

**Ensemble coding of faces occurs in children and develops dissociably from
coding of individual faces**

Gillian Rhodes¹, Markus Neumann¹, Louise Ewing^{1,2}, Samantha Bank¹, Ainsley
Read¹, Laura M. Engfors¹, Rachel Emiechel¹ & Romina Palermo¹

¹ARC Centre of Excellence in Cognition and its Disorders, School of Psychology,
University of Western Australia

²School of Psychology, University of East Anglia

Word Count: 3980 (excluding refs)

Please address correspondence to:

Professor Gillian Rhodes

School of Psychology

University of Western Australia

35 Stirling Highway

Crawley, WA 6009

AUSTRALIA

phone: +61-8-6488-3251

fax: +61-8-6488-1006

email: gillian.rhodes@uwa.edu.au

Running head: development of ensemble coding

Keywords: face processing, ensemble coding, set averaging, development

Research Highlights

- Ensemble coding of face identity is present by 6-8 years of age
- This ensemble coding of social groups increases from 6-18 years
- Its development is dissociable from improvements in individual face coding
- Children may use ensemble coding to access gist information about social groups

Abstract

Ensemble coding allows adults to access useful information about average properties of groups, sometimes even in the absence of detailed representations of individual group members. This form of coding may emerge early in development with initial reports of ensemble coding for simple properties (size, numerosity) in young children and even infants. Here we demonstrate that ensemble coding of faces, which provides information about average properties of social groups, is already present in 6-8 year old children. This access to average information increases with age from 6 to 18 years and its development is dissociable from age-related improvements in the coding of individual face identities. This dissociation provides the first direct evidence that distinct processes underlie ensemble and individual coding of face identity, evidence that has been lacking from adult studies. More generally, our results add to the emerging evidence for impressively mature sensitivity to statistical properties of the visual environment in children. They indicate that children have access to gist information about social groups that may facilitate adaptive social behaviour.

Ensemble coding of faces occurs in children and develops dissociably from coding of individual faces

We subjectively experience a rich and detailed visual world despite very limited visual attention and memory capacity. Some aspects of this subjective experience are illusory, as illustrated by phenomena such as change blindness, where substantial changes can go unnoticed (Noë, Pessoa, & Thompson, 2000). A potential contributor to this rich subjective experience may be the ability to access summary information about group properties that can bypass limitations on processing of individual group members (for reviews see Alvarez, 2011; Whitney, Haberman, & Sweeny, 2014). For example, information about average properties of a set of items can be available, even when participants have little explicit memory for the items themselves (e.g., Ariely, 2001; Haberman & Whitney, 2007). This “ensemble coding” may provide useful information about the visual environment and contribute to our subjective experience of a rich visual world, in the absence of a detailed representation.

Ensemble coding has been reported for many simple visual features, such as size, orientation and direction of motion (for a review, see Whitney et al., 2014). It also occurs for more complex stimuli, such as faces, with ensemble coding reported for identity, expression, gender, attractiveness and gaze direction of groups of faces (e.g., de Fockert & Wolfenstein, 2009; Haberman & Whitney, 2007; Kramer, Ritchie, & Burton, 2015; Neumann, Schweinberger, & Burton, 2013; Sweeny & Whitney, 2014; Walker & Vul, 2014). In some cases, the ensemble information is available in the absence of information about individual faces (e.g., Haberman & Whitney, 2007) whereas in other cases both types of information are available (e.g., Kramer et al., 2015; Neumann et al., 2013). In both instances, a hallmark of ensemble coding is the

erroneous “recognition” of faces with average group properties that have not been seen. This ensemble coding of faces appears to tap higher-level coding mechanisms, as it does not correlate with ensemble coding of low-level features (Haberman, Brady, & Alvarez, 2015) and is generally reduced for inverted faces (Haberman & Whitney, 2009; Sweeny & Whitney, 2014).

An intriguing aspect of ensemble coding is that it appears able to bypass, to some extent, capacity limitations that apply to the coding of individual items, such as the working memory capacity of 3-5 items (Luck & Vogel, 1997). It is facilitated when attention is distributed globally rather than focused on individual items (Chong & Treisman, 2005) and can even occur for items in the impaired hemifield of neglect patients (Pavlovskaya, Soroker, Bonneh, & Hochstein, 2015). The precise mechanisms underlying ensemble coding are unknown, but it may reflect the pooling of partial information across items (for a review see Whitney et al., 2014). Such pooling can potentially yield representations of average properties that are more accurate than the representations of individual properties, by averaging across uncorrelated (e.g., random) noise associated with individual representations (Alvarez, 2011). Another suggestion is that average group information is available in the initial feedforward of information to high-level cortical representations, with access to detailed properties of individual group members requiring additional feedback and focal attention to lower-level level cortical representations (Hochstein, Pavlovskaya, Bonneh, & Soroker, 2015).

Little is known about the development of ensemble coding. Do young children also have access to ensemble information about average group properties or is this ability relatively slow to mature? Children have substantial processing capacity limitations that could limit the detailed processing of individual group members

(Cowan et al., 2005; Riggs, McTaggart, Simpson, & Freeman, 2006; Simmering, 2012; for a review see Sweeny, Wurnitsch, Gopnik, & Whitney, 2015), but these should be no impediment to ensemble coding, which is supposed to largely bypass such limitations. Consistent with this suggestion, ensemble coding has been reported for very simple properties in young children (circle size, Sweeny et al., 2015) and even infants (numerosity, Zosh, Halberda, & Feigenson, 2011). Therefore, children may have access to useful ensemble information.

Early maturing ensemble coding would allow early access to any functional benefits of such coding, such as the rapid deployment of appropriate behavioural responses towards groups. Sensitivity to collective characteristics has also been tentatively linked with ‘the grand illusion’, i.e., (erroneous) perception that we hold a detailed representation of our visual world (e.g., Whitney et al., 2014). Thus, ensemble impressions could also contribute to our stable experience of visual completeness across the lifespan, ensuring that the world outside of our highly constrained attentional focus does not instantly disappear from awareness (for discussion see Sweeny et al., 2015).

Here we focus on the ensemble coding of face identity. Identity is a complex visual property that is important for social interaction, and for which specialized neural and computational machinery is available (Kanwisher, 2000; Kanwisher & Barton, 2011; Rhodes, 2011). We ask whether children show ensemble coding for this important face property. Ensemble coding of face identity could provide useful gist information about social groups, not only about identity per se, but also about other important attributes, such as trustworthiness, dominance and competence, that may be associated with identity and are inferred from facial appearance (Willis & Todorov, 2006). Ensemble coding for face identity has been reported in a small

group of older children and adolescents (9-14 year-olds) (Rhodes, Neumann, Ewing, & Palermo, 2014). Here we ask whether younger children show ensemble coding of face identity.

We also ask whether ensemble coding of face identity increases with age, and if so, whether this development is dissociable from improvements in the coding of individual face identity. A dissociation would constitute novel and important evidence that (at least partly) distinct processes underlie the two forms of coding, evidence that is currently lacking in the case of face identity coding. Ensemble information is sometimes available in the absence of detailed information about group members, a pattern that is highly suggestive of dissociable processes (Haberman & Whitney, 2007; Whitney et al., 2014). However, this pattern has not been found for face identity (Neumann et al., 2013; Rhodes et al., 2014). Thus, if we find a developmental dissociation between the two forms of coding it would provide important new evidence for distinct processes.

We measured ensemble coding and individual coding of face identity in participants ranging in age from 6 to 18 years, using a child-friendly immediate memory task. The task was presented as a game in which participants see teams (sets of four faces) and must decide whether or not test faces (individual faces or set averages) were present in those teams and therefore eligible for a prize.

To summarize, we aim to better understand the visual coding capabilities of children and adolescents (hereafter referred to as children), asking whether young children demonstrate an adult-like ability to abstract ensemble information from groups of faces, whether this capacity increases with age, and whether any such developmental increase is dissociable from improvements in the coding of individual identities. The results will enrich our understanding of children's visual experience

and provide a novel test of influential claims that ensemble and individual face identity coding rely on distinct processes.

Method

Participants

We recruited 105 participants (66 male), ranging in age from 6 to 18 years. Participants aged 6 to 15 years ($N = 52$, 27 male) were recruited from local schools and community groups and participants aged 17 to 18 years ($N = 53$, 39 male) were recruited from the University of Western Australia.¹ The lower age bound was determined by pilot testing, which indicated that our task was too difficult for younger children. We included young adults because face identity recognition continues to improve into adulthood (e.g., Germine, Duchaine, & Nakayama, 2011; Susilo, Germine, & Duchaine, 2013). The distribution of ages within this 6-18 year range reflected our desire to test a substantial number of children aged 8 years or younger, to determine whether ensemble coding of identity occurs in children this young (the only previous study of ensemble coding of face identity in children tested a small sample of 9-14 year-olds, Rhodes et al., 2014), as well as the availability of participants during our sampling period.

¹ The data from seven of the participants was also reported as part of the typically-developing control group ($n = 9$) in Rhodes et al (2014).

Task

We measured ensemble and individual coding using a child-friendly immediate memory task adapted by Rhodes et al (2014) from Neumann et al.'s (2013) original adult task (see Rhodes et al., 2014 for full details). Participants are instructed to remember faces in "winning teams" that competed at a "wacky sports carnival". Their job is to identify whether a subsequent face is part of the previous team, in order to determine whether they are eligible for a prize. In the first half of the experiment, participants must match the exact face (same-image condition), and in the second half they must match the person across different images (different-image condition). One example was given at the beginning of each condition, before the practice trials, using face stimuli with which the at children were familiar ("The Wiggles").

The basic trial structure is as follows: participants saw a fixation cross for 500 ms, followed by a study set of four faces (one randomly assigned to each screen quadrant), the winning team, for 2000 ms, followed by a single test face (in the centre of the screen) for 500 ms, followed by a prompt asking whether that test face (same-image condition) or that person (different-image condition) was a member of the team they had just viewed (see Figure 1). Participants used labelled keys on a keyboard to respond either "Yes" or "No" and initiated the next trial by pressing the space bar. The test faces were either set averages or individuals, from either the studied set (matching condition) or another set (mismatching condition).




Figure 1. A sample trial showing the trial structure (top) and the test conditions (below). Test faces could be either set averages or exemplars. In each case, the test face could be (exemplar), or be made from (set average), the study images (same-image condition) or different images of those identities (different-image condition). Test faces could be taken either from the study set shown on the trial (match condition) or from another study set (non-match condition).

Trials were blocked by image condition, with the same-image condition completed before the different-image condition. In the same-image individual condition, identical images were used as study and test faces. In the different-image individual condition, test faces depicted different images of the identities shown in the study sets. In the same-image average condition, the test faces were averages constructed from the images used in the study sets. In the different-image average condition, they were averages constructed from different images of the identities shown in the study sets. The different-image condition provides a more demanding test of identity recognition, by reducing the availability of pictorial cues. The participants were explicitly informed that the task in the different-image condition was more difficult, and were warned that, "It can be tricky, because it's not the same exact photo". The task was constructed using 8 different images of each of 10 unfamiliar male identities that were collected from various Internet sources. As a result of sourcing images from the Internet, images varied in lighting, viewpoint, expression, etc. We created 10 face sets ("teams" in the cover story) from different combinations of the 10 identities. Each set consisted of four different identities, and each identity occurred in four different sets. All were presented in grey scale and displayed inside oval masks that covered most of the hair.

In each image condition (same-image, different-image), each set was presented four times, once followed by a test face that represented: i) an identity from the previous set (matching individual), ii) a different identity from another set (mismatching individual), iii) the average of the previous set (matching set average), or iv) the average of a different set which had no overlap in identity with the previously seen set (mismatching set average) (40 trials). A different random trial order was used for each participant. Breaks were provided after 20 trials, and each

block started with instruction screens and four practice trials, using different characters from the cartoon show, “The Simpsons”. During the practice trials, participants received feedback in the form of a golden star that appeared on the monitor for correct responses. No feedback was provided during the experiment proper. The task took approximately 15 minutes.

General Procedure

Children and adolescents completed the task on a 15” Macbook Pro laptop computers running E-Prime 2.0 Professional (Psychology Software Tools, Inc., Sharpsburg, PA, USA). They were tested individually (either at home, school, after-school care facility or UWA psychology lab) and the experimenter monitored engagement and provided verbal encouragement. They received a certificate and sticker, or a certificate, movie ticket and small toy or chocolate, depending on the length of their test session. The 17 and 18 year-olds completed the task on standard PCs running E-Prime 2.0 Professional. They were tested individually and received course credit or payment. All participants completed the task as part of a larger battery.² The study was approved by the Human Research Ethics Committee at the University of Western Australia and all participants, along with the parents of children, provided written consent.

²The 17 and 18 year-olds also repeated the task, to provide additional data for a different adult individual differences project. These additional data were excluded from our analyses, to match task length across ages. There were no significant differences in unbiased recognition scores between the first and second halves in any condition, t 's < 1.66 , p 's $> .104$, and the results do not change if the full datasets are used for the 17 and 18 year-olds. Here we conservatively report results using only the length-matched datasets.

Results

Table 1 shows the mean proportion of “present” responses for set average and individual test faces on same-image and different-image trials, in match and mismatch conditions averaged across all ages and also for 6-8 year-olds separately (because we examine whether ensemble coding is present in this youngest group). These scores were used to calculate unbiased recognition scores, by subtracting the proportion of “present” responses on mismatch trials from the proportion of “present” responses on match trials (Figure 2). For individual test faces, recognition scores index accuracy of individual coding. For set average test faces, recognition scores index strength of ensemble coding of identity (actually incorrect responses, because averages were never shown). Positive scores indicate that participants used identity information from the study sets to make their decisions. Zero represents chance performance and negative scores represent below-chance performance (ie., fewer “present” responses on match than mismatch trials). Recognition scores and age showed significant deviations from normality on K-S tests, but skew and kurtosis were acceptable for parametric statistics ($-0.64 < \text{skew} < .25$, $-0.38 < \text{kurtosis} < 0.47$) (Stuart & Kendall, 1958). We present results separately for the same-image and different-image conditions, because floor effects are likely to contaminate performance in the different-image condition (Rhodes et al., 2014).

Table 1. Mean proportion of “present” responses for individual and set average test faces on same-image and different-image trials, in match and mismatch conditions, for all ages (6-18 years) and for 6-8 year-olds.

<i>Test Face</i>	<i>Image Condition</i>	<i>Match/Mismatch</i>	<i>Mean</i>	<i>SE</i>	<i>95% CI</i>	
					<i>LB</i>	<i>UB</i>
Ages 6-18 (N = 105)						
Individual	Same	Match	.736	.017	.703	.770
		Mismatch	.258	.016	.228	.289
	Different	Match	.539	.019	.501	.577
		Mismatch	.341	.017	.308	.374
Set Average	Same	Match	.634	.017	.600	.669
		Mismatch	.308	.018	.273	.343
	Different	Match	.550	.019	.513	.586
		Mismatch	.371	.021	.331	.412
Ages 6-8 (N = 26)						
Individual	Same	Match	.635	.035	.563	.706
		Mismatch	.262	.020	.220	.303
	Different	Match	.496	.034	.427	.565
		Mismatch	.392	.037	.316	.469
Set Average	Same	Match	.589	.028	.530	.647
		Mismatch	.389	.038	.311	.466
	Different	Match	.573	.040	.490	.656
		Mismatch	.492	.034	.422	.563




Figure 2. Mean (SEM) recognition scores (proportion “present” responses on match trials minus proportion “present” responses on mismatch trials) for individual faces and set averages for studied face sets in the same-image and different-image conditions.

Is ensemble coding present in young children (6-8 year-olds)?

We wanted to determine whether ensemble coding was already present in our youngest participants. Because there were relatively few 6 and 7 year-olds but a substantial number of 8 year-olds, we examined performance in the 6-8 year old range. Set average recognition scores were significantly above chance (zero) for the same-image condition, $t(25) = 4.91$, $p < .0001$, Cohen's $d = 1.96$ (Figure 2). They were also significantly above chance in the different-image condition, $t(25) = 2.36$, $p = .026$, Cohen's $d = 0.94$ (Figure 2), although many individual children were performing around chance (Figures 3, 4). These results indicate that ensemble coding of face identity occurs in young children. Moreover, the presence of significant ensemble coding in the different-image condition indicates some sensitivity to higher-level properties related to face identity, not just simple image-based properties. A limitation is that most of the children in our 6-8 year age range were actually 8 years of age (20/26), with only five 6 year-olds, and one 7 year-old, but we can conservatively conclude that ensemble coding of face identity is present by 8 years of age.

Is there an advantage for recognition of set averages over individual faces?

There was no advantage for recognition of set averages over individual faces in either same-image or different-image conditions for either the full sample (ages 6-18 years) or 6-8 year-olds (Figure 2). Scores were numerically *higher, not lower*, for recognition of individual faces than set averages in all four cases. These results replicate other face identity findings with adults and extend them to children (Kramer et al., 2015; Neumann et al., 2013). Therefore, an advantage for recognition of set averages, which is sometimes seen for non-face stimuli (e.g., Ariely, 2001), and which is taken as evidence for distinct processes for ensemble and individual coding of other attributes, is not found for face identity. This result makes the developmental dissociation, examined below, critical for determining whether distinct processes are involved.

Does ensemble coding increase with age and is any increase dissociable from increases in individual face recognition?

Same-image condition (Figure 3). Age correlated significantly with recognition of set averages, $r = .254, p = .009, N = 105, (95\%CI = .058, .430)$ indicating a developmental increase in ensemble coding. As expected, age also correlated significantly with recognition of individual faces, $r = .492, p < .0001, N = 105 (95\%CI = .356, .611)$, replicating the well-known developmental improvement in face recognition performance. Most importantly, the correlation between age and recognition of set averages remained significant after controlling individual recognition scores, partial $r = .205, p = .037, df = 102 (95\%CI = .006, .392)$. Thus the developmental increase in ensemble coding of face identity can be dissociated from

improvements in the coding of individual face identities. Moreover, the correlation between age and individual recognition also remained significant after controlling for set average recognition scores, partial $r = .473$, $p < .0001$, $df = 102$ (95%CI = .331, .595), providing further evidence for dissociability.

Different-image condition (Figure 4). Age correlated significantly with recognition of both set averages, $r = .369$, $p < .0001$, $N = 105$ (95%CI = .203, .518), and individual faces, $r = .381$, $p < .0001$, $N = 105$ (95%CI = .211, .530). As in the same-image condition, this correlation remained significant after controlling for individual recognition scores, partial $r = .351$, $p < .0001$, $df = 102$ (95%CI = .172, .501). Thus even in this more challenging version of the task, which taps higher-level face processing, the developmental increase in ensemble coding of face identity can be dissociated from improvements in individual face recognition. Moreover, the correlation between age and individual recognition again remained significant after controlling set average recognition scores, partial $r = .363$, $p < .0001$, $df = 102$ (95%CI = .178, .528), providing further evidence for dissociability.




Figure 3. Scatterplots showing the association of age with set average recognition (a) and individual recognition (b) in the same-image condition. Best-fitting regression lines are shown. Recognition is calculated as proportion “present” responses on match trials minus proportion “present” responses on mismatch trials. $N = 105$ (multiple participants occupy some data points).

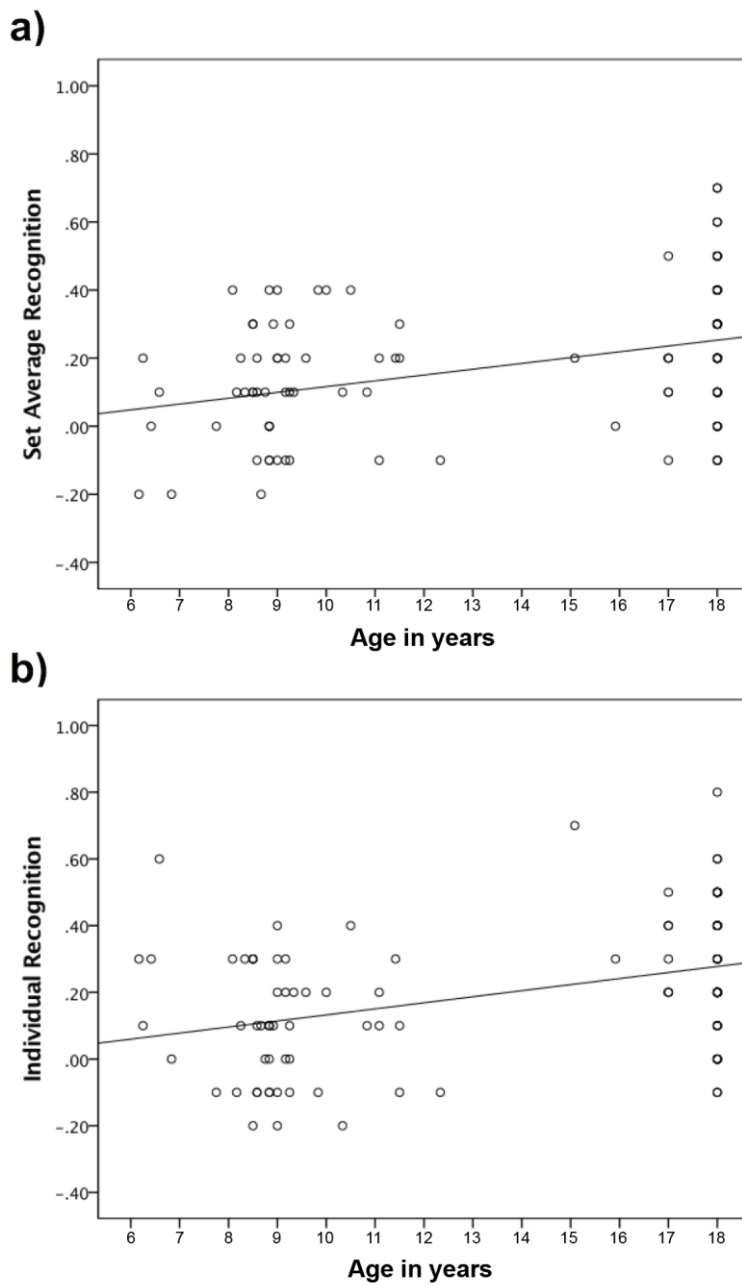


Figure 4. Scatterplots showing the association of age with set average recognition (a) and individual recognition (b) in the different-image condition. Best-fitting regression lines are shown. Recognition is calculated as proportion “present” responses on match trials minus proportion “present” responses on mismatch trials. $N = 105$ (multiple participants occupy some data points).

General Discussion

Our results add to emerging evidence that sensitivity to average properties of groups is present early in development. Ensemble coding of face identity was present by 8 years of age, and possibly as early as 6 years of age. Some ensemble coding was observed even when the average test faces were made from different images of studied individuals, indicating sensitivity to higher-level face properties, not just lower-level image properties. These results demonstrate that children's ensemble coding extends beyond simple object features (Sweeny et al., 2015; Zosh et al., 2011) to more complex properties.

A capacity to abstract average information at a glance may be especially useful for children given the substantial limitations upon their emergent attention and working memory capacities. It is interesting to note that other phenomena with low capacity requirements, such as 'pop out' effects in visual search, are also present early in development (Adler & Orprecio, 2006). Thus, despite their many visual immaturities (Maurer & Lewis, 2001), young children may readily access information about both simple object features and more complex group properties that could contribute to an adult-like subjective experience of a richly detailed visual environment (Whitney et al., 2014).

The ability to abstract average information for social groups could serve to guide appropriate and adaptive behavioural responses to those groups. It could potentially provide gist information about many attributes that are related to identity, and inferred from faces, such as trustworthiness, dominance, competence and attractiveness, because averages can preserve any consistent bias on such dimensions. For example, averages made from sets of individual faces that are all judged to be high (or low) on a dimension such as trustworthiness do indeed look high (or low) on

that dimension (e.g., Sutherland et al., 2013). Thus the ensemble coding capabilities shown here could potentially play an important role in guiding children's social behaviour towards groups.

The abstraction of average face properties may also play a role in the representation of individual faces. Both adults and children code faces relative to averages that function as a perceptual norms and are continuously updated by experience (Jeffery et al., 2010; Jeffery, Read, & Rhodes, 2013; Leopold, O'Toole, Vetter, & Blanz, 2001; Rhodes et al., 2005). Ensemble coding may contribute to these norms, which children can also abstract from sequentially viewed faces (e.g., in prototype abstraction studies) (de Haan, Johnson, Maurer, & Perrett, 2001; Gastgeb, Rump, Best, Minshew, & Strauss, 2009; Gastgeb, Wilkinson, Minshew, & Strauss, 2011; Rubenstein, Kalakanis, & Langlois, 1999; Strauss, 1979; Walton & Bower, 1993). The sequential presentation and (relative) absence of time pressure in these previous investigations contrasts powerfully with the conditions of the current study: where children were shown groups of faces simultaneously, with insufficient time for detailed and deliberate coding. An interesting question for future research is whether distinct processes underlie the abstraction of average information in these two paradigms.

Ensemble coding was present in our youngest children, but it also increased with development. So too did individual coding, replicating the well-known developmental improvement in face recognition (Carey, De Schonen, & Ellis, 1992; Germine et al., 2011; Weigelt et al., 2014). Critically, however, the developmental increase in ensemble coding was dissociable from this developmental increase in individual recognition. This dissociation cannot be attributed to task differences between ensemble and individual coding, because the task was identical in both cases,

and it provides important direct evidence that the processes underlying ensemble coding of face identity are (at least partially) distinct from those that deliver detailed individual representations of faces. In the absence of any evidence, either here or elsewhere, that ensemble face identity information is available in the absence of information about individual items, this developmental dissociation provides crucial evidence for distinct processes.

The distribution of ages was not entirely uniform in our sample. The 6-8 year range was deliberately over-represented, so that we could test whether ensemble coding was present this early. In contrast, some teenage years (12-16 years) were relatively under-represented. Individual recognition of adult faces can dip slightly in this period, due to changes in pubertal status and social goals (Carey, Diamond, & Woods, 1980; Picci & Scherf, 2016). However, the dip is small relative to the strong monotonic improvement in face recognition performance that occurs from early childhood to adulthood (Germine et al., 2011). Our sample was sensitive to that developmental improvement, with recognition of individual faces and set averages both increasing significantly with age. Critically, the non-uniform age distribution did not prevent evidence for dissociable processes emerging, with a clear developmental dissociation between the recognition of individual faces and set averages.

Future studies are needed to determine whether the developmental dissociation in ensemble coding and individual coding of face identity seen here generalizes beyond face identity, to other aspects of faces such as expression, and to other kinds of objects and properties. If children's ensemble coding is less affected than individual coding by capacity limits that decrease during development, as claimed, then we suggest that the dissociation observed here may be a very general

phenomenon. Another interesting future direction will be to determine just how early in development the ensemble coding of face properties, like identity and expression, emerges. This work would require new tasks, as the one used here proved too difficult for children younger than six years of age. Given the early presence of ensemble coding for simpler visual properties (Zosh et al., 2011), and early sensitivity to average properties of faces in (sequential presentation) prototype abstraction studies (de Haan et al., 2001), this highly efficient ability to ‘read’ social information from a crowd may well be present very early.

Acknowledgements

This research was supported by the Australian Research Council Centre of Excellence in Cognition and its Disorders (CE110001021), an ARC Professorial Fellowship to Rhodes (DP0877379) and an ARC Discovery Outstanding Researcher Award to Rhodes (DP130102300). We thank Eleni Avard, Nichola Burton, and Francis Caulfield for assistance with testing, and Nichola Burton for assistance with preparation of figures. Ethical approval was granted by the Human Research Ethics Committee of the University of Western Australia.

References

- Adler, S. A., & Orprecio, J. (2006). The eyes have it: visual pop - out in infants and adults. *Developmental Science*, 9(2), 189-206.
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15(3), 122-131.
- Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science*, 12(2), 157-162.
- Carey, S., De Schonen, S., & Ellis, H. D. (1992). Becoming a Face Expert [and Discussion]. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 335(1273), 95-103.
- Carey, S., Diamond, R., & Woods, B. (1980). Development of face recognition: A maturational component? *Developmental Psychology*, 16(4), 257-269.
- Chong, S. C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics*, 67(1), 1-13.
- Cowan, N., Elliott, E. M., Saults, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. A. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, 51(1), 42-100.
- de Fockert, J., & Wolfenstein, C. (2009). Rapid extraction of mean identity from sets of faces. *The Quarterly Journal of Experimental Psychology*, 62(9), 1716-1722.
- de Haan, M., Johnson, M. H., Maurer, D., & Perrett, D. I. (2001). Recognition of individual faces and average face prototypes by 1-and 3-month-old infants. *Cognitive Development*, 16(2), 659-678.
- Gastgeb, H. Z., Rump, K. M., Best, C. A., Minshew, N. J., & Strauss, M. S. (2009). Prototype formation in autism: can individuals with autism abstract facial prototypes? *Autism Research*, 2(5), 279-284.
- Gastgeb, H. Z., Wilkinson, D. A., Minshew, N. J., & Strauss, M. S. (2011). Can individuals with autism abstract prototypes of natural faces? *Journal of Autism and Developmental Disorders*, 41(12), 1609-1618.
- Germine, L. T., Duchaine, B., & Nakayama, K. (2011). Where cognitive development and aging meet: Face learning ability peaks after age 30. *Cognition*, 118(2), 201-210.
- Haberman, J., Brady, T. F., & Alvarez, G. A. (2015). Individual differences in ensemble perception reveal multiple, independent levels of ensemble representation. *Journal of Experimental Psychology: General*, 144(2), 432.
- Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), R751-R753.
- Haberman, J., & Whitney, D. (2009). Seeing the mean: ensemble coding for sets of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3), 718.
- Hochstein, S., Pavlovskaya, M., Bonneh, Y. S., & Soroaker, N. (2015). Global statistics are not neglected. *Journal of Vision*, 15(4), 7-7.
- Jeffery, L., McKone, E., Haynes, R., Firth, E., Pellicano, E., & Rhodes, G. (2010). Four-to-six-year-old children use norm-based coding in face-space. *Journal of Vision*, 10(5), 1-19.

- Jeffery, L., Read, A., & Rhodes, G. (2013). Four year-olds use norm-based coding for face identity. *Cognition*, *127*(2), 258-263.
- Kanwisher, N. (2000). Domain specificity in face perception. *nature neuroscience*, *3*, 759-763.
- Kanwisher, N., & Barton, J. J. S. (2011). The functional architecture of the face system: Integrating evidence from fMRI and patient studies. In A. J. Calder, G. Rhodes, M. H. Johnson & J. V. Haxby (Eds.), *The Oxford Handbook of Face Perception* (pp. 111-130). Oxford: Oxford University Press.
- Kramer, R. S. S., Ritchie, K. L., & Burton, A. M. (2015). Viewers extract the mean from images of the same person: A route to face learning. *Journal of Vision*, *15*(4), 1-1.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *nature neuroscience*, *4*(1), 89-94.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279-281.
- Maurer, D., & Lewis, T. L. (2001). Visual acuity and spatial contrast sensitivity: Normal development and underlying mechanisms. *Handbook of developmental cognitive neuroscience*, 237-250.
- Neumann, M. F., Schweinberger, S. R., & Burton, A. M. (2013). Viewers extract mean and individual identity from sets of famous faces. *Cognition*, *128*(1), 56-63.
- Noë, A., Pessoa, L., & Thompson, E. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, *7*(1-3), 93-106.
- Pavlovskaya, M., Soroker, N., Bonneh, Y. S., & Hochstein, S. (2015). Computing an Average When Part of the Population Is Not Perceived. *Journal of Cognitive Neuroscience*.
- Picci, G., & Scherf, K. (2016). From caregivers to peers: Puberty shapes human face perception. *Psychological Science*, *in press*.
- Rhodes, G. (2011). Face recognition. In D. Reisberg (Ed.), *The Oxford Handbook of Cognitive Psychology* (pp. 46-68). New York: Oxford University Press.
- Rhodes, G., Neumann, M. F., Ewing, L., & Palermo, R. (2014). Reduced set averaging of face identity in children and adolescents with autism. *The Quarterly Journal of Experimental Psychology*(ahead-of-print), 1-13.
- Rhodes, G., Robbins, R., Jaquet, E., McKone, E., Jeffery, L., & Clifford, C. W. G. (2005). Adaptation and face perception: How aftereffects implicate norm-based coding of faces. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the Mind to the World: Adaptation and Aftereffects in High-Level Vision* (pp. 213-240). Oxford: Oxford University Press.
- Riggs, K. J., McTaggart, J., Simpson, A., & Freeman, R. P. J. (2006). Changes in the capacity of visual working memory in 5- to 10-year-olds. *Journal of Experimental Child Psychology*, *95*(1), 18-26.
- Rubenstein, A. J., Kalakanis, L., & Langlois, J. H. (1999). Infant preferences for attractive faces: a cognitive explanation. *Developmental Psychology*, *35*(3), 848.
- Simmering, V. R. (2012). The development of visual working memory capacity during early childhood. *Journal of Experimental Child Psychology*, *111*(4), 695-707.

- Strauss, M. S. (1979). Abstraction of prototypical information by adults and 10-month-old infants. *Journal of Experimental Psychology: Human Learning and Memory*, 5(6), 618.
- Stuart, Alan, & Kendall, Maurice George. (1958). *Advanced theory of statistics (Vol 1)*. London: Charles Griffin & Co.
- Susilo, T., Germine, L., & Duchaine, B. (2013). Face recognition ability matures late: Evidence from individual differences in young adults. *Journal of Experimental Psychology: Human Perception and Performance*, 39(5), 1212.
- Sutherland, C. M., Oldmeadow, J. A., Santos, I. M., Towler, J., Burt, D. M., & Young, A. W. (2013). Social inferences from faces: Ambient images generate a three-dimensional model. *Cognition*, 127(1), 105-118.
- Sweeny, T. D., & Whitney, D. (2014). Perceiving Crowd Attention Ensemble Perception of a Crowd's Gaze. *Psychological Science*, 25(10), 1903-1913.
- Sweeny, T. D., Wurnitsch, N., Gopnik, A., & Whitney, D. (2015). Ensemble perception of size in 4-5 - year - old children. *Developmental Science*, 18(4), 556-568.
- Walker, D., & Vul, E. (2014). Hierarchical Encoding Makes Individuals in a Group Seem More Attractive. *Psychological Science*, 25(1), 230-235.
- Walton, G. E., & Bower, T.G.R. . (1993). Newborns form "prototypes" in less than 1 minute. *Psychological Science*, 4(3), 203-205.
- Weigelt, S., Koldewyn, K., Dilks, D. D., Balas, B., McKone, E., & Kanwisher, N. (2014). Domain - specific development of face memory but not face perception. *Developmental Science*, 17(1), 47-58.
- Whitney, D., Haberman, J., & Sweeny, T. D. (2014). From textures to crowds: multiple levels of summary statistical perception. *The new visual neurosciences*, 695-710.
- Willis, J., & Todorov, A. (2006). First Impressions: Making Up Your Mind After a 100-Ms Exposure to a Face. *Psychological Science*, 17(7), 592-598.
- Zosh, J. M., Halberda, J., & Feigenson, L. (2011). Memory for multiple visual ensembles in infancy. *Journal of Experimental Psychology: General*, 140(2), 141.