Pilot Study Findings

Tables 3.10 through 3.21 show the average, median, and standard deviations for the overlap estimate, intersection sketch, and single-target sketch measures.

Questions One and Two: How do participants perform with multi-layer DDS images compared to single-layer DDS images shown side-by-side? Is there a point where the multitude of additional layers causes enough visual interference that the task is more accurately performed looking at the targets side-by-side?

This analysis sought to answer the question of how task performance for overlaid DDS images compared to performance for the side-by-side DDS images. In the analysis of performance for the *Side-by-Side* condition is compared to performance with *C7*, the seven-distractor condition. The linear regression analysis predicts that *C7* will have the highest error percentages for the overlap estimation task and the lowest sketch scores for the sketch task, thus, *C7* represents worst-case performance for overlaid DDS images. Because performance for the *Side-by-side* condition was always worse on average than for *C7* and because *C7* represents worst-case performance, it is possible to statistically compare the *Side-by-side* condition with just *C7* and draw conclusions for all the overlaid DDS conditions. All analyses were performed using SPSS for Windows statistical software [SPSS, 2001].

By making fewer statistical comparisons, the statistical power of the results is maintained. The more comparisons one makes with the same data, the more the chance for a Type II error increases. Because performance for the *Intersection* was always near to the performance for *C0*, and because the *Intersection* and *C0* have essentially the same semantic interpretation (neither have distractors) I decided not to analyze the results for the *Intersection* condition.

For the overlap estimation task, a comparison of means shows that performances for *C7* is significantly better than for the *Side-by-Side* view ($p = 0.01$). *Target Display Type* is not a significant factor, nor is there an interaction between *Display Condition* and *Target Display Type*.

For the intersection sketch task a comparison of means shows that performances for the *Side-by-Side* and *C7* are not significantly different ($p = 0.877$). *Target Display Type* is not a significant factor, nor is there an interaction between *Display Condition* and *Target Display Type*. Figures 3.35 through 3.40 show the results of the analysis.

	Mean	Median	Std Deviation
Overall	.085	.054	.10

Table 3.10: Overall mean, median, and standard deviation of participant error in the overlap estimation task. Error is defined as the absolute value of the difference between the estimated area and the actual area. – *Display Condition* levels *C0-C7*.

	Mean	Median	Std Deviation
Color-Color	.063	.045	.06
Color-Bump	.078	.053	.08
Bump-Bump	.115	.066	.12

Table 3.11: Overall mean, median, and standard deviation of participant error by *Target Display Type* averaged across all eight levels of *Display Condition* for the overlap estimation task– *Display Condition* levels *C0-C7*. Notice that the mean performance for the *Bump-Bump* group is nearly twice that of the *Color-Color* group, but that the median is not. The standard deviation is also larger for the *Bump-Bump* group.

	Mean	Median	Std Deviation
C0	.058	.044	.071
C1	.059	.044	.057
C2	.061	.050	.050
C3	.080	.047	.104
C4	.098	.056	.109
C5	.084	.055	.100
C6	.130	.094	.123
C7	.112	.068	.127

Table 3.12: Overall mean, median, and standard deviation of participant error by *Display Condition* averaged across all three *Target Display Type* levels for the overlap estimation task– *Display Condition* levels *C0-C7*. Notice that the mean performance for *C7* is nearly twice that of *C0*, but that the median value is not. The standard deviation is also larger for *C7*.

	Color-Color			*Color-Bump*			*Bump-Bump*		
Display Condition	Mean	Median	Std Deviation	Mean	Median	Std Deviation	Mean	Median	Std Deviation
C0	.042	.028	.034	.053	.032	.049	.077	.051	.106
C1	.054	.044	.045	.052	.044	.048	.070	.051	.072
C2	.059	.045	.053	.055	.046	.048	.068	.056	.049
C3	.064	.036	.113	.064	.048	.056	.110	.064	.125
C4	.066	.047	.060	.070	.044	.073	.158	.120	.148
C5	.064	.050	.059	.069	.055	.084	.119	.069	.134
C6	.085	.056	.089	.138	.100	.127	.168	.133	.137
C7	.069	.055	.059	.118	.103	.108	.151	.089	.175
S-S	**.175**	**.128**	**.154**	**.142**	**.110**	**.109**	**.186**	**.160**	**.142**
I	.054	.035	.056	.055	.044	.049	.073	.058	.056

Table 3.13: Error statistics for the overlap estimation task by *Target Display Type* and *Display Condition*. The bottom two rows are the *Side-by-side* and *Intersection* views. Note that the values for the *Side-by-side* view are over twice those for *C7* in the *Color-Color* group.

	Mean	Median	Std Deviation
Overall	.648	1.000	.577

Table 3.14: Overall mean, median, and standard deviation of the sketch scores for the intersection sketch task – *Display Condition* levels *C0-C7*. Lower sketch scores indicate poorer performance, and a score of 1 indicates a correct sketch. In all cases the median score is equal to one, showing that over half of the sketches were judged to be correct.

	Mean	Median	Std Deviation
Color-Color	.719	1.000	.503
Color-Bump	.662	1.000	.574
Bump-Bump	.563	1.000	.637

Table 3.15: Overall mean, median, and standard deviation of the sketch scores by *Display Condition* averaged across all eight levels of *Display Condition* for the intersection sketch task– *Display Condition* levels *C0-C7*.

	Mean	Median	Std Deviation
C0	.784	1.000	.431
C1	.780	1.000	.437
C2	.713	1.000	.508
C3	.677	1.000	.549
C4	.601	1.000	.594
C5	.541	1.000	.655
C6	.523	1.000	.672
C7	.563	1.000	.652

Table 3.16: Overall mean, median, and standard deviation of the sketch scores by *Display Condition* averaged across all three *Target Display Type* levels for the intersection sketch task– *Display Condition* levels *C0-C7*.

	Color-Color			*Color-Bump*			*Bump-Bump*		
Display Condition	Mean	Median	Std Deviation	Mean	Median	Std Deviation	Mean	Median	Std Deviation
C0	.930	1.000	.226	.920	1.000	.234	.860	1.000	.286
C1	.970	1.000	.157	.940	1.000	.193	.840	1.000	.293
C2	.950	1.000	.182	.900	1.000	.267	.780	1.000	.337
C3	.870	1.000	.282	.930	1.000	.226	.630	1.000	.438
C4	.940	1.000	.193	.900	1.000	.267	.570	.500	.440
C5	.900	1.000	.267	.890	1.000	.291	.630	1.000	.438
C6	.810	1.000	.363	.700	1.000	.429	.490	.500	.434
C7	.850	1.000	.323	.680	1.000	.426	.490	.500	.468
S-S	.730	1.000	.394	.810	1.000	.333	.610	.750	.432
I	.970	1.000	.120	.960	1.000	.170	.830	1.000	.296

Table 3.17: Statistics for the intersection sketch task by *Target Display Type* and *Display Condition*. The bottom two rows are the *Side-by-side* and *Intersection* views.

	Mean	Median	Std Deviation
Overall	.930	1.000	.23

Table 3.18: Overall mean, median, and standard deviation of the sketch scores for the single-target sketch task – *Display Condition* levels *C0-C7*.

	Mean	Median	Std Deviation
Color-Color	.975	1.000	.14
Color-Bump	.978	1.000	.13
Bump-Bump	.839	1.000	.33

Table 3.19: Overall mean, median, and standard deviation of the sketch scores by *Display Condition* averaged across all eight levels of *Display Condition* for the single-target sketch task– *Display Condition* levels *C0-C7*.

	Mean	Median	Std Deviation
C0	.987	1.000	.081
C1	.977	1.000	.106
C2	.980	1.000	.098
C3	.920	1.000	.246
C4	.917	1.000	.268
C5	.940	1.000	.216
C6	.873	1.000	.296
C7	.850	1.000	.346

Table 3.20: Overall mean, median, and standard deviation of the sketch scores by *Display Condition* averaged across all three *Target Display Type* levels for the single-target sketch task– *Display Condition* levels *C0-C7*.

	Color-Color			*Color-Bump*			*Bump-Bump*		
Display Condition	Mean	Median	Std Deviation	Mean	Median	Std Deviation	Mean	Median	Std Deviation
C0	1.000	1.000	.000	1.000	1.000	.000	.960	1.000	.137
C1	1.000	1.000	.000	.990	1.000	.071	.940	1.000	.164
C2	.990	1.000	.071	1.000	1.000	.000	.950	1.000	.152
C3	.980	1.000	.141	.970	1.000	.157	.810	1.000	.348
C4	1.000	1.000	.000	.970	1.000	.157	.780	1.000	.406
C5	1.000	1.000	.000	.970	1.000	.157	.850	1.000	.323
C6	.900	1.000	.267	.990	1.000	.071	.730	1.000	.394
C7	.930	1.000	.248	.930	1.000	.248	.690	1.000	.451
S-S	1.000	1.000	.000	1.000	1.000	.000	.960	1.000	.137
I	1.000	1.000	.000	1.000	1.000	.000	.950	1.000	.152

Table 3.21: Statistics for the single-target sketch task by *Target Display Type* and *Display Condition*. The bottom two rows are the *Side-by-side* and *Intersection* views.

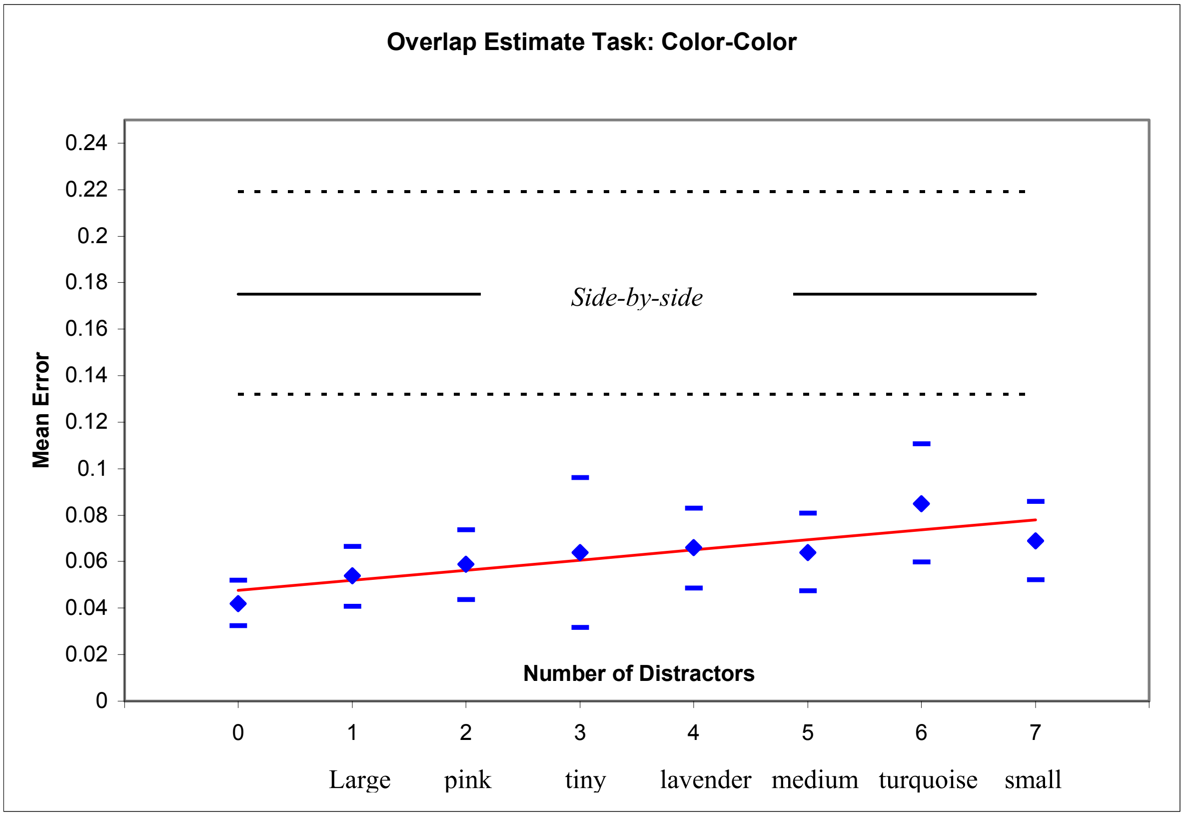

Figure 3.35: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the *Color-Color* group. The blue diamonds mark the average error for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the top of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

People performed significantly worse when the target images were displayed side-by-side than when the target images were displayed in one image with seven other distractor shapes. This illustrates the power of overlaying images when performing spatial correlation tasks. Both *C7* and *C6* are significantly different from the *Side-by-Side* view ($p = 0.011$) for *C7* and ($p = 0.045$) for *C6*. The remaining conditions are also significantly different.

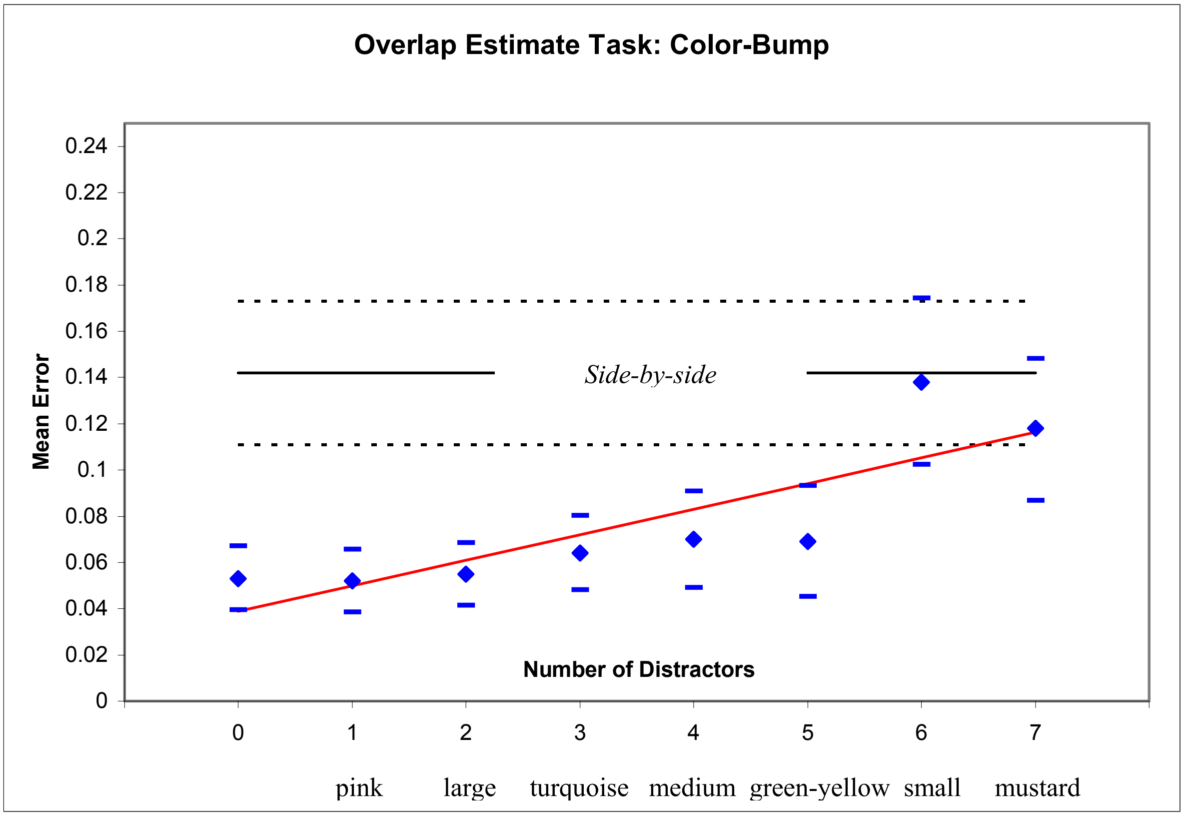

Figure 3.36: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the *Color-Bump* group. The blue diamonds mark the average error for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the top of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

Neither *C6* or *C7* are significantly different from the *Side-by-Side* view. However *C5* is ($p = 0.002$), as are the remaining levels.

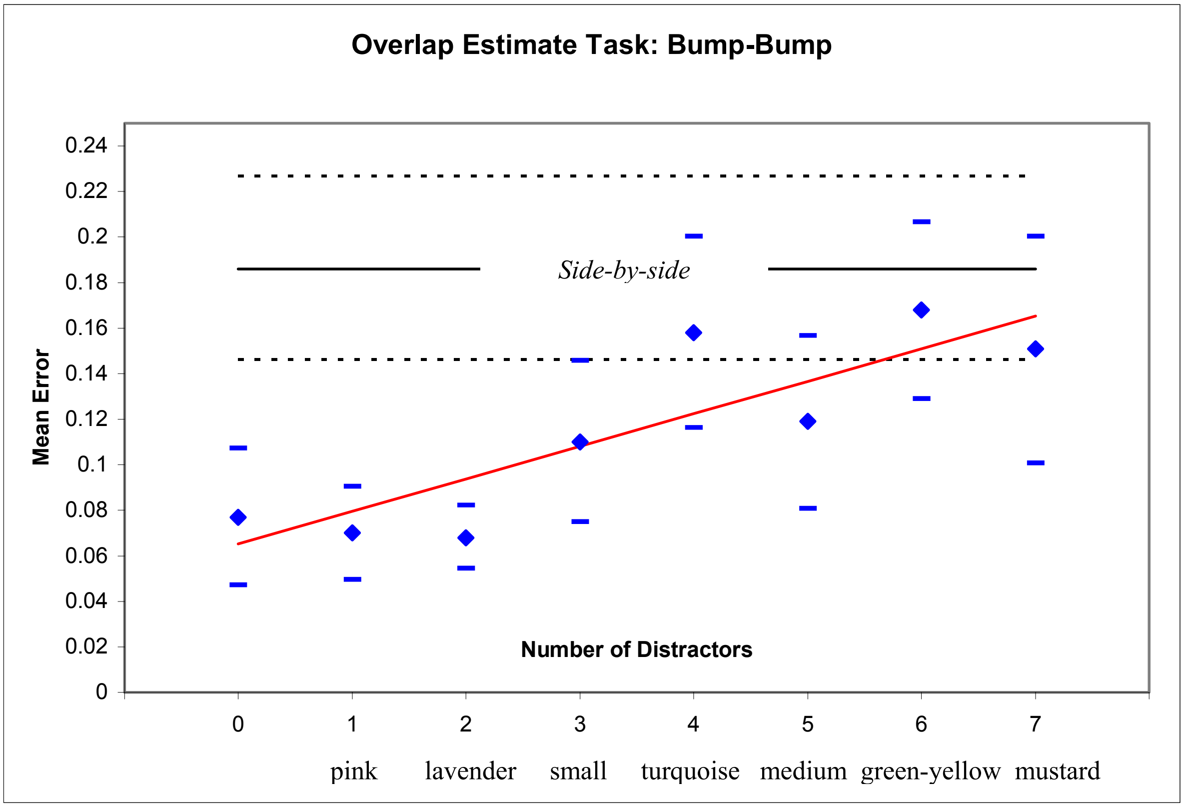

Figure 3.37: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the *Bump-Bump* group. The blue diamonds mark the average error for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the top of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

Only levels *C0* (p = 0.004), *C1* (p = 0.003), and *C2* (p = 0.002) are significantly different from the *Side-by-side* view.

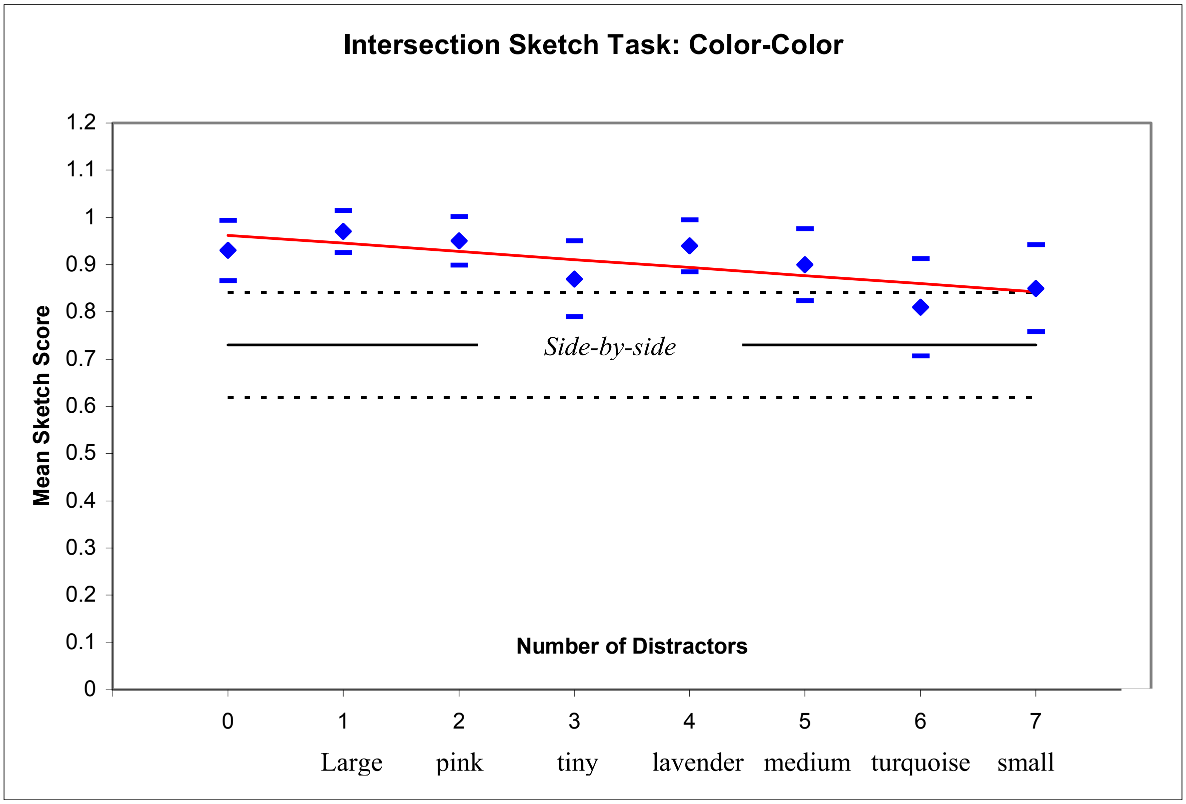

Figure 3.38: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the *Color-Color* group. The blue diamonds mark the average sketch score for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the bottom of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

C4 ($p = 0.027$), *C2* ($p = 0.012$), *C1* ($p = 0.004$), and *C0* ($p = 0.017$) are all significantly different from the *Side-by-side* view.

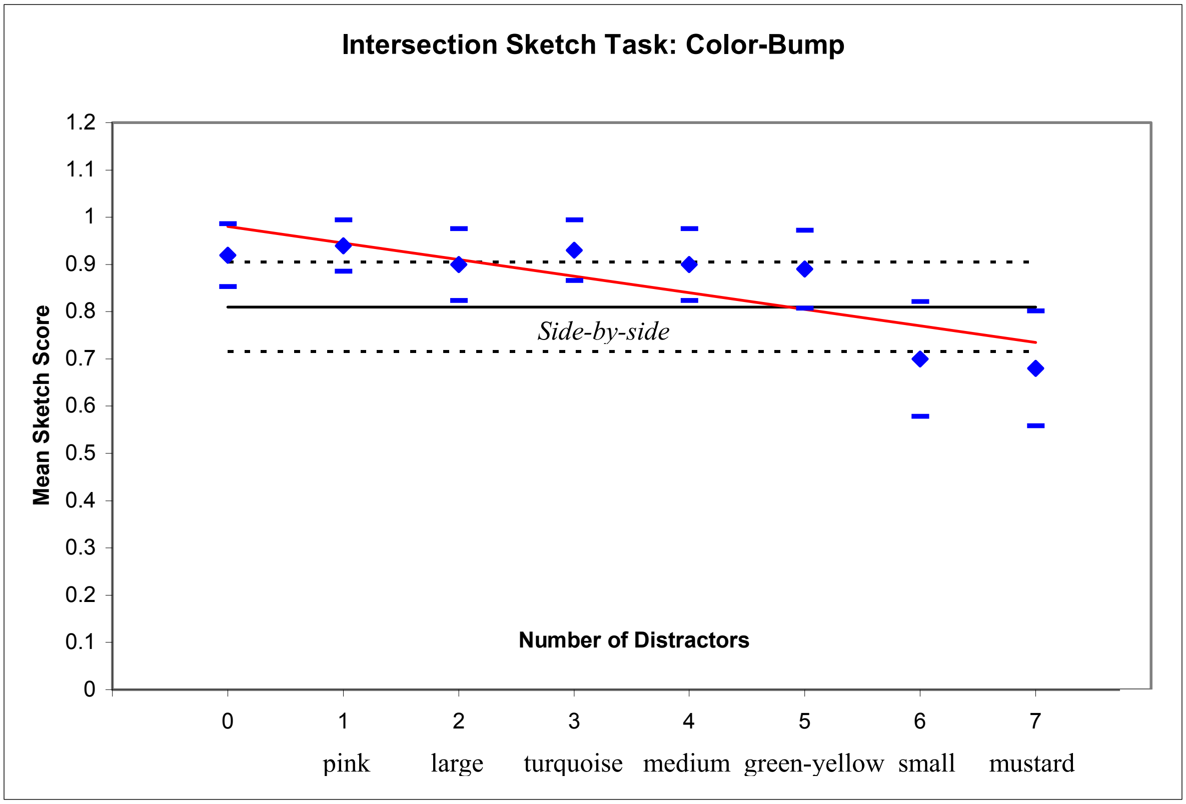

Figure 3.39: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the *Color-Bump* group. The blue diamonds mark the average sketch score for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the bottom of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

Only *C1* is significantly different from the *Side-by-side view* ($p = 0.045$).

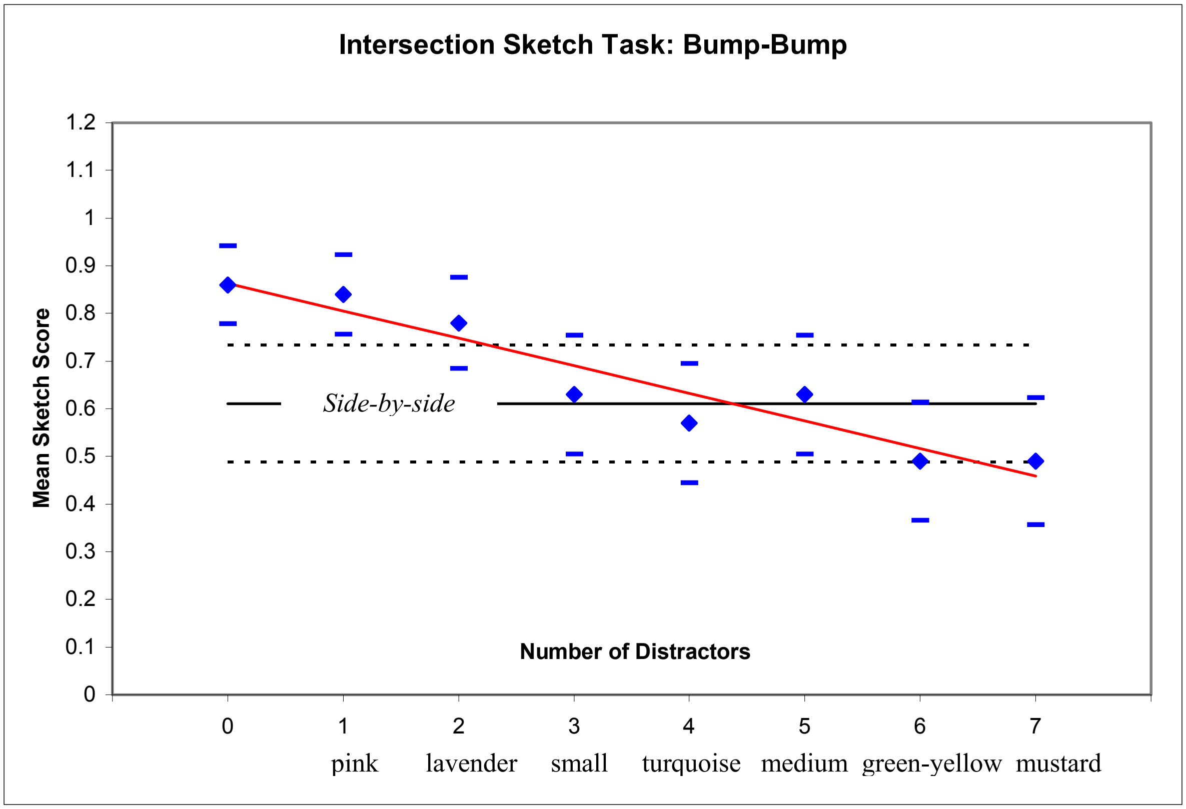

Figure 3.40: Performance for the Overlaid DDS conditions compared to the *Side-by-Side* condition for the Bump-Bump group. The blue diamonds mark the average sketch score for each *Display Condition* level *C0-C7*. The blue horizontal bars above and below the diamonds mark the standard error, which is a measure of the amount of variation in the data. The red line is the linear fit. The solid black horizontal line at the bottom of the graph shows the average performance for the *Side-by-side view*. The black dashed lines above and below it show the standard error for the *Side-by-side view*. Distractor type is listed at the bottom of the graph.

Only *C0* ($p = 0.03$) and *C1* ($p = 0.034$) are significantly different from the *Side-by-side* view.

Question Three: Are DDS alpha-blended layers more visually salient than DDS bump-mapped layers in a multi-layer visualization?

This analysis looks at participant performance on the three outcome measures, error in overlap estimation and sketch shape accuracy, and relates the outcome measures to the number of distractors in the test images, thus only *C0* through *C7* are under consideration. The analysis also considers differences in *Target Display Type* groups and whether performance on any of the three tasks was different for the *Color-Color*, *Color-Bump*, or *Bump-Bump* groups.

A regression F-test is performed to test for significant slope due to the number of distractors in the trial images. SAS 8.2 for Windows statistical software was used in the analysis [SAS, 2002]. The mixed procedure was used with participant ID as a random factor, *Display Condition* (levels *C0-C7*) and *Target Display Type* as fixed factors. Both participant ID and *Target Display Type* were considered nominal variables. Because only the levels of *Display Condition* that represent the number of distractors were included in this analysis, the *Display Condition* variable was analyzed as a continuous, quantitative variable. This model enabled the software to produce parameter estimates for the intercepts and slopes with the number of distractors as the independent variable and the outcome measure as the dependent variable. The analysis looked at all three *Target Display Type* groups together and individually.

The main effects looked at were the effect of *Display Condition*, and *Target Display Type*. In addition the interaction effect was also examined. In the SAS output below the effect of *Display Condition* is labeled DC, the effect of *Target Display Type* is labeled TDT, and the interaction effect is DC*TDT.

Tables 3.22 through 3.24 show the SAS output for the regression test. Plots of the results of are shown in Figures 3.41 through 3.43. For all three measures the number of distractors significantly affected performance: for the overlap estimation and intersection sketch tasks ($p < 0.001$) and ($p = 0.007$) for the single-target sketch task.

Type 3 Tests of Fixed Effects

Effect	Num DF	Den DF	F Value	Pr > F
DC	1	27	21.87	<.0001
TDT	2	1140	2.96	0.0522
DC*TDT	2	1140	1.92	0.1466

Table 3.22: SAS output for the main effects of *Display Condition*, and *Target Display Type*, and the interaction effect of *Display Condition * Target Display Type* for the overlap estimation task. The number of observations, N, is 1200 as only conditions *C0-C7* are included in the analysis.

Type 3 Tests of Fixed Effects

Effect	Num DF	Den DF	F Value	Pr > F
DC	1	27	41.67	<.0001
TDT	2	1140	2.53	0.0801
DC*TDT	2	1140	4.29	0.0140

Table 3.23: SAS output for the main effects of *Display Condition*, and *Target Display Type*, and the interaction effect of *Display Condition * Target Display Type* for the intersection sketch task. The number of observations, N, is 1200 as only conditions *C0-C7* are included in the analysis.

Type 3 Tests of Fixed Effects

Effect	Num DF	Den DF	F Value	Pr > F
DC	1	27	14.80	0.0007
TDT	2	1140	0.54	0.5850
DC*TDT	2	1140	4.13	0.0164

Table 3.24: SAS output for the main effects of *Display Condition*, and *Target Display Type*, and the interaction effect of *Display Condition * Target Display Type* for the single-target sketch task. The number of observations, N, is 1200 as only conditions *C0-C7* are included in the analysis.

Target Display Type was found to be borderline significant – performance was different when both targets were displayed with DDS alpha-blending versus DDS bump-mapping ($p = 0.05$) for the overlap estimation task, and for the intersection sketch task ($p = 0.08$), but was not significant for the single-target sketch task ($p = 0.59$). The interaction term, *Display Condition * Target Display Type* was significant for the overlap estimation task ($p = 0.01$), for the intersection sketch task ($p = 0.01$), and for the single-target sketch task ($p = 0.016$). A significant interaction effect indicates that the slopes for each *Target Display Type* group were significantly different in direction, as is shown in the graphs. The linear models based on the estimated intercepts and slopes from the regression are given below:

Overlap Estimation Task:

When analyzed by *Target Display Type* the effect of the number of distractors is significant: *Color-Color* ($p = 0.02$); *Color-Bump* ($p = 0.0007$); *Bump-Bump* ($p = 0.03$).

$$\text{Error}_{\text{Color-Color}} = 0.048 + 0.004 * \text{Number of Distractors} \quad (3.1)$$

$$\text{Error}_{\text{Color-Bump}} = 0.039 + 0.011 * \text{Number of Distractors} \quad (3.2)$$

$$\text{Error}_{\text{Bump-Bump}} = 0.065 + 0.014 * \text{Number of Distractors} \quad (3.3)$$

Intersection Sketch Task:

When analyzed by *Target Display Type* the effect of the number of distractors is significant: *Color-Color* ($p = 0.0129$); *Color-Bump* ($p = 0.0006$); *Bump-Bump* ($p = 0.0045$).

$$\text{Sketch}_{\text{Color-Color}} = 0.963 - 0.017 * \text{Number of Distractors} \quad (3.4)$$

$$\text{Sketch}_{\text{Color-Bump}} = 0.980 - 0.035 * \text{Number of Distractors} \quad (3.5)$$

$$\text{Sketch}_{\text{Bump-Bump}} = 0.863 - 0.058 * \text{Number of Distractors} \quad (3.6)$$

Single-target Sketch Task:

When analyzed by *Target Display Type* the effect of the number of distractors is significant: *Color-Color* ($p = 0.012$); *Color-Bump* ($p = 0.055$); *Bump-Bump* ($p = 0.0197$).

$$\text{Sketch}_{\text{Color-Color}} = 1.014 - 0.011 * \text{Number of Distractors} \quad (3.7)$$

$$\text{Sketch}_{\text{Color-Bump}} = 1.002 - 0.007 * \text{Number of Distractors} \quad (3.8)$$

$$\text{Sketch}_{\text{Bump-Bump}} = 0.975 - 0.039 * \text{Number of Distractors} \quad (3.9)$$

The *Color-Color* group has the shallowest slope, followed by *Color-Bump* and *Bump-Bump.* The line for the *Color-Bump* group crosses that of the *Color-Color* group between 1 and 2 distractors for both outcome measures. The results for the single-target sketch task show the *Color-Color* and *Color-Bump* groups to have the same slope – this is most likely due to target A being displayed with DDS alpha-blending in both cases.

Magnitude of Effect for Overlap Estimation Task

The difference in performance for zero distractors and seven distractors predicted by the analysis for the *Color-Color* group is 3 percentage points error for the overlap estimation task. For the overlap estimation task an increase of 3 percentage points error is not of practical importance. There is enough variability within a single participant's performance for any given level of distractors, variability much larger than 3 percentage points, that to predict a 3 percentage point change in performance does not hold much meaning on an individual per-participant basis. That the participants as an overall group should see a 3 percentage point increase in error, likewise is of little practical importance.

The difference in performance for zero distractors and seven distractors predicted by the analysis for the *Color-Bump* group is 7 percentage points. At zero distractors the linear fit is 3.9 and at seven distractors it is 11.6, nearly a tripling in percentage point error predicted by the model. The difference in performance for zero distractors and seven distractors predicted by the analysis for the *Bump-Bump* group is also high – 10 percentage points. The practical importance of this result is more meaningful.

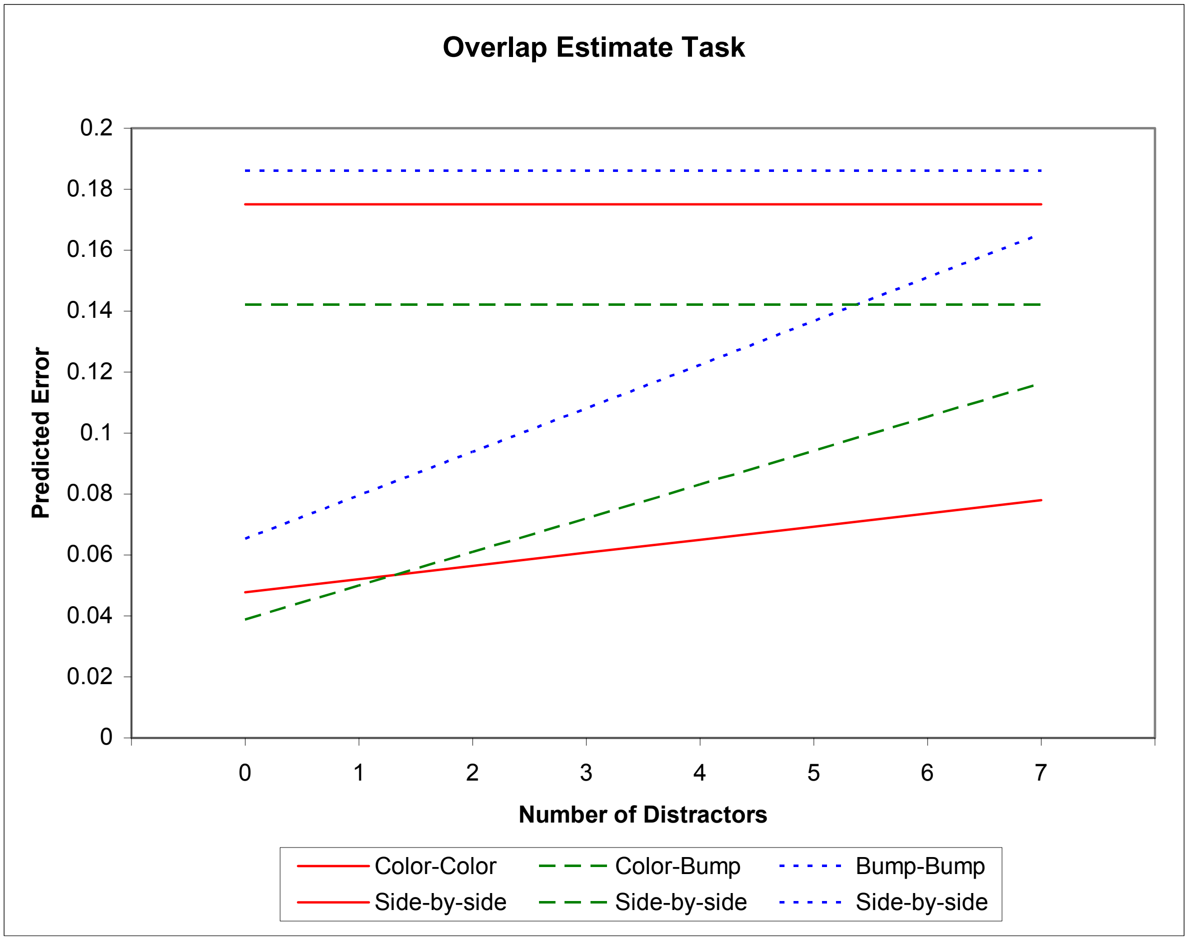

Figure 3.41: Plots of the linear fits for the overlap estimation task. The *Color-Color* group (red line) is least affected by added distractors, and the *Bump-Bump* group is the most affected.

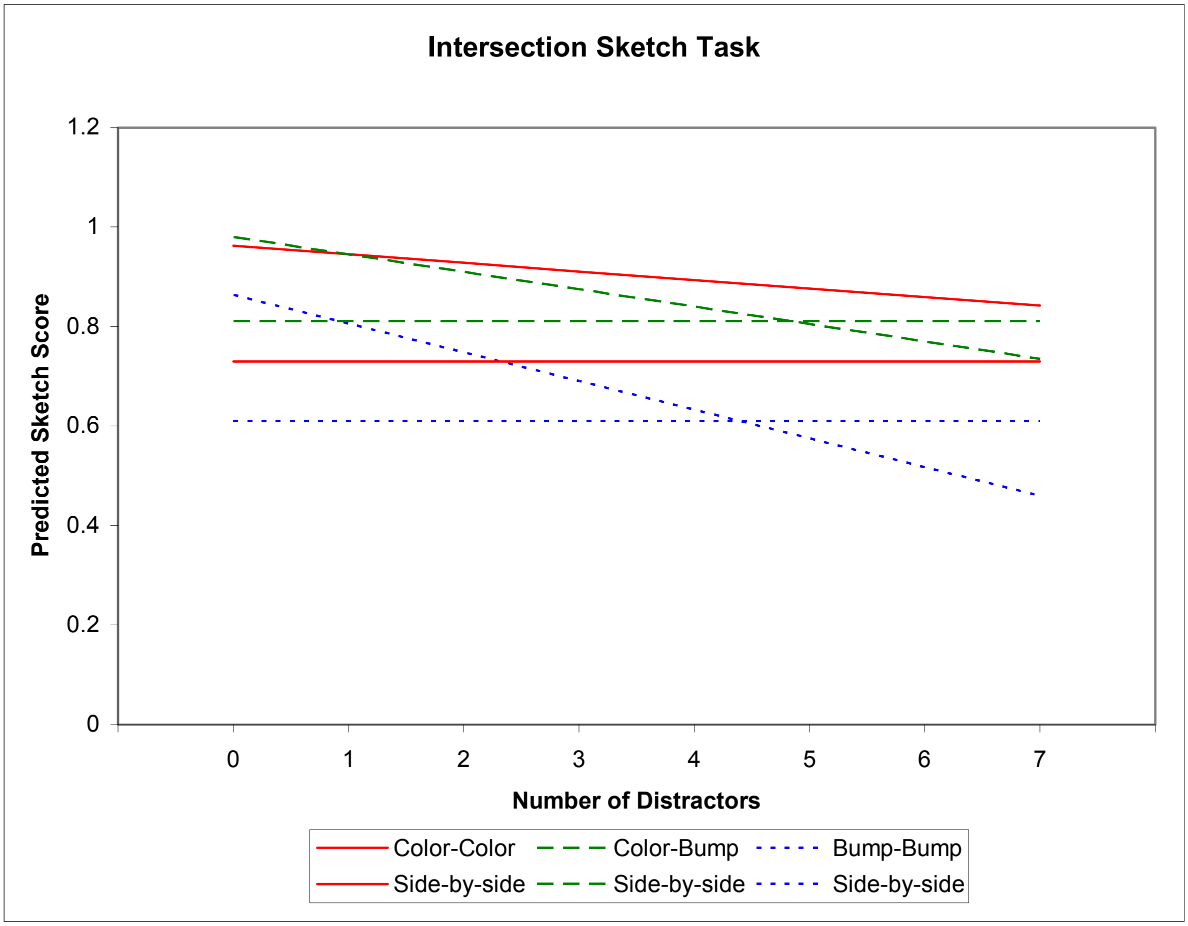

Figure 3.42: Plots of the linear fits for the intersection sketch task. The *Color-Color* group (red line) is least affected by added distractors.

Figure 3.43: Plots of the linear fits for the single-target sketch task. The *Bump-Bump* group is most affected by added distractors, whereas performance for the *Color-Color* and *Color-Bump* groups is nearly identical. This is most likely due to the fact that the single-targets for both conditions were DDS alpha-blended layers.

The results of the analysis show that the *Color-Color* group was much less affected by distractors than the *Color-Bump* and *Bump-Bump* groups. In Figure 3.35 performance for the *Color-Color* group is near-flat and well below that of the *Side-by-side view.* However, in Figure 3.36, performance for the *Color-Bump* group approaches that of the *Side-by-side view* at six distractors, and in Figure 3.37, performance for the *Bump-Bump* group approaches that of the *Side-by-side view* at three distractors.

Several participants in the *Color-Bump* and *Bump-Bump* groups indicated they had confused one of the target bump layers with a distractor bump layer. The two bump sizes closest in appearance are 0.25 and 0.14. In the *Color-Bump* group, target B is displayed with size 0.14 bumps and the interfering shape added in *C6* is displayed with size 0.25 bumps. In the *Bump-Bump* group, target A is displayed with size 0.14 bumps and the interfering shape added in *C3* is displayed with 0.25 bumps. The effect of this confusion can clearly be seen in the plots of subject performance for the *Color-Bump* and *Bump-Bump* groups (Figures 3.36 and 3.37). For the *Color-Bump* group the break in performance occurs at *C6*; for the *Bump-Bump* group the break occurs at *C3*. For the *Color-Color* group there is no equivalent break in performance.

In the regression test both *Target Display Type* and the interaction term *Target Display Type * Display Condition* were either borderline significant or significant. This result indicates that both the intercepts and slopes of the linear fits of performance for the three groups are different. Figure 3.41 shows a clear difference between the linear fits for the *Color-Color* and *Bump-Bump* groups, with the *Color-Color* group being least effected by distractors.

Lessons Learned from Pilot Study

In the main study the number of bump layers is reduced from four to three and the differences in bump sizes between the layers is increased. In the pilot study all DDS bump-mapped layers are assigned the same height value, whereas in the main study larger bumps, which appear to be under the smaller bump layers, are assigned a greater height value. Increasing the height value of the larger bumps in the main study is intended to counteract the observation that layering small bumps on top of large bumps has the effect of decreasing the apparent height of the large bumps.

In the pilot study the alpha-blended layer colors are pastel-like colors, selected because they have close to the same perceived brightness (Figure 3.44a). While running the pilot study, I learned that perceptually isoluminant colors are not easily distinguishable and lose coherence of motion; figure and ground are harder to separate when both are displayed with isoluminant colors; perceptual illusions are not seen as well with isoluminance [Hoffman 1998, Livingstone and Hubel 1988]. Perceptual isoluminance is discussed more in Chapter Two.

Performance for the *Color-Color* session for the pilot study did not show any negative effects due to the perceived isoluminance of the colors used for the alpha-blended targets and distractors. However, for the main study I decided to explore colors that differ in perceived luminance and changed to more saturated colors (Figure 3.44b).

Another difference between the pilot and main studies is the choice of spot sizes for the alpha-blended layers. In the pilot study, of the five DDS alpha-blended layers only three different spots sizes are used, whereas in the main study each DDS alpha-blended layer has a unique spot size. In the pilot study, two pairs of layers share spots sizes – the green-yellow and mustard layers, and the lavender and turquoise layers. Because the green-yellow, mustard, and lavender layers are all targets in the experiment, discrimination of DDS alpha-blended layers for the pilot study is based on hue differences alone.

The visual layer numbers for the targets are also different between studies. In the pilot study, during the *Color-Color* session, the green/yellow and mustard layers are targets, corresponding to color layers four and five. In the *Color-Bump* group the lavender and tiny bump layers are targets, corresponding to color layer two and bump layer four. In the *Bump-Bump* group the tiny and large bumps are targets, corresponding to bump layers four and one.

(A) Pilot study colors. **(B)** Main study colors.

Figure 3.44: Colors used in the pilot and main study.

In the main study, during the *Color-Color* session, the red and green layers are targets, corresponding to color layers two and three. In the *Color-Bump* group the blue and medium bump layers are targets, corresponding to color layer one and bump layer two. In the *Bump-Bump* group the medium and large bumps are targets, corresponding to visual layers two and one. In the main study the color target layers are the three bottom-most layers, whereas in the pilot study the color targets in the *Color-Color* session are the two top-most layers. Increasing the depth of the target layers in the main study is meant to increase the difficulty of the discrimination task.

In the pilot study the independent variable *Target Display Type,* is a between-subjects factor, whereas in the main study, I decided to make *Target Display Type* a within-subjects factor. The 30 participants in the pilot study are divided into three separate groups of ten, each group is assigned to one level of *Target Display Type*, and each participant completes only one session. In the main study each participant completes all three levels of *Target Display Type*. 30 participants are divided into three groups of ten, and each group completes the three sessions (*Color-Color*, *Color-Bump*, *Bump-Bump*) in a different order. The total number of participants is the same for both experiments, but the main study has three times the number of observations. More data points for the same number of participants is one of the major advantages of a within-subjects design over a between-subjects design. Another advantage of the within-subjects design is that each subject can act as his own control, which increases the power of the statistical test for differences between experimental treatments; see Maxwell and Delaney [2000] for a detailed discussion of the benefits of within-subjects and between-subjects designs.