

Seeing What Is Not There: Learning Context to Determine Where Objects Are Missing Supplementary Material

1 Introduction

In this supplementary material, we present a detailed illustration of the retrieved missing curb ramp regions to help readers to qualitatively evaluate the effectiveness of our work. Two methods are compared: the SFC network with hard negative mining and the Spatial Prior baseline. We show the top 48 retrieved regions from each method and provide comments on some of the representative regions.

Please view this document in color.

2 SFC + Hard Negative Mining

Tables 1-3 show the top 48 retrieved missing curb ramp regions, with three color codes. A green box represents a correct region where curb ramps are missing. For example, in Table 1, regions 9 and 13 both have clear pedestrian crosswalks marked on the road, but curb ramps are missing from the area where the cross walk intersects the road curb. Note that the SFC network model accurately locates the candidate curb ramp area at the center of the image patch. Also, it apparently learns different styles of crosswalk markings: basic lines marking (9) and white stripes (13). A red box represents a region that does not need a curb ramp. For example, regions 2, 3 and 6 are from one side of the road and there is no visible marked crosswalk. However, they all show one similarity: there are road lamps nearby, which are visually similar to traffic light poles and thus serve as strong contextual information. In the training images, there are multiple intersections that indeed contain curb ramps for one side of the road with a traffic light. A blue box represents a region that already has curb ramps, but the object detector fails to find them. For example, regions 4, 8, and 10 all contain curb ramps that are occluded by pedestrians, which cause the detector to fail. In other cases, either the curb ramp is not visually salient (regions 14, 16) or its appearance is affected by illumination and shadows (regions 11, 12).

The same color coding is used throughout this document. Overall, the performance of the proposed approach is largely constrained by the accuracy of the detector as most of the incorrect regions are in blue (75%). Since our context model is detector-agnostic, this suggests that the performance of the current system can be further improved if a better curb ramp detector is available. In fact, with an ideal detector, the precision can be improved from 15% (7/48) to 41% (7/17).

3 Spatial Prior

The Spatial Prior learns general positions of curb ramps so the regions it predicts are mostly around the intersection between buildings/sky and the ground plane. Meanwhile, it is having difficulty in finding the scene-specific location of curb ramps. Compared to the SFC + Hard Negative Mining approach, the Spatial Prior (Tables 4-6) finds fewer true missing curb ramp regions and its retrieved regions are not well localized with (missing) curb ramps at the center.

SFC + Hard Negative Mining



Table 1: Retrieved missing curb ramp regions, ranked by context scores high to low. **Green** box: correct missing curb ramp region. **Red** box: region does not need a curb ramp. **Blue** box: region already has curb ramp, bad detector result.

SFC + Hard Negative Mining



Table 2: Retrieved missing curb ramp regions, ranked by context scores high to low. **Green** box: correct missing curb ramp region. **Red** box: region does not need a curb ramp. **Blue** box: region already has curb ramp, bad detector result.

SFC + Hard Negative Mining



Table 3: Retrieved missing curb ramp regions, ranked by context scores high to low. Green box: correct missing curb ramp region. Red box: region does not need a curb ramp. Blue box: region already has curb ramp, bad detector result.

Spatial Prior



Table 4: Retrieved missing curb ramp regions, ranked by spatial prior scores high to low. **Green** box: correct missing curb ramp region. **Red** box: region does not need a curb ramp. **Blue** box: region already has curb ramp, bad detector result.

Spatial Prior



Table 5: Retrieved missing curb ramp regions, ranked by spatial prior scores high to low. **Green** box: correct missing curb ramp region. **Red** box: region does not need a curb ramp. **Blue** box: region already has curb ramp, bad detector result.

Spatial Prior



33



34



35



36



37



38



39



40



41



42



43



44



45



46



47



48

Table 6: Retrieved missing curb ramp regions, ranked by spatial prior scores high to low. **Green** box: correct missing curb ramp region. **Red** box: region does not need a curb ramp. **Blue** box: region already has curb ramp, bad detector result.