Someone Like You? Modelling Face Memory Reveals Task-generalizable Representations

Jiayu Zhan^a, Oliver G. B. Garrod^a, Nicola van Rijsbergen^a & Philippe G. Schyns^{a, b}

^a Institute of Neuroscience and Psychology, University of Glasgow, Scotland G12 8QB, United Kingdom.

^b School of Psychology, University of Glasgow, Scotland G12 8QB, United Kingdom.

Corresponding author

Philippe G. Schyns

Institute of Neuroscience and Psychology, 58 Hillhead Street, Glasgow, Scotland, UK, G12 8QB.

Tel.: +44 (0) 141 330 4937

E-mail: philippe.schyns@glasgow.ac.uk

ABSTRACT

Current theories of cognition are cast in terms of information processing mechanisms that use mental representations. For example, consider the mechanisms of face identification that use mental representations to identify familiar faces under various conditions of pose, illumination and ageing, or to draw resemblance between family members. Providing an explanation of these information processing mechanisms thus relies on showing how the actual information contents of these representations are used. Yet, these representational contents are rarely characterized, which in turn hinders knowledge of mechanisms. Here, we address this pervasive gap by characterizing the detailed contents of mental representations of familiar faces using a new methodological approach. We used a unique generative model of face identity information combined with perceptual judgments and reverse correlation to model the 3D representational contents of 4 familiar faces in 14 participants. We then demonstrated the validity of these contents using everyday perceptual tasks that generalize face identity and resemblance judgments to new viewpoints, age and sex with a new group of participants. Our work highlights that such models of mental representations are critical to understanding generalization behavior and for characterizing the information processing mechanisms that must use them.

The cognitive mechanism of recognition is guided by mental representations that are stored in memory (1-4). The information contents of these representations must therefore be sufficiently detailed to enable accurate recognition (i.e. identifying 'Mary' amongst other people) and sufficiently versatile to enable recognition across diverse everyday tasks —e.g. identifying Mary in different poses, at different ages or identifying her brother based on family resemblance (5-7). To model the face identity information that is stored in memory and guides such tasks, we developed a novel methodology based on reverse correlation (see Figure 1A) and a new 3D face information generator (see Figure 1B).

On each experimental trial, we generated a set of 6 new 3D faces (see Figure 1A), each with a unique and randomly generated identity. Critically, each face shared other categorical face information (i.e. sex, age and ethnicity) with one of the four familiar faces—e.g. with the familiar target face of 'Mary.' To achieve this, we used a Generalized Linear Model (GLM) to decompose the familiar target face into a categorical component (e.g., for 'Mary' the average of all white females faces of 30 years of age) plus a residual component that defines the specific identity of the familiar face (see Figure 1B, and *Materials and Methods, Generative Model of Face Identity*). We then generated new random identities by keeping the categorical component of the target constant (e.g., white female, 30 years of age) and adding a random component of identity (see Figure 1B and *Materials and Methods, Reverse Correlation Experiment, Random Face Identities*). Participants saw these randomly generated faces in full frontal view and selected the one that most resembled the familiar target (e.g., 'Mary') and rated its similarity to the target on a 6-point Likert scale, ranging from not at all ('1') to highly similar ('6'). To resolve the task, participants must compare the randomly generated faces presented on each trial with their mental representation of the familiar target in full frontal view. Therefore, each face selected comprises a match to the participant's mental representation of the target, which is estimated by the similarity rating of that face (see Figure 1A for an example trial and *Materials and Methods, Reverse Correlation Experiment, Familiar Faces, Participants and Procedure*).

Figure 1. Reverse correlating mental representations of familiar faces. (A) Task. Illustrative experimental trial with 6 randomly generated face identities. We instructed participants to use their memory to select the face most similar to a familiar identity (here, 'Mary') and then to rate the similarity of the selected face (purple frame) to their memory of 'Mary' (purple pointer). (B) Generalized Linear Model (GLM) of face identity. In its forward computation flow (see identity modelling solid arrow), the General Linear Model (GLM) decomposes a 3D, textured face (e.g. 'Jane' or 'Tom') into a non-identity face shape average capturing the categorical factors of face sex, ethnicity, age and their interactions plus a separate component that defines the identity of the face (illustrated by the 3D shape decomposition; 2D texture, not illustrated, is independently and similarly decomposed). Heat

maps indicate the 3D shape deviations that define 'Jane' and 'Tom' in the GLM in relation to their categorical averages. In the reverse flow (see dashed arrow of identity generation), we can randomize the 3D shape identity component (and 2D texture component, not illustrated here), add the categorical average of 'Jane' (or 'Tom') and generate random faces, each with a unique identity that share all other categorical face information with 'Jane' and 'Tom.'

After many such trials, we used reverse correlation (8) to estimate the information content of the mental representation of each target familiar face ($N=4$, see Figure S1B) in each participant ($N=14$). Specifically, we build a statistical relationship between the information content of the faces that the participant selected on each trial with their corresponding similarity ratings. In a second stage, we tested with a new group of participants whether these modelled mental representations were sufficiently detailed to enable identification of each target familiar face and sufficiently versatile to enable resemblance judgments across diverse everyday tasks--i.e. generalization across new viewpoints, age and siblings.

RESULTS

1. Analysis of the Contents of Mental Representation of Familiar Faces

To reconstruct the information contents of mental representations, we used linear regression to compute the single-trial relationship between <similarity ratings, random face identity components> for each target familiar face and participant. Specifically, we computed separate regressions between the similarity ratings and each 3D shape vertex and each RGB texture pixel that comprise the face identity components. We then used the resulting Beta coefficients to model the 3D shape and texture identity components that characterize the participant's mental representation of each familiar face in the GLM (see Figure S2 and *Materials and Methods, Analyses, Linear Regression Model and Reconstructing Mental Representations*).

With this approach, we can formally characterize and then compare the participant's mental representation of a familiar face with the ground truth face—i.e. the objective identity component of the scanned familiar face, see Figure S1B. We focus only on 3D shape because there were few non-systematic relationships for texture (see *Materials and Method, Analyses, Vertex Contribution to Mental Representations, Step1* and Figure S3). To illustrate, grey faces on the x-axis of Figure 2A show the ground truth identity component in the GLM for Inward and Outward 3D shape deviations in relation to the categorical average (i.e., all white females of 30 years of age, like 'Mary'). For example, Mary's nose is objectively thinner than the average of white females of her age, and so these vertices deviate inward (darker grey tones indicate increasing deviations). Likewise, her more pouty mouth is shown as an outward 3D shape deviation. The y-axis of Figure 2A uses the same format to show the mental representation of Mary in one typical participant, where colors indicate increasing deviations. These contents reveal faithful representations of, for example, a thinner nose and a pouty mouth (see *Materials and*

Method, Analyses, Vertex Contribution to Mental Representations). A scatter plot visualizes the vertex by vertex fit between the mental representation (y-axis) and the ground truth 3D face (x-axis). The white diagonal line provides a veridical reference, where the identity component in the mental representation is identical to the ground truth face, for every single 3D vertex coordinates.

Our analyses reveal the specific vertices near the veridical line that faithfully represent 'Mary' in the mind of this participant ($p < 0.05$, two-tailed), as colored dots reported on the scatter and located on the y-axis faces in Figure 2A. In contrast, white vertices away from the veridical line did not faithfully represent the identity (see *Materials and Methods, Analyses, Vertex Contribution*). We repeated the analysis of represented contents for each participant ($N = 14$) and familiar face ($N = 4$). Figure 2B reports the collated group results, using the format of Figure 2A, where colors now indicate N, the number of participants who faithfully represented that identity in their mind with this particular 3D shape vertex. Figure 2B demonstrates that mental representations comprised similar information contents across the 14 individual participants. Most (10/14) faithfully represented 'Mary's' thin nose, 'John's' receding eyes and wider upper face (13/14), 'Peter's' prominent eyebrow and jawline (13/14), 'Stephany's' protruding mouth (13/14), see *Material and Methods, Analyses, Vertex Contribution to Mental Representations*.

Figure 2. Contents of mental representations of familiar faces. (A) Mental representation of 'Mary' (a typical participant). *Ground truth:* 3D vertex positions deviate both Inward (-) and Outward (+) from the categorical average to objectively define the shape of each familiar face identity (see Figure 1A). Grey scale values reported on the flanking faces color-code the normalized magnitudes of inward and outward deviations from the categorical average. *Mental representation:* Inward and Outward colored faces highlight the individual 3D vertices whose position faithfully deviate from the categorical average in the GLM. Blue to red colors

represent the normalized magnitudes of their deviations. *2D scatter plots*: Scatter plots indicate the relationship between each vertex deviation in the ground truth (rank-ordered on a normalized scale on the X-axis) and the corresponding vertex in the memory representation (also rank-ordered on a normalized scale on the Y-axis). The white diagonal line provides the reference of veridical mental representation in the GLM—i.e. a hypothetical numerical correspondence between each shape vertex position in the ground truth face and in the mental representation of the same face. White dots indicate vertices that were not faithfully represented. (B) Mental Representations (group results). Same caption as Figure 2A, except that the colormap now reflects the number of participants ($N = 14$) who faithfully represented this particular shape vertex.

Such convergence of represented contents across participants suggests that the face representations are multivariate (i.e. comprising contiguous surface patches rather than isolated vertices). As a final step, we extracted the main multivariate components of represented surface patches by applying across observers and familiar faces Non-negative Matrix Factorization (NNMF, (9)) to the faithfully represented 3D vertices (see *Materials and Method, Analyses, Components of Memory Representation*). Figure 3A shows the multivariate diagnostic components that faithfully represent ‘Mary’ (i.e. the red framed map of vertices in Figure 3A). Importantly, these diagnostic components of familiar face identity have complementary nondiagnostic components (i.e. the grey framed map of vertices in Figure 3A), which capture variable face surfaces that do not comprise the participants’ mental representations. Supplementary Figure S4 shows the diagnostic and nondiagnostic components of the other familiar faces.

2. Efficacy of the Information Contents of Mental Representations in New Participants and Tasks

Here, we develop the critical demonstration that the information contents of the mental representations we modelled are valid. That is, the contents enable accurate identification of each target face and they also enable resemble tasks that preserve their identity. We asked a new group of participants (called ‘validators’, see *Generalization Experiment, Participants*) to resolve a variety of novel resemblance tasks that are akin to everyday tasks of face recognition. Success on these tasks would demonstrate that the diagnostic components derived from the previous experiment comprise identity information that can be used in new generalization tasks. Therefore, although the components are extracted under one viewpoint (full-face), one age (for each identity) and one sex (that of the identity), here we tested the generalization of identification performance to new viewpoints, ages and sex.

For this demonstration, we synthesized new diagnostic (vs. nondiagnostic) faces that were parametrically controlled for the relative strength of the diagnostic multivariate components of identity vs. their nondiagnostic complement (see Figure 3A and *Generalization Experiments, Stimuli*). It is important to emphasize that both

diagnostic and nondiagnostic faces are equally faithful representations of the original ground truth. That is, their shape features are equidistant from the shared categorical average. However, whereas the diagnostic components deviate from the average with multivariate information extracted from the participants' mental representations, the nondiagnostic components do not. We hypothesized that, though equidistant from the categorical average, only the diagnostic components will impact performance on the novel tasks. For all synthesized faces, we changed their viewpoint (rotation of -30 deg, 0 deg and +30 deg in depth), age (to 80 years old), and sex (to opposite) using the generative model--see Figure S5.

In three independent resemblance tasks – changes of viewpoint, age and sex – we tested the identification performance of 12 validators using a 5 Alternative Force Choice task (i.e. responding one of four familiar identities plus a 'don't know' response, see *Materials and Methods, Generalization Experiments, Participants and Procedure*). In each task, we found a significantly higher global identification performance for diagnostic faces (see Figure 3B, red curves) than for nondiagnostic faces (black curves, $p < 0.05$, 1-tailed, see also *Materials and Methods, Generalization Experiments, Analyses*). Thus, the diagnostic contents of the mental representations we modelled do indeed contain the information that can resolve identity and novel resemblance tasks.

A. Diagnostic and Nondiagnostic Faces

B. Identification Performance of Diagnostic and Nondiagnostic Faces

Figure 3. Generalization of Performance Across Tasks. (A) Diagnostic and nondiagnostic Faces. *Left panel:* The red framed map shows the multivariate diagnostic components of faithful 3D shape representation of 'Mary'; the grey framed map shows the nondiagnostic complement (1 - diagnostic components). *Middle panel:* Faces synthesized with increasing amplification (0.33 to 1.67) of the diagnostic (top) vs. nondiagnostic (bottom) components. *Right panel:* For each synthesized face, we changed its viewpoint (30° left and 30° right), age

(80 years old) and sex, shown here for faces synthesized at amplification = 1. (B) Task Performance. For each condition of generalization (y-axis) and familiar identity (x-axis), 2D plots show the median identification performance computed across participants (y-axis) for faces synthesized with the diagnostic (red curves) and nondiagnostic (grey curves) faces, at different levels of amplification of the multivariate components (x-axis). Shadowed regions indicate median absolute deviations (MAD) of identification performance.

DISCUSSION

Mental representations stored in memory are critical to guide the information processing mechanisms of cognition. Here, with a novel methodology based on reverse correlation and a new 3D face information generator, we modelled the information contents of mental representations of 4 familiar faces in 14 individual observers. We showed that the contents converged across participants on a set of multivariate features (i.e. local and global surface patches) that faithfully represent 3D information that is objectively diagnostic of each familiar face. Critically, we showed that validators could identify new faces generated with these diagnostic representations across new resemblance tasks—i.e. changes of pose, age and sex—but performed much worse with equally faithful, but nondiagnostic features. Together, our results demonstrate that the modelled representational contents were both sufficiently precise to enable face identification within task and versatile enough to generalize usage of the identity contents to other resemblance tasks.

At this stage, it is worth stepping away from the results and emphasize that it is remarkable that the reverse correlation methodology works at all, let alone producing robust generalization across resemblance tasks. In the experiment, we asked observers to rate the resemblance between a remembered familiar face, and randomly generated faces, that by construction are very unlike the target face (never identical, and almost never very similar). And yet, our results show that the representational contents we modelled following such a task were in fact part of the contents that objectively (i.e. faithfully) support identity recognition. This raises a number of important points that we now discuss.

Relationships to Face Space

It is important to clarify that we modelled face identification information in a face space that belongs to the broad class of 3D morphable, Active Appearance Models of facial synthesis (AAMs, (10, 11)). These models contain full 3D surface and 2D texture information about faces and so superseded the former generation of 2D image-based face spaces ((12, 13) (14)). To synthesize faces, we used a GLM to decompose each face identity as a linear combination of components of shape and texture added to a local average (that summarizes the categorical factor of age, gender, ethnicity and their interactions, cf. Figure 1B). To model the mental representations of faces, we estimated the identity components of shape and texture from the memory of each observer. Thus, for both stimulus synthesis and mental representation estimation, we used the same AAM framework.

There is a well-known problem with using AAMs to model the psychology of face recognition. Perceptual expertise and familiarity are thought to involve representations of faces that enable the greater generalization performance that is widely reported (15-18). However, AAMs typically adopt a brute force approach to identity representation: a veridical (i.e. totally faithful) deviation of each physical shape vertex and texture pixel from an average. Thus, as AAMs overfit identity information, they appear as a priori weak candidate models to represent the perceptual expertise (14).

Our approach of studying the contents of mental representations offers a solution to this information processing conundrum. We showed that each observer faithfully represented only a proportion of the objective identity information that defines a familiar face identity. Our key theoretical contribution to face space is to formalize the diagnostic information as a reduced set of multivariate face features that can be construed as dimensions of the observer's face space. Observers develop these dimensions whenever they interact with the objective information that represents a new face identity in the real world. We modelled the objective information that is available to the observer for developing their face space dimensions via learning as the veridical shape and texture information of the AAM (14, 19, 20). Key to demonstrating the psychological relevance of our psychological face space dimensions is that they should comprise identity information sufficiently detailed to enable accurate face identification and sufficiently versatile to enable similarity judgments of identity in novel tasks. We demonstrated this potential when validators identified faces synthesized with the diagnostic dimensions in novel resemblance tasks. Thus, by introducing reduced faithful mental representations of identity information in the objective representations of AAMs we provide the means of modelling the psychological dimensions of face space.

Our practical contribution to face recognition is that we can now precisely track the development of the psychological dimensions of face space. AAMs enable a tight control of objective face information at synthesis, such as ambient factors of illumination, pose and scale, but also categorical factors of gender, sex, age and ethnicity and components of identity. Thus, it is now possible to tightly control the statistics of exposure to faces in individual observers, model the diagnostic dimensions of the psychological face space that are learned, and finally test the efficacy of the psychological face space as we did here. And when we understand how ambient and categorical factors influence performance, we can switch to understanding familiar face identification in the wild, where all ambient and categorical factors get naturally mixed up, and where the influence of each factor to identification performance becomes near impossible to disentangle, precluding a detailed information processing understanding of face identification mechanisms.

Mental Representations or Task Representations?

A general question with reverse correlation tasks is whether the resulting models represent a particular visual category (here, the visual identity of a face) or the task

from which the model was reconstructed (20-23). We contributed to this debate by showing that the identity information reconstructed in one task had efficacy in other tasks that involved identity. Importantly, the tasks were designed to test two classes of factors: ambient and categorical. For example, we showed that the identity component extracted in one ambient viewpoint (full face, 0 deg) could be used to generalize identification of the same face under two new ambient viewpoints (-30 and +30 deg of rotation in depth). We also showed that the identity component extracted for identities (all < 40 years of age) generalized to older age (80 years). Furthermore, we also showed that though extracted from a given sex, the identity component would generalize to another sex, a kinship task. Hence, we found no dramatic differences due to the effect of task of extraction of the identity component. Rather, the extracted representational basis is useful for all tasks tested, whether using ambient or categorical factors of face variance. This therefore suggests that we have tapped into some essential information about familiar face identity. However, we acknowledge that the generalizations we observe might still be a function of an interaction between the nature of memory and the similarity task from which we estimated the identity component. The component could have differed had the task been more visual than memory based (e.g. identification of the same face under different orientations, or a visual matching task) and we might not have derived an identity component that enabled such effective generalization.

Implications of Explicit Models of Mental Representations of Visual Categories

Our models of mental representation should be construed as the abstract information goals that the visual system predicts when identifying familiar faces. We term the faithful components as the ‘abstract information goal’ because it has to be broken down into global and local constituents according to the constraints of representation and implementation at each level of the visual hierarchy—or their analogues in a multi-layered deep convolutional network. To be clear, though we modelled the mental representation of a face identity in an AAM it is important to state that we do *not* assume that memory really represents faces in this way (i.e. as demarcations to an average, separately for 3D shape and 2D texture). AAM is only a state-of-the-art, mathematical modelling framework. We fully acknowledge there are many possible concrete implementations into a neural, or a neurally-inspired architecture that could deliver AAM-like performance without assuming an explicit AAM representation. What is clear is that whichever implementation, in whichever architecture, the abstract information represented will have to enable the performance characteristics our resemblance tasks demonstrated.

For example, we would hypothesize that the diagnostic identity components in Figure 3A are broken down, bottom to top, into the representational language of V1—i.e. as representation in multi-scale, multi-orientation Gabor-like, retinotopically mapped receptive fields (24, 25); at intermediate levels of processing, as the sort of local surface patches (26, 27) that we reveal, and at the top level as the combinations of surface patches that enable identification and resemblance

responses. Under a framework of top-down prediction (28, 29), the abstract information goal of a familiar face identity should trim, in a top-down manner, the fully-mapped but redundant information on the retina into the task-relevant features that are transferred along the occipital to ventral/dorsal visual hierarchy (30). Tracing the construction of such a reduced memory representation of face identity in the brain should enable an accurate and detailed modelling of the processing mechanism along the visual hierarchy (see also (31-34)). What our work critically provides is an estimate of the end goal of the hierarchy (i.e. the diagnostic component), which is also a prediction of what is important in the input. It is in this sense that mental representations guide task-specific information processing in the brain. Without knowing mental representations, we do not have even have an information needle to search in the fabled haystack of brain activity, let alone reconstruct the mechanisms that process its contents.

CONCLUDING REMARKS

We modelled the critical mental representations of that guide the processing of visual information of familiar face identities. In several resemblance tasks that require usage of face identity, we demonstrated the efficacy of the contents we modelled. Our approach and results open new research avenues for the interplay between visual information, categorization tasks and their implementation as information processing mechanisms in the brain.

REFERENCES

1. Bar M (2009) The proactive brain: memory for predictions. *Philos T R Soc B* 364(1521):1235-1243.
2. Bar M, et al. (2006) Top-down facilitation of visual recognition. *Proc Natl Acad Sci U S A* 103(2):449-454.
3. Ullman S, Assif L, Fetaya E, & Harari D (2016) Atoms of recognition in human and computer vision. *P Natl Acad Sci USA* 113(10):2744-2749.
4. Harel A, Kravitz DJ, & Baker CI (2014) Task context impacts visual object processing differentially across the cortex. *Proc Natl Acad Sci U S A* 111(10):E962-971.
5. O'Toole AJ (2011) Cognitive and Computational Approaches to Face Recognition *The Oxford Handbook of Face Perception*, eds Rhodes G, Calder A, Johnson M, & Haxby JV), pp 15 -30.
6. Tsao DY & Livingstone MS (2008) Mechanisms of face perception. *Annu Rev Neurosci* 31:411-437.
7. Rosch E & Mervis CB (1975) Family Resemblances - Studies in Internal Structure of Categories. *Cognitive Psychol* 7(4):573-605.
8. Ahumada A & Lovell J (1971) Stimulus Features in Signal Detection. *J Acoust Soc Am* 49(6):1751-&.
9. Lee DD & Seung HS (1999) Learning the parts of objects by non-negative matrix factorization. *Nature* 401(6755):788-791.

10. Cootes TF, Edwards GJ, & Taylor CJ (2001) Active appearance models. *Ieee T Pattern Anal* 23(6):681-685.
11. Blanz V & Vetter T (1999) A morphable model for the synthesis of 3D faces. *Comp Graph*:187-194.
12. Turk M & Pentland A (1991) Eigenfaces for recognition. *J Cogn Neurosci* 3(1):71-86.
13. Rhodes G & Jeffery L (2006) Adaptive norm-based coding of facial identity. *Vision Res* 46(18):2977-2987.
14. O'Toole AJ, Castillo CD, Parde CJ, Hill MQ, & Chellappa R (2018) Face Space Representations in Deep Convolutional Neural Networks. *Trends Cogn Sci*.
15. Young AW & Burton AM (2018) Are We Face Experts? *Trends in Cognitive Sciences* 22(2):100-110.
16. White D, Phillips PJ, Hahn CA, Hill M, & O'Toole AJ (2015) Perceptual expertise in forensic facial image comparison. *Proc Biol Sci* 282(1814).
17. Eger E, Schweinberger SR, Dolan RJ, & Henson RN (2005) Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *Neuroimage* 26(4):1128-1139.
18. Jenkins R, White D, Van Montfort X, & Burton AM (2011) Variability in photos of the same face. *Cognition* 121(3):313-323.
19. Gosselin F & Schyns PG (2002) RAP: a new framework for visual categorization. *Trends Cogn Sci* 6(2):70-77.
20. Schyns PG (1998) Diagnostic recognition: task constraints, object information, and their interactions. *Cognition* 67(1-2):147-179.
21. Liu ZL (1996) Viewpoint dependency in object representation and recognition. *Spatial Vision* 9(4):491-521.
22. Schyns PG, Goldstone RL, & Thibaut JP (1998) The development of features in object concepts. *Behav Brain Sci* 21(1):1-17; discussion 17-54.
23. Mangini MC & Biederman I (2004) Making the ineffable explicit: estimating the information employed for face classifications. *Cognitive Sci* 28(2):209-226.
24. Kay KN, Naselaris T, Prenger RJ, & Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452(7185):352-U357.
25. Smith FW & Muckli L (2010) Nonstimulated early visual areas carry information about surrounding context. *P Natl Acad Sci USA* 107(46):20099-20103.
26. Peirce JW (2015) Understanding mid-level representations in visual processing. *J Vis* 15(7):5.
27. Kubilius J, Wagemans J, & Op de Beeck HP (2014) A conceptual framework of computations in mid-level vision. *Front Comput Neurosci* 8:158.
28. Friston KJ & Kiebel S (2009) Predictive coding under the free-energy principle. *Philos T R Soc B* 364(1521):1211-1221.
29. Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36(3):181-204.
30. Zhan J, Van Rijsbergen N, Ince RA, & Schyns PG (2018) Dynamic Construction of Reduced Representations in the Brain for Perceptual Decision Behavior. *bioRxiv*.
31. Gosselin F & Schyns PG (2003) Superstitious perceptions reveal properties of internal representations. *Psychol Sci* 14(5):505-509.

32. Smith ML, Gosselin F, & Schyns PG (2012) Measuring Internal Representations from Behavioral and Brain Data. *Current Biology* 22(3):191-196.
33. Nestor A, Plaut DC, & Behrmann M (2016) Feature-based face representations and image reconstruction from behavioral and neural data. *Proc Natl Acad Sci U S A* 113(2):416-421.
34. Nestor A, Plaut DC, & Behrmann M (2011) Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *P Natl Acad Sci USA* 108(24):9998-10003.

MATERIALS AND METHODS

Generative Model of Face Identity.

We designed a generative model to objectively characterize and control face identity variance, using a database of 355 3D faces (acquired with a 4D face capture system, see *Supplementary Information 3D Face Database*) that describes each face by its shape (with 3D coordinates for each one of 4,735 vertices) and its texture (with the RGB values of 800*600 pixels, see Figure S1A). It is critical to reiterate that the familiar faces were not part of the 3D face database.

To design the Generative Model, we first applied a high-dimensional General Linear Model (GLM), separately to each 3D vertex coordinate and 2D pixel RGB value, to model and explain away variations in face shape and texture that arise from the non-identity categorical factors of sex, age, ethnicity, and their interactions. The GLM therefore: 1) extracted as a non-identity face average the shape and texture face information explained by non-identity categorical factors; and also 2) isolated the residual information that defines the 3D shape and 2D texture identity information of each face--i.e. the identity residuals.

To further control identity information, we applied Principal Components Analysis (PCA) to the identity residuals of the 355 faces, separately for shape and texture. The PCA represented shape residuals as a 355-dimensional vector in a 355-dimensional space of multivariate components, and a separate PCA represented the texture residuals as a 355*5 (spatial frequency bands)-dimensional matrix in a space of 355*5 multivariate components. Two sets of PCA coordinates therefore represented the objective shape and texture information of each identity in the principal components space of identity residuals.

Furthermore, Figure S1B illustrates how the generative model controlled the non-identity and identity factors using the 4 familiar faces of our experiment. First, we scanned the four familiar faces of the experiment (2nd column). We fitted each into the GLM model to derive a ground truth face (the 3rd column), with minimal distortions (shown in the 1st column).

The model generates new 3D faces by adding a given set of identity residuals to different non-identity GLM averages, to change their age, sex or ethnicity separately, or jointly sex and ethnicity. The outcomes are older, sex swapped, ethnicity swapped and sex and ethnicity swapped versions of the same identity (the 4th to 7th column). We used these generative properties to derive the stimuli of the generalization experiment.

Figure S1. (A) 3D Face parameters. We parametrized the shape of a face with the 3D coordinates of 4,735 vertices and its texture with 800*600 RGB 2D pixels. (B) GLM control of the categorical averages and identity components. *Distortion.* Distortion quantifies, vertex per vertex the quality of the 3D GLM fit of the scanned familiar faces. *Changing the categorical averages.* In each column, the GLM controls the factors of sex, ethnicity and age using local averages, while the identity residuals are kept constant.

Reverse Correlation Experiment

Participants. We selected 14 Western Caucasians (7 females, mean age = 25.86 years, SD = 2.26 years) who were familiar with each familiar identity as work colleagues. We assessed familiarity on a 9-point Likert scale, from not at all familiar '1' to highly familiar '9'. Table S1 reports the familiarity ratings for each identity and participant. All participants had normal or corrected-to-normal vision, without a self-reported history or symptoms of synaesthesia, and/or any psychological, psychiatric or neurological condition that affects face processing (e.g., depression, autism spectrum disorder or prosopagnosia). They gave written informed consent and received £6 per hour for their participation. The University of Glasgow College of Science and Engineering Ethics Committee provided ethical approval.

Familiar Faces. We scanned four faces 'Mary' and 'Stephany' (white Caucasian females of 36 and 38 of age, respectively), and 'John' and 'Peter' (white Caucasian males of 31 and 38 years of age, respectively) who were familiar to all participants as work colleagues. As we will explain, we used these scanned faces to compare the objective and mentally represented identity information in each participant.

Random Face Identities. We reversed the flow of computation in the generative model to synthesize new random identities while controlling their non-identity factors (see Figure 1B, the reverse direction is indicated by the dashed line). We proceeded in three steps: First, we fitted the familiar identity in the GLM to isolate its non-identity averages, independently for shape and texture. Second, we randomized identity information by creating random identity residuals—i.e. we generated random coefficients (shape: 355; texture: 355*5) and multiplied them by the principal components of residual variance (shape: 355; texture: 355*5). Finally, we added the random identity residuals to the GLM averages to create a total of 10,800 random faces per familiar identity in the reverse correlation experiment.

Procedure. Each experimental block started with a centrally presented frontal view of a randomly chosen familiar face (henceforth, the target). On each trial of the block, participants viewed six simultaneously presented randomly generated identities based on the target, displayed in a 2 x 3 array on a black background, with faces subtending an average of 9.5° by 6.4° of visual angle. We instructed participants to respond on one of 6 buttons to choose the face that most resembled the target. The six faces remained on the screen until response. Another screen immediately followed instructing participants to rank the similarity of their choice to the target, using a 6-point Likert scale ('1' = not similar, '6' = highly similar) with corresponding response buttons. Following the response, a new trial began. The experiment comprised 1,800 trials per target, divided into 90 blocks of 20 trials each, run over several days, for a grand total of 7,200 trials per participant. Throughout, participants sat in a dimly lit room and used a chin rest to maintain a 76 cm viewing distance. We ran the experiment using the Psychtoolbox for MATLAB R2012a.

Analyses

Linear Regression Model. For each participant and target face, each trial produced two outcomes: one matrix of 4,735*3 vertex (and 800*600 RGB pixel) parameters corresponding to the shape (and texture) residuals of the chosen random face on this trial, and one corresponding integer that captures the similarity between the random identity parameters and the target. Across the 1,800 trials per target, we linearly regressed (i.e. RobustFit, Matlab 2013b) the 3D residual vertices (separately for the X, Y and Z coordinates) and residual RGB pixels (separately for R, G and B color channel) with the corresponding similarity rating values. These linear regressions produced a linear model with coefficients Beta_1 and Beta_2 vectors for each residual shape vertex coordinate and residual RGB texture pixel, for each familiar face and participant. Figure S2A illustrates the linear regression model for the 3D vertices of 'Mary.' Henceforth, we focus our analyses on the Beta_2 coefficients because they quantify how shape and texture identity residuals deviate from the GLM categorical average to represent the identity of each familiar face in the memory of each participant.

Reconstructing Mental Representations. Beta_2 coefficients can be amplified to control their relative presence in a newly synthesized 3D face. Figure S2B1 illustrates such amplification for one participant's Beta_2 coefficients of shape and texture of 'Mary.' Following the reverse correlation experiment, we brought each participant back to fine-tune their Beta_2 coefficients for each familiar face, using the identical display and viewing distance parameters as in the reverse correlation experiment. In a self-adaptive procedure, we initialized Beta_2 amplification with equally spaced values between 0 and 50, with 10 unit increments. We then narrowed the amplification range to participant's responses until convergence, keeping the same total number of stimuli (i.e. 6 faces) per trial. Figure S2B2 illustrates the adaptive procedure.

The experiment comprised one session per familiar face, with familiar face order randomized across participants. Each session started with the screen presentation of the front view of one familiar face target to instruct participants as to the target of the session. On each trial, 6 faces initially amplified between 0 and 50 appeared on the screen, randomly positioned in a 2 by 3 array against a black background. We instructed participants to choose the face that best resembled the familiar identity by pressing one of six response buttons. The 6 faces remained on the screen until response, immediately followed by the next trial. We repeated the trial 5 times, with the same 6 faces in different random array positions, to determine the next amplification range. We narrowed the amplification range every 5 trials by finding the minimum and maximum values that bound the participant's 5 choices. With this new range, we produced 6 new faces by evenly sampling the amplification values and again tested the participant over 5 new trials. We iteratively repeated sequences of 5 testing trials, updates of the amplification range, until it stabilized—i.e. remained constant over three blocks of 5 trials. We used the median of the final

amplification range as value to generate the fine-tuned Beta_2 coefficients that we call *mental representation* in our analyses (see Figure S2B2).

Figure S2. Reverse-correlating the visual information contents of familiar face representations. (A) Estimating Beta_2 Coefficients. We linearly regressed the 3D vertices of shape (separately for the X, Y and Z coordinates, texture not illustrated) with similarity

judgments of the selected random identities (illustrated here for 'Mary'). For each vertex, Beta_2 coefficients are color-coded according to their magnitude. Yellow-to-red indicates an outward change from the categorical average; turquoise-to-blue indicates an inward change from the categorical average. (B) Fine-tuning Beta_2 Coefficients. (B.1) *Amplifying Beta_2 coefficients.* Illustration of the amplification of Beta_2 coefficients. (B.2) *Illustration of the fine-tuning experiment.*

Vertex Contribution to Mental Representations. Vertices, whether in the ground truth face or in the participant's mental representation can deviate inward or outward in 3D from the corresponding vertex in the common categorical average of their GLM fits (cf. Figure 1B). Thus, we can compare the respective deviations of their 3D vertices in relation to the common GLM categorical average. To evaluate this relationship, we rank-ordered ground truth vertices from most Inward (-1) to most Outward (+1) in normalized deviation on the X-axis of a 2D scatter plot; we also reported the corresponding rank of each vertex of the mental representation on the Y-axis (as shown Figure 2A). If ground truth and mental representations were identical, their vertex-by-vertex deviations from the GLM categorical average (i.e. Euclidean distance) would be identical and would form the veridical diagonal straight white line provided as a reference in the scatter plot of Figure 2A.

Using this veridical line as a reference, for each participant and familiar face representation, we proceeded in three steps to classify each vertex as either 'faithful' or 'not faithful', and to test whether the vertices in mental representations deviated from the categorical average more than would be expected to occur by chance.

Step 1: We constructed a permutation distribution by iterating our regression analysis 1,000 times with random permutations of the choice response across the 1,800 trials. To control for multiple comparisons, we selected maximum (vs. minimum) Beta_2 coefficients across all shape vertices (and texture pixels), separately for the X, Y and Z coordinates (RGB color channels) from each iteration. We used the resulting distribution of maxima (and minima) to compute the 95% confidence interval of chance-level upper (and lower) Beta_2 value and classified each Beta_2 coefficient as significantly different from chance ($p < 0.05$), or not. We consider the vertex (or pixel) as significant if the Beta_2 coefficient of any coordinate (or color channel) was significant. There were very few significant pixels, with almost no consistency across participants (see Figure S3), so we excluded texture identity residuals from further analyses.

Step 2: We used the chance-fit Beta coefficients in Step 1 and the Beta_2 amplification value derived in **Reconstructing Mental Representation** to compute the equation $GLM + \beta_1 + \beta_2 * amplification\ value$ (cf. Figure S2B). As a result, we built a distribution of 1,000 chance fit faces

Step 3: To classify whether each significant 3D vertex in the mental representation of a participant is more similar to ground truth than we would expect by chance, we computed D_{chance} , the mean Euclidean distance between the 1,000

chance fit faces and the veridical line, and D_{memory} , the distance between the same mental representation vertex and the veridical line. If $D_{\text{memory}} < D_{\text{chance}}$, this significant vertex is 'faithful' because it is significantly closer to the veridical line than chance (and we plot it with blue to red colors in Figure 2A); if $D_{\text{memory}} > D_{\text{chance}}$, the vertex is not faithful (and we plot it in white in Figure 2A, together with the nonsignificant vertices).

To derive group results, we counted across participants the frequency of each faithful vertex and used a Winner-Take-All scheme to determine group-level consistency. For example, if 13/14 participants represented this particular vertex as 'faithful,' we categorized it as such at the group level and reported the number of participants as a color indicating 13 participants. If there was no majority for a vertex, we color-coded it as white (see Figure 2B).

Figure S3. Beta_2 coefficients of texture. Yellow colored overlays on each familiar face illustrate the significant Beta_2 coefficients for RGB texture pixels in each participant (labelled S1-S14). Dark purple pixels represent non-significant RGB coefficients.

Components of Memory Representation. The purpose of the following analysis was to find common diagnostic components (multivariate features) that emerged in the group-level memory representation of each face identity. To do so, we factorized with Non-negative Matrix Factorization (NNMF) the total set of memory representations across familiar identities and observers.

For each participant, we recoded each vertex in the identity residuals of each familiar face as ‘faithful’ = 1, ‘not faithful’ or not significant = 0, resulting in a 4735-d binary vector. We pooled 56 such binary vectors (across 4 targets x 14 observers = 56) to create a 4735 by 56 (i.e. vertex-by-model) binary matrix to which we applied NNMF to derive 8 multivariate components that captured the main features that faithfully represent familiar faces in memory across participants (Figure S4A shows each NNMF component).

To determine the loading (i.e. the contribution) of each NNMF component in the group-level mental representation of each familiar face identity, we computed the median loading of this component on the 14 binary vectors representing this identity in the 14 observers. We applied a 0.1 loading threshold (> 73 percentile of all 8 components $\times$ 4 identities median loadings) to ascribe a given component to a familiar face representation. The colored boxplot in Supplementary Figure S4A represents the loading of each NNMF component at the group-level representation, showing that at least 2 above-threshold NNMF components represent each familiar identity.

We then constructed the diagnostic component of a familiar identity representation as follows: for each vertex we extracted the maximum loading value across the NNMF components representing it, and normalized the values to the maximum loading across all vertices. This produced a 4735-d vector V_d that weighs the respective contribution of each 3D vertex to the faithful representation of this familiar identity that we call the “diagnostic component.” The heat maps in the left column of Figure 4SB represent the diagnostic component of each familiar identity.

Crucially for our validation experiment, we were then able to define a nondiagnostic component as the complement of the diagnostic component $V_n = 1 - V_d$. It is important to emphasize that we adjusted the total deviation magnitude of the diagnostic and nondiagnostic components from the categorical average—i.e. by equating the total sum of their deviations. This ensures that diagnostic and nondiagnostic components are both equidistant from the average face in the objective face space. The right column of Figure 4SB shows the nondiagnostic component of each familiar identity representation.

Figure S4. NNMF multivariate and compact representations. A. NNMF representations of faithful 3D vertices across the mental representations of participants. The x-axis presents each NNMF component, where colors indicate the relative weight of each shape vertex in the component (normalized by maximum weight across components). Boxplots on the y-axis show the loading of each NNMF component on the faithful representations ($N = 14$, one per participant) of each familiar identity ($N = 4$ familiar identities), with colored boxes indicating above 0.1 threshold loading for NNMF components. B. Diagnostic and nondiagnostic components for each familiar identity. Heat maps in the left column show the group-level diagnostic component for each familiar identity; heat maps in the right column show the complementary nondiagnostic components.

Generalization Experiments

Participants. We selected 11 Western Caucasians (7 females) and 1 East Asian (female), with mean age = 28.25 years and SD = 4.11 years. Each is familiar with the identities as work colleagues and also assessed familiarity on a 9-point Likert scale (see Table S2). All participants had normal or corrected-to-normal vision, without a self-reported history or symptoms of synaesthesia, and/or any psychological, psychiatric or neurological condition that affects face processing (e.g., depression, autism spectrum disorder or prosopagnosia). They gave written informed consent and received £6 per hour for their participation. The University of Glasgow College of Science and Engineering Ethics Committee provided ethical approval.

Stimuli. For each familiar identity, we synthesized new 3D faces that comprised graded levels of either the diagnostic or the nondiagnostic shape

components. Specifically, we used the normalized diagnostic component V_d and its nondiagnostic complement V_n to synthesize morphed faces with shape information of each target identity as follows:

$$\text{Diagnostic Faces} = \text{Ground Truth} \times V_d \times \alpha + \text{Categorical Average} (1 - V_d \times \alpha)$$

$$\text{Nondiagnostic Faces} = \text{Ground Truth} \times V_n \times \alpha + \text{Categorical Average} (1 - V_n \times \alpha)$$

with amplification factor $\alpha = 0.33, 0.67, 1, 1.33, 1.67$, to control the relative intensity of diagnostic and nondiagnostic shape changes. We rendered all these morphed shapes with the same average texture. The first rows of Figure S5-1 to S5-4 show the morphed faces for each familiar identity. We added as filler stimuli the grand average face (for both shape and texture) of the 355 database faces.

We also changed the viewpoint, age and sex of all of these synthesized faces. Specifically, we rotated them in depth by -30 deg, 0 deg and $+30$ deg and using the generative model, we set the age factor to 80 years/swapped the sex factor, keeping all other factors constant (see *Generative Model of Face Identity* in Figure 1B and Figure S1B).

Procedure. The experiment comprised 3 conditions (viewpoint, age and sex) that all participants accomplished in a random order, with one condition per day. In the Viewpoint condition, participants ran 15 blocks of 41 trials (5 repetitions of 123 stimuli). Each trial started with a centrally displayed fixation for 1s, followed by a face on a black background for 500ms. We instructed participants to name the face as ‘Mary,’ ‘Stephany,’ ‘John’ or ‘Peter,’ or respond ‘other’ if they could not identify the face. Participants were required to respond as accurately and as quickly as possible. A 2s fixation separated each trial. Participants could break between blocks. In the Age and Sex conditions, participants ran 5 blocks that repeated 42 trials. They were instructed to respond “Old ‘Mary,” “Old Stephany,” “Old John,” “Old Peter” or “Other” in the age condition, and “Mary’s brother,” “Stephany’s brother,” “John’s sister,” “Peter’s sister” or “Other” in the sex condition.

Analyses. For each participant and generalization condition, we computed the percent correct identification of diagnostic and nondiagnostic faces. In each of the three Viewpoints, we paired the diagnostic and nondiagnostic stimuli of each participant and identity, at each level of feature intensity. In Age and Sex, we paired the diagnostic and non-diagnostic stimuli of each participant and identity at each intensity level. Then, in each of the 5 tasks, we used a Wilcoxon sign ranked test that compared performance between diagnostic vs. nondiagnostic component trials and obtained the Wilcoxon sign ranked statistics W . We determined the statistical significance of W ($p < 0.01$) using a bootstrapped null distribution of W ($N = 1,000$ iterations). On each iteration, we randomly shuffled the observers’ responses across all trials within a task, and computed the W between two random sets of trials of equal cardinality (i.e. 5 levels $\times$ 4 identities $\times$ 12 validators = 240 trials in each random set). We corrected significance across the 5 tasks—i.e. $p < 0.05$.

Figure S5-1. Diagnostic (left) and nondiagnostic (right) faces of 'Mary'. Each row presents the main conditions of stimulus synthesis (i.e. 3 viewpoints, age and sex). Each column presents a level of diagnostic (vs. nondiagnostic) component amplification in the face.

Figure S5-2. Diagnostic (left) and nondiagnostic (right) faces of 'Stephany'. Same caption as in Figure S5-1.

Figure S5-3. Diagnostic (left) and nondiagnostic (right) faces of 'John'. Same caption as in Figure S5-1.

Figure S5-4. Diagnostic (left) and nondiagnostic (right) faces of 'Peter'. Same caption as in Figure S5-1.

SUPPLEMENTARY INFORMATION

3D Face Database. The face database comprised 197 females, 158 males, 233 Western Caucasian, 122 East Asian, age between 16 and 86, SD = 15.06, scanned in-house with a Di4D face capture system, at a high resolution in shape (4,735 3D vertex coordinates) and texture (800*600 RGB pixels, see Figure S1A). All 3D models were in full color with hair removed, posing with a neutral facial expression.

Table S1. Identity familiarity ratings of 14 participants in the reverse correlation experiment.

Participants	Mary	Stephany	John	Peter
1	6	4	3	7
2	7	8	9	6
3	7	9	7	3
4	8	8	7	8
5	3	4	4	4
6	3	3	4	5
7	5	6	9	4
8	5	6	9	5
9	8	7	7	4
10	8	6	8	10
11	7	7	8	7
12	7	6	6	7
13	5	3	9	7
14	9	9	9	9
Mean	6.29	6.14	7.07	6.14
SD	1.86	2.03	2.09	2.07

Note: Ratings are from 1 (not familiar at all) to 9 (highly familiar).

Table S2. Identity familiarity ratings of 12 participants in the generalization experiment.

Validators	Mary	Stephany	John	Peter
1	9	7	6	6
2	7	5	6	4
3	9	9	9	9
4	7	8	9	6
5	7	6	5	4
6	6	4	5	4
7	8	6	9	8
8	9	6	9	9
9	8	5	7	6
10	6	9	3	4
11	9	9	7	9
12	9	5	9	7
Mean	7.83	6.58	7	6.33
SD	1.19	1.78	2.04	2.06

Note: Ratings are from 1 (not familiar at all) to 9 (highly familiar).