

Patch2Vec: Supplementary Material



Figure 1: Single-click selection. Given an input image (a), we managed to associate pixels from the same texture with similar embeddings (shown after dimensionality reduction in (b)). Notice how, e.g., the background contains pixels with different colors (ranging from black to yellow) and is partly shadowed, yet we associate all background pixels with a similar embedding. (c) The user clicks a single pixel (cyan circle); we show  $l_2$  distances between the selected pixel's embedding and all other pixels'. Thresholding the distances (c) produces a single-click selection result (d). However, some edges might be inaccurate (blue arrows), which can be fixed by an optional step that snaps the selection to edges, producing the final result (e).




Figure 2: Masking algorithm comparison. See main text for details.




Figure 3: Embedding consistency across frames. The user clicks on a single point in the first frame (white circle, top left). For each other frame we show the location which is closest, in our embedding space, to the clicked pixel (magenta circle). Notice how the magenta circles track the biker, even though we used trivial  $l_2$  distance to a single pixel's embedding.



Figure 4: Embedding result, training set part 1. For visualization purposes, we only show a 3D representation (false-color images) of our 128D embedding.



Figure 5: Embedding result, training set part 2.



Figure 6: Embedding result, validation set.



Figure 7: Embedding result, test set part 1.





Figure 8: Embedding result, test set part 2.




Figure 9: Super-pixel creation results on the NYU-depth V2 dataset.