Appendix I: Looking a little closer at discriminating faces and bisected color disks

1.1 Introduction

Throughout the early chapters of this thesis where the dual-task paradigm was used to study attention, we stressed the importance of verifying that the central letter discrimination task was effective at engaging the attentional resource. To make certain that attention was made unavailable to the periphery, we used a bisected red-green discrimination task and demonstrated that in the dual-task condition, performance on this task was dramatically impaired. In contrast, subjects achieved high levels of performance on face-processing tasks in the near-absence of attention. From a computational perspective, this result is very surprising because a red-green bisected disk discrimination task (telling a red-green disk from a green-red one) can be easily accomplished with a few lines of Matlab code. On the other hand, for current machine vision algorithms, face or object recognition is still a very difficult problem.

Given the efficiency of natural stimulus processing and the dismal failure of the visual system at discriminating the red-green colored disks, our question still remains: For natural stimuli, where does the ability to efficiently process them break down? In the earlier chapters of this thesis, we had tried to identify the limits of pre-attentive processing by making the face tasks more difficult (an approach that admittedly is based on rather subjective criteria for difficulty). In this section, we instead try to make the face discrimination task more similar to disk discrimination task—the rationale being that, at some point in the continuum between face discrimination and disk discrimination, face task performance will have to break down in the absence of attention. In what follows, I will present some very preliminary data from two experiments that were conducted in an attempt to address this question.

In the first experiment, we asked whether the symmetry of the disk-discrimination task was in part responsible for the low levels of behavioral performance. In other words, did the task’s difficulty lie in the fact that the visual system is not adept at distinguishing between two objects that are mirror images of each other? To test this hypothesis, three subjects performed a dual-task experiment with the same setup as described in Chapter 2. The only difference in this experiment was that subjects were presented with the 45 degree viewpoints of faces (relative to an axis in the plane of the image; in other words, faces were seen somewhat in profile), and instead of reporting the gender of the faces, they had to report if the face was facing leftwards or rightwards. Thus this task involved discriminating between symmetrical views of the face stimuli.

The normalized performance from this experiment is presented in Figure A1 below. The figure shows that all three subjects achieved high performance on discriminating face viewpoints (mean normalized performance = 99.8 ± 7.8%). The performance of each subject was also not significantly different between single and dual-task conditions (p=.07, p=.09, p=1 for the three subjects). The results of this experiment demonstrate that in contrast to bisected color disks, making symmetry judgments of face stimuli is relatively effortless in the near-absence of attention.

What then is the key factor that differentiates disk processing from face processing? Is it that the red-green disks are defined in the red-green color domain, while the faces are characterized by natural colors and strong luminance variations? To address this question, I will present data from a second experiment in which we carried out two orthogonal manipulations: black-white bisected disk discrimination and a red-green face orientation judgment.

In our second experiment, we attempted to reduce the information of the face stimuli to the domain of red and green colors, of luminance comparable to the red-green disks that have been used in several experiments described in this thesis (Chapters 2, 3, and 4). This was accomplished by converting all color information from the original faces to grayscale values and then converting gray levels such that black was mapped to red and white to green. Examples of the face stimuli and mask we used are shown in Figure A2a. In separate sessions, subjects also discriminated between vertically bisected black-white and white-black disks.

Two subjects performed left-right view-discrimination task with these face stimuli in the dual-task paradigm. They were trained for 2 hours on this task before data was collected. Their normalized behavioral performance on the red-green face and white-black disk discrimination tasks is shown in Figure A2b. In comparison to the data presented in Chapters 2 and 3, subjects did show significant drops in performance on red-green face discrimination in the dual-task condition ($p<.007$). However, it must be noted, that the performance with these faces (mean normalized performance = 85.0 ± 7.0%) was still substantially higher than performance on the black-white disk discrimination task (mean normalized performance = 51.7 ± 1.0%).

It is important to emphasize again that the data presented in this appendix is preliminary, and I do not intend to make any decisive conclusions about it. However, these preliminary results do indicate that low-level differences between the disk and face tasks are not responsible for the amazing discrepancy observed in behavioral performance. Red-green disks are not more difficult for the visual system because they involve vertical orientation judgments between symmetric alternatives (Figure A.1), nor because they are defined by colors targeting primarily the parvocellular pathways, rather than luminance contrasts (optimized for the magnocellular pathways) (Figure A.2). Instead, the data presented here lends support to the notion that the critical difference between the two tasks lies in how “meaningful” these stimuli are to the visual system. Once in the red-green domain, the main difference between faces and disks is that for the faces, the distribution of red and green pixels defines something meaningful for the visual system, but the same is not true for the disk stimuli. At the neuronal level, this might imply that stimuli, which are explicitly represented at the neuronal level, escape stringent attentional requirements.

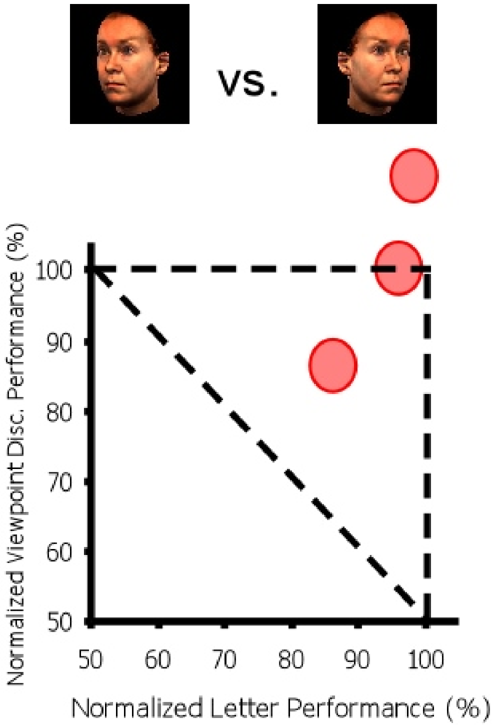

Figure A.1 Performance on view-point discrimination with faces. The data for 3 subjects on this task is shown here. Subjects achieve a high level of performance on this task in the dual-task condition (mean normalized performance = 99.8 ± 7.8%), demonstrating that discriminating between symmetric views of faces in the near-absence of attention is not a difficult problem for the visual system. These subjects had been previously trained in the face-gender discrimination task and performed this task without any prior training.

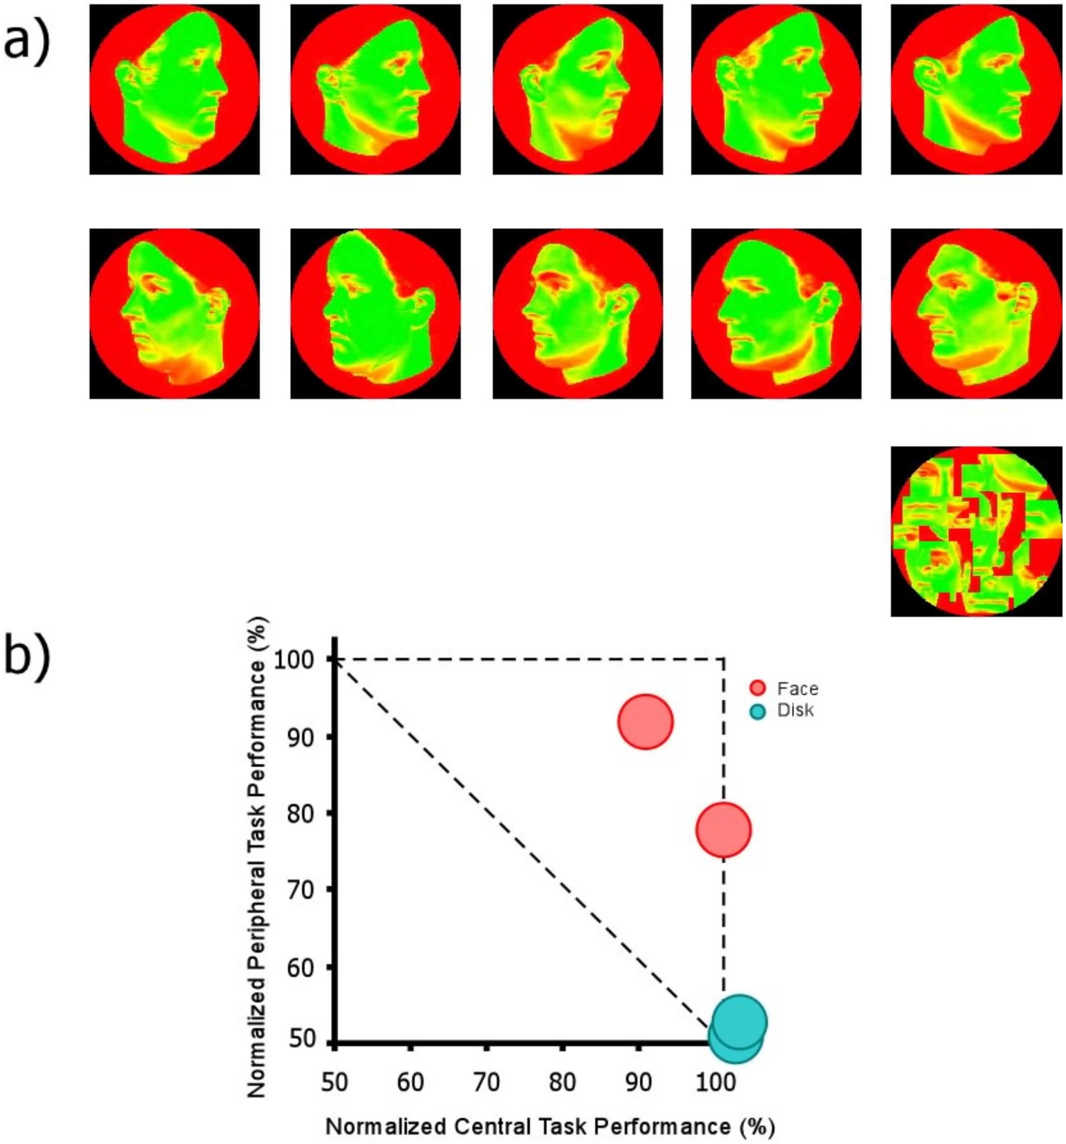

Figure A.2 Performance on red-green face and black-white disk discrimination. a) Examples of the face stimuli and the mask used for this experiment. The original faces (e.g. Figure 2.2) were converted to grayscale values and then gray levels were converted such that black mapped to red, and white to green. b) The data for 2 subjects on this task is shown here. Subjects achieve a high level of performance on the red-green face viewpoint discrimination task in the dual-task condition (mean normalized performance = 85.0 ± 7.0%). However, performance on a vertically bisected black-white disk discrimination task is at chance (mean normalized performance = 51.7 ± 1.0%). These two subjects, who had been previously trained on the face-gender discrimination task, were trained on the two tasks shown here for 2 hours. The data shown here was collected over 5 hours (15 blocks of each condition for each task).