Correlation Filters for Object Alignment (Supplementary Material)

Vishnu Naresh Boddeti
Carnegie Mellon University
naresh@cmu.edu

Takeo Kanade
Carnegie Mellon University
tk@cs.cmu.edu

B.V.K. Vijaya Kumar
Carnegie Mellon University
kumar@ece.cmu.edu

1. Comparing Appearance Models

We compare the performance of the appearance models (using the RANSAC BPSI shape model) on a per image basis by plotting the sorted error of the proposed vector correlation filter (VCF) as a reference along with the corresponding error for random forest (RF) and support vector machine (SVM) in Fig. 1 and Fig. 2 respectively. Notice that for most of the images, the RMSE of the VCF (blue line) is well below the RMSE of both RF (red spikes above blue line) and SVM (red spikes above blue line).

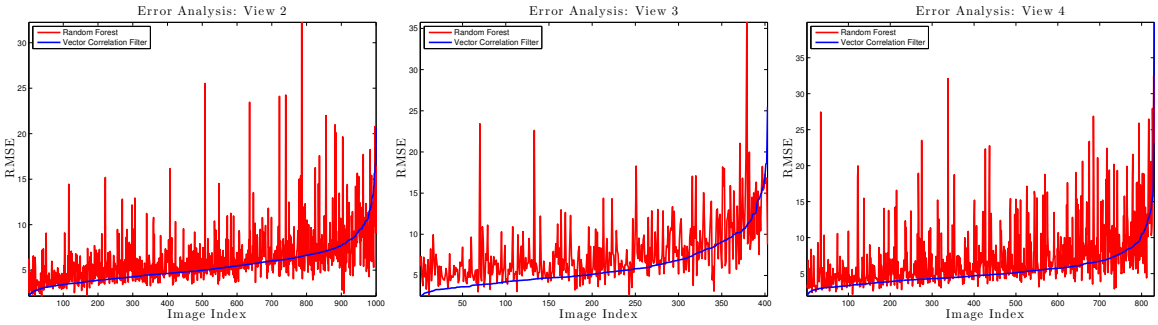

Figure 1. One-on-One comparison of the RMSE between the proposed vector correlation filter and random forest appearance models.

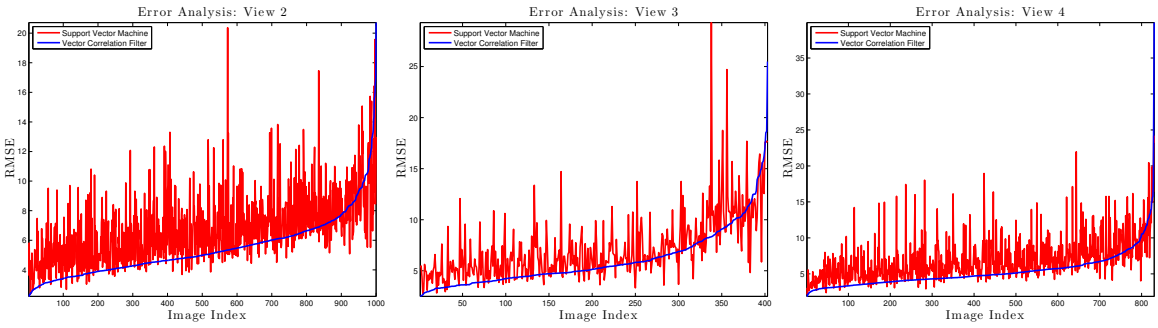

Figure 2. One-on-One comparison of the RMSE between the proposed vector correlation filter and SVM appearance models.

2. Comparing Shape Models

We compare the performance of the shape models (using the VCF based appearance model) on a per image basis by plotting the sorted error of Greedy BPSI as a reference along with the corresponding error for RANSAC BPSI and Oracle BPSI in Fig. 3 and Fig. 4 respectively. Notice that RANSAC BPSI and Oracle BPSI has more images with larger RMSE

(blue spikes above the red line). We further remind that Greedy BPSI can test fewer shape hypothesis compared to RANSAC BPSI for similar alignment performance.

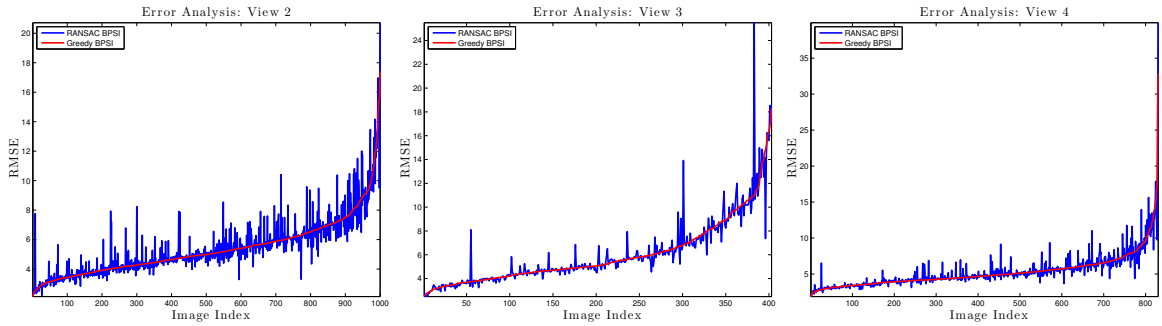

Figure 3. One-on-One comparison of the RMSE between the Greedy BPSI and RANSAC BPSI shape models.

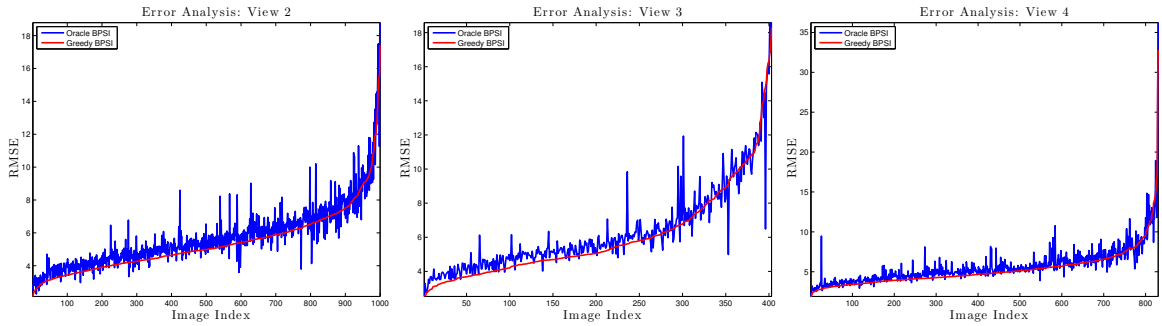

Figure 4. One-on-One comparison of the RMSE between the Greedy BPSI and Oracle BPSI shape models.

3. Robustness to Initialization Errors

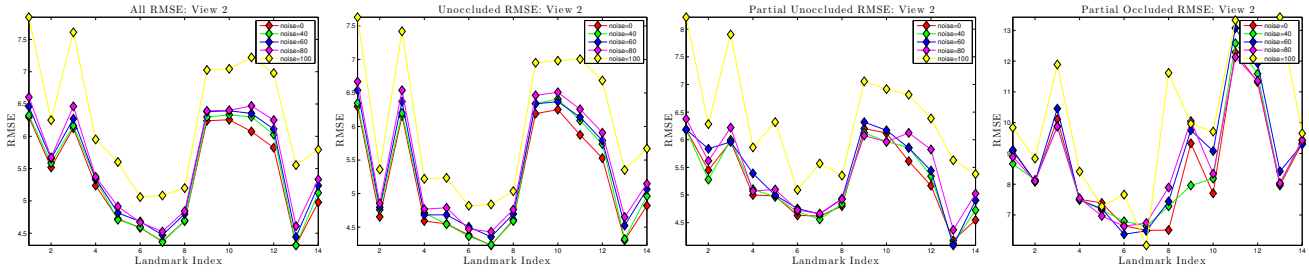

Figure 5. RMSE with varying levels of noise in the initialized shape.

To investigate the robustness of our algorithm to initialization errors, we initialize the car shape by a randomly perturbed mean shape with varying levels of noise. Fig. 5 shows the RMSE for different noise levels for the half-frontal pose and the trends for other poses are similar. While the performance of our alignment model drops as the noise level increases, it is however surprisingly robust with an average increase in error of less than 1 pixel even with noise levels of 100 pixel random shifts added to the initial shape. This is due to the robustness of the landmark detector as well as the fact that Greedy BPSI (and RANSAC BPSI) relies only on a minimal subset of features to generate a hypothesis and hallucinate the full shape.

4. Qualitative Results

Figure 6. Car Alignment: The three columns correspond to 1) VCF + Greedy BPSI, 2) VCF + RANSAC BPSI, 3) RF + RANSAC BPSI