova@eng.ox.ac.uk),
(2) Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, UK
(dkuzin1@sheffield.ac.uk), (3) Department of Physics, University of Lancaster, Lancaster, UK (b.simmons@lancaster.ac.uk)

We co-create machine learning products with UK and in-country commercial, GO and NGO partners to ensure the machine learning algorithms address appropriate user needs whether for tactical decision making or evidence-based policy decisions. In one particular case, we developed and deployed a novel algorithm, BCCNet, to quickly process large quantities of satellite imagery in response to natural disasters. Crowdsourcing provides an efficient mechanism to generate labels to prime machine learning algorithms for large scale data analysis. However, these labels are often imperfect with qualities varying among different citizen scientists, which prohibits the labels' direct use with many state-of-the-art machine learning techniques. BCCNet simultaneously aggregates biased and contradictory labels from the crowd and trains an automatic classifier to process new data. It integrates a convolutional neural network with a Bayesian probabilistic classifier combination algorithm ("Dynamic Bayesian combination of multiple imperfect classifiers", E. Simpson et al., 2013). In our work LeNet-5 with the Adam optimiser was chosen as the base neural network. A small amount of data labelled by hand through crowdsourcing platforms like Zooniverse (<https://www.zooniverse.org>) is used to train the neural network.

Our algorithm has been developed and deployed in collaboration with Rescue Global, a UK based not-for-profit, to generate damage heatmaps for disaster responders immediately following Hurricanes Irma and Maria (2017) and earlier versions following earthquakes in Nepal (2015) and Ecuador (2016). These heatmaps were passed to the UN, FEMA and over 60 NGOs during the response phase of Irma and Maria in a timely manner.

We analysed crowdsourced labels of damage from Digital Globe high resolution (30cm) optical satellite imagery of Dominica before and after Hurricane Maria in 2017. Crowd members were presented with a subset of satellite sub-images and asked to draw bounding boxes around all buildings and also mark building damage. We obtained 32,932 objects labelled by 13 volunteers (6 volunteers labelled each object on average). We extracted image patches from both 'before' and 'after' imagery corresponding to the bounding boxes as input for the neural network. Before and after image patches formed different channels of the neural network input layer. We defined ground truth as the crowd consensus output inferred using the whole dataset. We then divided the dataset in the ratio 70-10-20% into training, validation and test datasets. Our results were obtained from the trained neural network on the held-out test datasets over 30 Monte Carlo runs with random initialisation. The classification accuracy was $83 \pm 1\%$ overall, with $91 \pm 1\%$ correctly identified as background, $82 \pm 2\%$ correctly identified undamaged buildings and $91 \pm 2\%$ damaged buildings (see "BCCNet: Bayesian classifier combination neural network", O. Isupova et al., 2018 for further details).

This work was funded by the Alan Turing Institute's Data Centric Engineering programme and also through the UK Space Agency's International Partnerships Programme. The authors would like to thank Digital Globe, Planet and ESA for ongoing satellite data provision.