

Supplementary Material: Retrieving Visually Overlapping Images Through Interpretable Non-Metric Embeddings

Anonymous ECCV submission

Paper ID 4294

1 Interpretable Queries

Further to Section 4.2 in our main work, we demonstrate the interpretability of our method on additional test queries from the MegaDepth dataset. We show three different types of figures:

- Interpretability plots: For a query image from the test set what do the predicted enclosure and concentration of retrievals from the **training** set images tell us about their relationship to one-another?
- Generalization plots: Does our embedding generalize to the images from the **test** set and are pairwise relationship interpretations valid?
- Relative scale plots: Given pairs of images from the test set and their box representations, can we estimate their relative scale difference?

1.1 Interpretability Plots

These plots (Figures 4 5 6 7 8 9 10 11 12) illustrate the interpretability that we gain when using box embeddings trained with normalized visual overlap as the world-space measure. The key takeaway is that we can qualitatively observe the relationship between the query and each of the retrieved images. The relationships can be grouped into four categories: Given a query image, a retrieval can be a:

- zoom-out
- crop-out or oblique-out
- close-up
- clone-like

We show three different examples for the scenes *Venice* and *Florence* and two examples for *BigBen* and *NotreDame* additionally to Figure 5 in the main work. For each figure a single query and up to 36 retrieved images from the training set are shown. They are placed into one of six buckets according to the **predicted** normalized box overlap $\text{NBO}(\mathbf{b}_x \mapsto \mathbf{b}_y)$ and $\text{NBO}(\mathbf{b}_y \mapsto \mathbf{b}_x)$. The vertical axis describes the enclosure, so $\text{NBO}(\mathbf{b}_x \mapsto \mathbf{b}_y)$, or "how much surface from the query image is visible in the retrieved image". The horizontal

axis describes the concentration $\text{NBO}(\mathbf{b}_y \mapsto \mathbf{b}_x)$, in other words "how much surface of the retrieved image is visible in the query image".

The numbers below each image are enclosure and concentration estimated with box representations as well as ground-truth values estimated with semi-dense depth maps.

1.2 Generalization Plots

The interpretability plots retrieve images from the training set of our box embeddings. Next, we demonstrate qualitatively that the learned representations generalize to the images in the test set. Here the test set are images in MegaDepth that do not have dense depth information. Hence, we can only report qualitative results for this larger test set.

So, we now retrieve images from the test set using a random query image from the test set. Here, we plot retrieved images on a 2D grid, using enclosure and concentration of the retrieved image as 2D coordinate, where the x-coordinate denotes the concentration, and the y-coordinate denotes the enclosure.

These plots (Figures 13 14 15 16) also provide interpretability, as one can observe different clusters of images as zoom-outs, crop-outs, close-ups and clones of the query in similar "quadrants" as Interpretability plots.

1.3 Relative Scale Plots

Relative scale plots are further qualitative examples similar to Figure 1 and Figure 7 in the main paper.

We illustrate that we can estimate geometric relationships between two images from the test set using our box embeddings. For two images from the test set we can estimate the relative scale of the first image in the second image. Hence, we plot two test images and a rescaled version of the first image such that any geometric verification between the two images is now easier to do due to matching scale.

This relative scale estimate is relatively accurate if the images have zoom-in/zoom-out relationship. If the images are in crop-out/oblique-out to one-another, then the rescaling is not necessarily going to make matching easier. So, image pairs that seem to be failure cases in terms of estimated relative scale often have low enclosure value ($<80\%$), which means that these image pairs can be detected and treated accordingly. To demonstrate that this filtering approach is effective, we show failure cases with an enclosure of at least 80%.

2 Further Localization Evaluations

In this section we provide further evaluations of different embeddings for localization task.

2.1 7-Scenes

Figure 1 shows sorted rotation errors for each scene, similarly to Figure 6 in the main paper, for different embeddings. Again, rotation error is computed per query, when matched against 10-th and 30-th nearest neighbor from the training set (gallery).

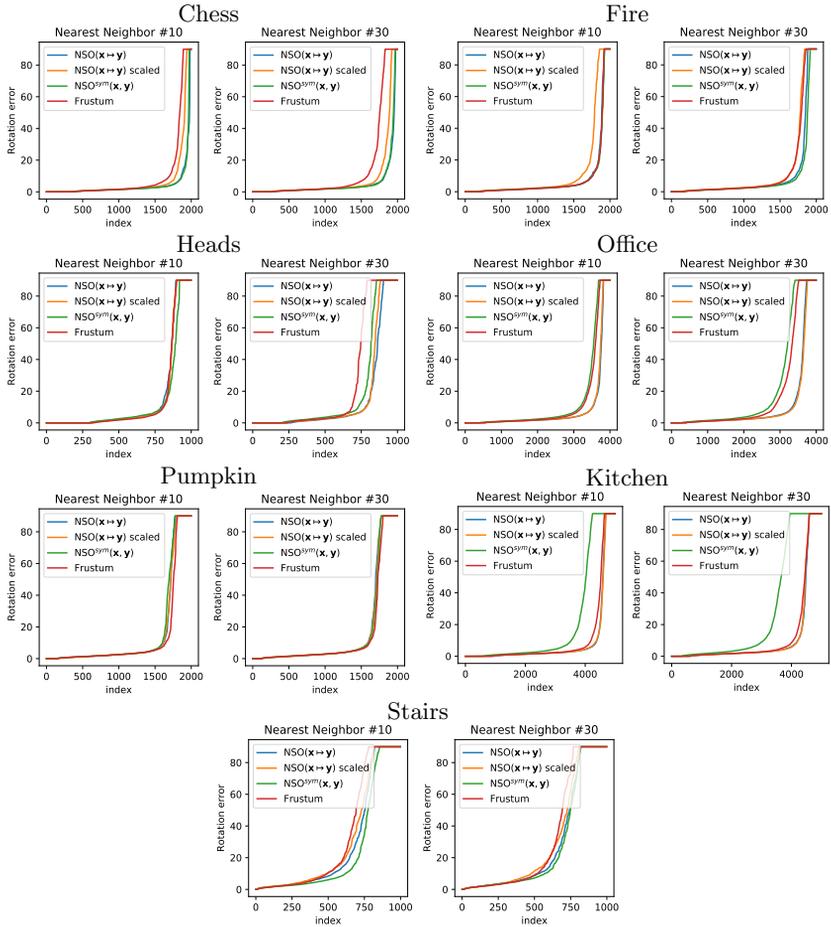


Fig. 1. 7-Scenes pose rotation error. Each plot shows (sorted) rotation error (capped at 90°) when each test image is matched against 10-th and 30-th closest retrieved image for pose estimation. As we can see, box embeddings with surface overlap measure tend to outperform alternatives, especially when rescaling images according to estimated relative scale.

When solving for the pose, retrieving 10-th nearest neighbor for matching seems to be sufficient to estimate good pose for most of the scenes. However,

even in this setting we can see that Frustum overlap is under-performing compared to embeddings trained with surface overlap. Only in Pumpkin scene with 10-th retrieved image, the frustum overlap is marginally better. Between surface overlap-based embeddings the performance is quite comparable. There are no systematic improvements nor deteriorations with relative scale correction as captured scenes are all rooms with limited scale variation.

2.2 Megadepth

Figure 2 shows evaluation results for Venice and Florence scenes to complement Figure 6 of the main paper.

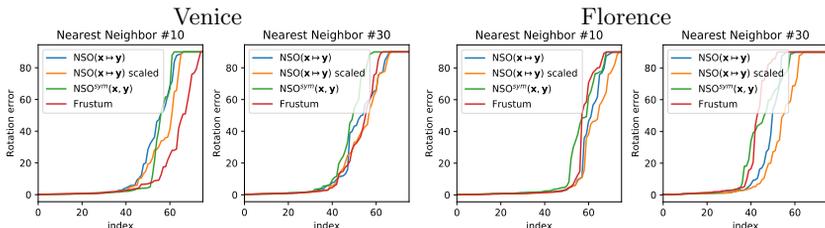


Fig. 2. Additional results to Fig. 6 on MegaDepth. Each plot shows (sorted) rotation error (capped at 90°) when each test image is matched against 10-th and 30-th closest retrieved image for pose estimation. As we can see, box embeddings with surface overlap measure tend to outperform alternatives, especially when rescaling images according to estimated relative scale. Results evaluated for 100 images from the test set.

In Table 1 we report median error translation and rotation error of estimated pose for 100 images of the test set (that have corresponding depth maps) similar to Table 2 in the main paper. The Megadepth scenes are not metric, so the scale factor of translation errors is not known. Furthermore, all these errors are relatively low, corresponding to accurate pose, so it is difficult to draw conclusions from these results.

Hence, we also evaluated different embeddings with a larger test set which consists of images in Megadepth that do not have depth maps. This results in 3165, 1931, 2255 and 1159 test set images for Big Ben, Notre-Dame, Venice and Florence, respectively. Table 2 show median errors for the larger test set. As can be seen, correcting for scale using Box embeddings is superior to alternatives on 3 scenes.

Similarly, (sorted) rotation error evaluated for the larger test set could be seen in Figure 3. Here the error is computed against 1st nearest neighbor retrieval. These plots indicate a similar conclusion. The surface overlap based embeddings are outperformed by Frustum overlap embedding for Florence scene. Florence scene has images that capture a large area with complex narrow streets, however the training set consists of only 1471 images. We suspect that our CNNs need more training data to learn generalizable surface overlaps.

Training \mathcal{L}	Box NSO($\mathbf{x} \mapsto \mathbf{y}$)	Box Scaled NSO($\mathbf{x} \mapsto \mathbf{y}$)	Vector NSO ^{sym} (\mathbf{x}, \mathbf{y})	Vector Frustum
Retrieval func.	NBO($\mathbf{b}_x \mapsto \mathbf{b}_y$)+ NBO($\mathbf{b}_y \mapsto \mathbf{b}_x$)	NBO($\mathbf{b}_x \mapsto \mathbf{b}_y$)+ NBO($\mathbf{b}_y \mapsto \mathbf{b}_x$)	Eucl. dist.	Eucl. dist.
Notre-Dame	.038, 0.79°	.038, 0.87°	.036 , 0.83°	.047, 1.05°
Big Ben	.067 , 0.87°	.089, 0.98°	.070, 0.87°	.096, 0.83°
Venice	.096, 1.01°	.098, 1.23°	.088 , 0.85°	.085, 0.91°
Florence	.081, 1.08°	.079, 1.10°	.054, 0.87°	.048 , 0.68°

Table 1. Comparison of rotation and translation errors on the MegaDepth dataset, where boxes learn surface overlap asymmetrically while vectors are trained symmetrically. The first entry of each cell denotes the translation error up to scale, the second entry is the rotation error in degrees. Results evaluated for 100 images from the test set.

Training \mathcal{L}	Box NSO($\mathbf{x} \mapsto \mathbf{y}$)	Box Scaled NSO($\mathbf{x} \mapsto \mathbf{y}$)	Vector NSO ^{sym} (\mathbf{x}, \mathbf{y})	Vector Frustum
Retrieval func.	NBO($\mathbf{b}_x \mapsto \mathbf{b}_y$)+ NBO($\mathbf{b}_y \mapsto \mathbf{b}_x$)	NBO($\mathbf{b}_x \mapsto \mathbf{b}_y$)+ NBO($\mathbf{b}_y \mapsto \mathbf{b}_x$)	Eucl. dist.	Eucl. dist.
Notre-Dame	0.84, 15.1°	0.81 , 13.8°	1.43, 38.2°	0.98, 25.2°
Big Ben	2.91, 58.0°	2.85 , 53.9°	2.81, 54.0°	3.30, 69.8°
Venice	3.20, 68.9°	3.24 , 58.6°	3.13, 70.8°	2.70, 65.6°
Florence	1.44, 35.0°	1.33, 31.6°	1.67, 42.2°	0.75 , 10.6°

Table 2. Comparison of rotation and translation errors on the MegaDepth dataset for test set without depth images, where boxes learn surface overlap asymmetrically while vectors are trained symmetrically. The first entry of each cell denotes the translation error up to scale, the second entry is the rotation error in degrees. Total number of images in these test sets (in Figure order): 3165, 1931, 2255 and 1159.

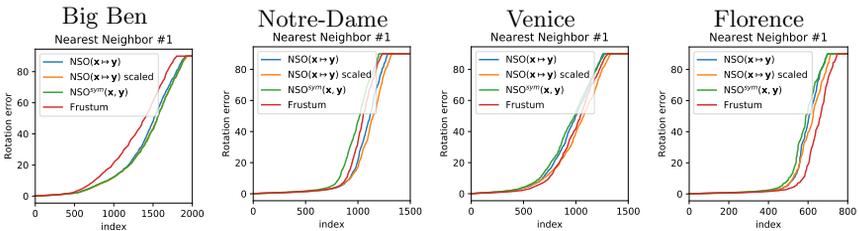
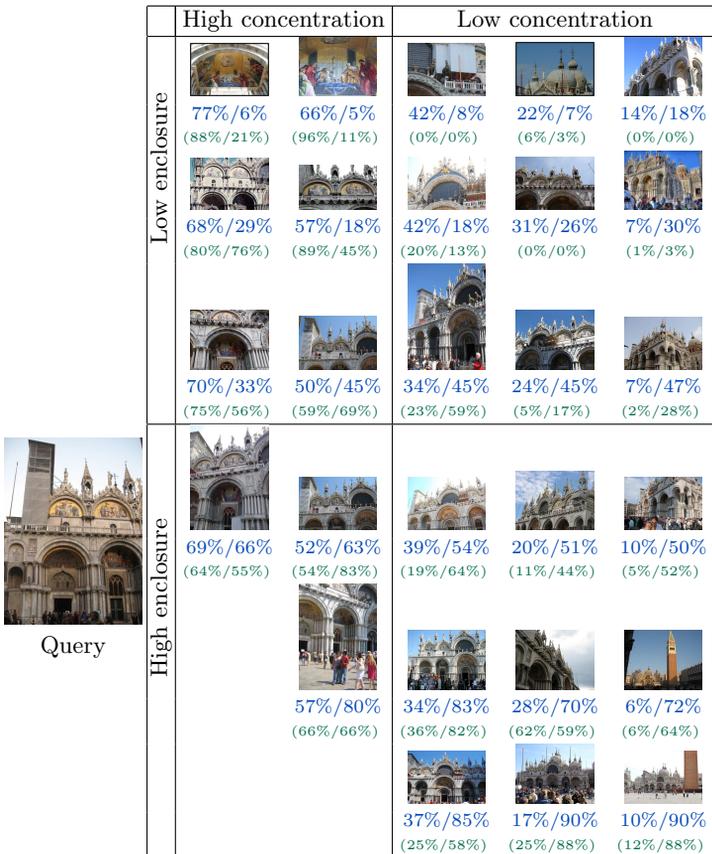


Fig. 3. Results on images without depth (MegaDepth) Each plot shows (sorted) rotation error (capped at 90°) when each test image is matched against the closest retrieved image for pose estimation. As we can see, box embeddings with surface overlap measure tend to outperform alternatives, especially when rescaling images according to estimated relative scale. Total number of images in these test sets (in Figure order): 3165, 1931, 2255 and 1159.



Results of predicted and ground-truth enclosure and concentration relative to the query image on the left. The numbers below each image indicate the predicted and ground-truth concentration/enclosure. It can be observed that the images in the upper left quadrant are close-ups of the query. The images in the lower left quadrant are clones of the query. The retrieved clones preserve the normals of the surfaces in the query. Images in the lower right quadrant are zoom-outs of the scene in the query. The images in the upper right corner are mostly oblique-outs and show the scene in the query from different angles. *This caption applies to all Interpretability plots unless otherwise stated.*

Fig. 4. Interpretability plot: Venice No. 1.

		High concentration			Low concentration			
270 271 272 273 274 275 276 277	Low enclosure							
		85%/5% (84%/13%)	75%/13% (73%/20%)	65%/9% (65%/15%)	48%/5% (74%/41%)	29%/16% (18%/20%)	15%/8% (6%/8%)	
		82%/21% (82%/30%)	78%/25% (74%/39%)	62%/34% (25%/29%)	43%/18% (47%/20%)	25%/16% (18%/22%)	12%/20% (8%/13%)	
		86%/49% (84%/80%)	82%/38% (70%/41%)	54%/37% (77%/74%)	44%/47% (77%/77%)	33%/44% (41%/68%)	11%/45% (45%/68%)	
278 279 280 281 282 283 284 285 286 287 288 289 290 291	High enclosure							
		84%/64% (77%/89%)	67%/63% (58%/85%)	67%/63% (58%/85%)	49%/55% (67%/71%)	31%/61% (59%/75%)		
		87%/66% (74%/89%)	67%/73% (57%/83%)	64%/71% (67%/90%)	37%/74% (31%/77%)	27%/78% (43%/76%)		
		89%/85% (78%/84%)	71%/84% (76%/89%)	65%/89% (45%/78%)				

Fig. 5. Interpretability plot: Venice No. 2.



	High concentration			Low concentration		
Low enclosure						
	88%/15% (87%/20%)	69%/16% (69%/21%)	66%/11% (75%/7%)	40%/5% (0%/0%)	19%/10% (1%/1%)	10%/6% (0%/0%)
	86%/32% (71%/25%)	83%/22% (82%/21%)	66%/22% (56%/17%)	50%/23% (48%/22%)	27%/17% (12%/12%)	11%/33% (10%/25%)
	89%/49% (83%/53%)	75%/41% (70%/41%)	55%/40% (64%/51%)	37%/36% (27%/42%)	29%/46% (31%/40%)	9%/38% (8%/29%)
High enclosure						
	90%/51% (83%/55%)	72%/50% (58%/37%)	58%/63% (50%/62%)	44%/54% (48%/55%)	25%/61% (27%/57%)	16%/55% (11%/46%)
	83%/71% (79%/63%)	83%/78% (76%/74%)	59%/76% (57%/74%)	43%/68% (44%/61%)	17%/79% (32%/82%)	
		68%/87% (70%/76%)	66%/87% (70%/82%)	49%/88% (65%/85%)	25%/85% (35%/84%)	

Fig. 6. Interpretability plot: Florence No. 1.

		High concentration			Low concentration		
360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404	Query	Low enclosure					
			71%/31% (79%/40%)	53%/29% (79%/32%)	44%/32% (69%/25%)	7%/25% (17%/21%)	
							
			81%/47% (79%/62%)	60%/39% (76%/36%)	46%/38% (67%/41%)		
		High enclosure					
			83%/62% (77%/63%)	77%/54% (79%/68%)	56%/55% (66%/36%)	14%/57% (30%/81%)	
							
			88%/79% (75%/65%)	74%/75% (78%/66%)	56%/77% (82%/70%)	49%/81% (74%/77%)	6%/77% (10%/75%)
							
			88%/85% (78%/77%)	83%/88% (76%/76%)	63%/92% (76%/78%)	25%/85% (54%/80%)	15%/86% (24%/80%)

Fig. 7. Interpretability plot: Florence No. 2.

		High concentration			Low concentration			
	Low enclosure							
		84%/16% (74%/17%)	80%/13% (76%/18%)	66%/17% (61%/24%)	42%/14% (60%/16%)	22%/5% (23%/4%)	13%/7% (5%/9%)	
								
		89%/19% (79%/31%)	68%/20% (71%/35%)	63%/22% (55%/32%)	39%/25% (43%/33%)	34%/21% (39%/34%)		
								
		86%/42% (73%/42%)	72%/42% (74%/42%)	65%/48% (71%/50%)	36%/36% (52%/34%)	33%/48% (34%/49%)	16%/45% (37%/31%)	
		High enclosure						
			89%/62% (73%/75%)	82%/64% (71%/56%)	58%/62% (48%/64%)	42%/58% (44%/38%)	19%/55% (35%/32%)	
								
	87%/78% (75%/72%)		68%/74% (64%/71%)	58%/78% (56%/73%)	41%/73% (43%/65%)			
								
	87%/84% (75%/76%)		74%/93% (76%/73%)					

Fig. 8. Interpretability plot: Florence No. 3.

		High concentration			Low concentration		
 <p>Query</p>	Low enclosure						
		88%/17% (89%/21%)	73%/12% (71%/6%)	65%/17% (61%/10%)		11%/8% (0%/0%)	
							
		90%/30% (90%/30%)	72%/28% (68%/24%)	66%/20% (62%/10%)	38%/31% (38%/29%)		
							
		93%/45% (94%/49%)	75%/38% (73%/36%)	62%/36% (64%/24%)	49%/40% (48%/41%)	17%/44% (16%/35%)	
	High enclosure						
		91%/61% (93%/65%)	82%/56% (79%/62%)	61%/61% (68%/56%)	38%/58% (35%/53%)	30%/65% (35%/78%)	9%/66% (7%/74%)
							
		93%/77% (89%/84%)	74%/77% (74%/78%)	53%/83% (50%/81%)	49%/78% (48%/79%)	24%/79% (22%/81%)	9%/79% (11%/81%)
							
		95%/91% (94%/88%)	67%/91% (80%/92%)	54%/88% (52%/88%)	39%/90% (36%/90%)	22%/86% (23%/88%)	13%/85% (14%/83%)

Fig. 9. Interpretability plot: Big Ben No. 1.



Query

	High concentration			Low concentration			
Low enclosure							
	85%/9% (82%/14%)	73%/14% (69%/16%)	65%/8% (57%/4%)	38%/16% (36%/16%)	20%/7% (43%/22%)	16%/8% (37%/24%)	
	92%/23% (94%/28%)	76%/21% (69%/22%)	54%/20% (53%/22%)	49%/18% (46%/21%)	28%/33% (26%/30%)		
	88%/45% (85%/40%)	81%/47% (81%/36%)	59%/47% (60%/46%)	50%/45% (49%/41%)	29%/44% (24%/36%)	14%/43% (11%/34%)	
	High enclosure						
		93%/62% (90%/53%)	80%/60% (81%/49%)	64%/62% (81%/86%)	46%/64% (46%/55%)	27%/57% (30%/53%)	16%/64% (13%/49%)
91%/73% (87%/52%)		83%/68% (86%/54%)	62%/68% (80%/84%)	39%/83% (43%/73%)	20%/80% (29%/71%)	6%/78% (8%/75%)	
91%/91% (92%/86%)		74%/85% (81%/66%)	59%/91% (68%/84%)	38%/87% (55%/78%)	33%/87% (43%/83%)		

Fig. 10. Interpretability plot: Big Ben No. 2.

		High concentration			Low concentration			
540 541 542 543 544 545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561 562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582 583 584	 Query	Low enclosure						
			85%/9% (87%/8%)	83%/6% (83%/5%)	64%/8% (80%/12%)	36%/9% (48%/8%)	18%/6% (23%/6%)	9%/11% (8%/4%)
								
			87%/23% (84%/21%)	73%/33% (75%/34%)	56%/30% (57%/33%)	42%/33% (49%/37%)	33%/29% (37%/30%)	12%/33% (18%/35%)
								
			87%/43% (82%/46%)	80%/35% (75%/39%)	59%/34% (62%/37%)	35%/38% (37%/39%)	30%/48% (30%/52%)	12%/35% (20%/38%)
High enclosure								
	90%/57% (88%/56%)	80%/58% (76%/58%)	63%/59% (66%/58%)	47%/53% (48%/53%)				
								
	93%/79% (91%/77%)	82%/75% (79%/77%)						
								
	90%/84% (88%/83%)	83%/88% (77%/87%)						

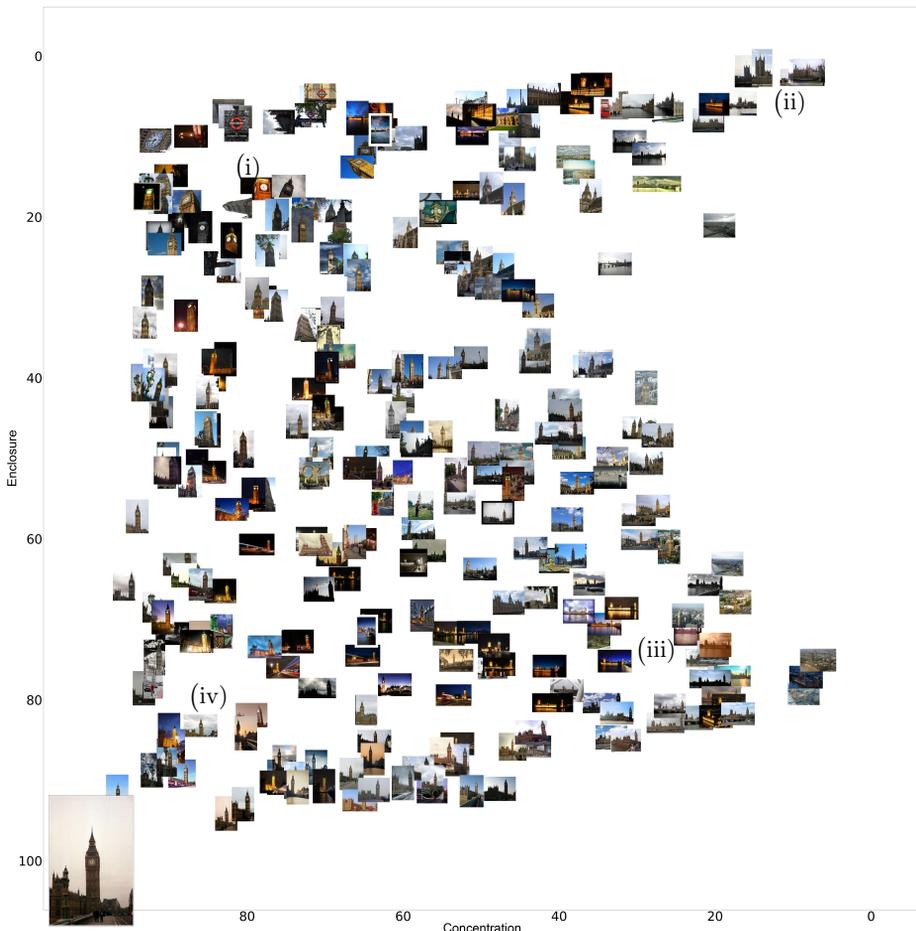
Fig. 11. Interpretability plot: Notre Dame No. 1.



Query

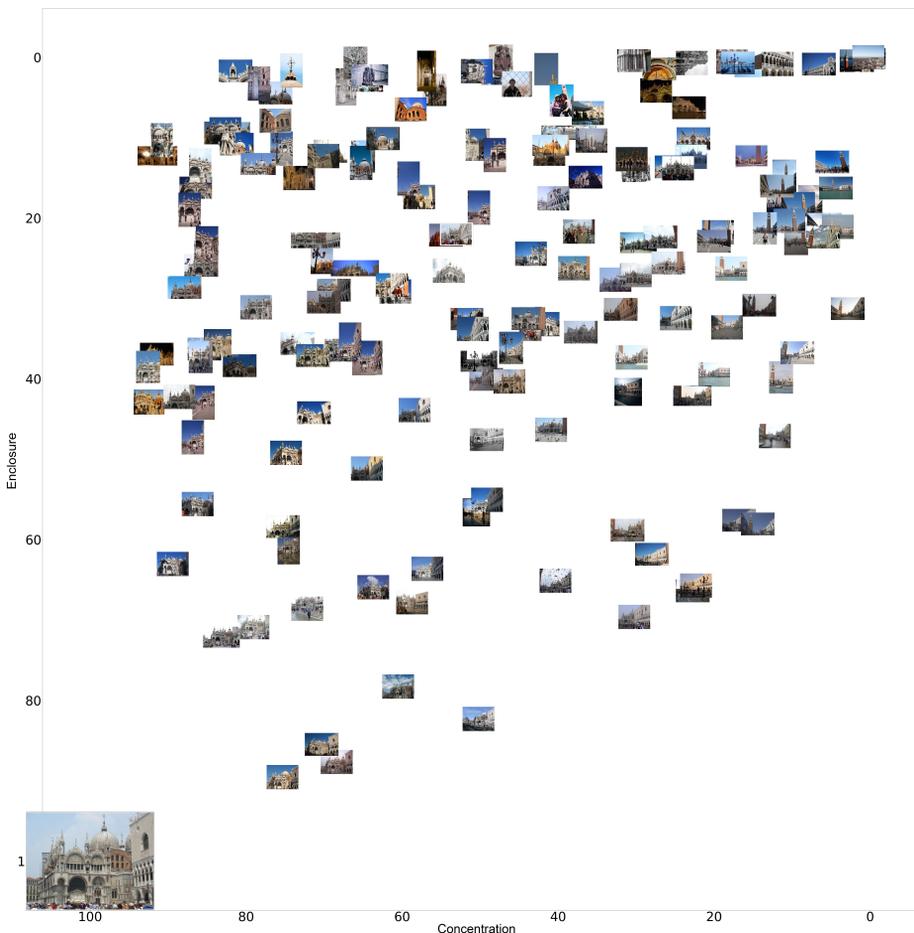
	High concentration	Low concentration	
Low enclosure	 83%/5% (88%/9%)	 74%/6% (80%/6%)	 57%/7% (56%/6%)
	 34%/15% (33%/14%)	 21%/8% (29%/17%)	 7%/16% (30%/9%)
	 70%/28% (83%/27%)	 51%/30% (64%/34%)	 36%/18% (45%/9%)
	 32%/26% (34%/17%)	 12%/24% (25%/15%)	
	 74%/47% (68%/42%)	 59%/45% (75%/35%)	 35%/36% (77%/72%)
	 31%/38% (34%/22%)	 15%/38% (37%/25%)	
High enclosure	 76%/62% (59%/43%)	 64%/49% (84%/50%)	 40%/54% (51%/55%)
	 30%/62% (45%/54%)	 76%/81% (89%/79%)	 55%/82% (84%/61%)
	 48%/82% (82%/89%)	 28%/67% (28%/72%)	 52%/96% (82%/81%)
	 41%/93% (81%/89%)	 29%/90% (76%/88%)	

Fig. 12. Interpretability plot: Notre Dame No. 2.



In this caption we expand on details and explanations of Fig.5 (Right) from the main paper. We show a query image from the test set (lower left corner) and the concentration and enclosure between randomly sampled test images from the MegaDepth SfM data set for which no depth maps are provided. The query image shows Big Ben from the view of the Westminster Bridge. (i) It can be observed that close-ups on the tower clock are clustered around the coordinates (80,15) which is consistent with our terminology of retrievals with high concentration and low enclosure. (ii) The images in the upper right corner show the waterfront side of Westminster Palace. These are crop-outs of the query image. In fact, the tower in the lower left corner of the query is one of the two towers that mark the corners of the water-front side of the palace. The retrievals in the upper right quadrant of the cluster therefore extends the view of the query. (iii) The images in the lower right area of the cluster clearly show zoom outs, with the pointy bell tower visible in all images. (iv) Lastly, one can observe that the images in the *clone - like* category are in fact similar views on Big Ben. Note that some of the retrievals are rotated images and sometimes cause outlier predictions. *This caption applies to all generalization plots unless otherwise stated.*

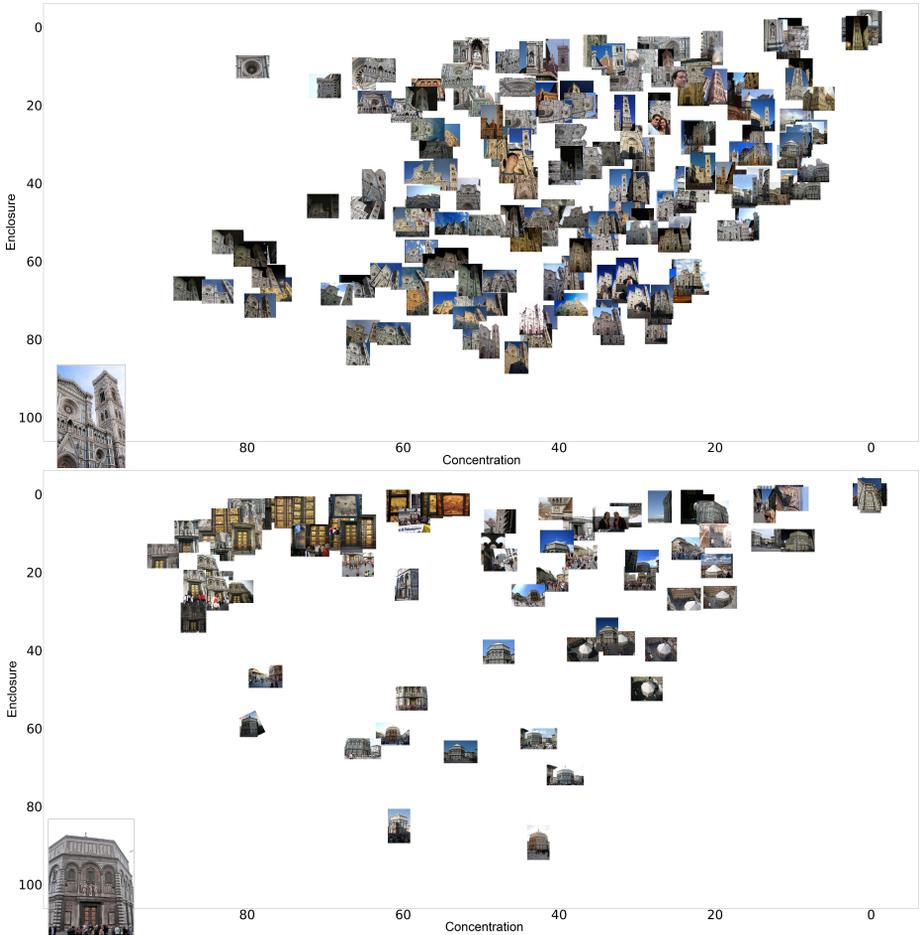
Fig. 13. Generalization plot: Big Ben.



We show a query image from the test set (lower left corner) and the concentration and enclosure between randomly sampled test images from the MegaDepth SfM dataset for which no depth maps are provided. The query image shows the side view on Saint Mark's Basilica. One can observe the front of the Basilica from a very oblique angle in the left-most fifth of the image.

The images in the left upper corner show images with high concentration and low enclosure. According to our classification these are close-ups. Especially around the coordinates (80, 10) one can clearly observe zoomed in views on the side of the Basilica. The right upper quarter of the cluster consist of the crop-outs and oblique-outs. Note the images around the coordinate (10,25). These are mostly front views on the Basilica, and correspond to the left-most part of the query image—from a very different angle. Lastly, observe that images in the left lower corner are similar to the query, and images around (20,60) are zoom-outs.

Fig. 14. Generalization plot: Venice.



Because the scenes is very complex, streets are narrow, and there are not many images of the same view we show two scenes from Florence.

Fig. 15. Generalization plot: Florence.

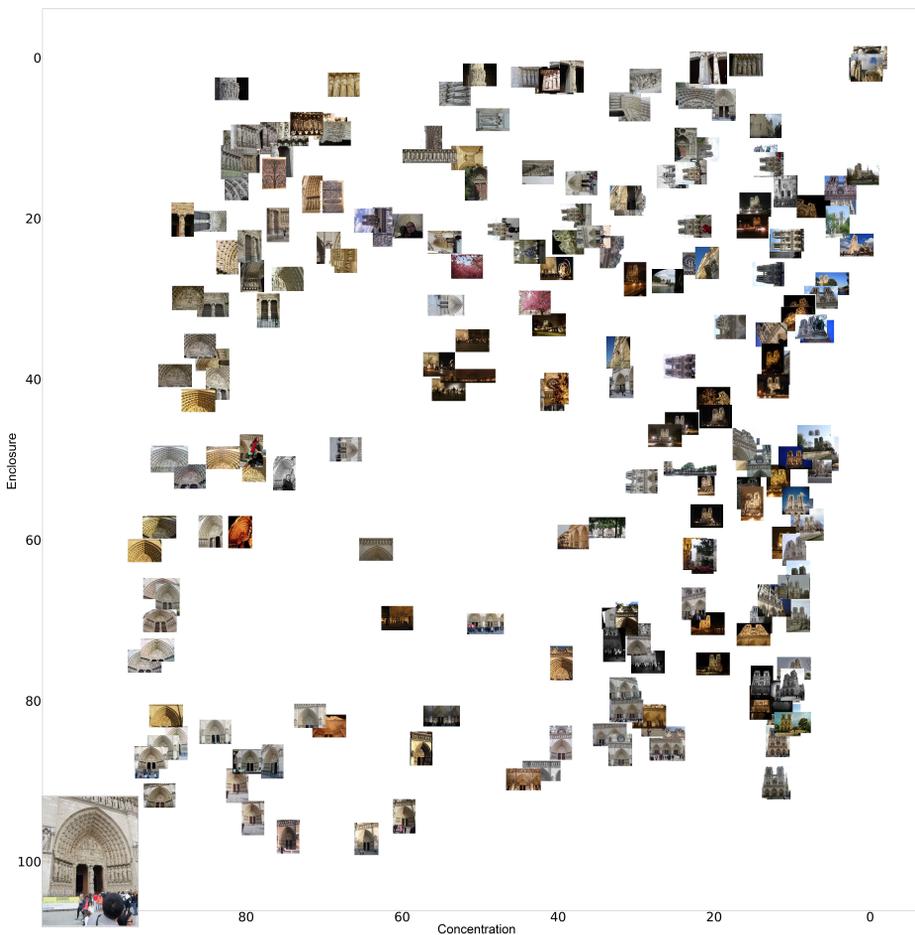
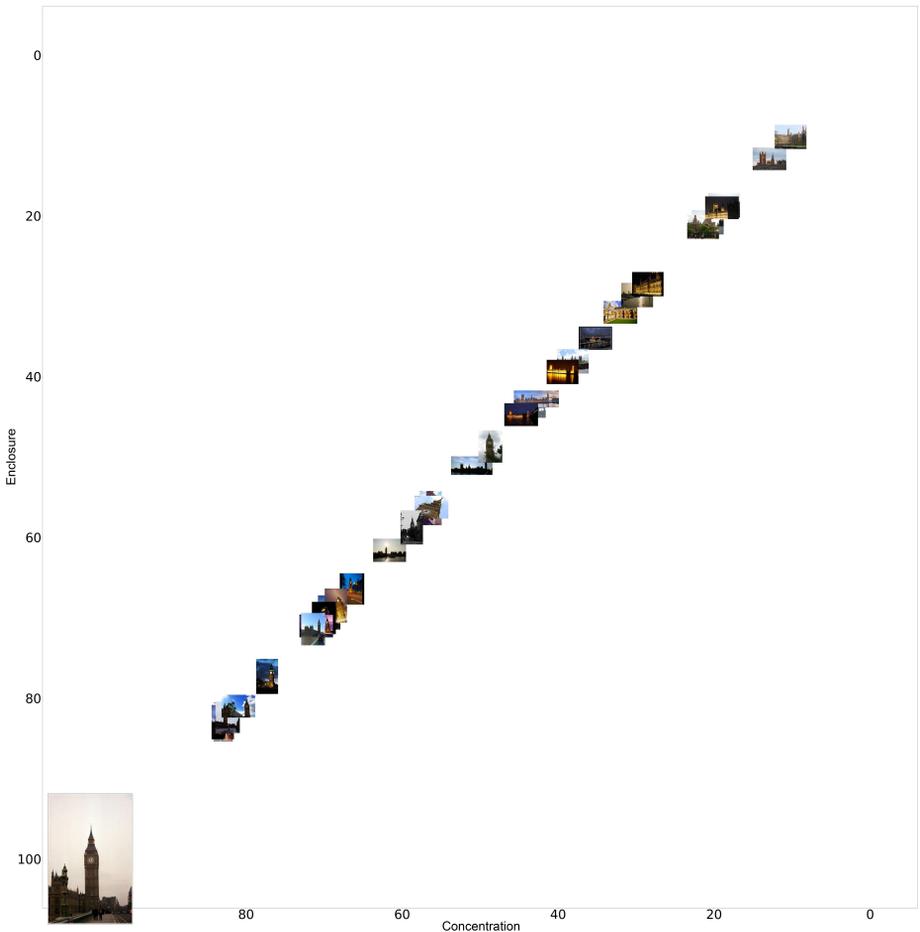


Fig. 16. Generalization plot: Notre Dame.



We show a query from the test set and report the predicted symmetric normalized surface overlap on a subset of test images. Because the embedding space measure is symmetric concentration and enclosure are equal for a given image. It can be observed that nearby images show similar views on the scene. However, the distance between the retrievals is not interpretable.

Fig. 17. Generalization plot: Big Ben NSO^{sym} .



87 % / 29%



98 % / 5%



83 % / 5%



79 % / 5%



82 % / 5%



97% / 33 %



65% / 30 %



73% / 14 %



76% / 27 %



93% / 39 %

855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899

Illustrated are several examples of how our method can estimate geometric relationships between images. For each pair the enclosure and concentration are calculated from which the relative estimated scaled can be derived. Based on that scale, the first image is resized and shown in the third position. The resized images match the scale of the scene in the first image to the scale in the second image. Note, that the resized images are sometimes very small, and the reader is encouraged to zoom into the images. The two numbers below each image pair show the estimated enclosure and concentration. Note that although some scale estimates are inaccurate, overwhelmingly the rescaling does not increase the scale difference between the two images, but only reduces it. *This caption applies to all Relative scale plots.*

Fig. 18. Relative scale plot: Venice No. 1.



95 % / 27%



69 % / 30%



63 % / 11%



95 % / 13%



88 % / 9%



69% / 18 %



94% / 30 %



93% / 9 %



95% / 36 %



85% / 29 %

Fig. 19. Relative scale plot: Venice No 2.

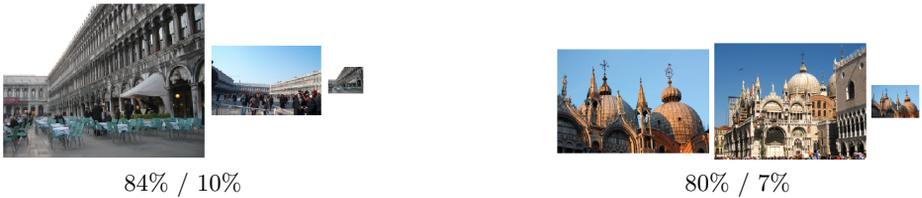


Fig. 20. Relative scale plot: Unsuccessful cases for Venice scene (test image pairs here were found by querying the database for images that had enclosure > 0.6 and $0.05 < \text{concentration} < 0.4$).

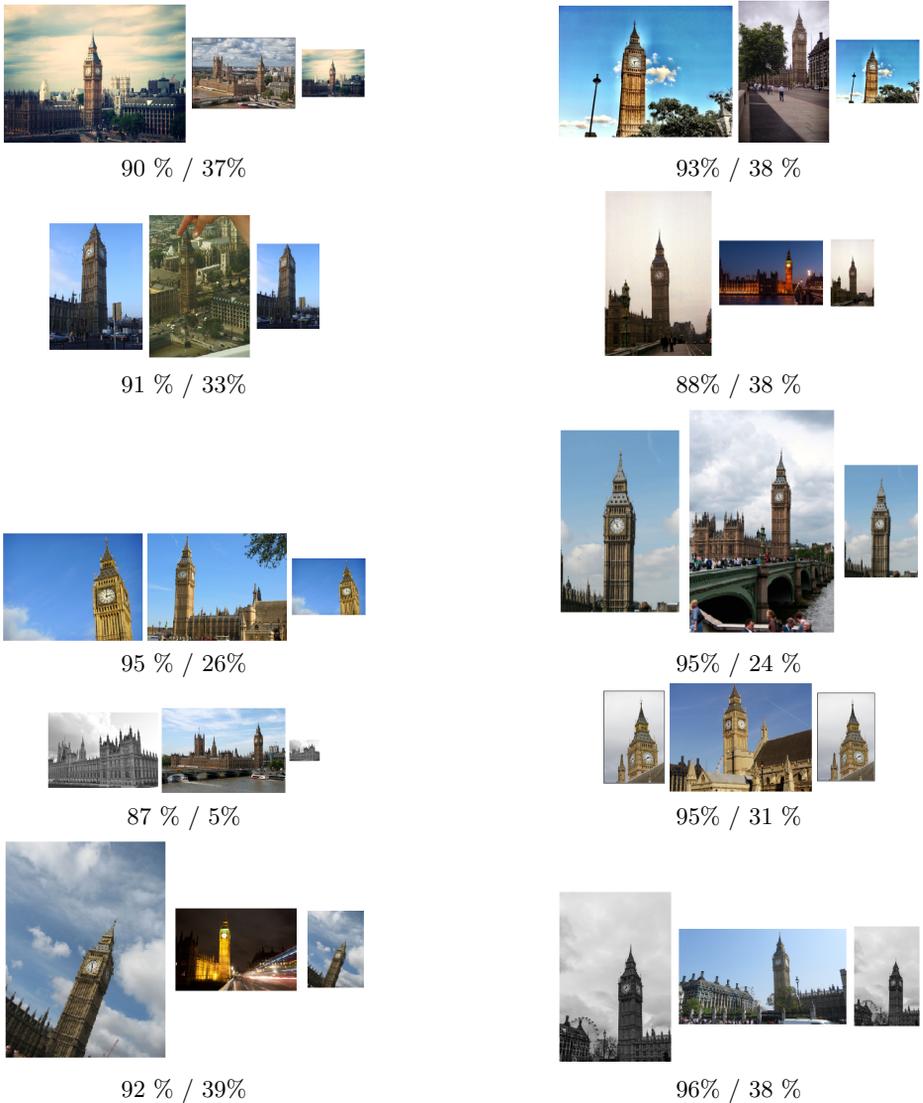


Fig. 21. Relative scale plot: Big Ben No. 1.

1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079



95 % / 24%



92 % / 24%



89 % / 21%



89 % / 21%



95 % / 24%



95% / 24 %



87% / 5 %



95% / 32 %



96% / 38 %



97% / 28 %

Fig. 22. Relative scale plot: Big Ben No. 2.



Fig. 23. Relative scale plot: Less successful cases Big Ben (18 out of 95 pairs shown in the document).

1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169

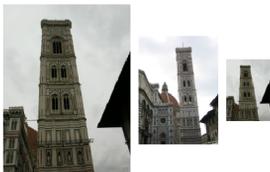
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169



88 % / 38%



91% / 39 %



82 % / 26%



81% / 5 %



92 % / 6%



87% / 32 %



71 % / 6%



70% / 10 %



81 % / 37%



91% / 39 %

Fig. 24. Relative scale plot: Florence No. 1.



65 % / 32%



81% / 25 %



69 % / 22%



64% / 33 %



88 % / 36%



66% / 38 %



82 % / 34%



74% / 16 %



74 % / 37%



83% / 37 %

Fig. 25. Relative scale plot: Florence No. 2.



82 % / 34%



89 % / 29%



88 % / 36%



82 % / 26%



87% / 32 %



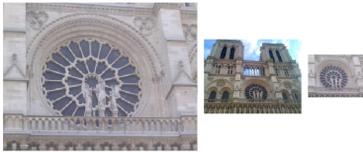
81% / 37 %



80% / 39 %



Fig. 26. Relative scale plot: Less successful cases Florence (7 out of 57 pairs shown in the document).



91 % / 30%



80% / 5 %



91 % / 5%



93% / 37 %



77 % / 36%



81% / 15 %



85 % / 39%



94% / 34 %



94 % / 5%



77% / 39 %

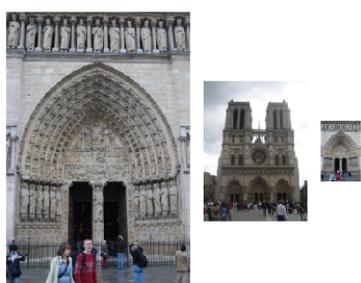
Fig. 27. Relative scale plot: Notre Dame No. 1.



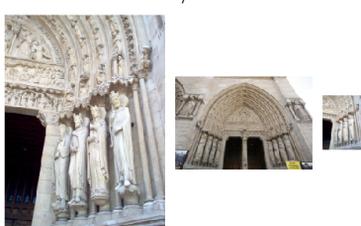
92 % / 10%



96 % / 28%



87 % / 14%



78 % / 13%



83 % / 38%



68% / 31 %



89% / 8 %



78% / 17 %



71% / 37 %



87% / 9 %

Fig. 28. Relative scale plot: Notre Dame No. 2.

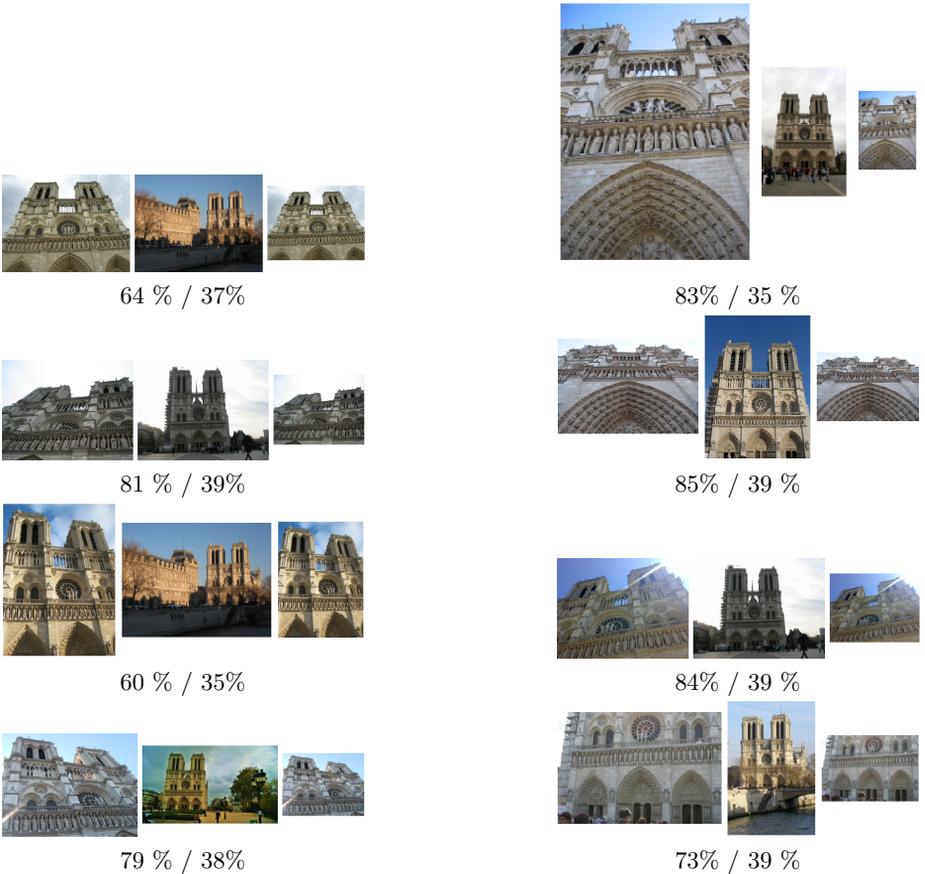


Fig. 29. Relative scale plot: Less successful cases Notre Dame (11 out of 61 pairs shown in the document).