

Average of Synthetic Exact Filters: Supplemental Material

David S. Bolme Bruce A. Draper J. Ross Beveridge

Colorado State University
Fort Collins, CO 80521, USA

`bolme@cs.colostate.edu`

Abstract

This is supplemental material submitted to CVPR 2009 to accompany Average of Synthetic Exact Filters. This document contains multiple figures that were too large for the paper. Included in this document are results for both the left and the right eye. In addition, sample output from the algorithms should give the reader a better idea of how ASEF compares to other methods.

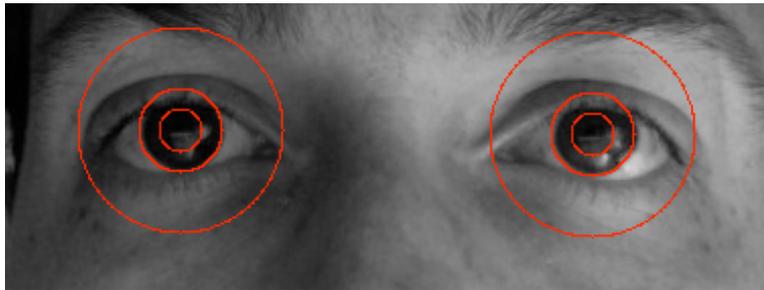


Figure 1. This figure shows three different eye localization success criteria overlaid on a sample image (from largest to smallest): $D < 0.25$, $D < 0.10$, $D < 0.05$. This gives some intuition as to the size of the targets that count as success. The criteria $D < 0.10$ was the primary operating point used for this work and is a target roughly the size of the iris.

Training Size	Experiment 1			Experiment 2		
	ASEF $\sigma = 1.0$	UMACE $\alpha = 0.9$	OTF $\alpha = 0.4$	ASEF $\sigma = 4.0$	UMACE $\alpha = 0.7$	OTF $\alpha = 0.1$
1						
2						
4						
8						
16						
32						
64						
8192						

Figure 2. This figure shows the effect of training set sizes on the filters. As the training set size increases, OTF over fits the training data. The averaging process in ASEF and UMACE prevents this effect. This figure also illustrates the effect of changing σ for ASEF training. Here you can see that $\sigma = 1.0$ emphasizes much higher frequencies than $\sigma = 4.0$. The higher frequencies in Experiment 1 allows ASEF to improve the accuracy of the localization, while the lower frequencies in Experiment 2 help to reduce false responses to other parts of the face. The parameter α has a similar effect which cannot be as easily seen in this figure. Because ASEF is trained on the full image, it can learn features that are not available to the other filters such as the shadow of the nose and eye brow that can be seen on the bottom row.

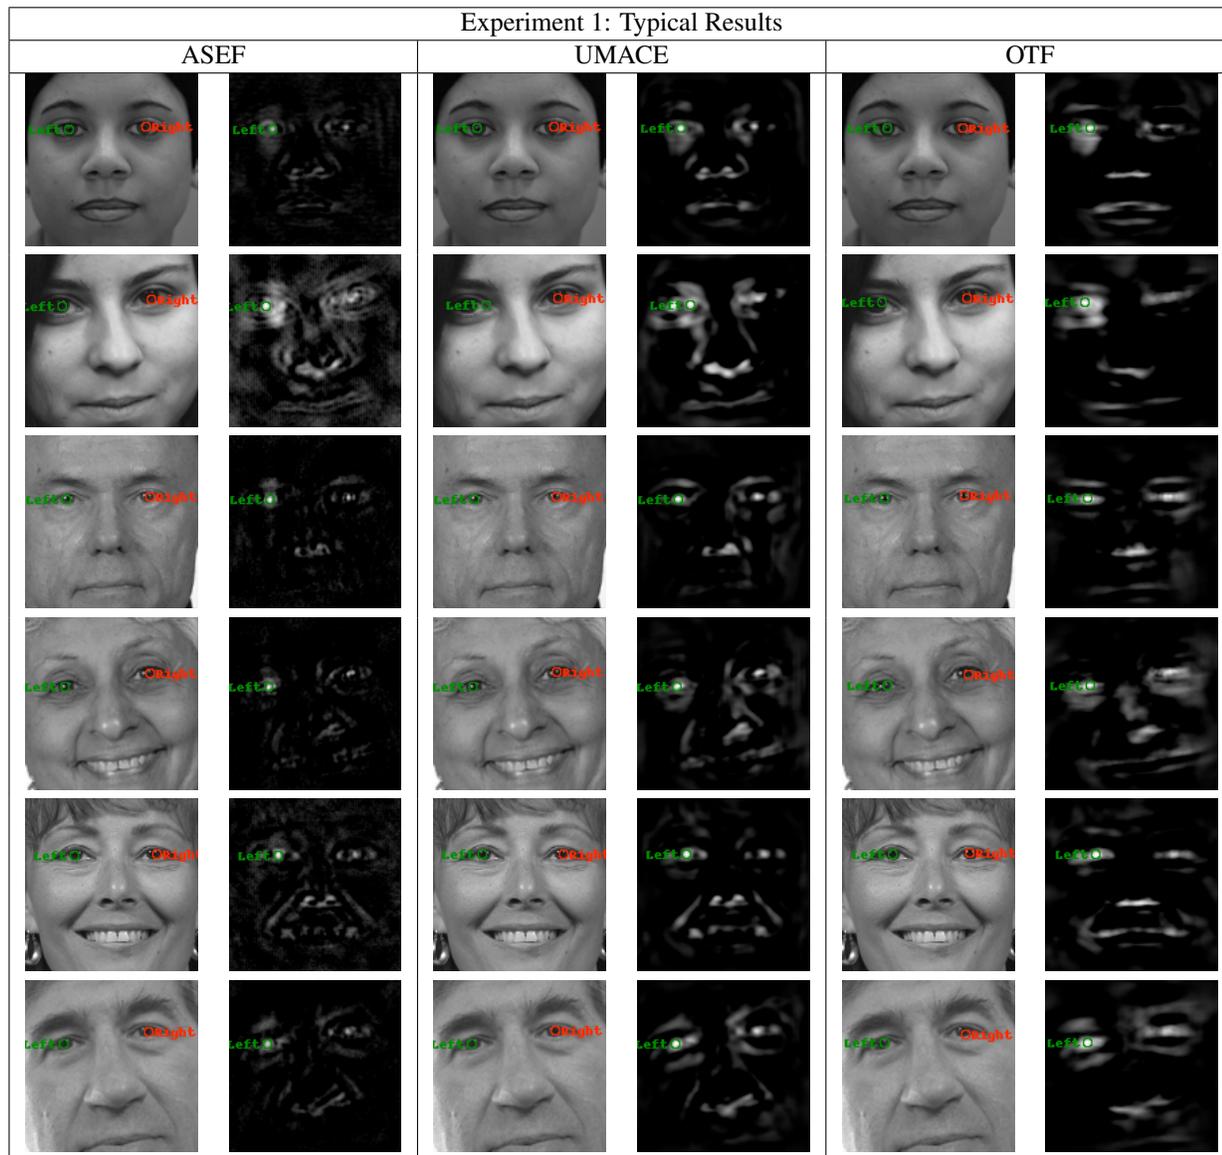


Figure 3. Experiment 1: This figure shows some typical responses for ASEF, UMACE, and OTF filters on the testing set. This is the easier experiment that is constrained to find maxima near the expected location of the eye. Here ASEF focuses on high frequency information that allow the filter to produce nice sharp peaks. ASEF also seems to do a better job suppressing the response to the other features of the face.

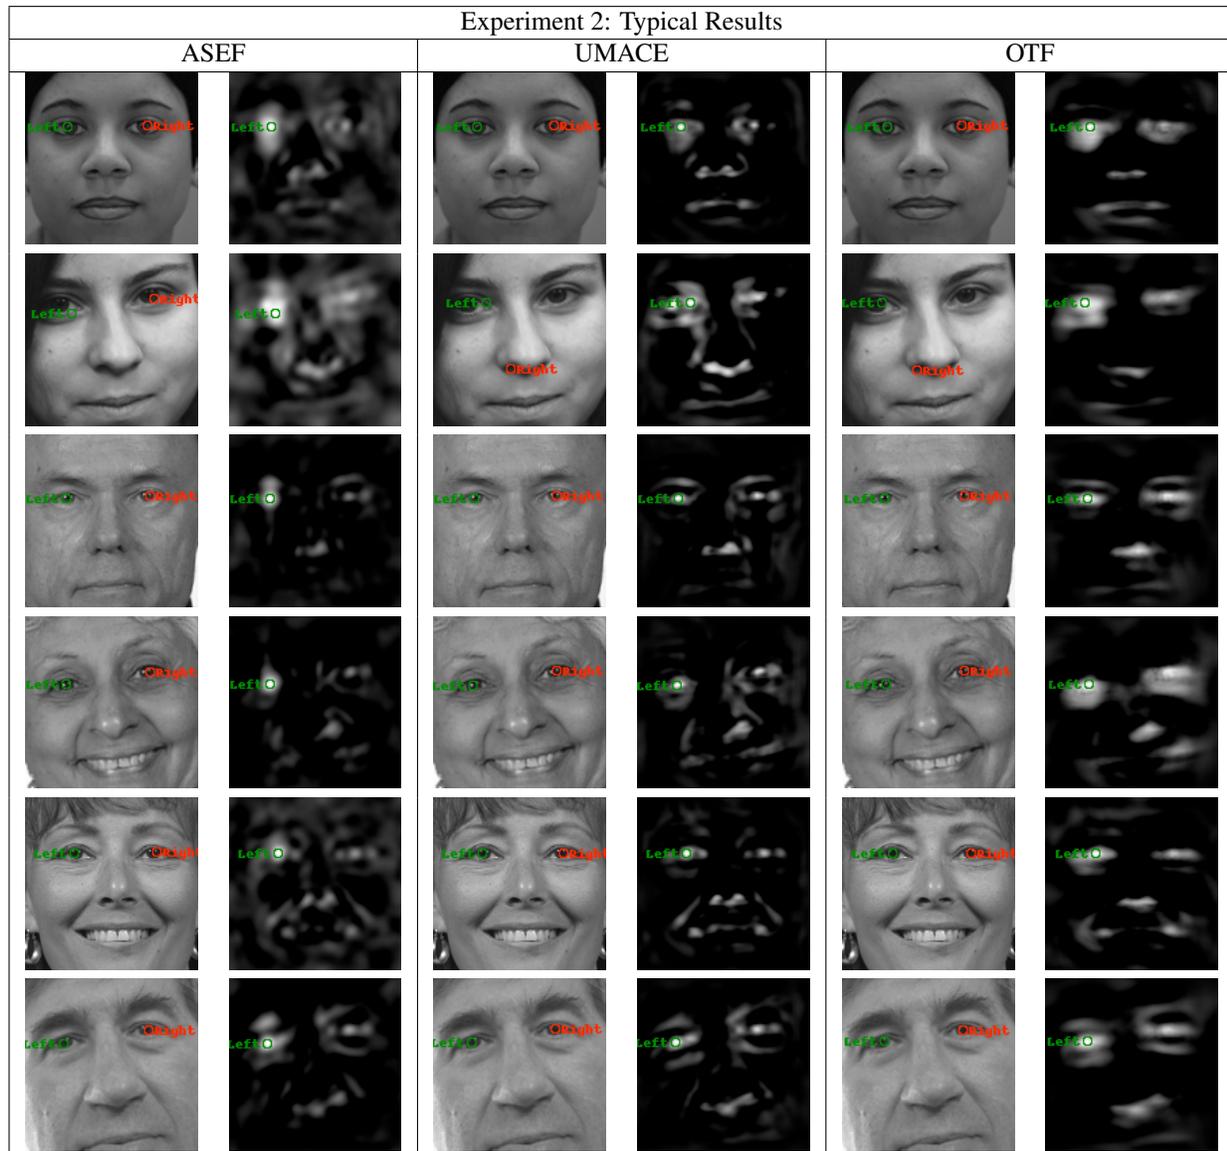


Figure 4. Experiment 2: This figure shows some typical responses for ASEF, UMACE, and OTF on the more difficult experiment. Because the filters are now looking for large responses any where in the image, ASEF uses a larger σ which helps to suppress responses to other parts of the face but does make the peak less sharp. All three of the algorithms are having difficulty accurately locating the eye in the second row.

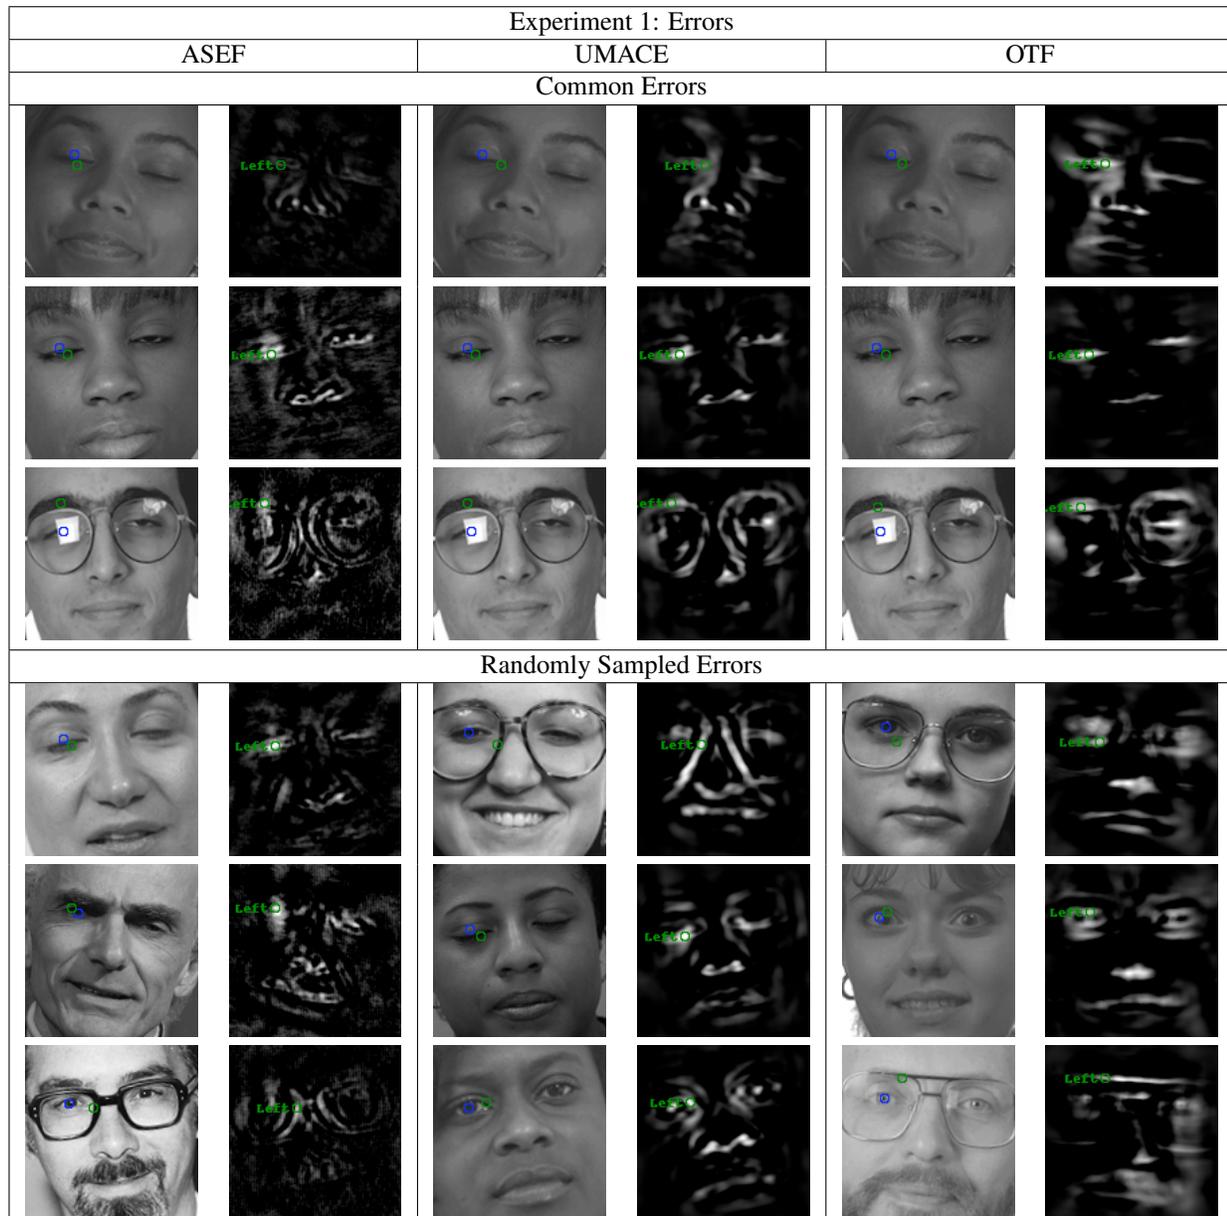


Figure 5. Experiment 1: This figures show some of the errors produced by the correlation filters on the easier experiment. The blue dot is the manually selected coordinate and the green dot is the coordinate selected by the correlation filter. In this case most errors for are caused by obvious problems in the testing imagery such as closed eyes, glasses with a lot of glare, or dark rimmed glasses. In many cases the filters all fail on the same images (see the first three rows).

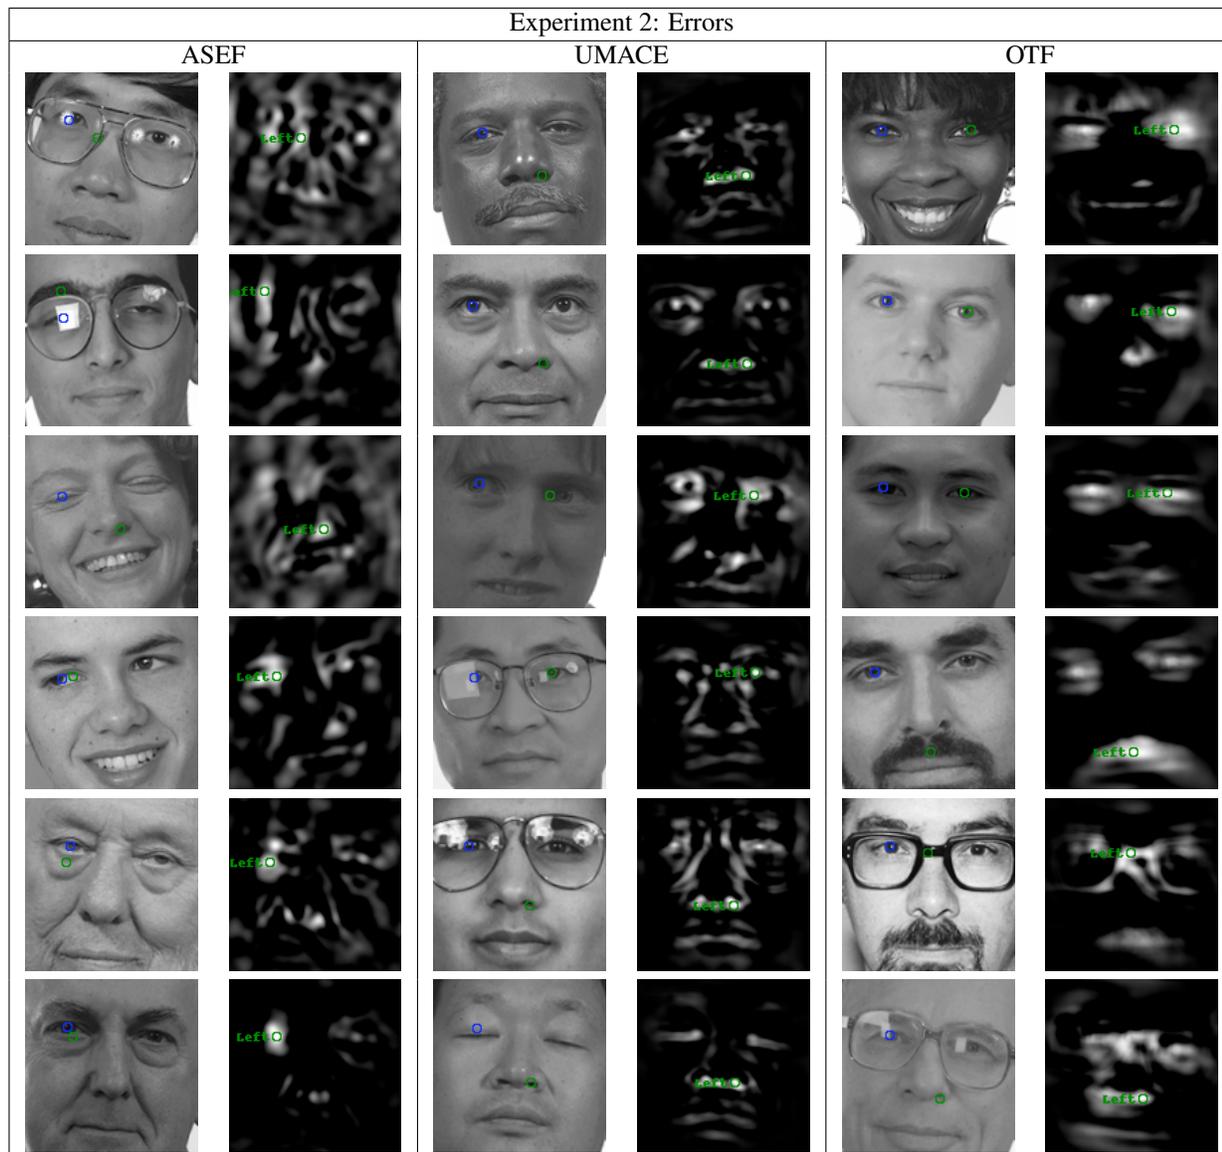


Figure 6. Experiment 2: These are some randomly sampled errors from the more difficult experiment. Because the algorithms now find the maximum of the entire image, it is possible for the algorithms to respond better to other features of the face. You can see that the majority of the errors for UMACE and OTF are responses to the "wrong" eye or to the nose. ASEF make very few of that type of error.

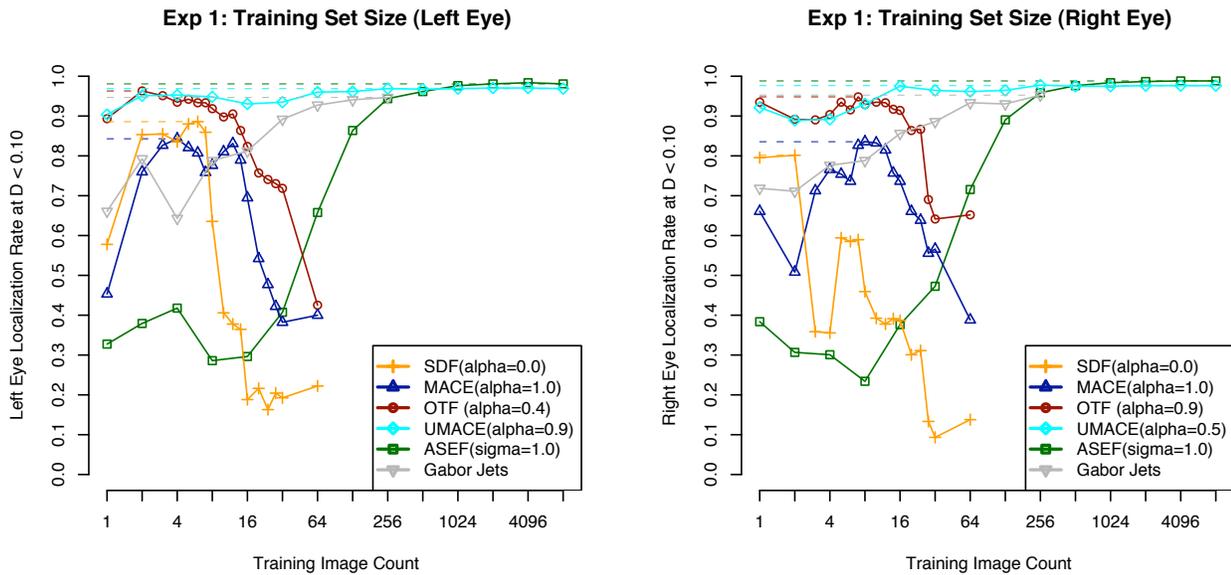


Figure 7. Experiment 1: These are the training plots for both eyes. The left eye plot is shown in the full paper. See the paper for a description.

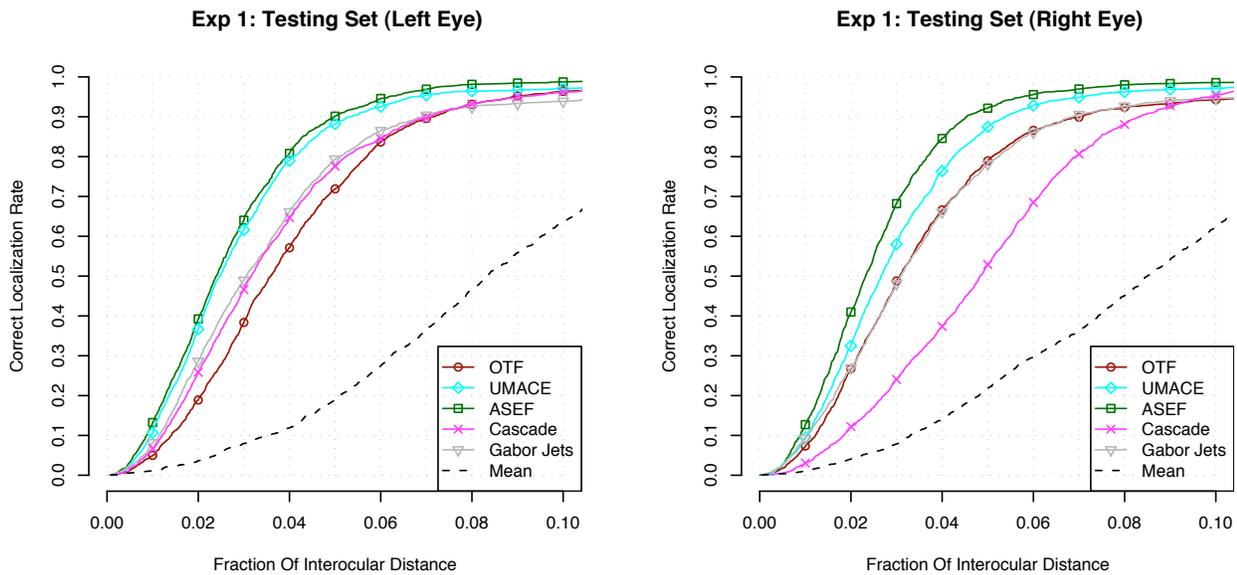


Figure 8. Experiment 1: These are the testing plots for both eyes. The left eye plot is shown in the full paper. See the paper for a description.

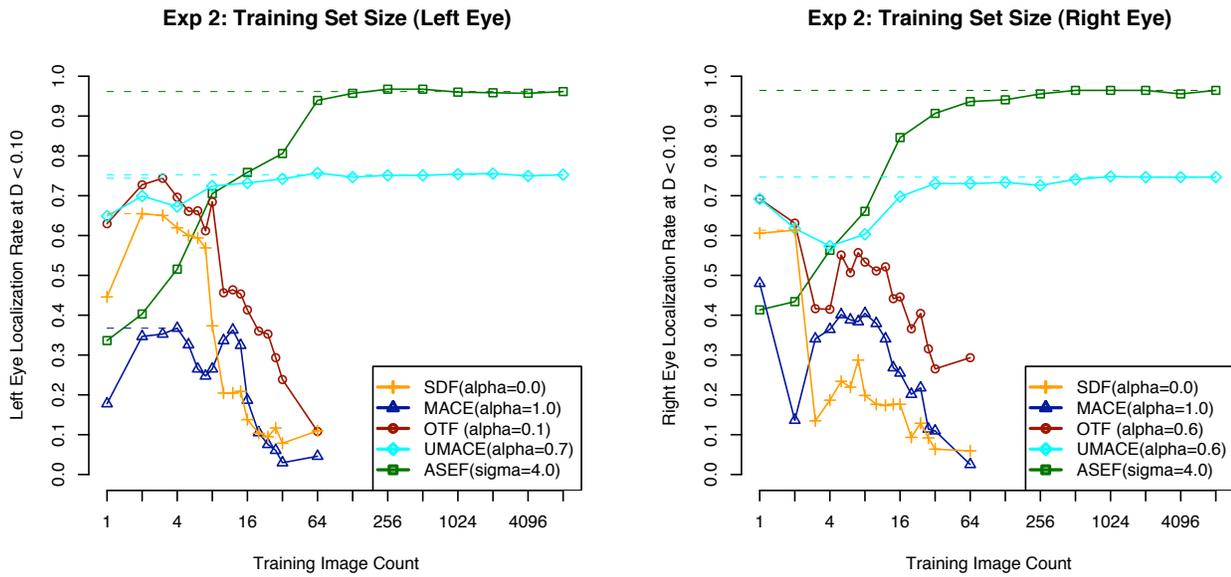


Figure 9. Experiment 2: These are the training plots for both eyes. The left eye plot is shown in the full paper. See the paper for a description.

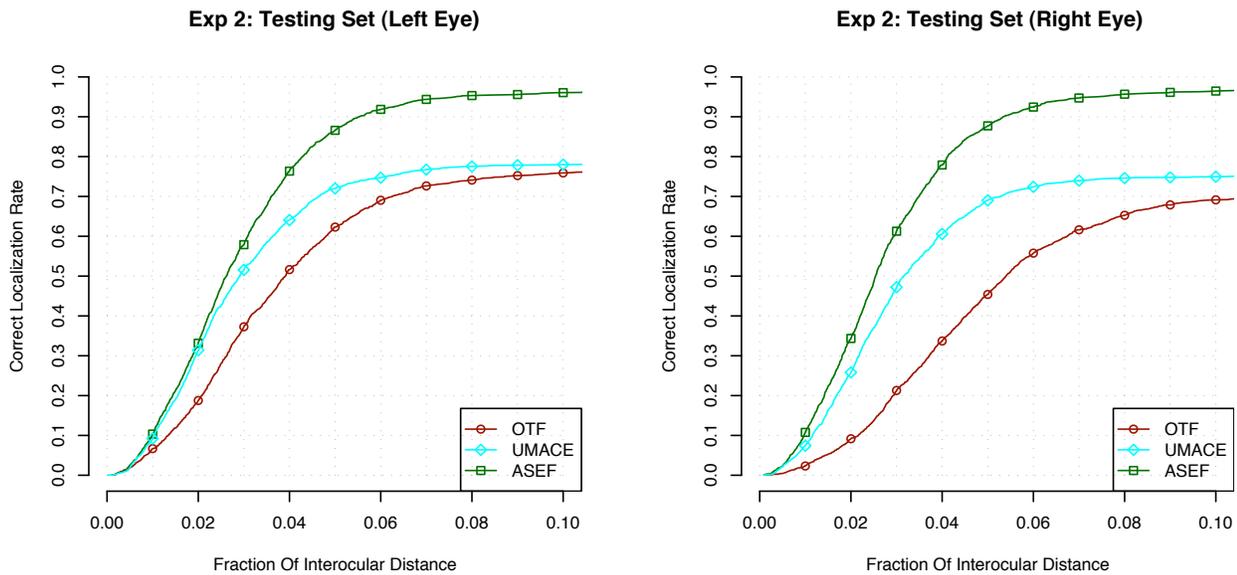


Figure 10. Experiment 2: These are the testing plots for both eyes. The left eye plot is shown in the full paper. See the paper for a description.

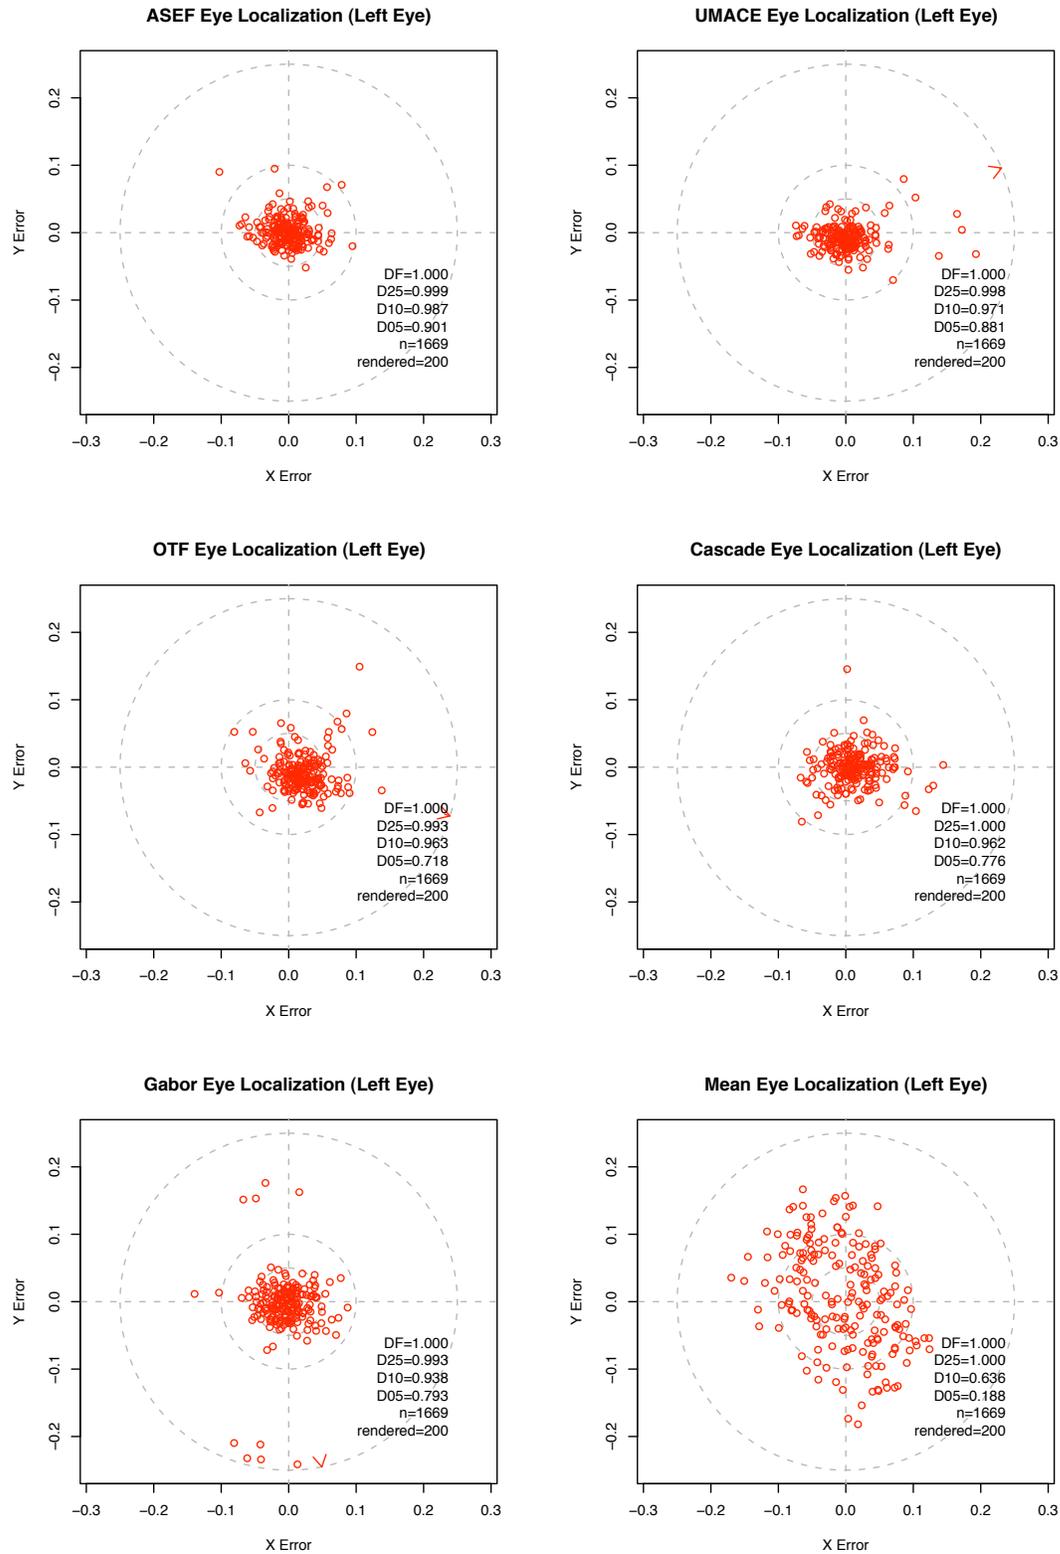


Figure 11. Experiment 1: This figure shows the output of the different algorithms relative to the manually selected coordinate. The manually selected coordinate is always at location (0,0). A red dot indicates the output for a single training image. Only 200 dots are rendered corresponding to the first 200 of the 1669 testing images. Arrows indicate the direction and distance (length) to outliers. Units are in the same terms as D and the dashed gray circles indicate $D = 0.25$, $D = 0.10$, and $D = 0.05$. All algorithms in this plot produces nicely clustered points centered on on the manual coordinate (0, 0). The mean plot indicates the variance of the coordinates in the testing set.

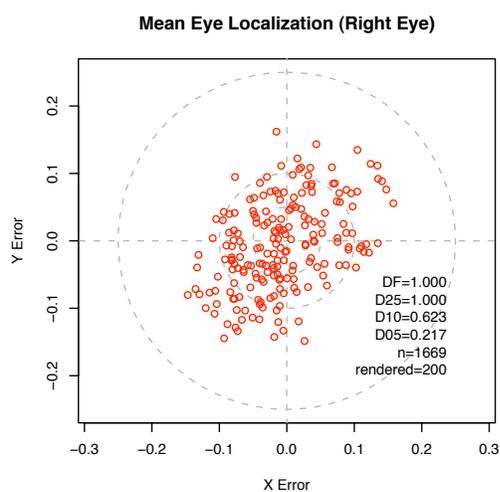
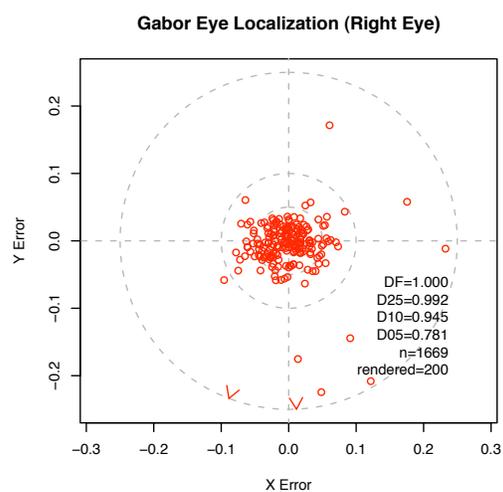
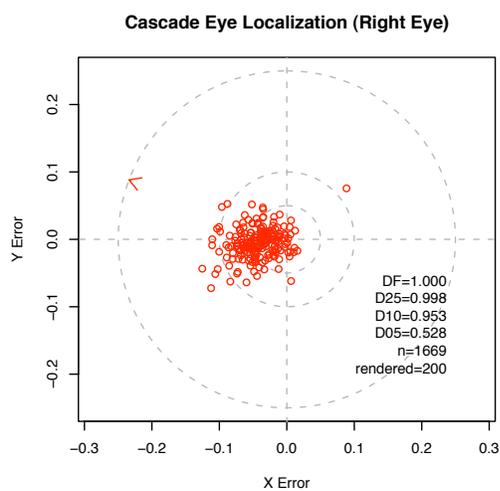
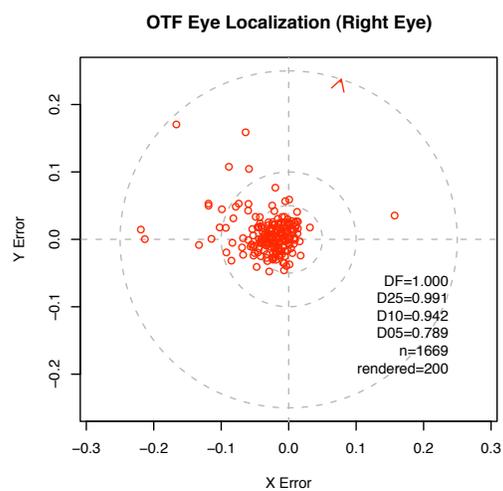
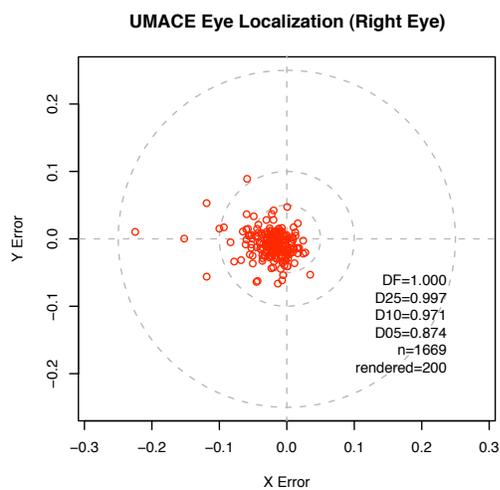
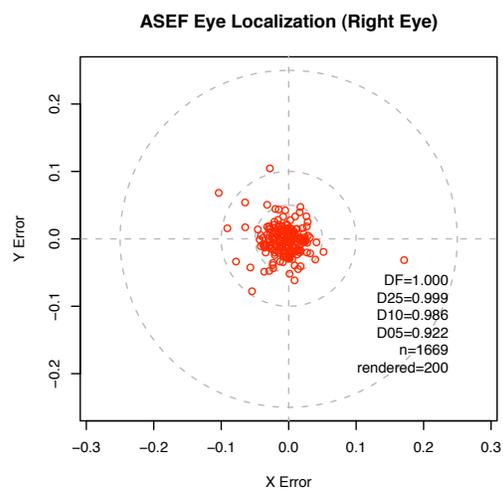


Figure 12. Experiment 1: See Figure 11 for a description.

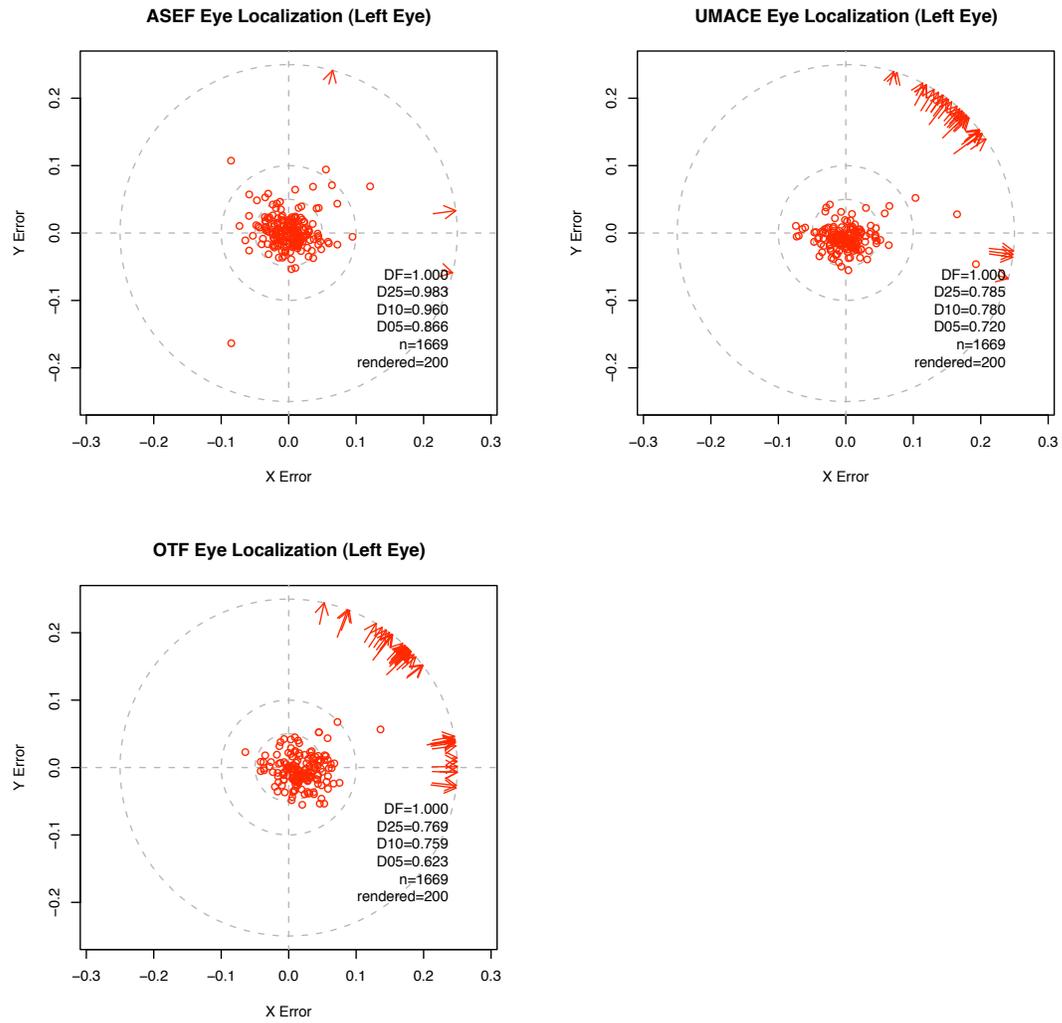


Figure 13. Experiment 2: See Figure 11 for a description. This figure shows that ASEF has far fewer errors than UMACE and OTF. The arrows in UMACE and OTF plots indicate that there are many false responses to the “wrong” eye or nose (the y-axis is reversed relative to the image y-axis). ASEF produces very few such errors.

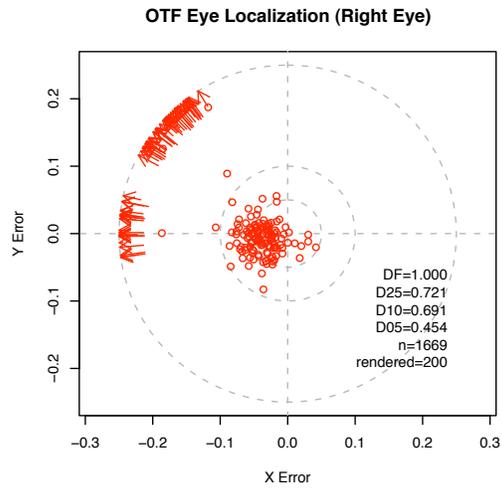
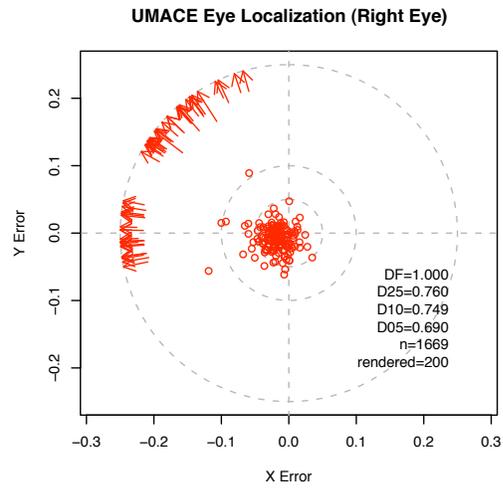
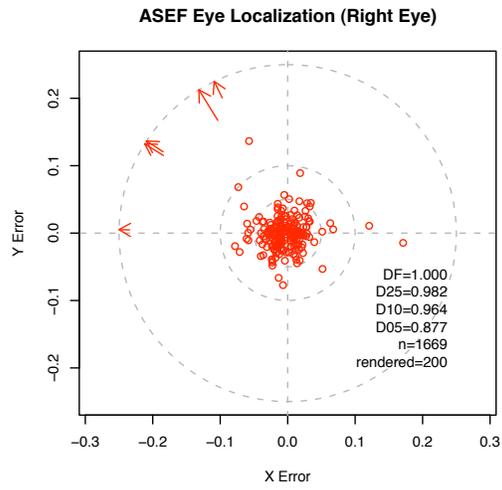


Figure 14. Experiment 2: See Figure 11 and Figure 13 for a description.