

On Rendering Synthetic Images for Training an Object Detector

Supplementary Material

Anonymous ECCV submission

Paper ID 1421

1 Influences of the Similarity Measures

As discussed in the main submission, the similarity measure based on the same family of image features as the detection framework yields the best performances. Table 1 of the main submission illustrates the influence of different similarity measures on the detection accuracy of various classifiers. We provide the corresponding plots here.

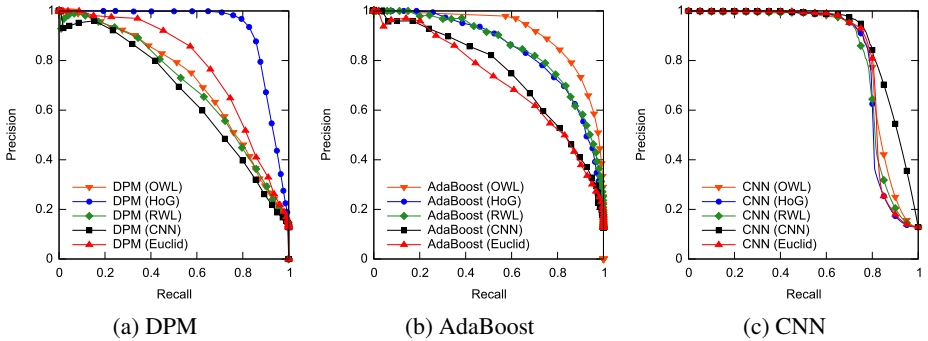


Fig. 1. The similarity measure has a strong influence on the final performance. We achieve better performance when using the similarity measure that relies on the same family of image features as the detection framework (best seen in color)




Fig. 2. Performances when using the Euclidean distance as a similarity measure. (best seen in color)

Fig. 1 shows the precision-recall curves for all possible combinations of the detection methods and the similarity measures we consider, and confirms that using the similarity measure that relies on the same family of images features as the detection framework yields better performance.

As shown in Fig. 2, relying on the Euclidean distance as a similarity measure to optimize the rendering parameters actually *degrades* the final performances for *all* the detectors.

2 Importance of the Rendering Parameters

To check whether the rendering effects have all a positive influence, and the importance of optimizing synthetic data generation parameters, we performed a set of evaluations in addition to those presented in Section 5.4 and Fig. 8 of the main submission.

We fixed all the capture parameters in Θ , setting their values to 0, which effectively means completely discarding the influence of the corresponding effect. We then varied only one of them and repeated this experiment with different values of the parameter, and for each of the capture parameters.

Classification method	Average precision			
	No effects	Boundaries blurring:		
		$\sigma^s = 1$	$\sigma^s = 1.5$	$\sigma^s = 2$
DPM	0.78	0.77	0.84	0.75
AdaBoost	0.65	0.73	0.79	0.79
CNN	0.86	0.89	0.85	0.86
	No effects	Motion blurring:		
		$\sigma_u^m = 0.3$ $\sigma_v^m = 0.3$	$\sigma_u^m = 0.5$ $\sigma_v^m = 0.5$	$\sigma_u^m = 1$ $\sigma_v^m = 1$
DPM	0.78	0.84	0.81	0.79
AdaBoost	0.65	0.72	0.79	0.79
CNN	0.86	0.87	0.88	0.89
	No effects	Random noise:		
		$\sigma^n = 0.5$	$\sigma^n = 0.9$	$\sigma^n = 1.1$
DPM	0.78	0.83	0.81	0.83
AdaBoost	0.65	0.75	0.79	0.75
CNN	0.86	0.89	0.86	0.85
	No effects	Material properties:		
		$w^d = 0.5$	$w^d = 1$	$w^d = 2$
DPM	0.78	0.83	0.84	0.86
AdaBoost	0.65	0.70	0.76	0.66
CNN	0.86	0.88	0.89	0.81

Table 1. Influence of various post-processing effects on the detection accuracy of different detectors.

The results are shown in Table 1 and Fig. 3. This shows that all the classifiers benefit from the application of every single post-processing effect.

To further highlight the effectiveness of the rendering parameters, in Fig. 4 we also compared the performance of every detector, trained on the real and synthetic data that was generated without using any of the post-processing steps with the ones when synthetic data was generated using all the post-processing steps with the parameters, optimized using appropriate similarity measures.

3 Importance of the Optimization over the Rendering Parameters

To show the importance of optimizing over the rendering parameters Θ , in Fig. 5 we compared the final performance obtained using optimized parameters with the final performance obtained with random parameters drawn from a uniform distribution. The minimum and maximum values for the random parameters were taken as the minimum and maximum values of the optimised parameters.

4 Rendering Parameters Distribution

Our method computes the capture parameters for each available real image. To show that these parameters are correlated in practice, Fig. 6 shows the distribution of each possible pair of parameters, for all the similarity measures we consider.

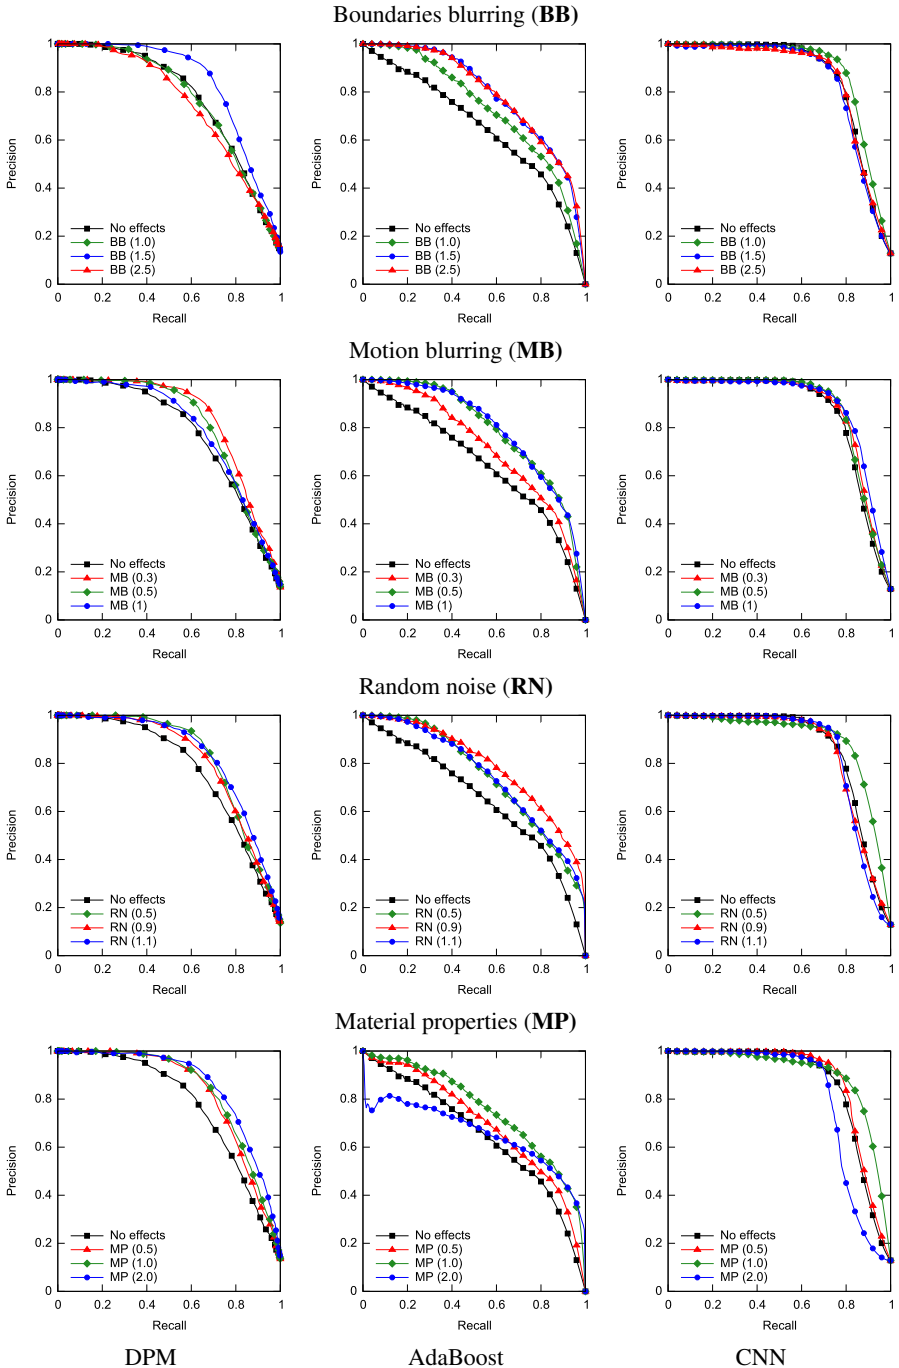


Fig. 3. Evaluation of the synthetic data generation effects. We fixed one capture parameter in Θ and then optimized the other parameters using the best similarity measure for each classification method. Each effect has clearly a positive influence of the quality of the synthetic data, however the impact is different for every classification method. (best seen in color)

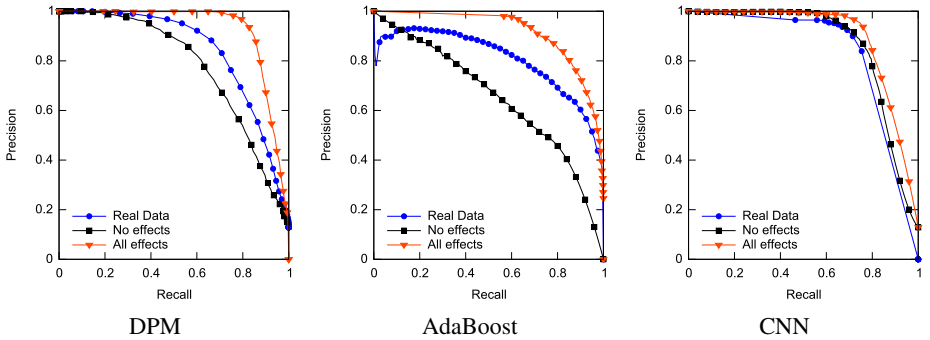




Fig. 4. Performances when using real data only, synthetic data without any post-processing effects and synthetic data with all the introduced effects, with the Θ parameters optimised according to the appropriate for every detector similarity measures. (best seen in color)




	Using real images only	Random parameters	Optimised parameters
Classification method:	Average precision:		
DPM	0.84	0.82	0.93
AdaBoost	0.80	0.82	0.92
CNN	0.85	0.87	0.89

Fig. 5. Comparison of the performances of different classifiers trained on real and synthetic data generated using corresponding similarity measures with those where the capture parameters are randomly selected. The optimized parameters always yield better performance. (best seen in color)


$d_{HoG}(\cdot, \cdot)$




$d_{WL}^L(\cdot, \cdot)$



$d_{WL}^R(\cdot, \cdot)$



$d_{CNN}(\cdot, \cdot)$



225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269




Fig. 6. Joint distributions of each possible pair of capture parameters, optimised using different similarity measures. The different parameters are clearly correlated, in a complex way. (best seen in color)