RUNNING HEAD: Children's sensitivity to intermediate complexity fragments

Children are sensitive to mutual information in intermediate-complexity face and non-
face features

Benjamin Balas, Amanda Auen, Alyson Saville, Jamie Schmidt & Assaf Harel

Author Note:

Benjamin Balas, Department of Psychology, Center for Visual and Cognitive Neuroscience, North Dakota State University. This research was supported by NSF grant BCS-1348627 awarded to BB.

Correspondence regarding this article should be addressed to Benjamin Balas, Department of Psychology, North Dakota State University, Fargo, ND 58102

Contact: Benjamin.balas@ndsu.edu

Word Count = 8000 words.

Keywords: Face detection; Visual Development; Object Recognition.

Abstract

Uncovering when children learn to use specific visual information for recognizing object categories is essential for understanding how experience shapes recognition. Research on the development of face recognition has focused on children's use of low-level information (e.g. orientation sub-bands), or on children's use of high-level information, namely, configural or holistic information. Do children also use intermediate complexity features for categorizing faces and objects, and if so, at what age? Intermediate-complexity features bridge the gap between low- and high-level processing: they have computational benefits for object detection and segmentation, and are known to drive neural responses in the ventral visual system. Here, we asked when children develop sensitivity to diagnostic category information in intermediate-complexity features. We presented children (5-10 years old) and adults with image fragments of faces (Experiment 1) and cars (Experiment 2) varying in their mutual information, which quantifies a fragment's diagnosticity of a specific category. Our goal was to determine whether children were sensitive to the amount of mutual information in these fragments, and if their information usage is different from adults'. We found that despite better overall categorization performance in adults, all children were sensitive to fragment diagnosticity in both categories, suggesting that intermediate representations of appearance are established early in childhood. Moreover, children's usage of mutual information was not limited to face fragments, suggesting the extracting intermediate-complexity features is a process that is not specific only to faces. We discuss the

RUNNING HEAD: Children's sensitivity to intermediate complexity fragments

implications of our findings for developmental theories of face and object recognition.

Introduction

Studying how early during development specific types of visual information are used for recognizing different object categories is an important way of understanding how learning and experience affect high-level visual representations (Nelson, 2001). The category of faces particularly stands out during development, both in terms of its specialized neural systems and dedicated modes of processing (Chung & Thomson, 1995; McKone, Crookes & Kanwisher, 2009). It thus follows that examining the time point during development in which face representations emerge can provide insight into how neural mechanisms for processing different objects become differentiated from one another as observers gain experience with their visual environment (Nelson, 2001).

One way that face recognition is “special” compared to other object categories is the visual information that face recognition is tuned to, with some visual features (by which we refer to any measurement one might use to characterize appearance) more useful for face recognition than others. Therefore, understanding how mechanisms for face recognition depend on different visual features is an important way to characterize how face recognition works, as it describes its unique properties in terms of specific computations. To describe face recognition in terms of information biases (i.e. differential reliance on distinct visual features) requires that we stipulate what class or classes of visual features we will use to do so. Because faces are complex, natural stimuli, they are highly multi-dimensional and as a result we have many options available to us along these lines. In large part, the literature regarding how information use for face recognition changes developmentally has focused on two levels of representation, which correspond to what we will refer to as *high-level* and *low-level* visual features. By *high-*

level visual features, we refer broadly to presumed mechanisms for face detection and recognition that use global features, which are typically described as reflecting “holistic” or “configural” face processing. The use of such features is almost always studied using specific behavioral effects (e.g. the inversion effect, the composite face effect) as proxies, largely because the field still lacks explicit computational models of what specific features contribute to this mode of processing (Maurer, Le Grand, & Mondloch, 2002). By *low-level* features, we refer to visual features that typically involve local processing of the image, and reflect basic properties of the image that can be measured at early stages of the visual system. For example, cells in primary visual cortex are known to be selective for orientation and spatial frequency (Valois & Valois, 1990), which immediately suggests the possibility of studying recognition performance as it depends on different sub-bands of information along these two dimensions. While there are measurable information biases favoring intermediate spatial frequencies (Nasanen, 1999) and horizontal orientation energy (Dakin & Watt, 2009) for face recognition, one theoretical limitation of characterizing face recognition in terms of low-level information use is the universality of the information: Local measurements of oriented contrast are not category-specific, (though the extent to which different sub-bands are weighted differently may be), and thus are often not sufficient to explain face-specific effects with face stimuli, which do not vary by these properties (e.g. Mooney faces or smiley faces).

In the present work, we adopt a third approach to characterizing the features used for face recognition, one that relies on applying tools for identifying meaningful intermediate-level features for face recognition that are on the one hand category-specific (unlike low-level features), and at the same time are easy to describe computationally

(unlike high-level features). By an intermediate-level feature, we refer to a description of the object that relies neither on purely local nor purely global information. Rather, intermediate-level features are partial representations of an object: They only provide information about some portion of the object. There are several reasons to consider the contribution that such features make to face recognition. First, while it might seem that information regarding the global configuration of an object's parts might require global features, overlapping partial configurations of smaller-scale features can be sufficient (Peterson, 2006). Indeed, Peterson & Hector (1996) have shown that object memory effects based on figure-ground assignment are explained more effectively by assuming that portions of contours and not the entire contours are being measured. Observers thus appear to use such intermediate-level features to encode simple objects. Second, for more complex objects, intermediate features offer distinct computational advantages. Ullman & Sali (2000) demonstrated that intermediate-level fragments were more useful for classifying faces and cars than global object views were, suggesting that these kinds of visual features are particularly important for categorization. Moreover, Ullman (2001) also demonstrated that similar features were also useful for segmenting objects from cluttered backgrounds, which is a notoriously difficult computational task. Intermediate-level fragments thus are superior to global features in some settings and it is computationally feasible to use them for real-world recognition tasks. Finally, physiological studies of the ventral visual system have demonstrated that some cortical areas are comprised of cells that appear to be sensitive to intermediate-level visual features. Area TE, for example, appears to be sensitive to proto-objects and portions of object contours (Tanaka, 1996), suggesting a code for objects in terms of intermediate-

level descriptors. Intermediate-level visual features are thus of clear computational value, appear to be measured in the ventral visual stream, and make contributions to behavioral visual recognition tasks. Notably, in spite of their putative role in object and face recognition, to date there have been effectively no attempts to determine how early during development intermediate complexity features emerge as a diagnostic source of information.

Compared to how very little is currently known about children's use of intermediate-level features for face recognition, there is a large literature describing how the use of high-level and low-level features for face recognition changes during childhood. With regard to developmental changes in how children use high-level visual features for face recognition, there are multiple examples of how performance in recognition tasks designed to reveal the use of different types of configural processing (Maurer, Le Grand & Mondloch, 2002) changes over the course of childhood. For example, changes in holistic processing are reflected in the extent to which the composite face effect (or CFE, Young et al., 1987) and the part-whole effect (Tanaka & Farah, 1993) are evident at different ages. The CFE appears to be measurable in 4-year-old children (de Heering et al., 2007), but only reaches adult levels around the age of 6 (Mondloch et al., 2007). The part-whole effect is also measurable in early childhood (3-4 years of age, Pellicano & Rhodes, 2003), as is the well-known face inversion effect (Yin, 1969; Schwarzer, 2000; Sangrigoli & de Schonen, 2004), suggesting that some aspects of configural processing may mature faster than others. One challenging aspect of all of these studies, however, is that though we can use these various tasks to identify when children begin to behave in an adult-like fashion with regard to each behavioral effect

RUNNING HEAD: Children's sensitivity to intermediate complexity fragments

(e.g. poorer performance for inverted faces), it remains a largely open question what specific visual features or computations lead to the behavioral effects that serve as proxies for putative “configural” mechanisms. That is, the field has largely accepted that the CFE (for example) reveals whether or not an observer is using holistic processing (see Richler, Palmieri & Gauthier, 2012), but we still don't really know what holistic processing is at a computational level. It probably involves obligatory measurement of face appearance at a large, possibly global, spatial scale, but the actual information used during holistic processing of a face remains underspecified, particularly in a developmental context.

Related studies describing how children develop sensitivity to the spacing between facial features are somewhat more computationally tractable, and thus may offer some direct insight into how high-level visual features for face recognition change during childhood. Inter-eye spacing is hard for even 8-year-old children to measure relative to adults, for example (Mondloch et al., 2004), and processing such metric relationships between features doesn't even seem to be adult-like at later ages approaching pre-adolescence (Mondloch, Le Grand & Maurer, 2002). These results (and others like them that use “2nd-order” face configuration to study development) link recognition performance to a specific candidate image property that can be clearly described and quantified (e.g. the distance between the two eyes), and so perhaps point more directly towards which high-level visual information emerges during development to process faces. The ecological validity of these specific features appears to be low, however, based on computational studies that demonstrate that geometric features like these are unreliable in natural images (Bruneli & Poggio, 1993; Gaspar, Bennett & Sekuler, 2007;

Taschereau-Dumouchel et al., 2010). Again, it seems to be the case that there is evidence that children's use of high-level visual information changes during childhood, but exactly what that information is and how early children start to use it is difficult to pin down.

If we consider low-level visual features that contribute to face recognition, there is also clear evidence that children use this information differently as they get older. Two nice examples of research along these lines are developmental studies of how spatial frequency use and the use of orientation sub-bands for face recognition each change during childhood. In the former case, adults tend to perform better at face recognition tasks when mid-range spatial frequencies are available (Costen, Parker, & Craw, 1996; Nasanen, 1999; Ruiz-Soler & Beltran, 2006). In the latter case, horizontal orientation energy tends to promote better recognition performance (Dakin & Watt, 2009; Goffaux & Dakin, 2010). Developmentally, children show different tuning for spatial frequency sub-bands than adults, suggesting that their use of SF channels is still changing. Leonard, Karmiloff-Smith, & Johnson (2010) demonstrated that while 9-10 year-old children showed adult-like biases for mid-band spatial frequencies that only applied to upright (not inverted) faces, younger children did not have the same preference for intermediate spatial frequencies. Similarly, Gao & Maurer (2011) showed that the tuning of identity and emotion recognition also showed some evidence of change between 10 and 14 years of age, again suggesting that the reliance of face processing mechanisms on specific spatial frequency channels is still changing late in childhood. Considering the information bias for horizontal orientation energy, an emerging body of work suggests that the contribution of these low-level features also changes during childhood. Infants as young as 3 months old show evidence of some preferential use of horizontal orientation energy

in face images (De Heering et al., 2016), but studies with children 5 years of age and older suggest that there is continual refinement of this tuning for horizontal orientation energy between the ages of approximately 6 and 10 years of age (Balas, Saville & Schmidt, 2015; Balas et al., 2015; Goffaux, Poncin & Schiltz, 2015). Indeed, studies with an extended developmental trajectory further indicate that there may be lifespan changes in the tuning of face recognition mechanisms to low-level features (Goffaux, Poncin & Schiltz, 2015; Obermeyer et al., 2012), which suggests the intriguing hypothesis that vocabularies for face recognition may remain plastic (or at least susceptible to change) throughout the life of an individual observer. Overall, these studies demonstrate that there is developmental change in how low-level features are used for face recognition, and provide a way to describe this development in terms of a specific vocabulary of features that we know are computed at early stages of cortical visual processing. While there have been several studies designed to reveal possible links between these low-level features and high-level behavioral effects (Goffaux & Rossion, 2006; Goffaux & Dakin, 2010), the key issue is that we don't know how simple, low-level features are integrated into more complex representations of face appearance. Without more ideas about how to build a bridge between low-level and high-level visual information, a comprehensive description of how information use changes developmentally in the context of face recognition will likely remain elusive.

Our goal, therefore, in the current study is to begin to build a bridge between low- and high-level features by investigating the emergence of sensitivity to visual features of *intermediate complexity*. Specifically, we use a computational measure first proposed by Ullman, Vidal-Naquet, & Sali (2002) to classify image fragments of face and other object

categories as a function of the *mutual information* they provide about their category (see Figure 1 for example of face fragments). Mutual information can be used to index the diagnosticity of an image fragment for basic-level categorization (i.e. determining that an image contains a face) by measuring the probability of how likely we are to find that pattern of intensity values in other face images and how likely we are to find that pattern in non-face images. A fragment will have high mutual information (or diagnosticity) if it is both highly likely to be found in a face image and unlikely to be found in non-face images, thus balancing the need for specificity and generalization. Mutual information values vary a great deal across different face fragments (Figure 1), indicating that some portions of a face are better indicators of category than others. Notably, fragments with high mutual information tend to be of intermediate complexity: They are neither very large or highly detailed fragments, nor are they very small and generic. These may therefore hint at the spatial scale and content of mid-level representations of face appearance, rather than low-level or high-level visual features. Moreover, they do not correspond to well-defined part descriptions, and thus avoid the existing theoretical tension between “parts” and “wholes” (for a discussion see Harel et al., 2007). In terms of behavioral responses, adult observers are faster and more accurate to categorize face fragments the higher mutual information they carry (Harel et al., 2007; 2011) indicating that this computational index is a perceptually meaningful characterization of how diagnostic a fragment is. Indeed, despite the fact that mutual information in itself is indifferent to categorical content, mutual information utilization differs for face and non-face fragments, with different RT functions for face and object (car) fragments. Face fragments displayed a monotonic relation between the mutual information level and RT,

whereas in car fragments this relation was in the shape of a step function, dividing the fragments into “low” or “high” in terms of their informativeness¹. The different RT curves may reflect the utilization of different kinds of diagnostic features, serving as unique “psychometric signatures” for different object categories. In terms of their neural correlates, face fragments with high mutual information also elicit more robust face-sensitive cortical responses than low mutual information face fragments (ERP: Harel et al., 2007; fMRI: Lerner et al., 2008), which further suggests that mutual information can serve as useful means to characterize diagnostic face information at an intermediate level of representation.

Our goal in the current study was therefore to examine how early during development children show sensitivity to mutual information level in face fragments and non-face object fragments ((Experiment 1 and Experiment 2, respectively). To do so we asked children between the ages of 5 and 10 years of age to carry out a simple detection task (a go/no-go judgment for each target category) using fragments that varied according to the MI they provided regarding target category. This detection task was used in a previous study examining the contribution of different orientation sub-bands to face recognition (Balas et al., 2015), as it is easier and more intuitive for young children than the standard 2AFC categorization task (see methods below). We predicted that sensitivity to mutual information level might be evident already during childhood, reflecting ongoing tuning of mid-level representations for facial appearance during this age range. Further, we hypothesized that the developmental trajectory associated with mutual information level in face fragments might differ from that of non-face fragments,

¹ The differences in RT curves were also found for car experts (Harel et al., 2010), strengthening the conclusion that these patterns may serve as a category-specific behavioral signature.

mirroring the emergence of other face-selective behavioral signatures during development (Balas et al., 2015; 2016; Gao & Maurer, 2011; Goffaux, Poncin & Schilz, 2015; Mondloch et al., 2004).

Experiment 1

In our first task, we examined children's and adult's sensitivity to mutual information level in a face detection task. In particular, we investigated whether or not children were sensitive to the higher diagnosticity of face fragments with a higher mutual information level, and if the degree of that sensitivity changed with age during middle childhood (5-10 years).

Methods

Participants. We recruited a total of 61 participants from the Fargo-Moorhead community to take part in this study. Since we did not have clear a-priori predictions regarding the specific time point in which sensitivity to mutual information will emerge, we divided our child participants into two groups based on age: 5-7 year-old children (N=19, Mean age=6;2) and 8-10 year-old children (N=22, Mean age=8;10 female). Adult participants (N=20) were recruited from the NDSU Undergraduate Psychology study pool. Children and their caregivers received compensation for their participation, while adult participants received course credit. Informed consent was obtained from adult participants and the legal guardians of child participants. Children 7 years and older also provided written assent to participate. All participants (or their caregivers) reported either normal or corrected-to-normal vision.

Stimuli. We presented our participants with gray-scale images of face and non-face (car) fragments. Both classes of fragments were previously used in behavioral and electrophysiological (Harel et al., 2007; 2011) studies designed to reveal how adult face processing is impacted by the extent to which image fragments are diagnostic of face/non-face categories. The most important feature of this stimulus set is that for each fragment, the *mutual information* provided by that fragment for category membership has been calculated. Briefly, this measure is a way to quantify how informative a given fragment is regarding the presence of a category (e.g. “face”). The mutual information associated with each fragment is calculated by determining how likely one is to find the fragment in images containing a face and how likely one is to find the fragment in images that do not contain faces. A high value implies that the fragment is very diagnostic: It appears frequently in images containing a face, and appears rarely in images containing non-faces. A detailed description of how mutual information is calculated can be found in (Ullman, Vidal-Naquet, & Sali, 2002), but for our purposes, the critical property of the stimuli is that we know how this value varies across our stimulus set. Critically, all of the fragments in our stimulus set are intermediate complexity fragments, but they vary in their mutual information level, and thus, their diagnosticity.

We selected 96 non-face fragments depicting portions of cars and 96 face fragments. All car fragments had a “medium” level of diagnosticity for their category. Face fragments were selected to include 32 “low” mutual information fragments, 32 “medium” mutual information fragments, and 32 “high” mutual information fragments. Within each of these three levels, fragment size varied, and the regions of the face

depicted within each fragment also varied. Examples of these fragments can be found in Figure 1.

Figure 1 – Examples of face fragments with varying levels of mutual information. We have included fragments of different sizes from all three levels of mutual information to demonstrate the variability within our stimulus set.

Procedure. We asked participants to complete a “go/no-go” judgment using the full set of face and non-face fragments, similar to prior work examining the contribution of different orientation sub-bands to face recognition (Balas et al., 2015). We chose this task primarily because it is easier for our youngest participants to understand and execute than a 2AFC categorization task. Young children don’t need to remember which button

goes with which response, for example, nor do they have to provide a response to stimuli that they have difficulty identifying as a member of either category.

Stimuli were presented in a pseudo-randomized order for 3 seconds each, (see Balas et al., 2015) and participants were instructed to press a large red button as quickly as possible if and only if the fragment depicted part of a face. Otherwise, participants were asked to withhold making a response and wait until the image disappeared. We included a 1-second inter-stimulus interval between trials and presented each image once.

Participants completed the task seated approximately 60cm away from a 1200x800 MacBook Pro laptop. All stimulus display and response collection routines were written using the Psychtoolbox extensions for Matlab (Brainard, 1997; Pelli, 1997). The task was administered in a dark, sound-attenuated room, and most participants completed the task in approximately 20 minutes.

Results

We examined three aspects of participants' performance: Their *hit rate* for correctly labeling face fragments as faces, their *response latency* for correctly labeled face stimuli, and their *false alarm rate* for incorrectly labeling non-face fragments as faces. For both the hit rate and response latency, we submitted the values obtained from each participant to a 3x3 repeated-measures ANOVA with participant age (5-7 years old, 8-10 years old, adults) as a between-subjects factor and face diagnosticity (low, medium, or high) as a within-subjects factor. To analyze false alarm rates, we carried out a 3x1 ANOVA with participant age as a between-subjects factor. We conducted separate analyses of the hit rates for each mutual information level and the false alarm rate because the false alarm

rate for each observer is shared across all three hit rates, making it not especially useful for converting these values to d' or other signal detection theory descriptors of performance. Specifically, differences in d' and response bias across conditions within one participant can only reflect hit rate differences given that there is only one source of false alarms.

Hit Rates. Our analysis of participant hit rates revealed a main effect of diagnosticity ($F(2,116)=99.66$, $p<0.001$, partial $\eta^2=0.63$), and a main effect age group ($F(2,58)=9.39$, $p<0.001$, partial $\eta^2=0.98$). The interaction between diagnosticity and age group did not reach significance ($F(4,116)=1.18$, $p=0.325$, partial $\eta^2=0.049$). The main effect of diagnosticity was the result of significant differences between all three levels of mutual information as revealed by Bonferroni-corrected post-hoc comparisons. The main effect of age was the result of a significant difference between 5-7 year-old performance ($M=0.68$, 95% CI = [0.62 0.74]) and both 8-10 year-old performance ($M=0.85$, 95% CI=[0.80 0.91]) and adult performance ($M=0.82$, 95% CI = [0.76 0.87]). No other pairwise comparisons between age groups reached significance.

Figure 2 – Average hit rate as a function of diagnosticity (mutual information) and age (5-7 years old, 8-10 years old, and adults) for face detection. Error bars indicate +/- 1 s.e.m.

Response Latency. Our analysis of response latency to correctly-detected face images revealed significant main effects of diagnosticity ($F(2,116)=54.23$, $p<0.001$, partial $\eta^2=0.48$) and age ($F(2,58)=43.55$, $p<0.001$, partial $\eta^2=0.60$). The main effect of diagnosticity was the result of significant differences between the response latency to images with “High” mutual information ($M=0.77$, 95% CI=[0.74 0.81]) and images with both “Low” mutual information ($M=0.95$, 95% CI=[0.89 1.00]) and “Medium” mutual information ($M=0.92$, 95% CI=[0.88 0.96]). The comparison between “Low” and “Medium” mutual information did not reach significance. The main effect of age was the result of significant differences in performance between all three age groups.

Besides these main effects, we also observed a significant quadratic contrast between face diagnosticity and age ($F(2,58)=3.62$, $p=0.033$, partial $\eta^2=0.11$). To examine the nature of this interaction, we carried out follow-up tests to determine how adjacent levels of mutual information affected performance within each group. Specifically, within each age category, we determined whether “Low” and “Medium” response latency differed significantly, and whether “Medium” and “High” response latency differed significantly using one-tailed, paired-samples t-tests (a directional test is appropriate here due to the expectation that higher mutual information should lead to faster response latencies). We found that while all three age groups exhibited a significant difference between “Medium” and “High” latencies (5-7 year olds: $p<0.001$; 8-10 year-olds, $p<0.001$; Adults, $p<0.001$), only adults exhibited a significant difference between “Low” and “Medium” response latencies (5-7 year-olds, $p=0.71$; 8-10 year-olds, $p=0.55$; Adults, $p=0.043$). We note that this latter result for Adult observers does not survive a Bonferroni correction for multiple comparisons, but nonetheless that it provides some insight into what drives this relatively small interaction effect.

Figure 3 – Average response latency for correctly labeled faces as a function of diagnosticity (mutual information) and age (5-7 years old, 8-10 years old, adults). Error bars indicate ± 1 s.e.m.

False Alarms. Finally, our analysis of false alarm rates to non-faces revealed a significant effect of age group ($F(2,58)=4.27$, $p=0.019$, partial $\eta^2=0.13$). Pairwise comparisons revealed that this main effect was the result of a lower false alarm rate in adults ($M=0.079$, 95% CI=[0.03 0.12]) relative to 5-7 year-olds ($M=0.18$, 95% CI=[0.13 0.23]). The false alarm rate observed in 8-10 year olds ($M=0.16$, 95% CI=[0.11 0.20]) did not significantly differ from either age group.

Discussion

The goal of Experiment 1 was to examine how early during middle childhood children show sensitivity to the information contained in intermediate-level visual features of

faces. Starting with our results from adult observers, we replicate and extend previous results obtained using mutual information as a criterion for defining face fragments that are more or less category-diagnostic (Harel et al., 2007; 2011). Specifically, face fragments with higher mutual information are more reliably categorized as faces in our go/no-go task by adults, and they are also categorized as such with a shorter response latency. This further demonstrates the utility of mutual information as a useful computational tool for identifying diagnostic face fragments.

Developmentally, our data demonstrates that both young children and older children are sensitive to mutual information contained in face fragments. In all age groups, Hit rates are larger for high mutual information face fragments, and response latencies shorter. This suggests that sensitivity to mid-level face features may be relatively mature early in childhood. In comparison, low-level biases for face recognition develop relatively slowly during childhood (Balas, Schmidt & Saville., 2015; Goffaux et al., 2015). We did not find an interaction between age group and mutual information in our hit rate data, which might suggest a general cognitive improvement during childhood in the extraction of intermediate-level diagnostic information, that is, an improvement that perhaps does not reflect any face-specific mechanisms (Crookes & McKone, 2009). Note, however, that we cannot argue that these effects do or do not reflect face-specific aspects of visual development without examining how children's sensitivity to mutual information level changes for a non-face category. To address this question, we conducted Experiment 2, in which we examined the extent to which children of the same age range as in Experiment 1 are sensitive to mutual information level contained in car fragments of varying levels of mutual information.

Experiment 2

In our second study, we examined the development of sensitivity to mutual information level in a non-face category (cars). Using the same methods as Experiment 1, we investigated how performance in car detection varied with age and mutual information level when face fragments were used as the distractor category.

Methods

Participants. We recruited a sample of 56 participants from the Fargo-Moorhead community to take part in Experiment 2. As in Experiment 1, the sample was comprised of three groups: 5-7 year-old children (N=20, Mean age=6;4), 8-10 year-old children (N=20, mean age=9;1) and adults (N=16). Recruitment and consent procedures were identical to those reported in Experiment 1, and none of the participants included in this sample had taken part in the first experiment.

Stimuli. The stimulus set used in this experiment was drawn from the same larger set of images used to select stimuli for Experiment 1. The key difference between this stimulus set and the previously described images is that in this case, we selected car fragments that varied according to the mutual information level provided by each image regarding the target category and chose face fragments from a fixed, intermediate, level of mutual information. Specifically, we selected 96 car fragments and 96 non-car (face) fragments for use in this task. The set of car fragments was comprised of 32 “low mutual information” car fragments, 32 “medium” fragments, and 32 “high” fragments. The set of

face fragments was comprised of 96 face fragments all with a “medium” level of mutual information. Within each category, image size varied substantially.

Procedure. We administered the same go/no-go task described in Experiment 1 to participants in this task. The only critical difference between tasks was that in this case, participants were instructed to press the response button only in response to car fragments and withhold responding if a face fragment was presented. Otherwise, all stimulus presentation, response collection, and display parameters were identical to those described in Experiment 1.

Results

As in Experiment 1, we examined participants' hit rate for responding correctly to car fragments, their response latency for correct responses to car fragments, and their false alarm rate in response to face fragments. We submitted the hit rate and response latency values to a 3x3 mixed-design ANOVA with age (5-7 years old, 8-10 years old, adults) as a between-subjects factor and mutual information level (low, medium, and high) as a within-subjects factor. We analyzed false alarm rates using a one-way ANOVA with participant age as a between-subjects factor.

Hit Rates. This analysis revealed only a main effect of mutual information level ($F(2,106)=121.5$, $p<0.001$, partial $\eta^2=0.70$). Neither the main effect of age ($F(2,53)=2.23$, $p=0.12$, partial $\eta^2=0.08$), nor the interaction between these factors ($F<1$) reached significance. The main effect of mutual information level was the result of lower hit rates

in response to low information fragments relative to medium (Mean difference = 0.28, 95% CI of the difference between means = [0.22 0.34] and high information fragments (Mean difference = 0.30, 95% CI of the difference between means = [0.23 0.36]). The average hit rates as a function of age and car diagnosticity are presented in Figure 4.

Figure 4 – Average hit rate as a function of diagnosticity (mutual information level) and age for car detection. Error bars indicate +/- 1 s.e.m.

Response Latencies. This analysis revealed main effects of age ($F(2,53)=5.19$, $p=0.009$, partial $\eta^2=0.94$) and mutual information level ($F(2,106)=32.47$, $p<0.001$, partial $\eta^2=0.38$). The interaction between these factors did not reach significance ($F<1$). The main effect of age group was the result of a significant difference between the performance of children in our youngest age group and adults (Mean difference = 0.25, 95% CI of the difference between means = [0.06 0.45]. Neither the difference between

adults and older children, nor the difference between younger children and older children reached significance. The main effect of mutual information level was the result of slower response latencies in response to low information fragments relative to medium (Mean difference = 0.21, 95% CI of the difference between means = [0.11 0.30] and high information fragments (Mean difference = 0.22, 95% CI of the difference between means = [0.14 0.31]). The average response latencies as a function of age and car diagnosticity are presented in Figure 5.

Figure 5 – Average response latency for correctly labeled cars as a function of diagnosticity (mutual information) and age (5-7 years old, 8-10 years old, adults). Error bars indicate +/- 1 s.e.m.

False Alarms. Finally, our analysis of false alarm rates (response errors made to face fragments) revealed a marginally significant effect of participant age ($F(2,53)=2.71$,

$p=0.076$, partial $\eta^2=0.093$). While the average proportion of false alarms did appear to decrease (5-7 year olds: $M=0.17$, 8-10 year-olds: $M=0.10$, adults, $M=0.049$) with age, no pairwise comparisons between age groups reached significance.

Discussion

Experiment 2 was conducted to assess the extent to which children in the middle childhood range are sensitive to the mutual information contained in a non-face category (cars) and not only to the face fragments. We found that participants across all age groups were sensitive to the diagnosticity of car fragments, as indicated by a main effect of MI level in both hit rate and response latency. Thus, children as early as the age of five years old can distinguish more diagnostic object features than less diagnostic object features. Note that while this may suggest at first that the sensitivity to mutual information is not face-specific, the patterns of information usage in the two experiments are not identical, suggesting a more complex pattern of results. We discuss this below.

General Discussion

The present study sought to determine how early during development intermediate complexity features emerge as a diagnostic source of information for the categorization of faces and objects. We tested this question in two separate experiments, using the objective and quantifiable measure of mutual information. In both of our experiments, we have found evidence that children as young as 5-7 years old are sensitive to category diagnosticity in image fragments of complex objects (faces and cars). Like adults, children in both of our target age groups showed better accuracy and faster response

latencies when categorizing fragments with high mutual information levels. This basic effect applied to both face and car categories, so this observation could be interpreted in terms of general visual development rather than any face-specific developmental trajectory that governs the recruitment of intermediate complexity features to support basic-level recognition. The present findings provide a first indication that children as early as five years of age are capable of extracting diagnostic features for categorization at an intermediate complexity level and not only at low- or high-level levels of visual representation.

Notably, in addition to the general effects of age and feature informativeness, we also observed critical differences in how participant age and mutual information level contributed to performance in our face and non-face detection tasks. In particular, the interaction we observed between age and mutual information level in the response latency data from Experiment 1 was not present in Experiment 2. This outcome could be interpreted to mean that for face recognition, there is ongoing development during middle childhood that leads to graded sensitivity to mutual information level in adulthood as opposed to the more non-linear pattern of results we observed in our younger age groups. To put this more simply, we suggest that middle childhood could be a period during which children are still learning how to use less diagnostic information for recognition, which means that the difference between “low” and “medium” amounts of mutual information is not meaningful to them. Adults, on the other hand, who have acquired the ability to make use of weakly diagnostic facial features, differ in that the increase in mutual information level from low to medium confers additional visual information that their visual system is capable of using for modest gains in recognition ability. This

outcome does resemble recent results describing the development of low-level orientation biases for emotion recognition in faces (Balas et al., 2015) and bodies (Balas et al., 2016), in that young children in both of these studies demonstrated particularly poor abilities with sub-optimal visual features. That is, while young children exhibited adult-like biases for orientation features in these emotion recognition tasks in terms of the direction of the effects, their absolute performance when the preferred features were not available was far worse than adults (in some cases, not different from chance). The current finding of lack of an RT difference between low and medium levels of mutual information in children but not in adults may reflect a similar developmental trajectory whereby optimal features are identified early in development and contribute to face representations quickly, but less optimal features are uniformly 'bad' at early stages of development and only gradually become useful. Further, the fact that we observed this effect for faces and not for cars may mean that this is a unique feature of how face recognition develops, either because children's exposure to faces relative to cars differs across middle childhood (Sugden, Mohamed-Ali & Moulson, 2014), or because there are differences in how cortical areas that support face and object recognition mature during this time period (Grill-Spector, Golarai, & Gabrieli, 2008). We emphasize however, that the critical interaction supporting this account is associated with a rather small effect size, and is only evident in response latencies and not in accuracy of performance. On balance, we therefore suggest that both the hit rate data and the response latency data suggest that information biases favoring higher amounts of mutual information in face fragments are adult-like at the earliest ages we tested, differing primarily in overall performance across levels with increasing age. And while the potential for face-specific development is

interesting to consider in light of other recent results, the current data does not make a strong case for this conclusion at this moment.

An important avenue of research to consider to help clarify the extent to which the current findings may reflect face-specific effects is the measurement of neural sensitivity to mid-level visual features during childhood. Adult observers' ERP responses are sensitive to mutual information in face and car fragments (Harel et al., 2007; 2010), but little is known about how either low-level or mid-level information biases for face recognition develop. An investigation of how ERP components like the P100 and the N170/N290 respond to varying levels of mutual information in face and non-face fragments would be an important step towards linking the computational and behavioral evidence for fragments of intermediate complexity as a mid-level representation of face appearance to real neural outcomes. Recent results examining the sensitivity of these components to the low-level visual information in faces (horizontal vs. vertical orientation energy) has revealed that face-sensitive ERP components change their tuning to orientation sub-bands during middle childhood in a category-selective way (Balas et al., 2017), supporting behavioral work indicating similar outcomes for response latencies (Balas et al., 2015). Applying these same methods to intermediate complexity features in faces and non-faces would thus provide an important complementary look at the development of visual recognition during this age range.

In summary, the present findings provide important and novel data demonstrating sensitivity to mid-level features for face and object recognition emerges as early as five years old of age. Complementing prior results examining children's use of visual information at low-level and high-level stages of face procession, the current study

demonstrates near-adult utilization of information contained in intermediate complexity level features. Our findings are significant not only in establishing the feasibility of intermediate-level features as a 'third way' for object representations, but also in providing explicit, objective, and quantitative means to describe what it means for a mid-level feature to be diagnostic. This will allow future investigations of how children develop 'vocabularies' of intermediate complexity visual features to form different object categories, while avoiding previous dichotomies between low and high accounts of face and object categorization. Deeper insights into how the content of face representations changes during development will depend on identifying and testing quantitative models like the current one, and our results offer a compelling example of how such computational models can be adapted for use with developmental populations.

References

- Balas, B., Auen, A., Saville, A. & Schmidt, J. (2016) Body emotion recognition disproportionately depends on vertical orientations during childhood, *International Journal of Behavioral Development*.
- Balas, B., Schmidt, J. & Saville, A. (2015) A face detection bias for horizontal orientations develops during middle childhood, *Frontiers in Developmental Psychology*, 6:772.
- Balas, B., Huynh, C., Saville, A. & Schmidt, J. (2015) Orientation biases for facial emotion recognition in early childhood and adulthood. *Journal of Experimental Child Psychology*, 140, 71-83.
- Balas, B., van Lamsweerde, A. E., Saville, A. & Schmidt, J. (2017) School-age children's neural sensitivity to horizontal orientation energy in faces. *Developmental Psychobiology*, 59, 899-909.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.
- Brunelli, R., & Poggio, T. (1993) Face recognition: features versus templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15, 1042-1052.

Chung, M-S. & Thomson, D.M. (1995) Development of face recognition. *British Journal of Psychology*, 86, 55-87.

Costen, N. P., Parker, D. M., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics*, 58, 602–612.

Crookes, K. & McKone (2009) Early maturity of face recognition: No childhood development of holistic processing, novel face encoding, or face-space. *Cognition*, 111, 219-247.

Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9(4):2, 1–10.

De Heering, A., Houthuys, S., & Rossion, B. (2007) Holistic face processing is mature at 4 years of age: Evidence from the composite-face effect. *Journal of Experiment Child Psychology*, 96, 57-70.

De Heering, A., Goffaux, V., Dollion, N., Godard, Or., Durand, K. & Baudouin, J-Y. (2016) Three-month-old infants' sensitivity to horizontal information within faces. *Developmental Psychobiology*, 58, 536-542.

De Valois, R. L. & De Valois, K. K. (1990) *Spatial Vision* (Oxford Univ. Press, New York).

Gao, X. & Maurer, D. (2011) A comparison of spatial frequency tuning for the recognition of facial identity and facial expressions in adults and children. *Vision Research*, 51, 508-519.doi: 10.1016/j.visres.2011.01.011

Goffaux, V. & Dakin, S.C. (2010) Horizontal information drives the behavioral signatures of face processing, *Frontiers in Psychology*, 1:143.

Gaspar, C., Bennett, P., & Sekuler, A. (2007) The ecological utility of interfeature distances for human face discrimination. *Journal of Vision*, 7(9), 500a.

Goffaux, V. & Rossion, B. (2006) Faces are “Spatial” – holistic face perception is supported by low spatial frequencies. *JEP:HPP*, 32, 1023-1039.

Goffaux, V., Poncin, A. & Schiltz, C. (2015) Selectivity of Face Perception to Horizontal Information over Lifespan (From 6 to 74 Year Old) *PLoS:ONE*, 10(9), e0138812.

Gosselin, F. & Schyns, P.G. (2001) Bubbles: A technique to reveal the use of information in recognition tasks. *Vision Research*, 41, 2261-271.

Grill-Spector, K., Golarai, G., & Gabrieli, J. (2008). Developmental neuroimaging of the human ventral visual cortex. *Trends in cognitive sciences*, 12(4), 152-162.

Harel, A., Ullman, S., Epshtein, B. & Bentin, S. (2007) Mutual information of image fragments predicts categorization in humans: Electrophysiological and behavioral evidence, *Vision Research*, 47, 2010-2020.

Harel, A., Ullman, S., Harari, D., & Bentin, S. (2011). Basic-level categorization of intermediate complexity fragments reveals top-down effects of expertise in visual perception. *Journal of Vision*, 11(8):18, 1–13.

Kelly, D.J., Quinn, P.C., Slater, A.M., Lee, K., Ge., L. & Pascalis, O. (2007) The other-race effect develops during infancy: Evidence of perceptual narrowing. *Psychological Science*, 18, 1084-1089.

Langner, O., Becker, E.S. & Rick, M. (2009) Bubbles Reveal Differential Use of Facial Information With Low Spatial Frequencies. *Psychological Science*, 20, 666-670.

Leonard, H.C., Karmiloff-Smith, A. & Johnson, M.H. (2010) The development of spatial frequency biases in face recognition, *Journal of Experimental Child Psychology*, 106, 193-207.

Lerner, Y., Epshtein, B., Ullman, S. & Malach, R. (2008) Class information predicts activation by object fragments in human object areas. *Journal of Cognitive Neuroscience*, 20, 1189-1206.

Macchi Cassia, V., Bulf, H., Quadrelli, E. & Prioetti, V. (2014) Age-related face processing bias in infancy: Evidence of perceptual narrowing for adult faces.

Developmental Psychobiology, 56, 238-248.

Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6, 255–260.

Mckone, E., Crookes, K., & Kanwisher, N. (2009). The cognitive and neural development of face recognition in humans. In M. S. Gazzaniga, E. Bizzi, L. M. Chalupa, S. T. Grafton, T. F. Heatherton, C. Koch, . . . B. A. Wandell (Eds.), *The Cognitive Neurosciences* (pp. 467-482). Cambridge, MA, USA.

Mondloch, C.J., Le Grand, R., & Maurer, D. (2002) Configural face processing develops more slowly than featural face processing. *Perception*, 31, 553-566.

Mondloch, C.J., Dobson, K.S., Parsons, J. & Maurer, D. (2004) Why 8-year-olds cannot tell the difference between Steve Martin and Paul Newman: Factors contributing to the slow development of sensitivity to the spacing of facial features, *Journal of Experimental Child Psychology*, 89, 159-181.

Mondloch, C.J., Pathman, T., Le Grand, R., Maurer, D., & de Schonen, S. (2007).

The composite face effect in six-year-old children: Evidence of adultlike holistic face processing. *Visual Cognition*, 15, 564-577.

Nasanen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, 39(23), 3824-3833.

Nelson, C.A. (2001) The development and neural bases of face recognition. *Infant and Child Development*, 10, 3-18.

Obermeyer, S., Kolling, T., Schaich, A. & Knopf, M. (2012) Differences between Old and Young Adults' Ability to Recognize Human Faces Underlie Processing of Horizontal Information. *Frontiers in Aging Neuroscience*, 4:3.

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-42.

Peterson, M.A. (1996) Overlapping partial configurations in object memory. In M.A. Peterson and G. Rhodes (Eds.), *Perception of Faces, Objects, and Scenes*. Oxford University Press, NY.

Pellicano, E., & Rhodes, G. (2003). Holistic processing of faces in preschool children and adults. *Psychological Science*, 14, 618–622.

Richler, J.J., Palmeri, T.J. & Gauthier, I. (2012) Meanings, mechanisms, and measures of holistic processing. *Frontiers in Psychology: Perception Science*, 3:553.

Ruiz-Soler, M., & Beltran, F. S. (2006). Face perception: An integrative review of the role of spatial frequencies. *Psychological Research*, 70, 273–292.

Sangrigoli, S., & de Schonen, S. (2004). Effect of visual experience on face processing: A developmental study of inversion and nonnative effects. *Developmental Science*, 7, 74–87.

Schwarzer, G. (2000) Development of Face Processing: The Effect of Face Inversion, *Child Development*, 71, 391-401.

Sugden, N. A., Mohamed-Ali, M. I., & Moulson, M. C. (2014). I spy with my little eye: Typical, daily exposure to faces documented from a first-person infant perspective. *Developmental Psychobiology*, 56, 249–261.

Tanaka, J.W. and Farah, M.J. (1993) Parts and wholes in face recognition. *Q. J. Exp. Psychol. Hum. Exp. Psychol.*, 46A, 225–245.

Tanaka, K. (1996) Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-139.

- Taschereau-Dumouchel, V., Rossion, B., Schyns, P.G. & Gosselin, F. (2010) Interattribute Distances do not Represent the Identity of Real World Faces. *Frontiers in Psychology*, 1:159.
- Ullman, S. (2007) Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Science*, 11, 58-64.
- Ullman, S. & Sali, E. (2000) Object classification using fragment-based representation. In S.-W. Lee, H.H. Bulthoff, & T. Poggio (Eds.), *Biologically motivated computer vision*, 1811, 73-87. Springer-Verlag, Berlin.
- Ullman, S., Vidal-Naquet, M. & Sali, E (2002). Visual features of intermediate complexity and their use in classification. *Nature Neuroscience* 5, 682–687.
- Yin, R.K. (1969) Looking at upside-down faces, *J. Exp. Psychol.*, 81, 141-145.
- Young, A.W. et al. (1987) Configurational information in face perception. *Perception*, 16, 747–759.