Conceptual rather than perceptual similarity enables generalization across
perceptually different appearances of familiarized faces

Maya Gotlieb¹, Naphtali Abudaraham¹, Yarden Shir¹ & Galit Yovel^{1,2}

¹School of Psychological Science

²Sagol School of Neuroscience

Tel Aviv University

Pre-registration of the two studies can be found in these links:

Study 1: <http://aspredicted.org/blind.php?x=6ti4u4>.

Study 2: <http://aspredicted.org/blind.php?x=83eh3z>.

The data files of these studies are posted in [OSF](#)

Correspondence regarding this work should be e-mailed to Galit Yovel,
gality@tauex.tau.ac.il

Abstract

Face recognition is a challenging categorization task, as in many cases the variability between different images of the same identity may be larger than the variability between images of different identities. Nevertheless, humans excel in this task, in particular for faces they are familiar with. What type of learning and what is the nature of the representation of the learned identity that support such remarkable categorization ability? Here we propose that conceptual learning and the generation of a conceptual representation of the learned identity in memory enables this classification performance. First, we show that humans learn to link perceptually different faces to the same identity, if faces are learned with the same conceptual information. Next, we show that this conceptual learning does not generate a single perceptual representation of the different appearances of each identity. Instead, perceptually dissimilar images of the same identity remain separated in the perceptual space and are linked conceptually rather than perceptually. This conceptual representation of face identity is advantageous, as it enables generalization across perceptually dissimilar images of the same identity/category, without increasing false recognition of perceptually similar images of different identities. A similar conceptual mechanism may also apply to other familiar categories such as familiar voices or objects of expertise that involve fine discrimination of a homogenous sets of stimuli that are linked to unique conceptual information. Overall these findings highlight the importance of studying the contribution of both cognition and perception to face recognition.

Recognition of a stimulus depends on matching its perceptual representation to a representation in memory that is defined by its categorical boundaries (Tanaka, Giles, Kremen, & Simon, 1998). One of the most challenging recognition tasks that humans perform on a daily basis is face recognition. Face recognition is computationally challenging because the perceptual variability of images of the same identity may sometimes exceeds the variability between identities (see Figure 1). Despite this challenge, humans excel in this task. This remarkable ability, however, is limited to familiar faces, highlighting the critical role of experience in face recognition (Jenkins, White, Van Montfort, & Mike Burton, 2011; Ritchie & Burton, 2017; Young & Burton, 2017, 2018). Thus, understanding how humans successfully perform this computationally complex task entails answering two main questions: First, what type of *experience* is needed to enable humans' exceptional recognition of familiar faces. Second, what is the nature of the representation that is generated, following this type of experience, that allows generalization across highly dissimilar images of the same identity and discrimination among different identities.

Figure 1: Example of face images of two women, each is shown at the center of each row. The images to the right are different appearances of the same identity and images to the left are different identities. This shows that perceptual variability within identity may be larger than between identities. Individuals who are familiar with these identities (an Israeli celebrity at the top row and a Brazilian celebrity at the bottom row) typically perform this challenging face matching task effortlessly (see Ritchie & Burton, 2015).

Empirical and theoretical accounts of familiar face recognition highlight the importance of experience with high variability of images of the same identity (Andrews, Jenkins, Cursiter, & Burton, 2015; Kramer, Jenkins, Young, & Burton, 2017; Kramer, Towler, Reynolds, & Burton, 2018; Ritchie & Burton, 2017). In the current study, we test two different representations that may be generated following this type of experience that may account for humans' familiar face recognition. The two types of representations are shown in Figure 2, in which faces are represented as dots in a multidimensional feature space (i.e. feature-vectors) (Valentine, 1999). Similarity between faces is represented by their face-space distance (i.e., distance between feature vectors). Each identity occupies a subspace (denoted by the black line) that includes its different learned appearances. An incoming image is considered familiar, if it falls within the

perceptual boundary of the identity subspace (Tanaka, Giles, Kremen, & Simon, 1998).

The dashed line in Figure 2 indicates the subspace boundary of a given identity. As shown in Figure 1, different appearances of the same individual may be perceptually different and therefore reside in distant locations in the face space. The two hypotheses propose different ways in which perceptually different appearances of the same identity may be linked to the same identity:

Figure 2: Two hypotheses that account for the generalization across perceptually different images of the same identity were tested: the perceptual single identity subspace hypothesis (left) and the conceptual multiple identity subspace hypothesis (right). Similarity among faces is depicted as distance in face space. The black line indicates the boundary of an identity subspace of a given learned identity (e.g., David). Faces that fall within the subspace belong to the same identity. The dashed line indicates identity boundary of perceptually similar appearances before learning. The two hypotheses make different predictions with respect to the generalization across perceptually different appearances of the learned identity. According to the perceptual single identity subspace hypothesis (left), the identity subspace expands following exposure to a new appearance (dotted line). According to the conceptual multiple identity subspace hypothesis (right), perceptually different appearances of the same identity are stored in separate subspaces and are linked conceptually. To decide between these hypotheses, following a supervised learning phase, in which the learned faces are presented with the same name label (see Figure 3), an average face (denoted by the red dot) that was not presented during learning is presented at test. The perceptual single subspace representation predicts that the average face will be included in the representation of the learned identity, whereas the conceptual multiple subspace representation predicts that the average face will not be included in the representation of the learned identity.

According to the *single identity subspace* hypothesis, each identity is represented by a single subspace (Figure 2, left). Following exposure to a perceptually different image

of the same identity, the boundary of the identity subspace is expanded to include the newly learned appearance of the same identity (dotted line). The updated identity subspace now occupies a larger area. Any face that falls within its boundaries, based on its perceptual representation, will be categorized as the same identity. This hypothesis is therefore named the *perceptual single identity subspace* hypothesis and may be consistent with previously suggested prototype-based representations (for review see, Minda & Smith, 2012; see also Burton et al., 2011; Kramer et al., 2015). The main disadvantage of this representation, is that it may include perceptually similar faces of different identities that fall within the relatively large identity subspace and lead to false identification. We therefore propose a new account for familiar face recognition, that we name the conceptual *multiple identity subspace* hypothesis. According to this hypothesis, learning a new appearance of the same identity, that is perceptually different, does not expand the boundaries of the identity subspace. Instead, a new identity subspace is formed. Thus, an identity may be represented by multiple identity subspaces that reside in distant locations and cannot be matched perceptually (Figure 2, Right). How are these separated subspaces then linked to the same identity? Because familiar faces are learned and then stored in memory with relevant conceptual information about a person (i.e., person knowledge), these multiple identity-subspaces can be matched based on their conceptual similarity rather than based on their perceptual similarity. The main benefit of such conceptual multiple identity subspace representation is that it enables generalization across perceptually different images of the same identity, without the cost of increasing false identification of perceptually similar faces of different identities. The main disadvantage of this representation is that it depends on specific supervised experience with the perceptually different appearances of the learned identities, and therefore does not generalize to other identities. However, familiar face recognition indeed involves supervised learning and does not generalize to other identities and is therefore consistent with such a representation.

To decide between these two hypotheses, we designed an experiment in which perceptually different faces, i.e. faces that reside in separate subspaces, were linked together to the same identity by giving them the same name label during a learning phase (i.e., conceptual/supervised learning). In Study 1, we first show that participants match perceptually different faces that are learned with the same name label to the same identity. Based on these findings, in Study 2 we examined how this conceptual

learning changed the representation of the learned identities. In particular we examined the boundaries of the identity subspace of the learned identities, by presenting participants with average images of the learned faces (see Figure 5) and examine whether they are included in the identity subspace of the learned identities. If a single subspace is generated for the learned identities, then the average face should be included in the representation. If the average face is not included in the representation, the learned identities are represented as separated subspaces (see Figure 2).

Study 1

To decide between the two hypotheses (Figure 2), we first had to assess whether learning faces that reside in different identity subspaces with the same conceptual label can link them to the same identity. In addition, in order to know whether faces reside in the same or different identity subspaces, we need to know what are the facial features that define the identity of a face (i.e., the dimensions of the face space). In previous studies, we discovered a subset of facial features that changing them moved faces outside of their identity subspace (i.e., changed the identity of the face). We therefore named them *critical features*, as they are critical for the identity of a face. In

Figure 3: A: Faces were changed by either modifying their critical features to move them outside of their identity subspace or modifying non-critical features to keep them within the same identity subspace (See also Figure S1). B: Design of Study 1: During the learning phase pairs of faces that differ in critical (shown in panel B) or non-critical features (not shown in panel B) are presented sequentially. Half of the subjects were presented with the faces with name labels and the other half with no name labels. Pairs of faces that belong to the same identity were given the same name labels. During the test phase the same identities (with a different facial expression) were presented in a simultaneously matching task.

contrast, changing another set of facial features did not change the identity of the face, and we therefore named them *non-critical features* (Abudarham, Shkiller, & Yovel, 2019; Abudarham & Yovel, 2016, 2018) (Figure 3A and Figure S1 and S2). In the learning phase, we presented pairs of images of the same identity that differed in critical features or in non-critical features. These face images were presented with name labels (supervised) to one group of participants and without labels (unsupervised) to another group of participants (Figure 3B shows only faces that differ in critical features). Following the learning phase, pairs of faces were presented side by side in a matching task without names and participants were asked to decide whether they are same or different people. We expect that when faces are learned without labels (unsupervised learning), they will be judged based on their perceptual similarity – that is, faces that differ in critical features will be judged as different identities and faces that differ in non-critical features will be judged as same identity faces. Here we examined whether perceptually different faces (i.e. faces that differ in critical features) that are learned with the same name label will be linked to the same identity.

Methods

Study 1

Pre-registration of methods and predictions of this study can be found here <http://aspredicted.org/blind.php?x=6ti4u4>.

Participants

A total of 50 psychology first year students participated in the experiment for course credit (12 males, mean age = 24). Half of the participants were assigned to the No name condition and the other half to the Name condition. The sample size was determined based on our previous face recognition studies that have used similar sample size (Schwartz & Yovel, 2016, 2018, 2019). The ethics committee of Tel Aviv

University approved the experiment and all participants signed a written informed consent.

Stimuli

Face stimuli were 45 color images of Caucasian adult male faces with neutral expression and 45 similar identities with a slight smile (See an example in Figure S3). The images were created with the FaceGen Modeller software, generating 3-dimensional faces with measurable features from 2-dimensional photographs. The images that were used to create the avatars were 15 different identities of unfamiliar male faces, originally taken from the color FERET database (Phillips, Martin, Wilson, & Przybocki, 2000), which were used by Abudarham & Yovel (2016) (Experiment 4).

The face stimuli included 15 different identities, 15 additional images of these identities with 5 critical features changes, and another 15 images of these identities with 5 non-critical features changes (see Figure 3 and supplementary Figure 1). The learning phase included 15 pairs of faces, in which 8 or 7 pairs were an original face and its corresponding identity that differ in critical features, and the other 7 or 8 remaining identities were an original face and its corresponding identity that differ in non-critical features. The pairs of identities that differ in critical or non-critical features were counterbalanced across participants. Pairs of faces that belong to the same identity were always presented one after the other.

The stimuli appeared at the center of the screen on a gray background. Labels were 15 common Israeli male names, written in Hebrew in a black, Times New Roman font above the face.

Procedure

Study phase

Participants went through a familiarization session, where they learned two different images of 15 face identities. Half of the participants were presented with name labels above the faces and the other half were presented with no labels. The name labels were randomly associated with face identities across participants. Each identity appeared ten times during the study phase – such that overall, the study phase included 150 trials.

Out of the 15 learned identities, half included the original and its modified version in which Critical features were changed, and the rest included the original and its modified version in which Non-critical features were changed. An identity learned with the Critical feature changed image in one version of the experiment, was learned with

the Non-critical feature changed image in the other version, and vice-versa. Each participant was randomly assigned to one of these two versions. Additionally, for half of the participants all faces in the study phase appeared with a neutral expression, and for the rest the same identities appeared with a smiling expression.

Each face image appeared on the screen sequentially for 1s. Each pair of faces of the same identity appeared one after the other. Participants were informed when reading the task instructions, that two faces of the same identity may look perceptually different and were given a real example of such a case (similar to Figure 1). In the *Name* condition, the name appeared above the face for the same duration and participants were asked to memorize the faces and names, as they will be tested on them later. In the *Name* condition, participants were told that face images that are presented with the same name belong to the same person. Such information was not given in the *No name* condition.

Test phase

During the test phase, a total of 30 face pairs were presented side by side, 15 same identity pairs (with either Critical or Non-critical feature changes as presented during the learning face) and 15 different identity pairs. Participants were asked to determine if the two images belong to the same identity or different identities by pressing one key for Same and another key for Different. Following each response, the participants were asked to rank how confident they are in their answer, on a scale of 1 (a little) to 5 (very much).

The facial expression of the faces presented in the study phase were different from the test phase. Half of the participants were presented with neutral expression faces in the learning phase and smiling in the test phase, and vice versa for the other half of the participants. This was done to avoid repeating exactly the same image during the study and test phase, such that performance reflects identity-based learning rather than image-based learning.

Data Analysis

Trials in which response times were longer than 10 seconds or shorter than 200 ms were excluded prior to analysis. There were no responses shorter than 200 ms. A total of 3.8% of the responses were longer than 10 seconds.

Results

Accuracy

We computed the proportion of correct responses for same identity pairs, for faces that differed in Critical and Non-critical features for each participant, and compared performance between participants that were presented with faces with or without labels during the learning phase. When faces were presented without name labels, we found a higher proportion of *same* responses for face pairs that differ in non-critical features than for faces that differ in critical features (Figure 4). These results confirm the feature manipulation that we performed, indicating that changing critical features moves the face outside of its identity subspace, whereas changing non-critical features moves the face within the identity subspace. Interestingly, when faces that differ in critical features were presented with the same name label, proportion of *same* responses increased significantly from a mean of 27% for faces presented without labels to 69% *same* responses for faces presented with labels, indicating that supervised learning can link faces from different identity subspaces to the same identity.

A mixed-ANOVA with label (Name/No name) as a between subject factor and type of feature change (Critical, Non-critical) as a within subject factor, revealed a significant interaction of Label and Feature Type [$F(1,48) = 11.73, p < .001, \eta^2_p = 0.2$]. The effect of Name was significant for Critical feature pairs ($t(48) = 6.04, p < .0001$, Cohen's $d = 1.7$), but not for Non-critical feature pairs ($t(48) = 1.51, p = .14$, Cohen's $d = 0.42$). Proportion of *same* responses were significantly higher for pairs of faces that differ in Non-critical features than Critical features, when faces were presented without names ($t(24) = 7.03, p < .001$, Cohen's $d = 1.4$) and also when they were presented with names, but to a smaller degree ($t(24) = 3.1, p < .005$, Cohen's $d = 0.61$).

We then ran the same analysis for matching different face pairs. We expected no effect of Label or Feature type for different identity faces. Indeed, no significant main effects of Feature Type ($F(1,48) = .6, p = 0.44, \eta^2_p = .012$), Label ($F(1,48) = 3.39, p = .07, \eta^2_p = 0.07$) and no interaction ($F(1,48) = 2.22, p = 0.14, \eta^2_p = .04$) were found for different identity pairs.

Figure 4: Performance for same identity pairs in a face matching task (Study 1). Pairs of same identity faces that differ in critical features (green) or non-critical features (yellow) were presented during learning either without or with name labels. When learned without names, faces that differ in critical features are perceived as different identities and faces that differ in non-critical features are perceived as same identity. When learned with names, proportion of same responses significantly increased for perceptually different faces, indicating that conceptual learning can link perceptually different faces to the same identity.

Reaction Times

Mean reaction times (RTs) were computed for each condition for correct responses only of same identity pairs. Seven subjects did not make any same response in the No name condition for faces that differ in Critical features, and were therefore not included in the analysis. This analysis therefore underestimates the effect of label for matching faces that differ in Critical features, but still shows the same pattern that we found with accuracy data. A mixed ANOVA with label (Name/No name) as a between subject factor and type of feature changes (Critical, Non-critical) revealed significantly faster RTs for pairs of faces that differ in Non-critical features than in Critical features ($F(1,41) = 5.28$, $p = 0.03$, $\eta^2_p = .11$). There were neither significant effect of label nor an interaction between Label and face type (see Table 1B for descriptive of RT data).

Confidence rating

Confidence rating was averaged across correct trials for each of the conditions for same faces. Similar to the reaction time data, the 7 subjects who did not make any correct response on same face pairs that differ in Critical features in the no-name condition were excluded from the analysis. A mixed ANOVA label (Name/No name) as a between subject factor and type of feature changes (Critical, Non-critical) revealed a higher confidence rating for matching pairs of faces that differ in Non-critical features than Critical features ($F(1,41) = 23.29$, $p < 0.001$, $\eta^2_p = .36$) and a marginally significant higher confidence for faces that were presented with than without names ($F(1,41) = 2.48$, $p = 0.053$, $\eta^2_p = .09$) (see Table 1C for descriptive of confidence rating data).

Table 1: Results of Study 1

A. Accuracy (Proportion Correct)

	Same Pairs				Different Pairs			
	No Name		Name		No Name		Name	
	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical
Mean	0.278	0.783	0.687	0.867	0.829	0.761	0.718	0.74
Std.	0.276	0.232	0.196	0.155	0.137	0.165	0.183	0.174
Minimum	0	0.14	0.29	0.57	0.5	0.5	0.25	0.38
Maximum	0.88	1	1	1	1	1	1	1

B. Reaction Time (ms.)

	Same Pairs				Different Pairs			
	No Name		Name		No Name		Name	
	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical
Mean	4274	3754	3485	2999	3502	3609	3602	3609
Std.	1993	1695	1252	1325	1141	1138	1311	1138
Minimum	1759	1547	1380	1177	1785	1421	1476	1421
Maximum	9616	8390	6070	7338	6329	6290	5858	6290

C. Confidence rating (1-5)

	Same Pairs				Different Pairs			
	No Name		Name		No Name		Name	
	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical	Critical	Non-Critical
Mean	3.4	4.1	3.8	4.2	4.2	4.0	3.7	3.9
Std.	0.7	0.6	0.6	0.7	0.5	0.5	0.7	0.6
Minimum	2.0	3.0	2.5	2.4	3.2	2.8	2.2	2.5
Maximum	5.0	5.0	5.0	5.0	5.0	5.0	5.0	5.0

Discussion

Results of Study 1 show that when faces were learned without labels, faces that differ in critical features were judged as different identities, whereas faces that differ in non-critical features were judged to belong to the same identity. These results confirmed our feature changing manipulation (Abudarham & Yovel, 2016) and show that when faces are learned with no conceptual labels, participants based their judgements on the perceptual similarity of the faces. Importantly, learning these same faces with name labels (supervised learning) can link faces that reside in different identity subspaces to the same identity, as indicated by a sharp increase in the proportion of *same identity* responses (Figure 4), almost to the level of faces that reside in the same identity subspace. These findings show that conceptual learning can link perceptually different faces to the same identity. However, the more interesting question is what is the nature of the representation that was generated for the learned identity. Study 2 was designed to answer this question and in particular to test the predictions of the *perceptual single identity subspace* and the *conceptual multiple identity subspace* hypotheses.

Study 2

To decide between the two hypotheses, in Study 2 we generated an averaged image between each original face and its changed version (Figure 5). The original face images and their changed versions, that differed in critical or non-critical features, were learned with labels in a face memory task, in the same way as was done in Study 1 (Figure 3). During the test phase, these learned faces were presented one at a time, together with their averaged images, which were not presented during the study phase, in an old-new face recognition task. We expect that the representation for faces that reside within the same identity subspace before learning, that is faces that differ in non-critical features, will also include their averaged image, despite the fact that it was not presented during learning. The two hypotheses make different predictions for faces that differ in critical features and reside in different identity subspaces (red dot in Figure 2): the *perceptual single identity subspace hypothesis* predicts that their averaged image will be recognized as an old face, similar to the learned faces, since they reside within the new boundaries of the same identity subspace; the *conceptual multiple identity subspace hypothesis* predicts that the averaged image is less likely to be recognized as an old face relative to the learned images, since it resides outside of their identity subspaces. We can therefore decide between these two hypotheses

by measuring the proportion of old responses to the averaged face of the faces that differ in non-critical features.

Methods

Pre-registration of methods and predictions of this study can be found here <http://aspredicted.org/blind.php?x=83eh3z>

Participants

A total of 25 psychology first year students participated in the experiment for course credit (7 males, mean age = 23). The sample size was determined based on our previous face recognition studies that have used similar sample size. The ethics committee of Tel Aviv University approved the experiment and all participants signed a written informed consent.

Stimuli

In addition to the stimuli presented in Study 1, we generated an average face for each identity and its corresponding Critical feature and Non-critical feature changed image (Figure 5) with FaceGen Modeler Tween tool, by setting the scale to the middle between original and its changed face. Thus, two average faces per identity were created by morphing the original image with the Critical image and with the Non-critical image of the same identity. For all the average images, a similar image with a smiling expression was also created. Thirty additional face identities were included as novel faces and were only presented in the test phase.

Validation of the average face stimuli: To assure that the average faces that were generated by FaceGen, in particular for faces that differ in critical features, were indeed perceived as the average of the original and the changed face, 10 participants were presented with the 3 images (Figure 5, top) of each of the 15 identities. The 3 faces of each identity were presented each on a separate screen - in one of 3 locations in a triangular arrangement. Participants were asked to move them on the screen with a mouse and arrange them along an imaginal horizontal line according to their perceptual similarity, such that the most dissimilar faces will be located at the ends of the line and the face that is similar to both faces, will be located in between them. Out of the 10 participants, 7 participants located all the average faces in between the original and the changed face for all 15 identities, 1 participant generated a different arrangement for 5 out of the 15 identities, 1 participant generated a different arrangement for 2 out of the 15 identities and one participant did not arrange any of the face triads correctly and therefore we assumed that they failed to understand the

instructions. These findings assure that overall the average faces that were generated based on faceGen parameters are also perceived as averages based on human perception.

Learning Phase

The Learning phase was identical to the learning phase with name labels used in Study 1.

Test Phase

The test phase was a recognition memory task, which included 75 trials (30 learned exemplars of 15 identity pairs (an original face and its modified version in which either Critical or Non-critical features were changed in two different version of the experiment), 15 never-before-seen average face images of these pairs, and 30 novel identity face. The face images appeared on the screen for 1000 ms. without the name labels. The faces presented in the test phase had a different expression from the ones seen in the study phase, such that participants who learned faces with a neutral expression were tested on smiling faces and vice-versa. The rationale was to assure that learning is not image-based, but is generalized to other images of the learned identity. The faces appeared in a pseudo-randomized order, to assure that two faces of the same identity are not presented consecutively.

Figure 5: Averaged faces were generated for the pairs of faces that differed in Critical and Non-Critical faces using FaceGen. The original and changed faces were presented during the study phase with name labels. The averaged faces were not presented during the study phase but only during the test phase.

Participants were asked to judge as quickly and accurately as possible for each face whether it was old (presented at study) or new, using one key for Old and another key for New. Following each response, the participants were asked to rank how confident they were in their decision, on a scale of 1 (not confident) to 5 (very confident).

Data Analysis

Trials in which response times were longer than 10 seconds or shorter than 200 ms. were excluded prior to analysis. There were no responses shorter than 200 ms. and only two responses that were longer than 10 seconds.

Results

Accuracy

We computed the proportion of *old* responses for original and their changed versions (learned faces) that were presented during the learning phase, as well as for their average face images (unlearned faces), that were not presented during the learning phase. A repeated measure ANOVA with Face Type (Learned, Unlearned) and Feature Type (Critical, Non-critical) as within-subject factors, revealed a significant interaction ($F(1,24) = 23.39$, $p < .001$, $\eta^2_p = 0.49$) of Face Type and Feature Type, indicating a very different pattern of response for faces that differ in critical and non-critical features. Post hoc tests, corrected for multiple comparisons revealed no difference in proportion of *old* responses for the learned and the averaged, unlearned faces ($t(24) = 0.22$, $p = 1$, $\text{cohen's } d = 0.05$) for faces that differed in non-critical features. For the critical feature changes there was a much lower proportion of *old* response for the unlearned, averaged faces than the learned faces ($t(24) = 5.86$, $p < 0.001$, $\text{Cohen's } d = 1.3$) (see Figure 6, Table 2A), consistent with the *conceptual multiple-identity subspace* hypothesis.

Proportion of correct responses for New faces was 89%, indicating that the participants' recognition performance was overall high.

Figure 6: Proportion of Old responses in the face recognition task (Study 2) for the learned faces (Original and Changed) and their averaged image (unlearned) for faces with Critical (Left) and Non-Critical (right) feature changes. The learned representation of faces that differ in critical features and are linked to the same identity do not include the average face. The learned representation of faces that differ in non-critical features does include the average face.

Reaction Time

A repeated measure ANOVA with Face Type (Learned, Unlearned) and Feature Type (Critical, Non-critical) as within-subject factors revealed no significant main effects or interactions. (see Table 2B)

Confidence Rating

A repeated measure ANOVA with Face Type (Learned, Unlearned (Average)) and Feature Type (Critical, Non-critical) as within-subject factors revealed a significant effect of feature type with higher confidence for Non-critical than Critical faces ($F(1,23) = 24.79$, $p < .001$, $\eta^2_p = 0.52$), and a significant interaction between Face Type and Feature type ($F(1,23) = 8.6$, $p = .01$, $\eta^2_p = 0.27$) that is similar to the pattern that was found for accuracy, with lowest confidence for the average face of the Critical pairs than the learned Critical faces and similar confidence for the average face of the Non-critical pairs and the learned Non-critical faces (see Table 2C).

Table 2: Results of Study 2

A. Accuracy: Proportion Old				
	Non-Critical		Critical	
	Learned	Unlearned	Learned	Unlearned
Mean	0.77	0.77	0.73	0.53
Std.	0.15	0.16	0.18	0.16
Minimum	0.50	0.38	0.38	0.25
Maximum	1.00	1.00	1.00	0.88

B. Reaction Time (Correct Old Response)				
	Non-Critical		Critical	
	Learned	Unlearned	Learned	Unlearned
Mean	1446	1442	1569	1503
Std.	387	447	421	428
Minimum	921	876	905	921
Maximum	2258	2528	2590	2221

C. Confidence Rating (1-5)				
	Non-Critical		Critical	
	Learned	Unlearned	Learned	Unlearned
Mean	4.2	3.8	4.0	4.2
Std.	0.5	0.6	0.5	0.5
Minimum	2.8	2.7	3.1	2.8
Maximum	5.0	5.0	5.0	5.0

Discussion

In Study 1 we found that faces that are learned with the same conceptual label can be linked to the same identity even if they are perceptually different. In Study 2 we examined what is the nature of the representation of the learned identity. In particular, we asked whether the perceptual boundary of the learned Identity is expended such that the perceptually different faces are included in the same identity subspace. In that case, we expect that an average face will be included in the single identity subspace of the learned identity (Figure 2, left). However, if perceptually different faces remain in separated identity subspaces, then the average face will fall outside of the identity subspace and will not be included in the representation of the learned identity (Figure 2, right). In this case, the perceptually different faces are linked conceptually rather than perceptually. We therefore named this alternative the *conceptual multiple identity subspace* hypothesis.

Results of Study 2 are consistent with the *conceptual multiple identity subspace* hypothesis. We found that despite the fact that participants learned to link perceptually different faces (i.e. faces that differ in critical features) to the same identity following conceptual/supervised learning (Study 1), the learned representation did not include their average face. The average faces were generated with a face generation software (FaceGen) and we therefore assured that they were indeed perceived as average faces of the learned faces in a perceptual judgement task performed by a different group of subjects.

These findings are therefore inconsistent with the generation of a single, wide identity subspace representation following conceptual learning. In contrast, the representation of face images that reside in the same identity subspace prior to learning (i.e. faces that differ in non-critical features), does include the average face of the learned identity, even though it was never presented during learning. Proportion of Old responses to the average, unlearned face was similar to the proportion of Old responses to the learned faces that differ in non-critical features. These findings indicate that supervised learning may expand the representation of face identity across separate identity subspaces, without generating a single identity subspace. Instead, the labels that are used for linking the faces to the same identity during learning, may link the perceptually different faces based on their conceptual similarity rather than based on perceptual similarity.

The role of labeling in face recognition has been shown in previous studies that reported that studying faces with different labels improves face recognition relative to faces that are presented with no labels or with the same label (Gordon & Tanaka, 2011; Schwartz & Yovel, 2016; Yovel et al., 2012). These studies used similar face images during study and test and attributed the effect of labeling to individuation, which encourages discrimination between different identities that are linked to different labels. Schwartz & Yovel (2016) further showed that this type of labeling is primarily beneficial for person-related labels (name, occupation) but not for person-unrelated labels (symbols, object names). These findings further highlight that important contribution of conceptual learning. not only for generalization across different appearances of the same identity as was demonstrated here, but also for individuation/discrimination of different identities.

The single and multiple identity subspace hypotheses studied here, may be similar to the prototype vs. exemplar-based representation, that were extensively

studied in the literature of category learning (e.g., Minda & Smith, 2012; Nosofsky & Zaki, 2002; Reed, 1972). The *perceptual single identity subspace representation* hypothesis is consistent with the prototypical representation, in that all exemplars of the same category share the same representation. The *conceptual multiple representation* hypothesis suggests an intermediate representation between the prototype and the exemplar-based representation (see also Love et al., 2004; Vanpaemel & Storms, 2008), in that it proposes a single representation for perceptually similar images, but separated representations for perceptually dissimilar images. This type of intermediate representation is therefore more economical than a representation that is based on separate representations for the many different appearances of the same category, as an extreme version of an exemplar-based representation proposes (e.g., Nosofsky & Zaki, 2002).

The conceptual multiple subspace representation assumes a linear transformation between the perceptual and the conceptual representations, such that perceptually different faces remain in distant locations. Another alternative that is a non-linear transformation from the perceptual to the conceptual representation such that the perceptually different subspaces are moved closer together according to their conceptual similarity and generate an identity-based representation. The current findings cannot rule out this type of identity-based conceptual representation. However, regardless of the nature of the conceptual representation, we propose that conceptual information plays a critical role in linking perceptually different faces to the same identity during learning and recognition. This conceptual learning generates an identity-based conceptual representation that links all the information that is unique to the familiar identity including its different perceptual appearances as well as all the non-facial knowledge (e.g., voice, body, semantic, emotional, episodic). We further discuss the nature of this conceptual representation in the general discussion.

General discussion

The main goal of the current study was to account for human's superb classification humans show for familiar faces. In particular, we asked what type of learning and what is the nature of the representation that enable humans to generalize across highly dissimilar faces of a familiar identity (Jenkins et al., 2011; Ritchie et al., 2015) (see Figure 1). Whereas previous studies have primarily emphasized the importance of exposure to high variability appearances of the same identity of a familiar person (Andrews et al., 2015; Kramer et al., 2017, 2018; Ritchie & Burton,

2017), here we highlight that this perceptual experience should be accompanied with conceptual information to enable the remarkable generalization that is reported for familiar identities. Accordingly, we propose that in order to link perpetually different faces that are perceived as different identities prior to learning (unfamiliar) to the same identity, faces need to be learned with the same conceptual label. In addition, this conceptual learning generates a conceptual representation that links perceptually different appearances to the same identity based on their conceptual similarity.

More broadly we propose that the face recognition system generates two different representations for face identity: an identity-general view-invariant perceptual

Figure 7: The face recognition system generates a feature-based, view-invariant perceptual representation that enables limited generalization across different head-views/expressions/illuminations for all faces. Generalization across perceptually different faces is enabled only for familiar faces based on conceptual (supervised) learning that links perceptually different images of the same identity to the conceptual representation of the familiar person. In this example we show 3 different images of the same identity of an Israeli celebrity (Ninet), two that are perceptually similar and one that is perceptually different.

representation that is used for all faces before the system determines whether they are familiar or not and an identity-based view-invariant conceptual representation in memory that is generated only for familiar faces (Figure 7). The perceptual representation enables generalization across relatively constrained variations of head views, expressions and illuminations. This view-invariant perceptual representation relies on view-invariant facial features that are shared for all faces such as the critical features we used in the current study (see also Abudarham & Yovel, 2016). The three

faces presented in Figure 7 all belong to the same identity but only two of them can be linked based on their perceptual similarity. If these faces are learned with the same conceptual label, this perceptual representation is then converted to an identity-based conceptual representation in memory. The conceptual identity-based representation links the different learned appearances of the familiar identities, as well as all other non-facial information that we have about that person. Here we suggest that it is this conceptual representation that enables generalization across perceptually different faces of the same identity for familiar but not unfamiliar faces.

The view-invariant perceptual representation that is described here, may be reminiscent of the structural code that was proposed by Bruce and Young (1986). However, Bruce and Young suggested a different structural code for familiar and unfamiliar faces. Here we propose that the view-invariant perceptual representation is generated for all faces regardless of their familiarity. This suggestion is based on our recent findings that show that the same view-invariant critical features are used for both familiar and unfamiliar faces (Abudarham et al., 2019). We therefore propose that the significant gap in performance between familiar and unfamiliar faces is mediated by the conceptual representation of the familiar face, which may correspond to Bruce and Young's face recognition unit (FRU), which unlike the structural code, is not purely perceptual, but links the perceptual and the conceptual knowledge about the familiar person.

Our proposal that the difference between familiar and unfamiliar faces emerges at the conceptual rather than the perceptual level is consistent with findings from neuroimaging and single unit recording studies of the primate face processing system. In a monkey single unit recording study, Chang and Tsao (2017) showed that neurons in the anterior face patch (AM), which is located at the end of the hierarchy of the visual face areas, generate a feature-based, view-invariant perceptual representation of face identity that enables generalization across different appearances of the same identity for relatively constrained unfamiliar faces (Chang & Tsao, 2017; Freiwald & Tsao, 2010). Importantly, however, a monkey fMRI study by Landi and Freiwald (2017) showed that familiar and unfamiliar faces generated different representations only in the temporal pole and perirhinal cortex, beyond high-level visual face areas, where perceptual and conceptual information are linked and memories are created. There was no evidence that the perceptual face areas, including the most anterior one (AM), generated different representations to familiar and unfamiliar faces (Landi & Freiwald,

2017). These findings are also consistent with humans' fMRI studies that overall show similar response of the ventral visual face areas (e.g., fusiform face area - FFA) to familiar and unfamiliar faces. The main difference between familiar and unfamiliar faces is typically found in non-visual, conceptual brain areas (Gobbini & Haxby, 2006, 2007; Natu & O'Toole, 2011). These findings are also consistent with single unit recordings from human medial temporal lobe that revealed single neurons that are highly selective to familiar identities. These neurons, known as concept cells, respond to highly variable images of the familiar person, even images in which the facial features are occluded but the identity can be still determined based on non-facial person knowledge (e.g., Halle Barry with a cat mask), as well as the name of the familiar person (Quiroga, Kreiman, Koch, & Fried, 2008; Quiroga, Reddy, Kreiman, Koch, & Fried, 2005). Finally, recent computational modelling that use deep convolutional neural network to account for the differences in performance for familiar and unfamiliar faces also supports such separation between an identity-general view-invariant perceptual representation used for all faces and an identity-based view-invariant conceptual representation used for familiar faces (Blauch, Behrmann, & Plaut, 2020). These findings are consistent with our suggestion that generalization of highly dissimilar images of the same identity is achieved only at the conceptual, identity-based representation, not at the preceding, view-invariant perceptual representation.

The effect of conceptual learning on the perceptual boundary of a category, has been investigated in previous studies of category learning for different types of categories (e.g., phonemes, sounds, objects, faces). These studies typically examine how learning visual stimuli with different labels determine the boundary between the learned categories. In particular, associating different labels to perceptually similar images, improves perceptual discrimination relative to similar labels or no labels (Gauthier, James, Curby, & Tarr, 2003; Goldstone, 1994; Goldstone, Lippa, & Shiffrin, 2001). The current findings are consistent with these findings and further emphasize the importance of studying the combined contribution of perceptual and conceptual information to the representation of the learned categories/identities. We further propose that when a system learns a new category of stimuli, it attempts to discover the critical features that are common to all stimuli to enable generalization and discrimination of the different appearances of the sub-categories. However, in many cases specific sub-categories/identities may vary in a unique manner that cannot be

generalized to other categories. In that case, a conceptual identity-based representation, in which the different unique variations are linked to the same conceptual label, enable conceptual generalization of perceptually different appearances. This representation is specific to each identity and is not generalized to other identities.

We conclude that resolving the challenging task of familiar face recognition, requires supervised learning with large variability of images and the generation of a conceptual representation for each familiar identity. This representation is specific for the learned identities as is indeed the case for familiar face recognition that does not generalize to unfamiliar faces. This proposed conceptual account for familiar face recognition may not be limited to familiar faces, but can account for any other familiar category that includes stimuli of higher variability within category than between categories, such as familiar voices (Lavan, Burston, & Garrido, 2018), musical tunes or objects of expertise (Harel et al., 2013).

References

- Abudarham, N., Shkiller, L., & Yovel, G. (2019). Critical features for face recognition. *Cognition*, 182. <https://doi.org/10.1016/j.cognition.2018.09.002>
- Abudarham, N., & Yovel, G. (2016). Reverse engineering the face space: Discovering the critical features for face identification. *Journal of Vision*, 16(3). <https://doi.org/10.1167/16.3.40>
- Abudarham, N., & Yovel, G. (2018). Same critical features are used for identification of familiarized and unfamiliar faces. *Vision Research*. <https://doi.org/10.1016/j.visres.2018.01.002>
- Andrews, S., Jenkins, R., Cursiter, H., & Burton, A. M. (2015). Telling faces together: Learning new faces through exposure to multiple instances. *Quarterly Journal of Experimental Psychology*. <https://doi.org/10.1080/17470218.2014.1003949>
- Blauch, N., Behrmann, M., & Plaut, D. (2020). Computational insights into human perceptual expertise for familiar and unfamiliar face recognition. *Cognition*.
- Burton, A. M., Jenkins, R., & Schweinberger, S. R. (2011). Mental representations of familiar faces. *British Journal of Psychology*. <https://doi.org/10.1111/j.2044-8295.2011.02039.x>
- Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. *Cell*. <https://doi.org/10.1016/j.cell.2017.05.011>
- Freiwald, W. A., & Tsao, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science*. <https://doi.org/10.1126/science.1194908>
- Gauthier, I., James, T. W., Curby, K. M., & Tarr, M. J. (2003). The influence of conceptual knowledge on visual discrimination. *Cognitive Neuropsychology*. <https://doi.org/10.1080/02643290244000275>
- Gobbini, M. I., & Haxby, J. V. (2006). Neural response to the visual familiarity of faces. *Brain Research Bulletin*. <https://doi.org/10.1016/j.brainresbull.2006.08.003>
- Gobbini, M. I., & Haxby, J. V. (2007). Neural systems for recognition of familiar faces. *Neuropsychologia*. <https://doi.org/10.1016/j.neuropsychologia.2006.04.015>
- Goldstone, R. L. (1994). The role of similarity in categorization: providing a groundwork. *Cognition*. [https://doi.org/10.1016/0010-0277\(94\)90065-5](https://doi.org/10.1016/0010-0277(94)90065-5)
- Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations

- through category learning. *Cognition*. [https://doi.org/10.1016/S0010-0277\(00\)00099-8](https://doi.org/10.1016/S0010-0277(00)00099-8)
- Gordon, I., & Tanaka, J. W. (2011). *The role of name labels in the formation of face representations in event-related potentials*. 5, 884–898.
<https://doi.org/10.1111/j.2044-8295.2011.02064.x>
- Harel, A., Kravitz, D., Baker, C. I., Campitelli, G., Cowan University, E., & Elinor McKone, A. (2013). Beyond perceptual expertise: revisiting the neural substrates of expert object recognition WHAT IS EXPERTISE AND WHY IS IT IMPORTANT TO STUDY IT? *Frontiers in Human Neuroscience*.
<https://doi.org/10.3389/fnhum.2013.00885>
- Jenkins, R., White, D., Van Montfort, X., & Mike Burton, A. (2011). Variability in photos of the same face. *Cognition*.
<https://doi.org/10.1016/j.cognition.2011.08.001>
- Kramer, R. S. S., Jenkins, R., Young, A. W., & Burton, A. M. (2017). Natural variability is essential to learning new faces. *Visual Cognition*.
<https://doi.org/10.1080/13506285.2016.1242522>
- Kramer, R. S. S., Ritchie, K. L., & Burton, A. M. (2015). Viewers extract the mean from images of the same person: A route to face learning. *Journal of Vision*.
<https://doi.org/10.1167/15.4.1>
- Kramer, R. S. S., Towler, A., Reynolds, M. G., & Burton, A. M. (2018). *Familiarity and Within- Person Facial Variability: The Importance of the Internal and External Features*. <https://doi.org/10.1177/0301006617725242>
- Landi, S. M., & Freiwald, W. A. (2017). Two areas for familiar face recognition in the primate brain. *Science*. <https://doi.org/10.1126/science.aan1139>
- Lavan, N., Burston, L. F. K., & Garrido, L. (2018). How many voices did you hear? Natural variability disrupts identity perception from unfamiliar voices. *British Journal of Psychology*. <https://doi.org/10.1111/bjop.12348>
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). *SUSTAIN: A Network Model of Category Learning*. 111(2), 309–332. <https://doi.org/10.1037/0033-295X.111.2.309>
- Minda, J. P., & Smith, J. D. (2012). Prototype models of categorization: basic formulation, predictions, and limitations. In *Formal Approaches in Categorization*. <https://doi.org/10.1017/cbo9780511921322.003>
- Natu, V., & O'Toole, A. J. (2011). The neural processing of familiar and unfamiliar

- faces: A review and synopsis. *British Journal of Psychology*.
<https://doi.org/10.1111/j.2044-8295.2011.02053.x>
- Nosofsky, R. M., & Zaki, S. R. (2002). Exemplar and Prototype Models Revisited: Response Strategies, Selective Attention, and Stimulus Generalization. *Journal of Experimental Psychology: Learning Memory and Cognition*.
<https://doi.org/10.1037/0278-7393.28.5.924>
- Phillips, P. J., Martin, A., Wilson, C. L., & Przybocki, M. (2000). An Introduction to evaluating biometric systems. *Computer*. <https://doi.org/10.1109/2.820040>
- Quiroga, R. Q., Kreiman, G., Koch, C., & Fried, I. (2008). Sparse but not “Grandmother-cell” coding in the medial temporal lobe. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2007.12.003>
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*.
<https://doi.org/10.1038/nature03687>
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*.
[https://doi.org/10.1016/0010-0285\(72\)90014-X](https://doi.org/10.1016/0010-0285(72)90014-X)
- Ritchie, K. L., & Burton, A. M. (2017). Learning faces from variability. *Quarterly Journal of Experimental Psychology*.
<https://doi.org/10.1080/17470218.2015.1136656>
- Ritchie, K. L., Smith, F. G., Jenkins, R., Bindemann, M., White, D., & Burton, A. M. (2015). Viewers base estimates of face matching accuracy on their own familiarity: Explaining the photo-ID paradox. *Cognition*, 141, 161–169.
<https://doi.org/10.1016/j.cognition.2015.05.002>
- Schwartz, L., & Yovel, G. (2016). The roles of perceptual and conceptual information in face recognition. *Journal of Experimental Psychology: General*, 145(11).
<https://doi.org/10.1037/xge0000220>
- Schwartz, L., & Yovel, G. (2018). Learning Faces as Concepts Rather Than Percepts Improves Face Recognition. *Journal of Experimental Psychology: Learning Memory and Cognition*. <https://doi.org/10.1037/xlm0000673>
- Schwartz, L., & Yovel, G. (2019). Independent contribution of perceptual experience and social cognition to face recognition. *Cognition*, 183.
<https://doi.org/10.1016/j.cognition.2018.11.003>
- Tanaka, J. (1998). Mapping attractor fields in face space: the atypicality bias in face recognition. *Cognition*. [https://doi.org/10.1016/s0010-0277\(98\)00048-1](https://doi.org/10.1016/s0010-0277(98)00048-1)

- Valentine, T. (1999). Face-Space Model of Face Recognition. In *Computational, geometric, and process perspectives of facial cognition: Contexts and challenges*.
- Vanpaemel, W., & Storms, G. (2008). In search of abstraction: The varying abstraction model of categorization. *Psychonomic Bulletin and Review*, 15(4), 732–749. <https://doi.org/10.3758/PBR.15.4.732>
- Young, A. W., & Burton, A. M. (2017). Recognizing Faces. *Current Directions in Psychological Science*. <https://doi.org/10.1177/0963721416688114>
- Young, A. W., & Burton, A. M. (2018). Are We Face Experts? *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2017.11.007>
- Yovel, G., Halsband, K., Pelleg, M., Farkash, N., Gal, B., & Goshen-Gottstein, Y. (2012). Can massive but passive exposure to faces contribute to face recognition abilities? *Journal of Experimental Psychology: Human Perception and Performance*, 38(2). <https://doi.org/10.1037/a0027077>

Figure S1:

In previous studies (Abudarham and Yovel, 2016, Abudraham et al, 2019), we discovered a subset of five features that changing them changed the identity of the face and we therefore named them Critical features. In contrast, changing a set of other five features did not change the identity of the face and we therefore named them Non-critical features (see also Figure S2). This figure shows an example of one familiar identity (left) and one unfamiliar identity (right) that were changed by replacing five critical features (top) or five non-critical features (bottom). The FaceGen images used in this paper were generated based on the set of 15 unfamiliar faces that were changed in the same way (Abudarham and Yovel, 2016). We used FaceGen images rather than these original images, in order to generate the average faces that were presented in Study 2.

Figure S2:

This figure shows the procedure of changing five critical and five non-critical features sequentially for the familiar identity (Ben Affleck) shown in Figure S1 (for more information see Abudarham & Yovel, 2016, Abudarham et al., 2019).

Figure S3:

For each identity we created two versions that differed slightly in facial expression, so the same exact image does not repeat at study and test. The figure shows an example of the neutral and smiling version of one identity – for the original, critical feature changed and non-critical feature changes.