

Supplemental Document 1: Results of Traffic Sign Detection

This supplemental document provides sample results for detection of traffic signs using the method in the paper entitled “Traffic Sign Detection and Classification in the Wild”. Each input image has resolution 2048×2048 . In each result row, the image on the left shows detected signs marked in a rectangle. Green rectangles indicate correctly detected signs, red rectangles indicate incorrectly detected signs (which are not actual traffic signs), and blue rectangles indicate signs undetected by our network. On the right is the corresponding probability heat map for the image. Blue corresponds to a low probability that a pixel contains a sign, while red corresponds to high probability. Signs are identified by setting a threshold on the probability map and using OpenCV’s efficient implementation of `groupRectangles`. We empirically set the threshold to 0.95, but carefully choosing this value could potentially further improve the results.

Below, we give 8 pairs of contrasting conditions. In each, we show two images illustrating ‘easier’ cases, and two ‘harder’ cases: for example, smaller signs are harder to correctly detect than larger signs. The results generally indicate that our methods work well on more difficult cases as well as easier ones. Nevertheless, we also give some cases in which the method has failed to deliver the correct result.

1 EASY AND DIFFICULT PAIRS

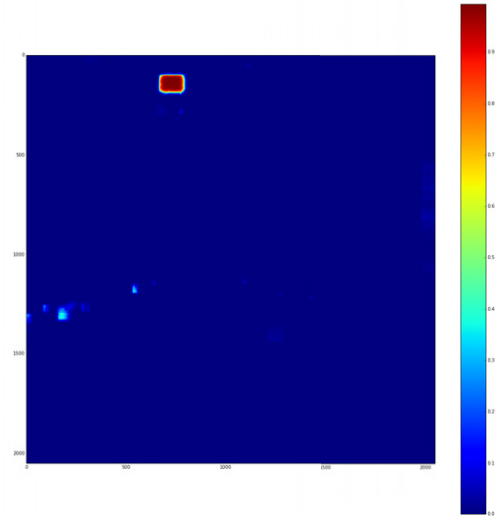
For each pair, we first give two ‘easier’ cases followed by two ‘harder’ cases.

1.1 Countryside and city

Images (a) and (c) were captured in the countryside, while (e) and (g) images were captured in the downtown region of a city. In the former, there are few other things that might be confused with traffic signs. In the latter, there are many advertisements, logos, other signs, etc., that might be confused with traffic signs, yet traffic signs are still correctly detected , and false positives rarely arise.



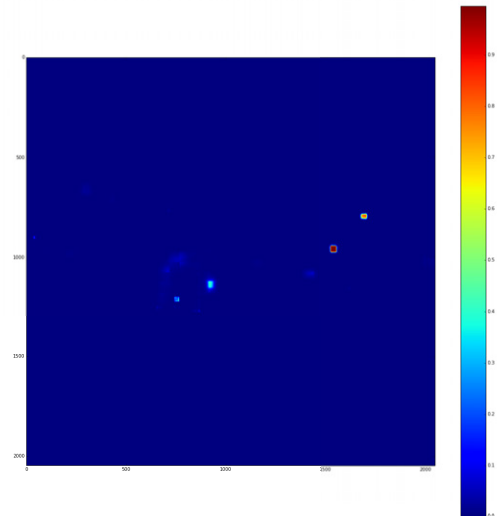
(a)



(b)



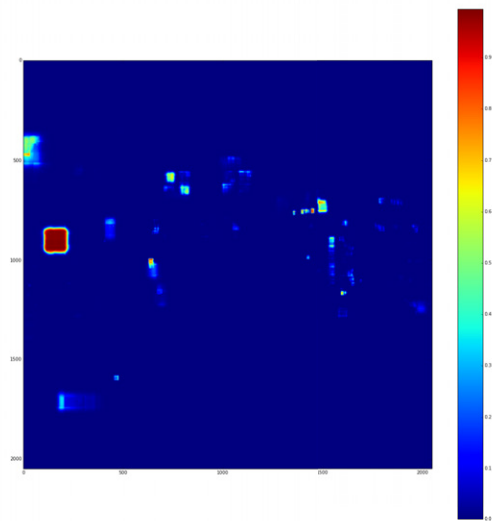
(c)



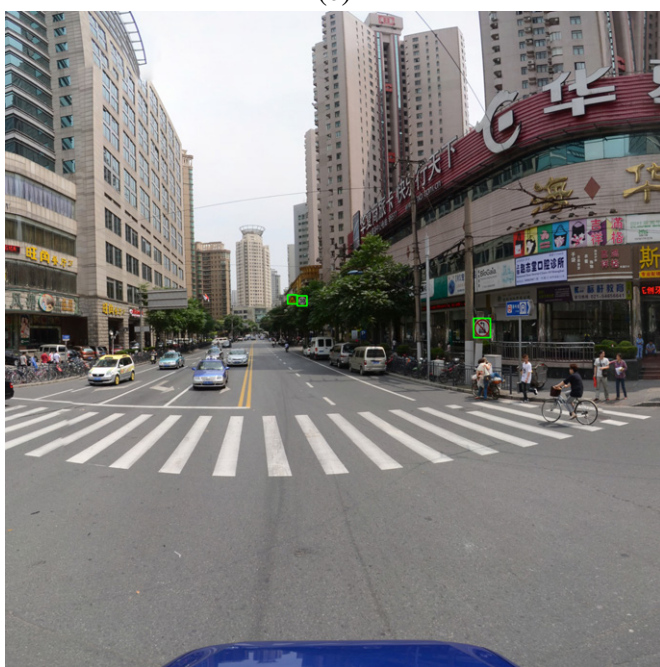
(d)



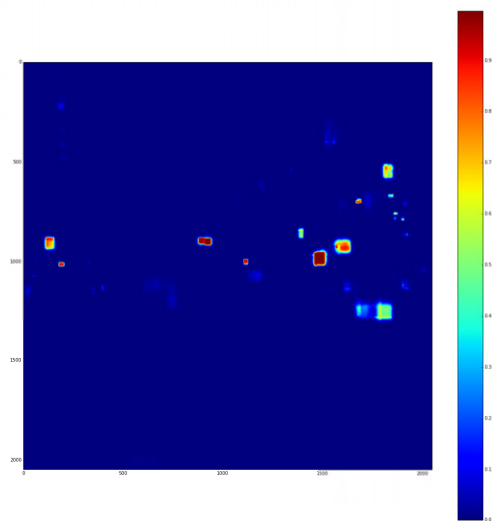
(e)



(f)



(g)



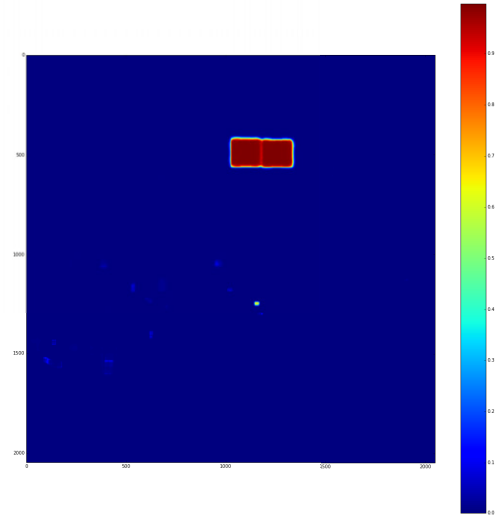
(h)

1.2 Large and small

Traffic signs in (a) and (c) are of large size so are easily detected. Those in (e) and (g) are much smaller, yet most signs of this size are still correctly detected. Only a very small mandatory sign in (e) is undetected—while the probability of that area containing a sign is high, it is discarded by bounding box regression step.



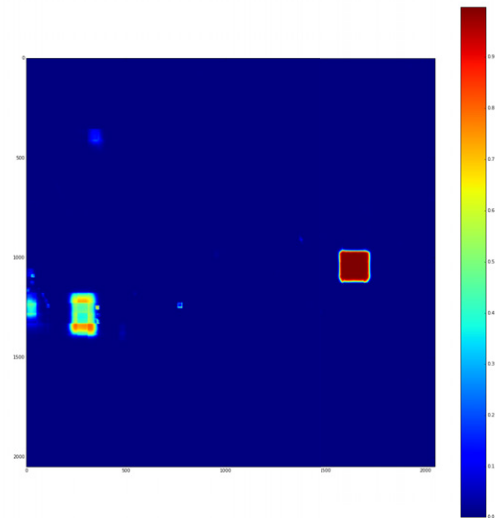
(a)



(b)



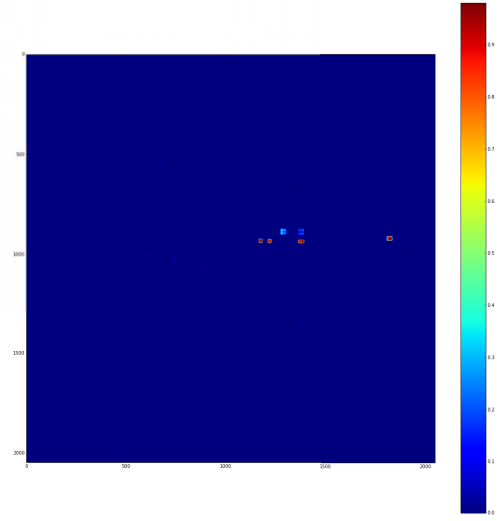
(c)



(d)



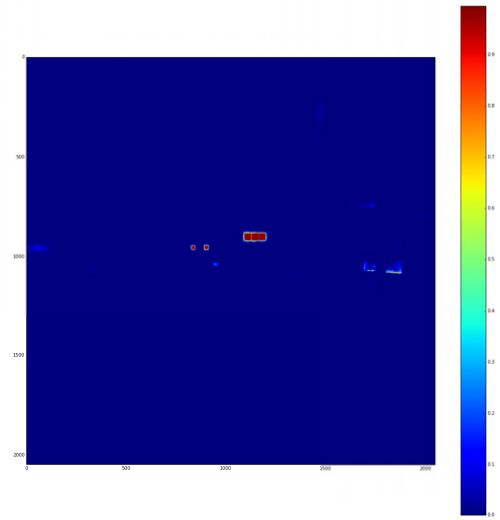
(e)



(f)



(g)



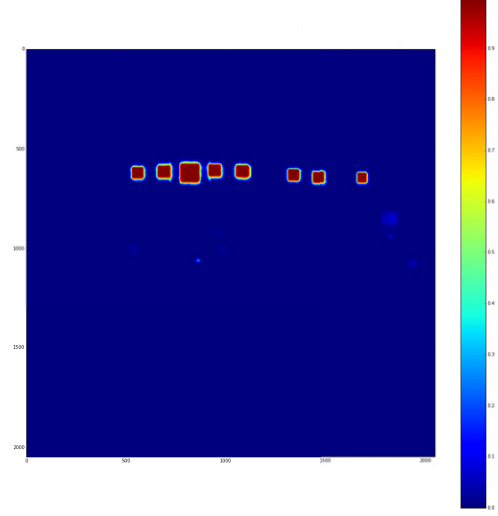
(h)

1.3 Frontal lighting and lighting from behind

Images (a) and (c) were captured under frontal lighting, while (e) and (g) were captured with lighting from behind. Images of the former kind generally have stronger colours and better contrast than the latter, Even under more challenging illuminance conditions, our network still performs well. In (g) the warning sign is not of a type in the standard list of signs, and remains undetected as many pixels have probability less than 0.95.



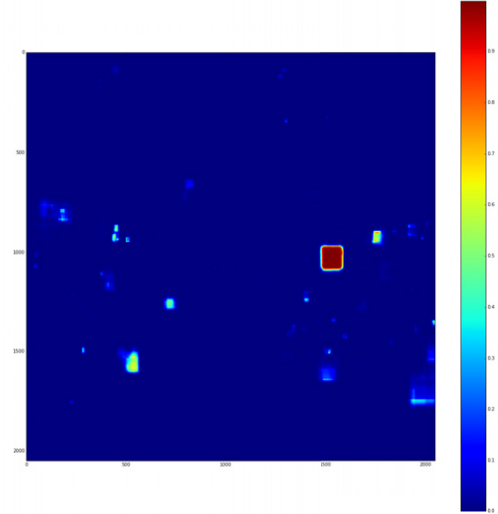
(a)



(b)



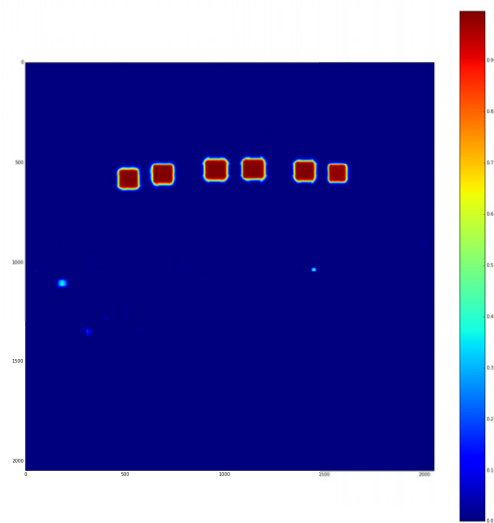
(c)



(d)



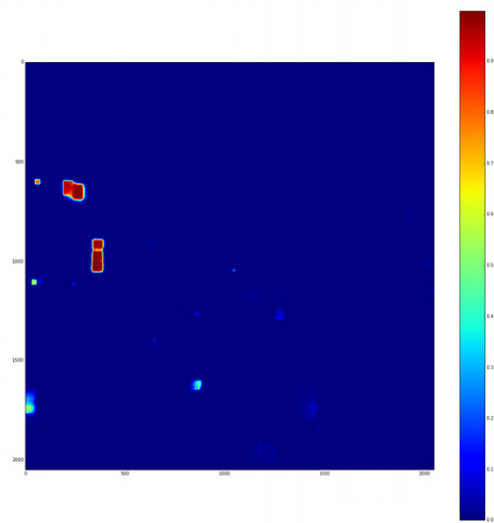
(e)



(f)



(g)



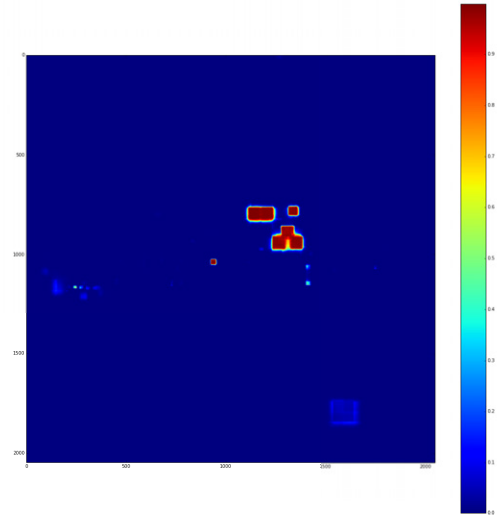
(h)

1.4 Orthogonal and oblique views

Traffic signs in (a) and (c) are more or less orthogonal to the direction to the camera, so our network has no difficulty in correctly detecting them. In (e) and (g), the signs are obliquely presented to the viewing direction. Nevertheless, our network still performs well in the latter cases.



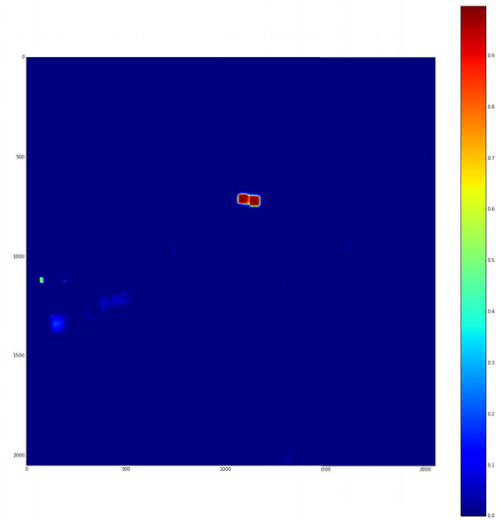
(a)



(b)



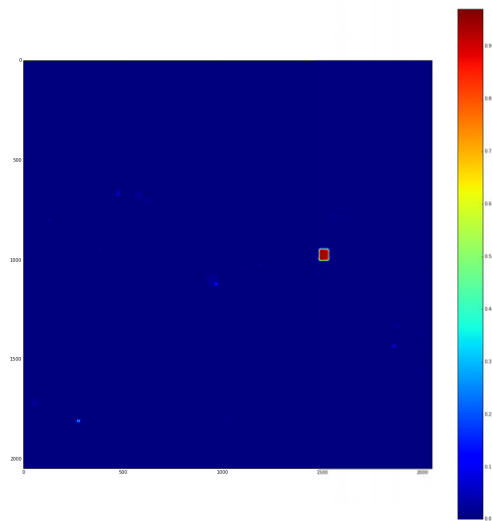
(c)



(d)



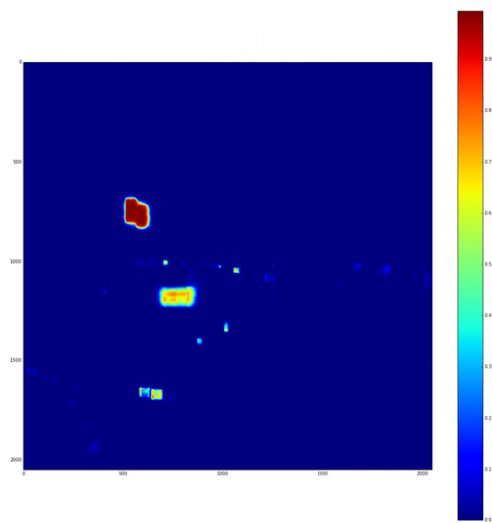
(e)



(f)



(g)



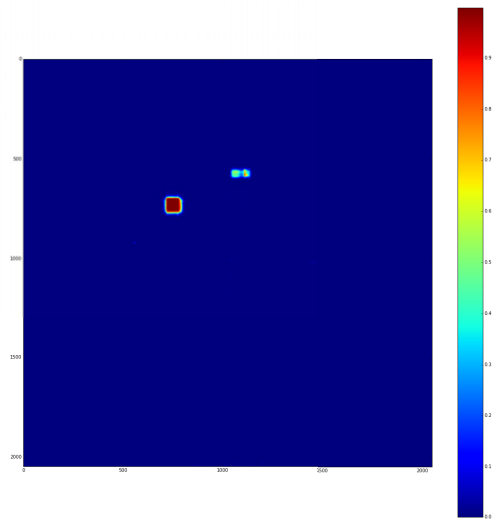
(h)

1.5 Single and multiple adjacent signs

There is only one traffic sign in each of images (a) and (c) while in (e) and (g) there are multiple traffic signs, some of which are almost touching. All traffic signs are correctly detected in all of these images—multiple adjacent traffic signs can be correctly handled.



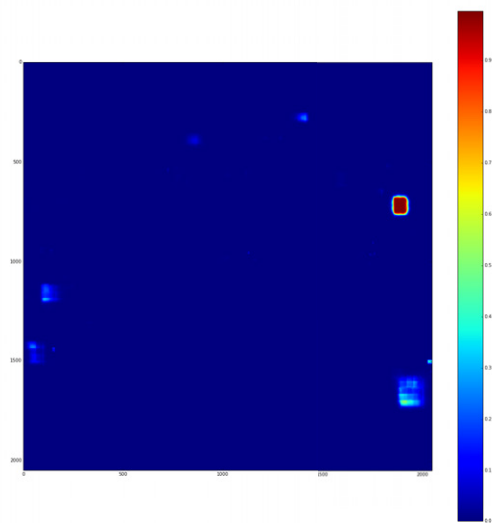
(a)



(b)



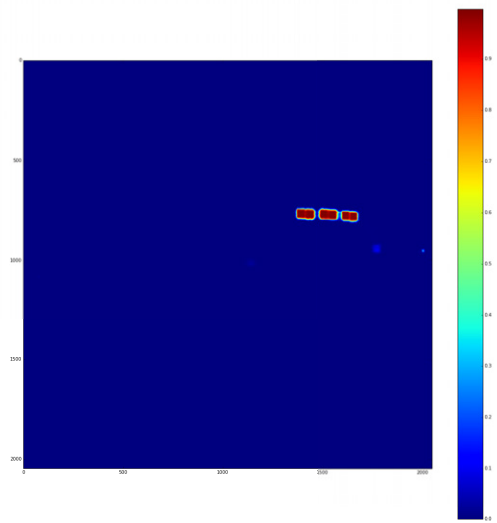
(c)



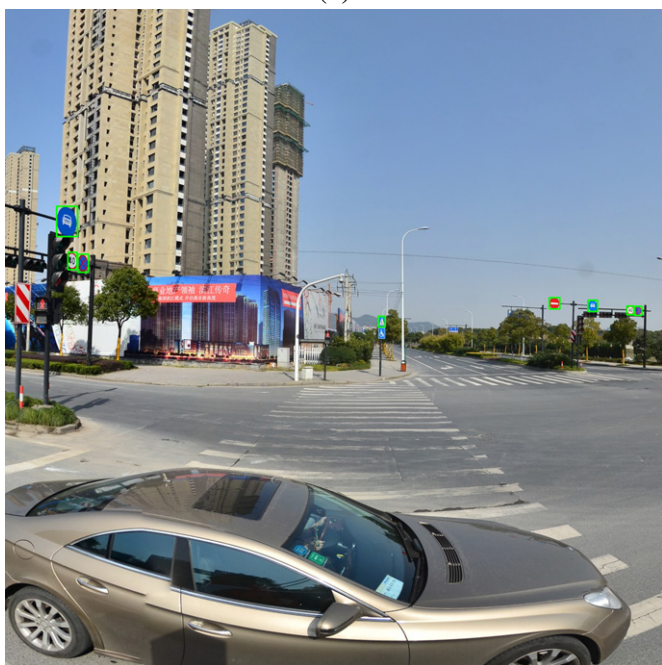
(d)



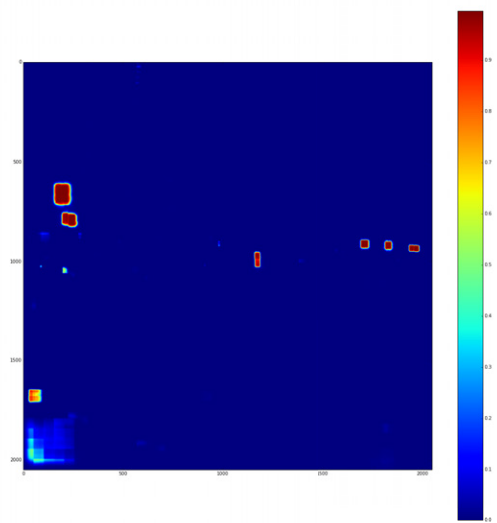
(e)



(f)



(g)



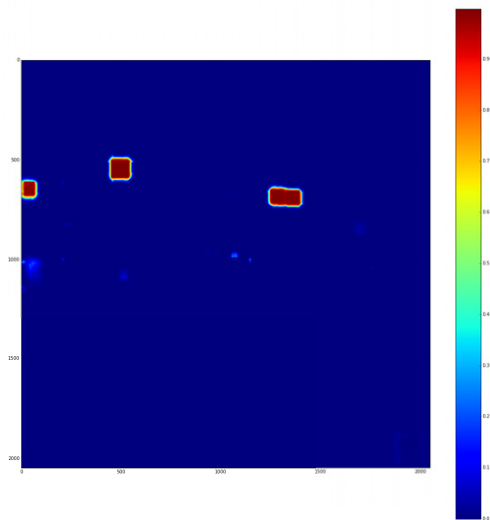
(h)

1.6 Signs with simple and complex backgrounds

Traffic signs in images (a) and (c) are suspended at height, and have simple backgrounds, making it easy to detect them. In images (e) and (g) the traffic signs are nearer the road level; such signs are more likely to have cluttered, complex backgrounds. In (e) and (g), the background has a richer texture but our network still performs well.



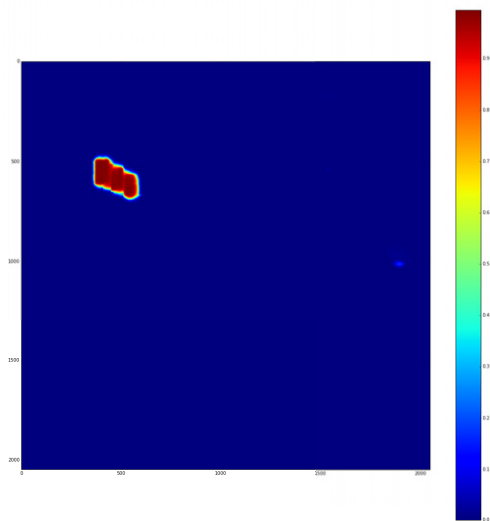
(a)



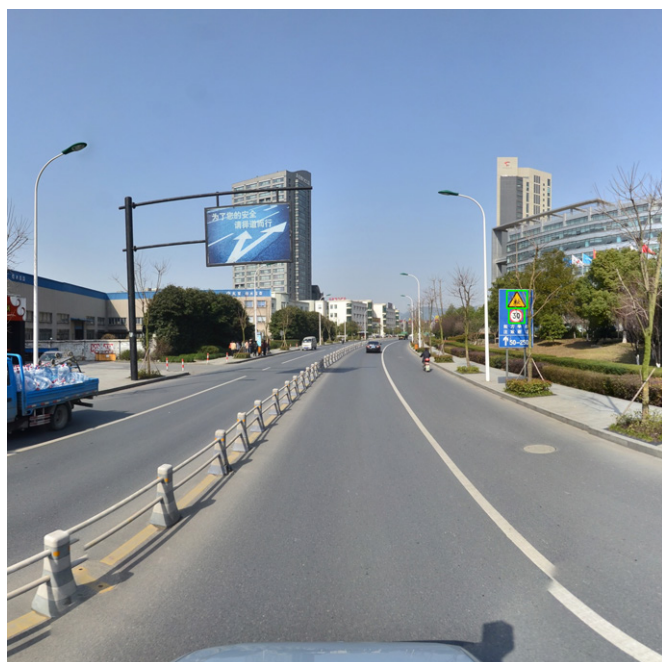
(b)



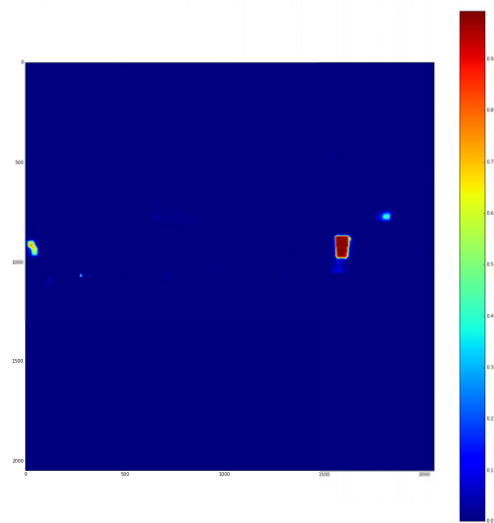
(c)



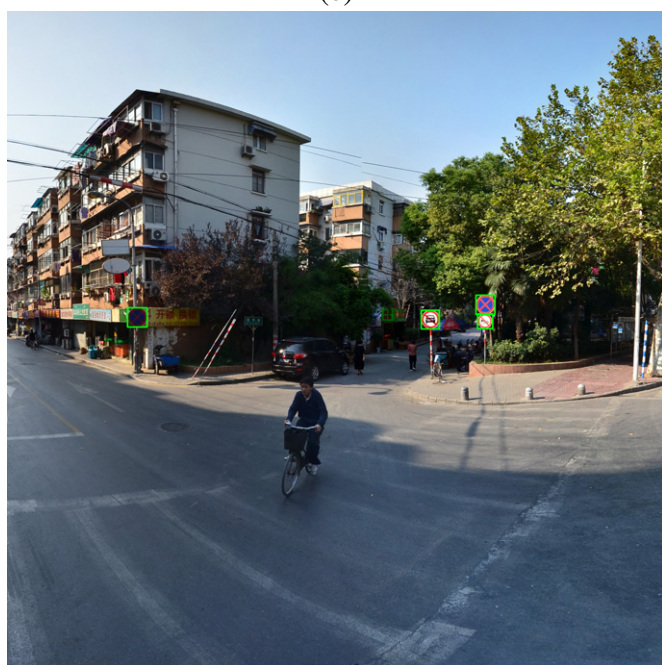
(d)



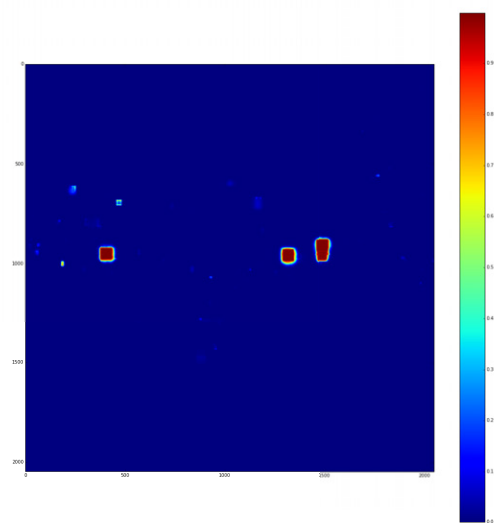
(e)



(f)



(g)



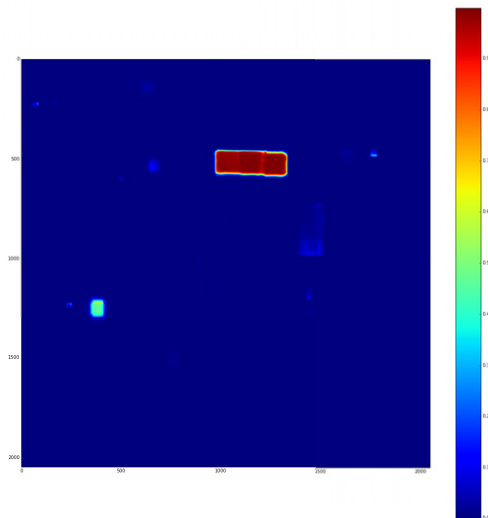
(h)

1.7 Occluded and unoccluded

There is no occlusion of the traffic signs in images (a) and (c), and our network performs well. In (e) the speed limit sign is partially occluded but our network still detects it. To the upper left of this speed limit sign is a prohibitory sign we have failed to detect, as it is of a kind not included in the training set. In (g) the bottom of the no parking sign and the top of the no bicycles sign are occluded. Both of them are correctly detected.



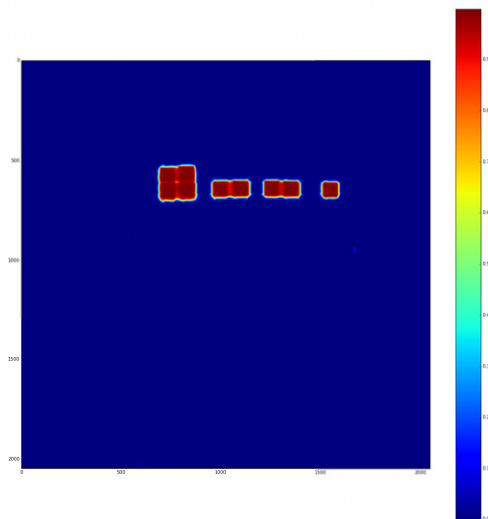
(a)



(b)



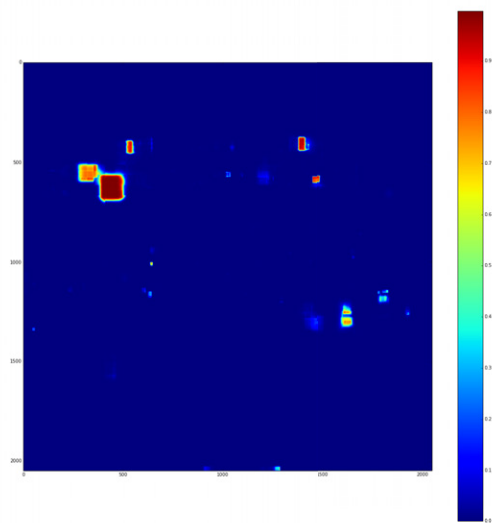
(c)



(d)



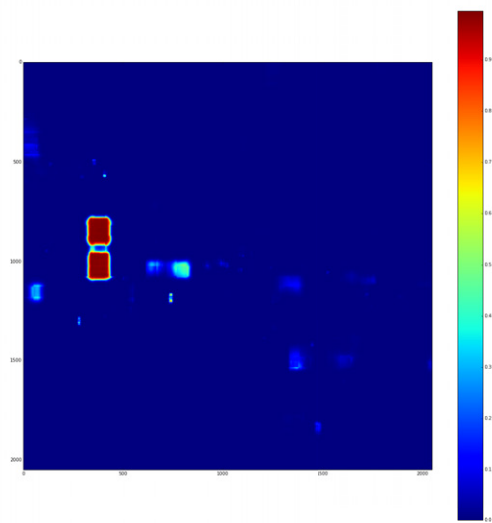
(e)



(f)



(g)



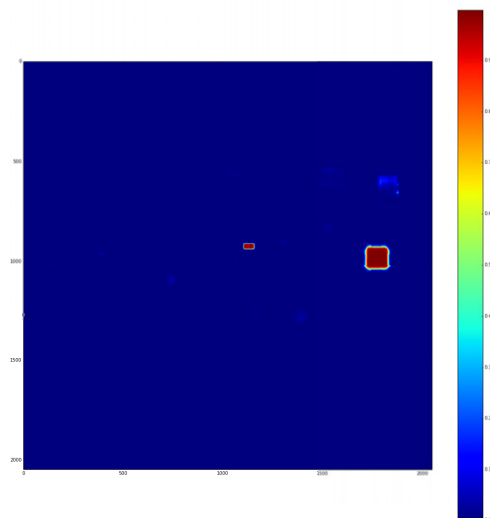
(h)

1.8 Cloudy and sunny

Images (a) and (c) were captured on cloudy days while (e) and (g) were captured on sunny days. The results show that our network is robust to variations in weather conditions which cause changes in illuminance.



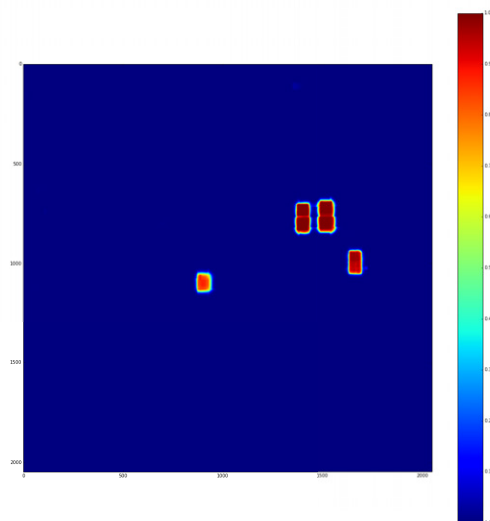
(a)



(b)



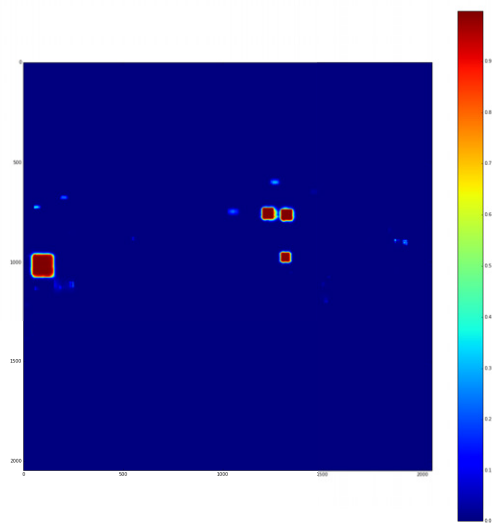
(c)



(d)



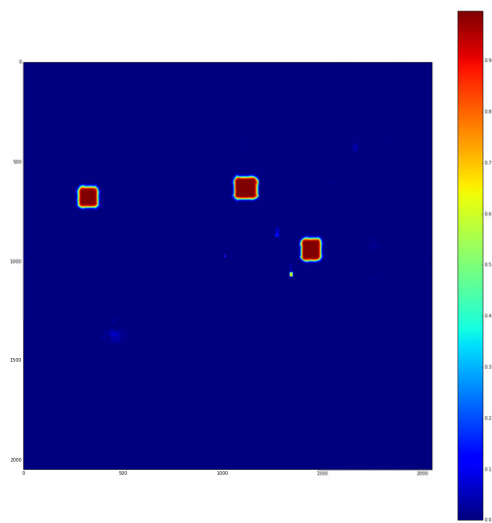
(e)



(f)



(g)



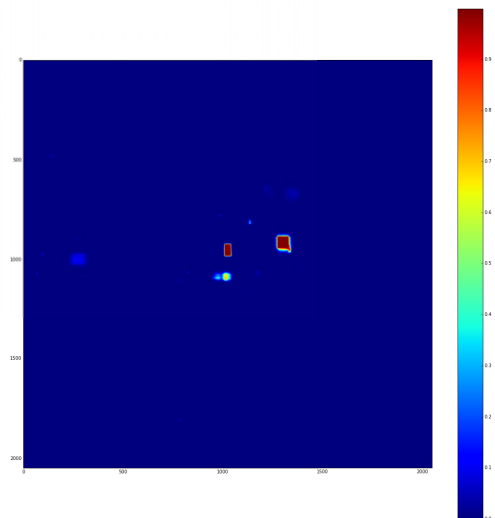
(h)

2 FAILURES

We now show some examples where our method has not worked well. In image (a) a no entrance sign is missed—it has a small size and considerable perspective distortion. In image (c) a no left turn sign is missed. In (e) the blue logo of a bank is misdetected as it has a similar blue circular shape to the mandatory sign class. In (g) a Volkswagen logo is misdetected.



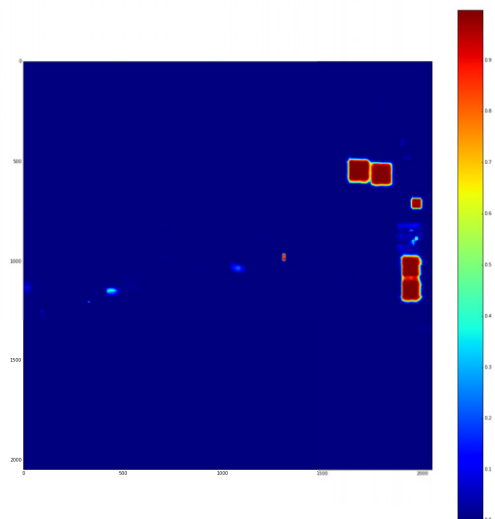
(a)



(b)



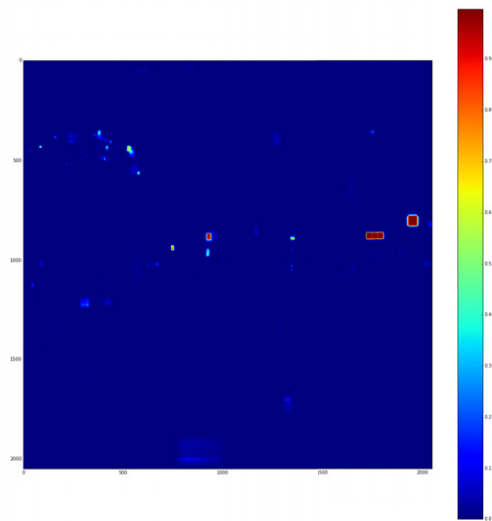
(c)



(d)



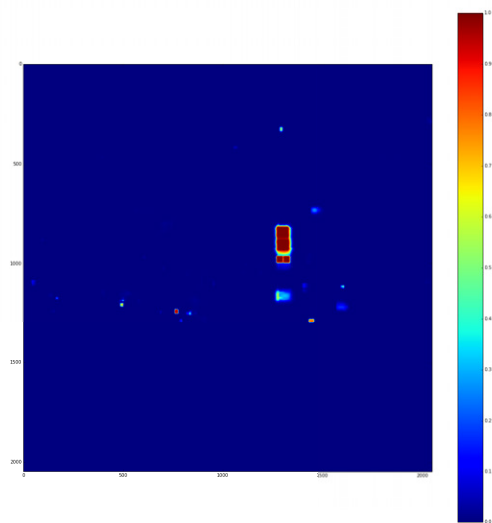
(e)



(f)



(g)



(h)