Internal Use Only
非公開

TR－H－096

0048

Exploring the Structure of Multidimensional Face Space

Robert A. JOHNSTON　Takashi KATO
Masami KANAZAWA　Masaomi ODA

1994. 8. 30

ATR人間情報通信研究所

〒619-02 京都府相楽郡精華町光台 2-2　☎07749-5-1011

ATR Human Information Processing Research Laboratories

2-2, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-02 Japan

Telephone: +81-7749-5-1011
Facsimile: +81-7749-5-1008

Exploring the Structure of Multidimensional Face Space.

Robert A. Johnston[1], Takashi Kato[2], Masami Kanazawa[2] and Masaomi Oda[2]

[1]University of Wales College of Cardiff, Cardiff, U.K. CF1 3YG
and
[2]ATR Human Information Processing Research Laboratories
2-2 Hikari-dai, Seiko-cho, Soraku-gun
Kyoto 619-02 JAPAN

Three experiments investigated the status of a model of multidimensional face space (Valentine, 1991a,b). Experiments 1 and 2 show how commonsense assumptions about similarities between faces due to age and gender can be accounted for by this type of model. A non-intuitive prediction of the model based on these findings is that face classification times will be slowest for adult male faces. This is demonstrated in Experiment 3. These findings and an unexpected discovery that face classification times are also slower to male child faces are used to explore possible structures of multidimensional face space.

Recently, a lot of attention has been given to how faces may be represented in memory. A current model, which is particularly influential, is Valentine's (1991a, 1991b) multidimensional face space. This model assumes faces to be stored as locations within a multidimensional space.

The heuristic of a multidimensional space has proven very efficient at understanding many of the phenomena which are observed in face recognition work.

Valentine (1991a) described the model at length so only a brief summary will be provided here. Essentially the model assumes that a location within a Euclidean multidimensional space furnishes a suitable metaphor for the mental representation of a face. The number and identity of the dimensions of the space have not yet been fully specified but they are assumed to be those physical and psychological attributes that underlies our superb ability to discriminate amongst many thousands, perhaps tens of thousands, of faces. There are many likely candidates for these dimensions such as face shape, face length, hair length, hair colour and perceived age to mention but a few. (See Shepherd, Ellis and Davies, 1977). The origin of the multidimensional space is the central tendency of the various dimensions and that the values of the feature dimensions for faces experienced will vary normally around this point. A consequence of this is that more typical faces will be seen (and represented) than distinctive faces and so the distribution of representations throughout the space will not be uniform. Representations of typical faces will be stored in areas with nearer neighbours than distinctive faces.

The differential density of similar representations for typical and distinctive faces has been used to explain many effects in face recognition research. These include findings from face recognition memory experiments (e.g., Light, Kayra-Stuart and Hollander, 1979; Bartlett, Hurry and Thorley, 1984; Valentine and Bruce, 1986a; Johnston and Ellis, in press), inversion experiments (e.g., Valentine and Bruce, 1986b), race-of-face experiments (e.g., Valentine and Endo, 1992) and face classification tasks (e.g., Valentine, 1991a; Johnston and Ellis, in press). This research is comprehensively treated elsewhere (e.g., Valentine, 1991b) so we restrict ourselves to an explanation of findings from work using the face classification task which directly relates to work reported in this study. Valentine and Bruce (1986a) devised a task which required subjects make a face classification decision to presented stimuli. The stimulus could either be a real face or a 'jumbled' face: a face where the gross internal features (eyes, nose and mouth) had been rearranged. Valentine and Bruce (1986a) determined that typical faces could be classified as faces faster than distinctive faces. This is accounted for by Valentine's (1991a) multidimensional space model in one of two ways depending on the appropriate framework for the space.

If the space is seen as a norm-based framework then the origin of the space will be a prototype face and as new faces are encountered they will be encoded in terms of how they deviate from this prototype. In this version, face representations are viewed as vectors from the origin of the space to the point representing the feature dimensions of a particular face. Distinctive faces will have longer vectors than typical faces and so when a face is presented for a subject to make a face classification judgement the longer vector will take longer to derive.

If the space is characterised as an exemplar-based framework the difference in classification times for distinctive and typical faces would still be predicted, but for a different reason. In this conception of Valentine's face space there is no norm or prototype face. The origin of the space has no special meaning, it is simply the central tendency of the feature dimensions. In this arrangement faces representations are usually considered as points within the space. In the face classification task when subjects have to judge if the presented stimulus is a face they form a representation of the target stimulus and the region around the stimulus is used to judge category membership, the greater the density of surrounding exemplars the faster the decision. This puts typical faces at an advantage since they are located in areas of higher exemplar density than distinctive faces. While the exemplar-based and norm-based models both predict the advantage shown for typical faces for separate reasons, it should be pointed out that these mechanisms will generally be correlated with one another: the further you move away from the origin the lower will be the density of exemplars in the local region of the face space.

Although, Valentine's model is proving to be particularly efficient in accommodating many disparate phenomena found in face recognition research, it is still far from fully understood. Almost exclusively, the work applied directly to test the model has employed unfamiliar adult male faces (e.g., Valentine 1991a; Valentine and Endo, 1992; Johnston and Ellis, in press; Johnston, Milne, Williams and Hosie, in revision; Johnston, Milne, Williams and Ellis, submitted). While there is work which comments on whether the space may be seen as an exemplar-based or norm-based framework, much of the nature of the space remains unknown and unexplored.

Valentine (1991a) emphasises that it is "implicit knowledge derived from a lifetime's experience with faces contributes to the normal distribution of faces within the multidimensional space" (p.166). The representations in the space include faces that are familiar in the sense of being celebrities or acquaintances and also faces of those people we would generally class as unfamiliar - people we pass in the street or who are less salient in our perception of them in the various media sources. Our intention in this study is to attempt to ascertain more details of the structure of this space.

We assume that the face space is used to represent both male and female faces. It is hard to see what alternative could reasonably be suggested. It would seem strange to propose that faces are first categorised according to gender and then represented in separate spaces. While it is not always clear what precisely makes a face male or female in our perception of it (Burton et al., 1993; Bruce et al, 1993) it seems safe to assume that the dimensions along which such a differentiation lies would be a useful way to discriminate faces within the multidimensional space. Indeed, this assumption is so straightforward that it would almost appear strange to raise this point, but this point is really intended to lead into a second important, but less obviously answered, question: Are children's faces also stored within the same multidimensional space as adult faces? The answer would seem to have to be affirmative. If the space is built up across a lifetime's experience of faces, many of the faces seen will be those of children. The number of children's faces experienced in proportion to the number of adult faces seen will vary over an individual's lifetime, but it is not easy to predict the full consequence of this. Nevertheless, in response to whether adults faces and children faces are stored within the same multidimensional space, again it is difficult to envisage how the conclusion could be anything other than affirmative. First, the proposal of two separate face spaces for young people and old people seems unreasonable - at what point would someone be in one group rather than the other? Are faces first classified this way before they are stored? Additionally, several researchers suggest that perceived age of face is probably a useful dimension within such a space (Shepherd, Ellis and Davies, 1977; Valentine, 1991a; Johnston et al., in revision).

The issue of the age and gender of faces suggests some interesting questions. First, consider gender, although normally there would be no difficulty in discriminating male from female faces it is not unreasonable to see this distinction as a continuum. It is a common experience to see masculine looking female faces and feminine looking male faces. How would such a dimension be represented within Valentine's space? If the label of the dimension is gender, is the central tendency a state of androgyny with absolute male and femaleness at opposite ends of this distribution? Alternatively, one of these states may represent the central tendency and the other could be represented in terms of how it deviates from this state? If this were the case it would seem intuitively more likely that females faces are closer to the central tendency than more differentiated male faces, however, the former arrangement with an androgynous centre seems more reasonable.

The possible dimension of age raises similar issues: what is the central tendency of such a dimension, middle age, young adulthood? The local distribution of the ages of faces experienced might be expected to vary as a function of an individual's age, presumably more children are seen by other children, but again the effects are difficult to predict. Valentine (1991a) has made it clear that the representations of all faces experienced will make up the population of the face space. As young adults (the age of most subjects in face recognition experiments) it is expected that most stored representations will be of faces encountered casually, in the streets or seen on TV. The number of these occurrences should far outweigh the number of faces of people known to be personally familiar (acquaintances). Is it possible to estimate what sort of central tendency on the age dimension this may give rise to?

By considering young children's faces as opposed to adult faces another possibility arises. In the case of gender, a continuous dimension seems plausible, but with the difference between adult and child faces a more abrupt differentiation between the two seems possible. Very young children's faces are not yet fully formed and also lack the later secondary sexual characteristics of adult faces and so perhaps children's faces do not vary in distinctiveness to the extent of adult faces. Is it possible to view the presence of children's faces within the multidimensional space in a manner similar to that which Valentine and Endo (1992) propose for other race faces? While wishing to be cautious in isolating faces in this way they are certainly qualitatively different from adult faces. What predictions would this make for the way they are represented? If children's faces are less differentiated than adult faces could they be regarded as 'tending towards typicality' in some way? (This description is at present necessarily vague.) If this is the case then the representations of these faces should be seen as generally towards the central tendency of the face space. Are they 'mixed in' (i.e., randomly distributed around the central tendency) or do they form a 'colony' like other-race faces, bunched together, perhaps near to the central tendency of the face space?

A second issue relates to the similarity of children's faces to female as opposed to male faces. Casual everyday language is more likely to refer to women's faces as being child-like than male faces. On a gross level, it can be noted that only adult male faces sport beard shadow but there are probably a number of other systematic differences (e.g., size of adam's apple, size of nose).

The purpose of Experiment 1 is to attempt to get to grips with this issue of similarity amongst female adult faces, male adult faces, female children's faces and male children's faces. The commonalties of these different groups will be explored by means of age and gender judgement tasks in order to try to obtain a more precise understanding about the arrangement of the age and gender dimensions. In Experiment 2 the relative locations of these faces within the face space will be explored by means of a face classification task.

EXPERIMENT 1

METHOD.

Subjects. Twelve Japanese volunteers who attended the ATR Human Information Research Laboratories in Kyoto to participate in paid experiments took part in this experiment. All had normal or corrected-to-normal vision.

Materials. Forty eight greyscale pictures of Japanese adults and children were used in this experiment. The stimuli were photographs which had been scanned into computer files and presented on the screen of a Macintosh Quadra 950 using the Superlab package. The faces were normalised so that the contrast was the same for all items and so that the ratio of interpupilary distance to other measures of the faces was kept constant. All faces were free from beards and glasses and were in a frontal neutral pose. The pictures were quasi-randomly selected from a larger set to provide four subgroups: female adults, male adults, female children and male children. The adult groups were comprised of pictures of individuals judged by the authors to be in their 20's and the child groups were comprised of individuals of about 5 or 6 years of age. There were twelve faces in each subgroup.

Design. This was a completely within subjects design. Subjects were shown the 48 faces in a separate random order and asked to make a particular decision to them. In one block they were asked to make an old/young decision and in another they were asked to make a male/female decision. The order of these tasks was counterbalanced across subjects. There were two within-subjects factors Age of face and Gender of face.

Procedure. Subjects were tested individually. Depending on which task s/he was to perform first, s/he was shown written instructions as to the nature of his/her response. At this stage s/he was not informed that there would be a subsequent task. After s/he had confirmed their understanding of the instructions, faces were presented on the screen for a maximum of one second and the subject had to decide if it was a young or old face (or, alternatively, male or female face) and make his/her response by pressing one of two keys on the computer keyboard. S/he was also told to respond as quickly and as accurately as possible. The displayed face disappeared when the subject made a response. The subjects engaged in a short practise session where faces prepared in the same manner as the experimental faces and comprising the same groups were presented. None of the faces used in the warm-up session reoccurred in the main experiment. On the completion of this warm-up session subjects immediately commenced the main experiment. The keypress recorded both response latencies and any errors made. All pictures had a presentation time of 1000 ms with an interstimulus interval of 500 ms. A response was considered to be an error if it was incorrect or if it had a latency in excess of 1500ms.

RESULTS.

Since the tasks employed in the experiment are very different they will be analysed separately and referred to as Experiment 1a (age) and Experiment 1b (gender).

The mean response latencies to classify correctly faces as young or old are shown in Table 1a. The overall error rate was less than 4% and so was not analysed further.

Table 1a: Mean response time to classify faces by age (msec).

	Female Adult	Male Adult	Female Child	Male Child
RT	850	747	733	744

The correct mean latencies for age classification were subjected to 2 x 2 ANOVA to examine the effects of age of face (adult or child; repeated measure) and gender of face (male or female; repeated measure). There was a significant main effect of age, $F(1,11) = 20.1$, $p < 0.001$, a significant main effect of gender, $F(1,11) = 7.76$, $p < 0.02$, and the age x gender interaction was also significant, $F(1,11) = 26.66$, $p =.0003$. Post hoc Tukey tests ($p = 0.01$) revealed that female adult faces took longer to classify as old than did all other types of faces.

The correct mean latencies for classifying faces according to gender can be seen in Table 1b. The number of errors was less than 4% and not analysed further.

Table 1b: Mean response time to classify faces by gender (msec).

	Female Adult	Male Adult	Female Child	Male Child
RT	817	774	953	982

A 2 x 2 ANOVA examined the effects of Age of face (old or young; repeated measure) and Sex of face (male or female; repeated measure). There was a significant main effect of Age, $F(1,11) = 55.49$, $p < 0.0001$. No other effects or interactions approached significance.

Clearly adult faces are easier to classify by gender than child faces. The mean time to classify female adult faces as female was longer than the mean time to classify male adult faces as male but not significantly so.

DISCUSSION.

The results of these two experiments showed effects that were predictable from common-sense intuitions. From Experiment 1a it can be seen that subjects are reliably slower in categorising female adult females as old compared with latencies for all other categories. A straightforward explanation for this would be that there is something about adult female faces that makes them similar to younger faces. It is important to note, however, that there seems to be no more uncertainty over child female or child male faces, an explanation for this will be offered later.

Examination of the latencies for gender decisions shows that adult faces are easily differentiated, but that subjects experience more difficulty with children's faces. This is not surprising since children's faces are not fully mature and perhaps lack important secondary sexual characteristics as cues to gender. It also appears that the source of resemblance between adult females and young faces on the age dimension does not cause difficulties when subjects are asked to judge female adult faces on a gender dimension.

An explanation for these effects will now be attempted in terms of the Valentine multidimensional face space. For the moment, this will ignore the difference between norm-based and exemplar-based models as the explanation will be the same in principle for both of them. When a subject is asked to judge the age (or gender) of a face s/he encodes the stimulus and compares it to faces which are nearby in the face space (close in distance for the exemplar-based model and close in angle for the norm-based model). Applying the same logic used to explain facedness decisions in the exemplar-based model, if there are many faces nearby all of one sort then the decision can be made very rapidly, and if there are fewer faces nearby the decision will take more time to be made. There is, however, a complication in these tasks not present in the facedness task. In the latter neighbouring faces can only be confirming, but in these present tasks there may be faces nearby to the representation which are disconfirming. There may be faces sufficiently (close) similar to be used by whatever mechanism performs the judgement, but some will be of the same category and some will be of a different category. In this event the decision process must also involve some weighing up of these contradictory sources of information.

Consider the situation for male adult faces in the age decision task. The stimulus face is encoded and compared to the general region of nearby representations, since these are almost exclusively old the decision can quickly be reached that this is an old face. Child male and child female faces must be located close to one another along the age dimension and so they act as confirming sources for one another. For the last group of faces the explanation must be rather more complicated. Female adult faces must be located somewhere between the child faces (both sexes) and the male adult faces. When a female adult face is presented for the age judgement, it is encoded and some representations near enough to take into account for the decision process will be young and some will be old. The two sets of conflicting information will need to be resolved and this requires extra time. This is revealed in the longer decision latencies to these faces. It is not an insoluble problem, subjects are making comparatively few errors, but they are being slowed down.

Turning now to the data from the gender judgements a similar explanatory mechanism can be constructed. Young and old faces are quite separate in the space along this dimension. When a target adult face is encoded it is compared with neighbouring representations in a region of space where there are few disconfirming faces and a decision can be rapidly achieved. There is more of a problem for performance with child faces on this gender assignment task. Since the child faces are not fully matured in terms of secondary sexual characteristics perhaps they are represented closer together along this dimension. When a target child face is encoded there are presumably many nearest neighbours in the space which belong to both sexes. This noise in the system causes uncertainty and results in the longer latency required to make this decision and the greater number of errors.

These findings allow a couple of tentative conclusions about the nature of the face space to be drawn. On the Age dimension, female adult faces will lie somewhere between male adult and both-gender child faces. What does this inform us about the nature of this dimension? As proposed earlier, a couple of options are possible: The central tendency of this dimension may be adult faces of 'middle age'[1] with children's faces at one end and faces of very old people at the other or alternatively the sheer lack of

[1] But not necessarily middle-aged, 'middle age' is intended to encompass the entire period from when the effects of sexual dimorphism on facial characteristics are complete until the period where old age begins to reduce the differentiation characteristic of adult faces.

differentiation of children's faces could place them at or near the origin of the space with older faces further spread out as the effects of sexual maturity takes its toll. The picture for the Gender dimension can be complimentary adult male and females are clearly differentiated, forming the alternate ends of this dimension, but the under developed faces of children lack sexual definition and are clustered together somewhere towards the centre of this dimension near the origin of the space.

Now although this appears to be a consideration of only two dimensions of what is unarguably a much higher dimensional space, there is an important point to bear in mind. The differences in faces which are present either because of sexual dimorphism or which occur through aging (we have to allow that these two effects are not entirely unrelated) are global differences which presumably affect many aspects of the face. These aspects will be those qualities which are also useful in discriminating amongst faces.

If this picture is correct it is possible to make another firm prediction with regard to the face classification task for adult and child faces. If children's faces represent one end of an 'age' dimension then face classification latencies should be made faster to adult's faces than to children's faces. If the former are clustered towards the origin of the space they should have many more representations near to them than the children's faces located out towards the periphery. If on the other hand the lack of sexual differentiation of children's faces means that they are clustered towards the centre, then it is to be expected that face classification decisions will be made more rapidly towards children's faces. Since in Experiment 1a it was determined that female adult faces were harder to judge as old than male adult faces there also arises the possibility that we may find longer latencies to male adult faces than female adult faces. It is entirely possible that the experiment is not sensitive enough to show such a difference, nevertheless the putative structure forbids the opposite outcome where longer latencies are witnessed to female adult faces. This is not just a trivial observation. The occurrence of longer face classification latencies to male adult faces would be predicted by Valentine's (1991a) model when coupled with the assumptions we have made above. However, this is by no means the only common-sense prediction from the data. In both previous tasks, male adult faces were responded to most rapidly. Presumably, this is a consequence of their suffering least interference by the other faces and thus it should follow, from Valentine's proposals that they will furnish slow response times in the face classification task. Nevertheless, there could be an argument to suggest that male adult faces are just easily and readily processed and this advantage will carry over to the face classification task also.

These predictions <u>can</u> hold true for both the exemplar-based and norm-based models, however, there are some subtle variations possible in the case on the exemplar-based model. First, the norm based model will be dealt with as its predictions are rather more straightforward.

In the norm-based model the time to complete a face classification task is determined by how closely the presented face resembles the prototype or norm face. In other words it will depend on the length of the vector which needs to be derived for the target face. On the basis of the discussion above, the prediction will be that the longest face classification latencies will be made to faces furthest from the face norm - furthest away from the central tendency of the face space.

In the exemplar-based model the prediction should be the same. It has been stressed that an assumption of Valentine's model is that there is a tendency for the exemplar density of the space to decrease as the distance from the central tendency increases. This should give rise to a similar effect to that proposed above. However, it is known from Valentine and Endo's(1992) work that this assumption does not have to hold everywhere in the face space. They showed that the axiom of decreasing density with increased distance from the origin of the space is violated in the case of other race faces. They assume that "the dimensions of the space are not optimum for the discrimination of faces drawn from a different population, and therefore it is reasonable to assume that other race faces would be more densely clustered in the space than own-race faces."(p.6). The important issue here, is that the prediction from an account relying solely on distance from the norm must necessarily be much stricter than one which depends on exemplar density which can be influenced by 'anomalous' local differences in exemplar density within the multidimensional space. The purpose of Experiment 2 is to determine whether the norm-based or exemplar-based versions of Valentine's framework can give a better account for the representation of faces.

EXPERIMENT 2

METHOD.

Subjects. Fourteen subjects who were paid volunteers in studies at the ATR Human Information Processing Research Laboratories acted as participants in this experiment. All had normal or corrected-to-normal vision. None had participated in Experiment 1.

Materials. These were the same 48 faces which had been employed in Experiment 1 and the same 48 faces made into 'jumbled' faces. In order to make the jumbled faces, the gross internal regions (eyes, nose and mouth) were divided up into roughly equal rectangular regions which were then reassembled in the order nose-mouth-eyes, mouth-nose-eyes and eyes-mouth-nose. In order to avoid any superficial differences between faces and non-faces the lines used to section features in the non-faces were left on the faces. Eight extra faces comprising two in each of the experimental conditions were treated in a similar fashion to be used as practise items.

Design. The same four sets of pictures (adult male, adult female, child male and child female) from Experiment 1 were shown to all subjects. These were accompanied by each of the 48 pictures presented in a 'jumbled' format. There were two within subject factors, the gender and age (child or adult) of the face. The dependent variable was the response time required to make a correct categorisation.

Procedure. Subjects were tested individually. At the start of the experiment they were presented with an example of an intact and a jumbled face and the nature of the task was explained to them. They were informed that they would see either faces or stimuli that looked like faces but which did not have a normal arrangement of features. Subjects were asked to try to decide as quickly, but as accurately, as possible whether each successive image was a face and to signal their decision by means of a key press. The key press recorded the response latency of the subject's response and also its accuracy. Prior to the start of the experiment subjects engage in a warm-up session where they performed this task to a smaller set of faces which again comprised the same four types of faces as the experimental block. None of the faces used in the warm-up session reoccurred in the main experiment. All pictures had a presentation time of 1000 ms with an interstimulus interval of 500 ms. A response was considered to be an error if it was incorrect or if it had a latency in excess of 1500ms.

RESULTS.

The mean times to classify correctly the intact stimuli as faces are shown in Table 2. The error rate was less than 2% and no further analysis was carried out.

Table 2: Mean time to classify correctly stimuli as faces (msec).

	Female Adult	Male Adult	Female Child	Male Child
RT	589	620	596	631

A 2x2 ANOVA was conducted to examine the effects of Gender of face (male or female; repeated measure) and age of face (adult or child; repeated measure). There was a significant main effect of Gender,

$F(1,13) = 6.96$, $p = 0.02$, No other main effects or interactions were significant [all F's < 1].

Subjects required more time to classify male faces as a face than they needed to classify female faces. This effect was independent of the age of the face.

DISCUSSION.

From the results of Experiment 2 it can be seen that male faces are taking longer to classify than female faces in a face classification task. This was expected in the case of the male adult faces, but it was not anticipated in the case of the male child faces. In Experiment 1 response latencies for both processing of these faces was not impaired by the presence of other faces in the experiment. In terms of Valentine's model this would be explained as follows. When a target male adult face is encoded and compared with its appropriate region within the multidimensional space, there are few representations which are disconfirming to the hypothesis that it is both male and adult. In order to explain the slower latencies for female adult faces in the age judgement task we proposed that representation for these face overlap on this dimension with younger faces. In the gender judgement task it is only the child faces which appear to interfere with one another which suggests that these two populations of faces are intermixed substantially along this dimension. This combination of factors suggested, according to Valentine's model, that subjects would demonstrate very long latencies to male adult faces in a facedness task. This is in contradiction to a common-sense interpretation which might have simply predicted that subjects were very good at processing male adult faces. The fastest latencies for female adult and female child faces was also anticipated as these faces will overlap in space with each other and, it is expected, that female child faces would overlap with male child faces thus explaining why latencies to the gender judgement task are extended for child faces. It is, however, unexpected that male child faces show such long latencies.

An explanation for this effect will be attempted with reference to the alternative norm-based and exemplar-based versions of Valentine's multidimensional space. Given the results from Experiment 1, a particular arrangement of old and young and male and female faces within the multidimensional space can be suggested. This arrangement is shown schematically in Figure 1 (it is very important to note that this diagram is intended only as an aid to discussion. It is unlikely that Age and Gender, if they turn out to be sensible dimensions, could ever be thought of as perpendicular to one another. Indeed, it seems most likely that they are related to one another since many changes which occur due to sexual maturation will also correlate with increasing age). In this arrangement faces which are both adult and male are thought to be clustered in the upper right quadrant, far removed from other types of faces, and clearly defined towards the masculine end of the gender dimension and the older end of the Age dimension. Female adults are not so isolated, they are well defined in terms of femininity, but share characteristics which make them overlap with female younger faces. Both female and male young faces are distributed to what we will call the young end of the Age dimension, but they are not clearly differentiated in terms of Gender and so fall just either side of the midpoint. This is a different account to those proposed earlier in this paper. A prior suggestion made in this account to describe the structure of the face space proposed that either adult faces that were fully sexually differentiated but not yet subject to reduced differentiation through the effects of old age would form the central tendency. Alternatively this area was proposed to be comprised of the very undifferentiated and hence tending-to-be-typical children's faces. Neither of these hypotheses sit comfortably with the data from Experiment 2. In view of the results from the Experiment 2, it is suggested that the central tendency along the age dimension falls some where between very young faces - typical through a lack of maturational differentiation - and adult faces which are seen so much more frequently - and hence are represented much more often within the face space. (Although in these particular experiments faces of adults in their twenties have been used, it is expected that this grouping includes all adult faces before an age where a decay of features leads to a decrement of differentiation.)

Figure 1. Schematic representation of Age and Gender Dimensions.

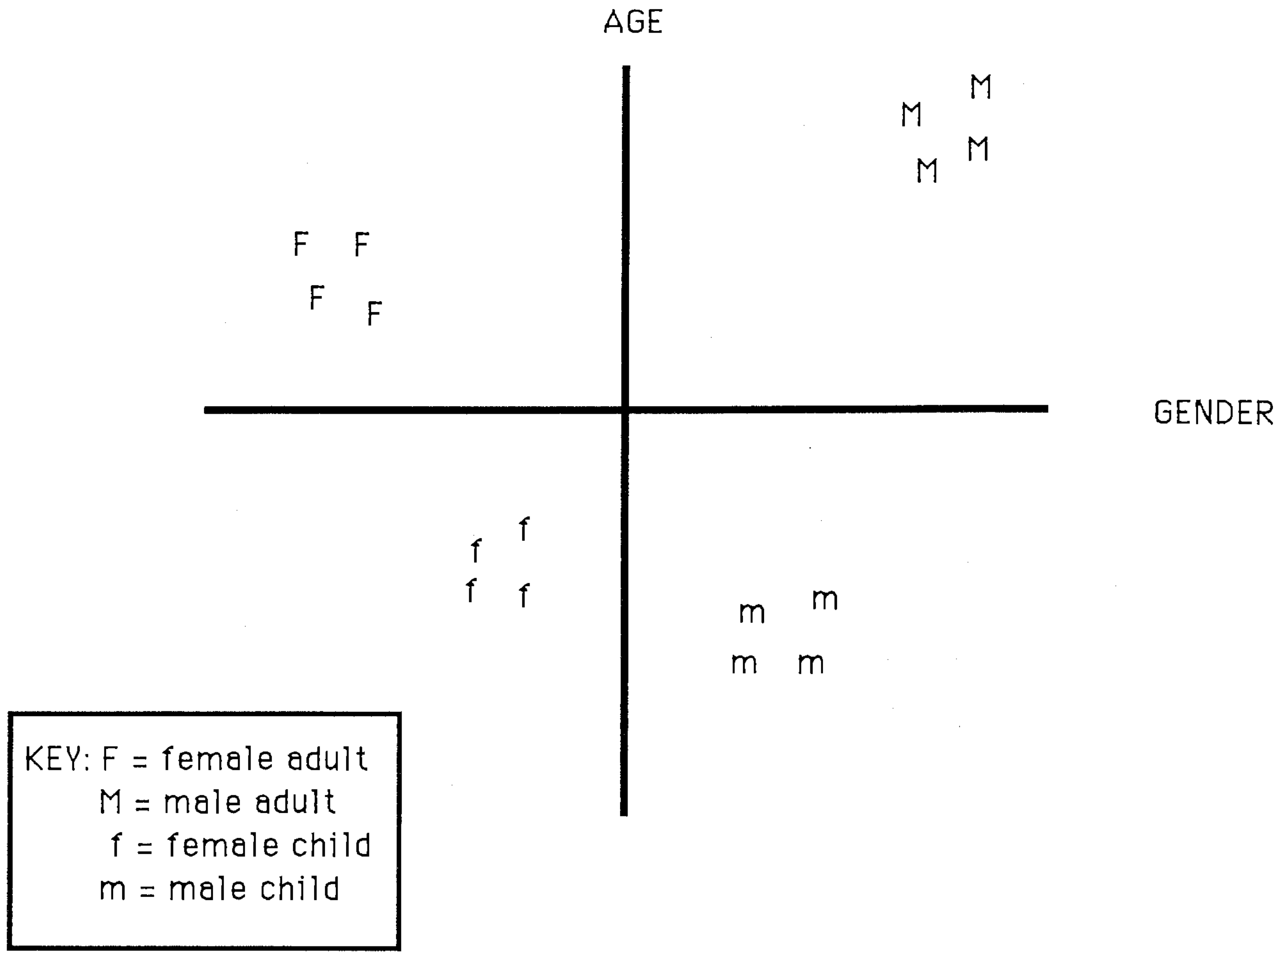

While it is, of course, impossible to make any hard quantitative judgements about latencies for the face classification decisions on the basis of this suggested arrangement, it is possible to suggest an ordinal scaling of RTs. The norm-based model would clearly rank RTs to male adult faces as slowest, but would be hard-pressed to discriminate amongst the other three groups of faces on the basis of the limited data from these experiments. The exemplar-based model would draw the same conclusion based on the assumption that exemplar density decreases with distance from the origin of the space. The discovery that long response latencies are also obtained with presentations of male child faces does not fit easily with these suggestions. It is difficult to determine how an explanation based on the length of derived vectors from a norm can accommodate this data, but an explanation based on an exemplar-based account provides a more flexible solution.

While in the simple case, the exemplar-based model assumes that exemplar density will be negatively correlated with distance from the central tendency of the multidimensional space, it is known that there are departures from this state of affairs (e.g., other race faces). It is suggested that the results of this experiment arise as a consequence of a similar phenomenon. It is proposed that the representations of child faces are relatively very sparse compared to adult faces - especially in the face spaces belonging to young adults (the subjects in these experiments). In the case of female child faces this effect is offset because they are both similar to male child faces along the gender dimension and female adult faces along the age dimension. Male child faces are considered to be very distinct from both type of adult face. The low exemplar density of the local space surrounding these face representations thus gives rise to comparatively lengthy response latencies in the face classification times. A similar effect has been shown by Valentine (1991a, Exp.. 5) where longer latencies were shown to other race faces in a face classification task.

GENERAL DISCUSSION.

Two experiments examined the possible structure of a multidimensional face space such as that proposed by Valentine (1991a, 199b). The first experiment explored similarities between four subsets of faces (male adults, female adults, male children and female children) and discussed possible structural arrangements of those face-types within the face space. Two options were suggested that made certain predictions about subjects performance with the different face-types in a face classification task.

Experiment 2 explored the prediction made by the alternative options and it was discovered that neither was a good explanation of the data collected. A third alternative, based on data from both experiments, was constructed. It is proposed that while the latter model can provide a reconciliation of all the data, it can do so only be appealing to the effects of local 'anomalies' in the exemplar density of the face space. This is not unprecedented as it has been proposed to account for recognition effects in the processing of other-race faces (e.g., Valentine and Endo, 1992). The most important consequence of such a decision is that it means the work reported here must add to the burgeoning accumulation of evidence that Valentine's face space should be viewed as having an exemplar-based framework (Valentine and Endo, 1992; Johnston et al., submitted: but see also Johnston and Ellis, In Press for alternative arguments). We have shown that a non-intuitive prediction of Valentine's model obtains: namely that face classification decisions are made more slowly to male adult faces although these faces can be classified by age and gender more rapidly than other groups of faces employed. The longer response times obtained with male child faces was unexpected and our future experimental work will attempt to understand how this effect arises.

Acknowledgements. The authors would like to acknowledge the support of President Y. Tokhura and Dr. S. Akamatsu for the opportunity to carry out these investigations. Bob Johnston is particularly grateful to his co-authors for opening up an exciting new avenue of research.

REFERENCES

Bartlett, J. C., Hurry, S. and Thorley, W. (1984). Typicality and familiarity of faces. Memory and Cognition, 12, 219-228.

Bruce, V., Burton, A.M., Hanna, E., Healy, P., Mason, O., Coombes, A., Fright, F. and Linney, A. (1993) Sex discrimination: how do we tell the difference between male and female faces? Perception, 22(2), 131-152.

Burton, A.M., Bruce, V. and Dench, N. (1993). What's the difference between men and women? Evidence from facial measurement. Perception, 22(2), 153 - 177.

Cohen, M. E and Carr, W. J. (1975). Facial recognition and the Von Restorff effect. Bulletin of the Psychonomic Society, 6, 383-384.

Ellis, H.D. (1992). The development of face processing skills. Philosophical Transactions of the Royal Society of London. B, 335, 105-111.

Goldstein, A. G. and Chance, J. E. (1980) Memory for faces and schema theory. Journal of Psychology, 105, 47-59.

Johnston, R.A. and Ellis, H.D. (In press) "Age effects in the processing of typical and distinctive faces.

Johnston, R.A., Milne, A.B., Williams, C.L. and Hosie, J. (submitted). Do distinctive faces come from outer space: An investigation of the status of a multi-dimensional face space.

Johnston, R.A., Milne, A.B., Williams, C.L. and Ellis, H.D. (In prep). Determining between two models of face space.

Light, L.L. Kayra-Stuart, F. & Hollander, S. (1979). Recognition memory for typical and unusual faces. Journal of Experimental Psychology: Human Learning and Memory, 5, 212-228.

Rhodes, G. (1988). Looking at faces: First order and second order features as determinants of facial appearance. Perception, 17, 43-63.

Valentine, T. and Bruce, V. (1986a). The effects of distinctiveness in recognising and classifying faces. Perception, 15, 525-535.

Valentine, T. and Bruce, V. (1986b). Recognising familiar faces: The role of distinctiveness and familiarity. Canadian Journal of Psychology, 40, 300-305.

Valentine, T. (1991a). A unified account of the effects of distinctiveness, inversion and race on face recognition. Quarterly Journal of Experimental Psychology, 43A, 161-204.

Valentine, T. (1991b). Representation and process in face recognition. In: Watt, R. (ed.) Vision and visual dysfunction. Vol. 14: Pattern recognition in Man and Machine. (series editor J. Cronley-Dillan). London: MacMillan.

Valentine, T. and Endo, M. (1992). Towards an exemplar model of face processing: The effects of race and distinctiveness. Quarterly Journal of Experimental Psychology,