

Supplementary Material for
RAID: A Relation-Augmented Image Descriptor

Paul Guerrero
KAUST, University College London

Niloy J. Mitra
University College London

Peter Wonka
KAUST

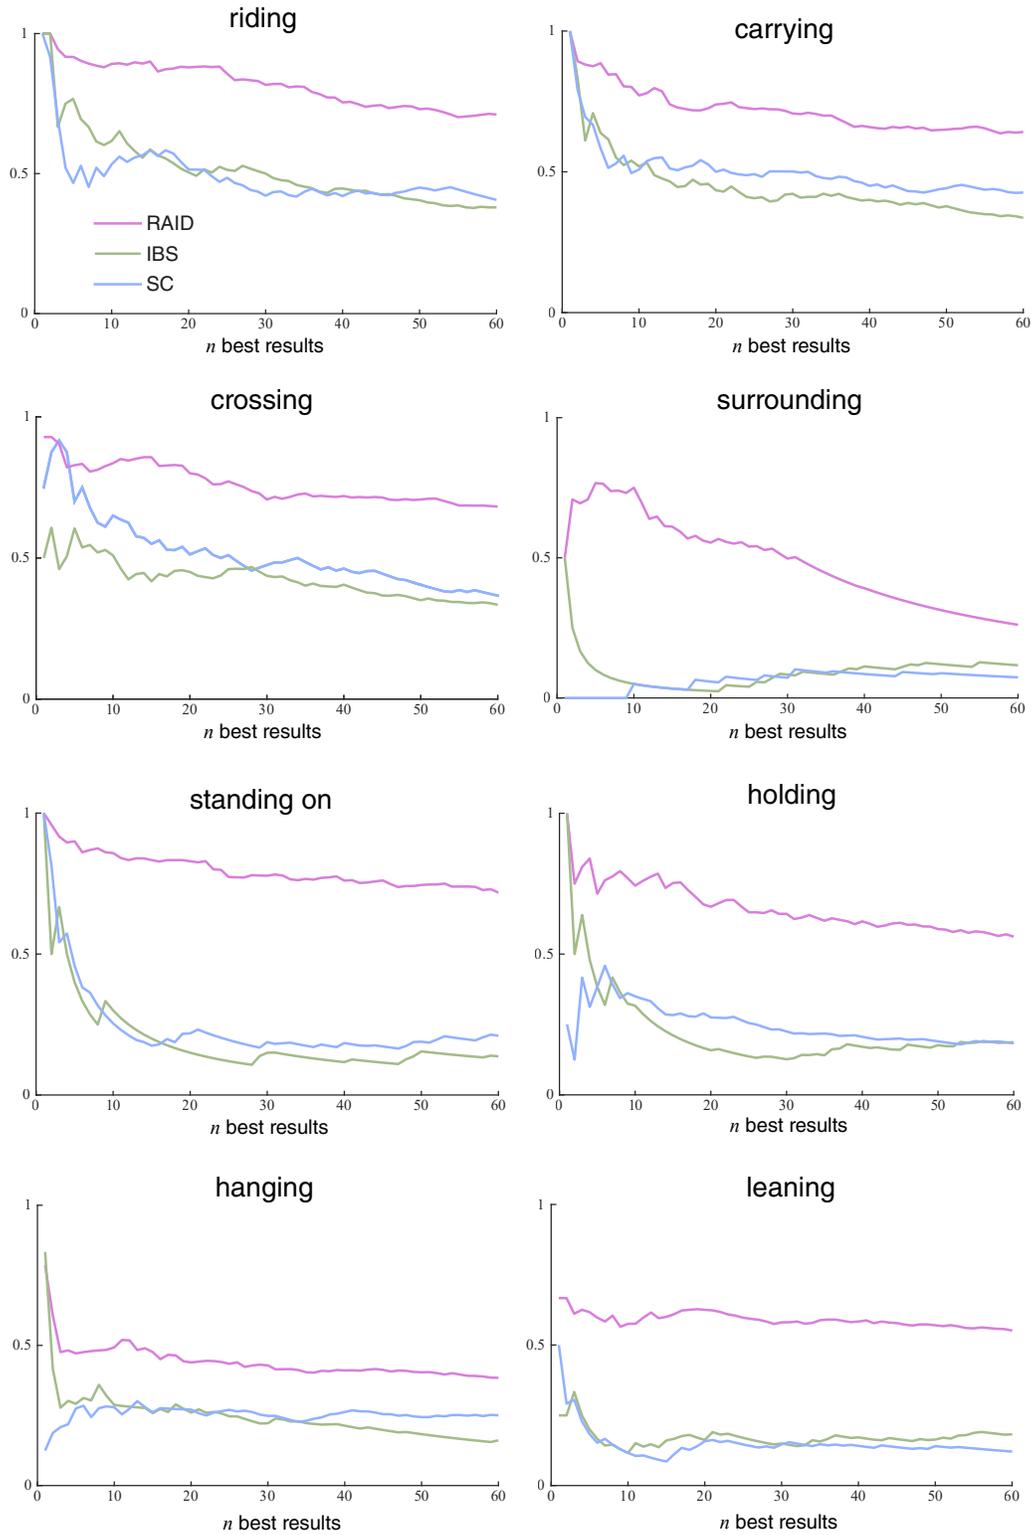
1 Synthetic Dataset

Below we show all 164 images of the synthetic dataset. Each image contains exactly one relationship between the orange source and the blue target region; 97 of these relationships are members of one or multiple relationship classes shown in Figure 3 of the paper.



3 Additional Image Retrieval Results

Here, we show extended results of the queries shown in the paper (the first 20 results for each query) and results for two additional queries, 'hanging' and 'leaning'. Both of these queries use synthetic query regions. The precision of the first n results of each query as a function of n is shown below, followed on the next pages by the results of each query.



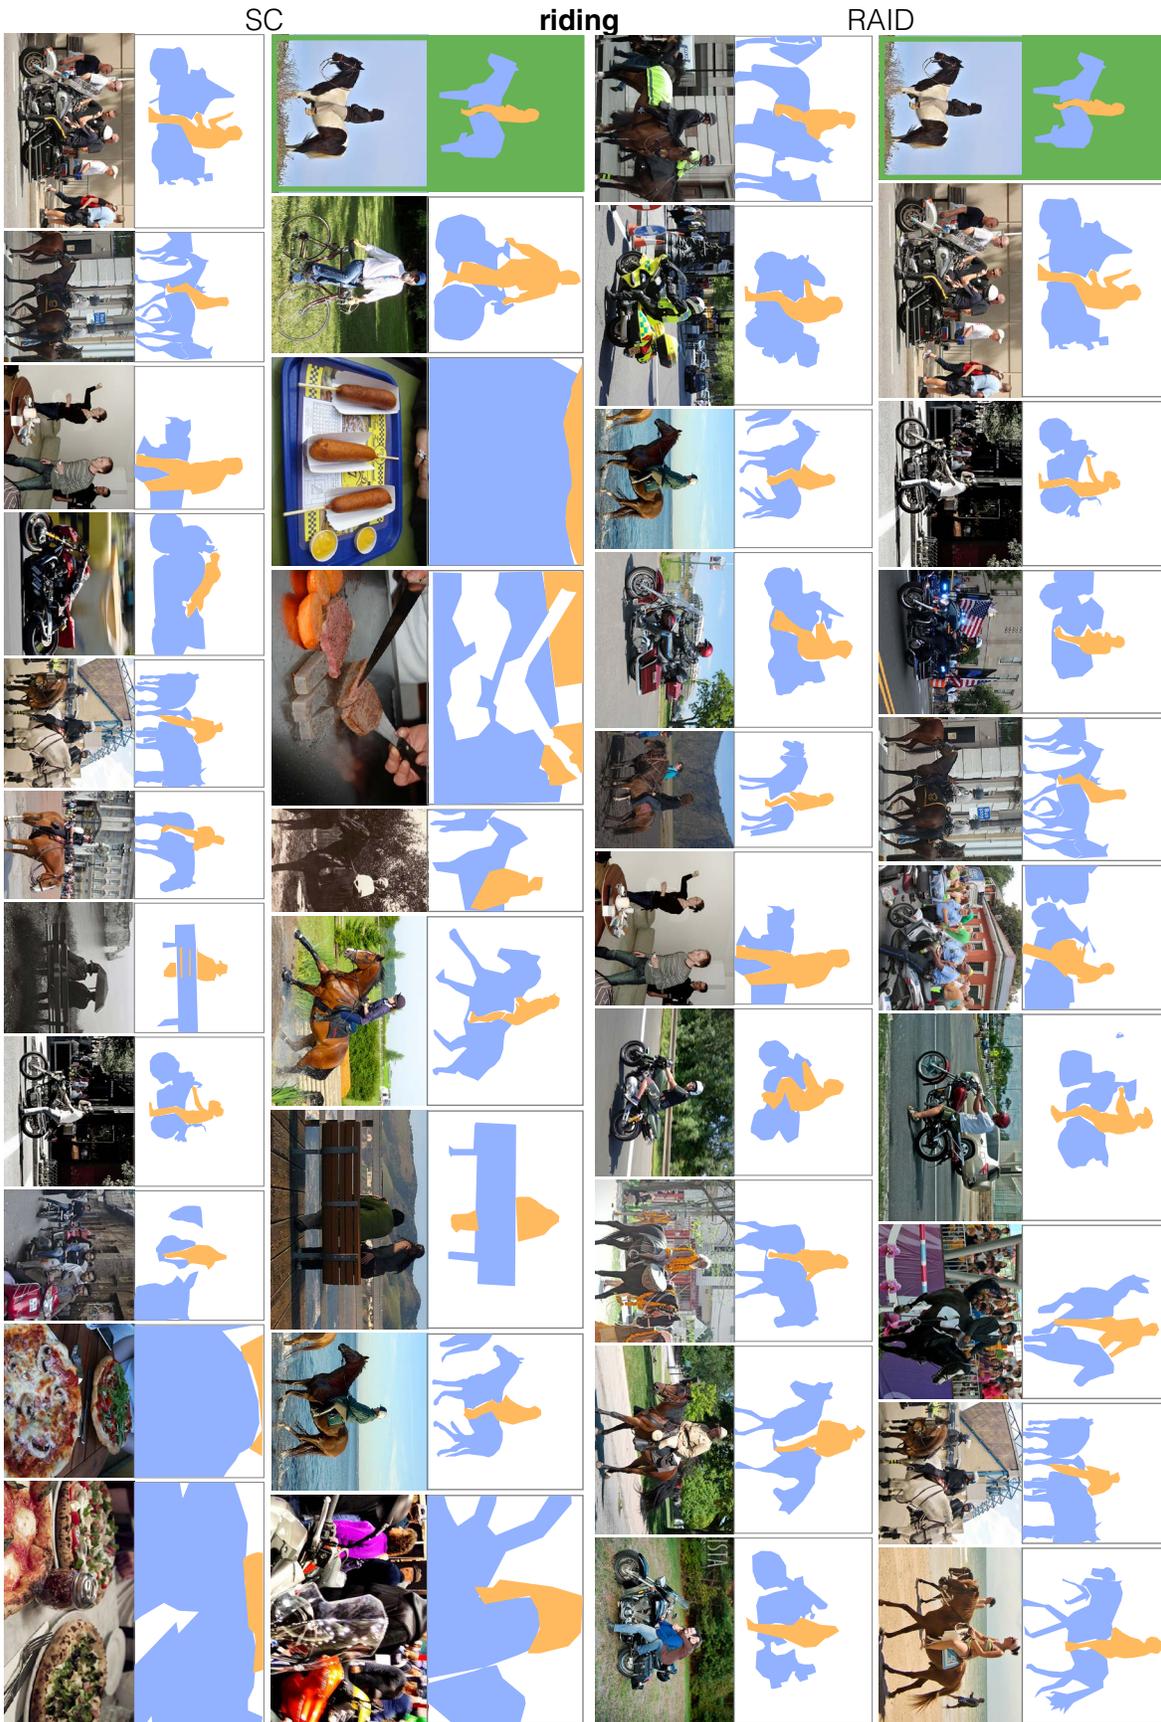


Figure 1: The 20 best results of the 'riding' query for both Shape Contexts (left) and RAID (right).



Figure 2: The 20 best results of the 'carrying' query for both Shape Contexts (left) and RAID (right).

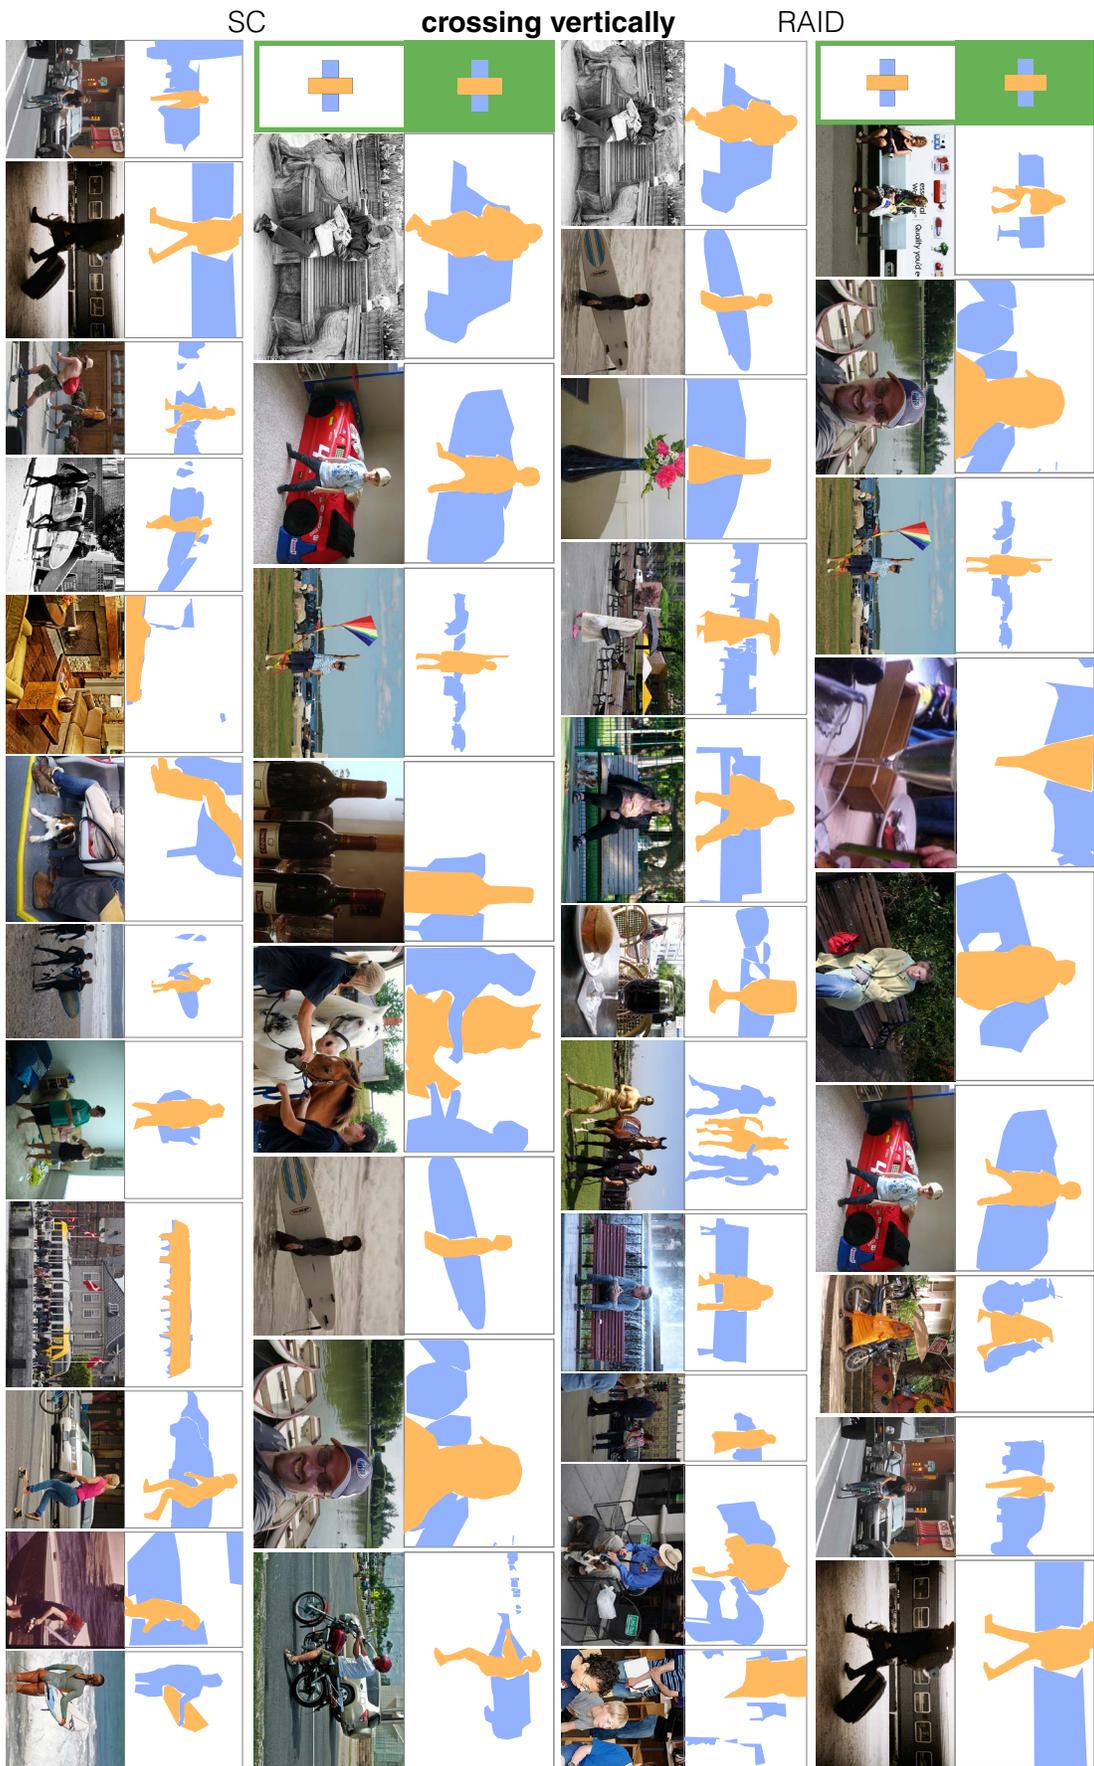


Figure 3: The 20 best results of the 'crossing vertically' query for both Shape Contexts (left) and RAID (right).





Figure 6: The 20 best results of the 'holding' query for both Shape Contexts (left) and RAID (right).



Figure 7: The 20 best results of the 'hanging' query for both Shape Contexts (left) and RAID (right).



Figure 8: The 20 best results of the 'leaning' query for both Shape Contexts (left) and RAID (right).