

# Exemplar learning reveals the representational origins of expert category perception

Elliot Collins<sup>a,b</sup> and Marlene Behrmann<sup>a,1</sup> 

<sup>a</sup>Department of Psychology and Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA 15213; and <sup>b</sup>School of Medicine, University of Pittsburgh, Pittsburgh, PA 15260

Contributed by Marlene Behrmann, March 2, 2020 (sent for review July 25, 2019; reviewed by James W. Tanaka and Alan C.-N. Wong)

**Irrespective of whether one has substantial perceptual expertise for a class of stimuli, an observer invariably encounters novel exemplars from this class. To understand how novel exemplars are represented, we examined the extent to which previous experience with a category constrains the acquisition and nature of representation of subsequent exemplars from that category. Participants completed a perceptual training paradigm with either novel other-race faces (category of experience) or novel computer-generated objects (YUFOs) that included pairwise similarity ratings at the beginning, middle, and end of training, and a 20-d visual search training task on a subset of category exemplars. Analyses of pairwise similarity ratings revealed multiple dissociations between the representational spaces for those learning faces and those learning YUFOs. First, representational distance changes were more selective for faces than YUFOs; trained faces exhibited greater magnitude in representational distance change relative to untrained faces, whereas this trained–untrained distance change was much smaller for YUFOs. Second, there was a difference in where the representational distance changes were observed; for faces, representations that were closer together before training exhibited a greater distance change relative to those that were farther apart before training. For YUFOs, however, the distance changes occurred more uniformly across representational space. Last, there was a decrease in dimensionality of the representational space after training on YUFOs, but not after training on faces. Together, these findings demonstrate how previous category experience governs representational patterns of exemplar learning as well as the underlying dimensionality of the representational space.**

visual expertise | perceptual learning | object recognition | category learning | mental representations

**R**esearch in vision science has uncovered several visual categories for which individuals evince highly skilled perception, including that of birds (1–3), dogs (4), cars (5), radiologic images (6), geological specimens (1), and computer-generated novel stimuli, such as greebles (7, 8), to name a few. The standard bearer for highly skilled visual perception, however, has been the category of faces for which observers demonstrate remarkable feats of recognition and individuation.

Several signatures of fine-grained visual perception have emerged as a result of these investigations of perceptual expertise and are used as diagnostic markers of visual skills. One such signature is the extent to which performance is adversely impacted by inversion (for example, 180° picture-plane rotation). That experts are more affected by inversion than novices is assumed to result from the increased sensitivity of the visual system to the typical arrangement of features in experts. Another characteristic of expert perception is that the discrimination of exemplars within the expert category is as good as the discrimination of exemplars between categories (2). This suggests that representations of individual objects of the expert category, even if similar visually, are easily encoded and matched to their stored counterparts: For example, whereas experts retrieve subordinate- and basic-level labels in their domain of expertise with

equal speed, they show the standard basic-level advantage outside this domain (3). A final criterion for ascribing expert-level competency concerns the ability to generalize across inputs. That is, individuals can rapidly discriminate or even individuate new exemplars that fall within the distribution of a previously established expert category, with little cost to processing efficiency (9).

The underlying theoretical assumption that brings together these different behavioral assays of expertise is the existence of a representational space in which exemplars are sufficiently differentiated to permit highly skilled perception. Face Space (10) is perhaps the best-established theory of a psychological similarity space. In Face Space, dimensions correspond to the different visual features across which faces vary, although these dimensions need not map onto easily verbalizable features (11). The Euclidean distance between unique exemplars within this space corresponds to perceptual similarity with typical faces occupying the centroid of the space and disproportionately atypical faces situated closer to the periphery of the space. Many of the predictions made by Face Space theory have been supported empirically (12), providing overwhelming support for the existence of such a psychological similarity space in visual processing.

## Significance

**Vision science has uncovered several behavioral signatures of highly skilled perception. Subsequent inferences assume that well-differentiated object representations permit fine-grained visual discrimination. However, it remains unknown how representations develop and what governing principles drive changes in representational space more generally. Here, we explore how previous category experience constrains change in representational space during novel exemplar learning. These experiments implicate three types of representational change by which individuals develop visual expertise and the manner in which subsequent exposure within the same domain influences representations. Our findings demonstrate that previous category experience influences the selectivity, locality, and change in dimensionality of representational space that occurs during learning. Together, this evidence reveals a representational mechanism by which highly skilled visual perception emerges.**

Author contributions: E.C. and M.B. designed research; E.C. performed research; E.C. contributed new reagents/analytic tools; E.C. analyzed data; and E.C. and M.B. wrote the paper.

Reviewers: J.W.T., University of Victoria; and A.C.-N.W., Chinese University of Hong Kong.

The authors declare no competing interest.

Published under the [PNAS license](#).

Data deposition: Data and experimental materials are publicly available at DOIs [10.1186/R1/11869524](#), [10.1186/R1/11991642](#), and [10.1186/R1/11991636](#).

<sup>1</sup>To whom correspondence may be addressed. Email: [behrmann@cmu.edu](mailto:behrmann@cmu.edu).

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1912734117/-DCSupplemental>.

First published May 4, 2020.

In the same way that Face Space offers a unifying theory of face perception, one can use the principles of a representational space to develop a representational theory of visual expertise. Across the many behavioral assays of visual expertise, some of which are discussed above, is the assumption that representations are sufficiently separated in space to permit rapid differentiation or individuation of visual objects. Additionally, the specific dimensions of an established representational space are sufficiently well established such that they are not dramatically perturbed as an individual encounters new exemplars within the existing distribution of the expert category. Indeed, these dimensions are so well established that manipulation of these features or their arrangement, such as in face inversion, results in an outsized decrement in processing in the expert compared to the novice perceiver. The obvious first question to ask is how an “expert representational space” develops with experience. How do representational spaces, in general, change as individuals move along the continuum of experience from novice to expert? One might predict that given differing amounts of previous experience, subsequent learning would differentially impact representational space. For example, when encountering a novel visual category, one might expect representational changes to be nonspecific, as dimensions of the space have not yet been established. In contrast, in individuals with substantial previous experience, one might expect representational changes to be privileged to certain areas within the representational space so as to permit differentiation among the most similar exemplars. Quantification of the space itself in the context of novel exemplar learning might itself serve as a behavioral assay of visual expertise. An expert representational space might simply correspond to a sufficiently well-established space, whose dimensions permit generalization to novel exemplars, akin to the “intrinsic manifold” described in the domain of motor learning (13, 14).

In the present study, we interrogate the changes in representational spaces at two different points along the continuum of visual expertise during exemplar learning. Doing so permits us to test several of the assumptions of the visual expertise literature, as well as harness a variety of behavioral assays under a single representational theory of visual expertise. We examine these predictions by adopting a microgenetic approach to characterize and quantify changes in representational space prior to and after a perceptual training paradigm.

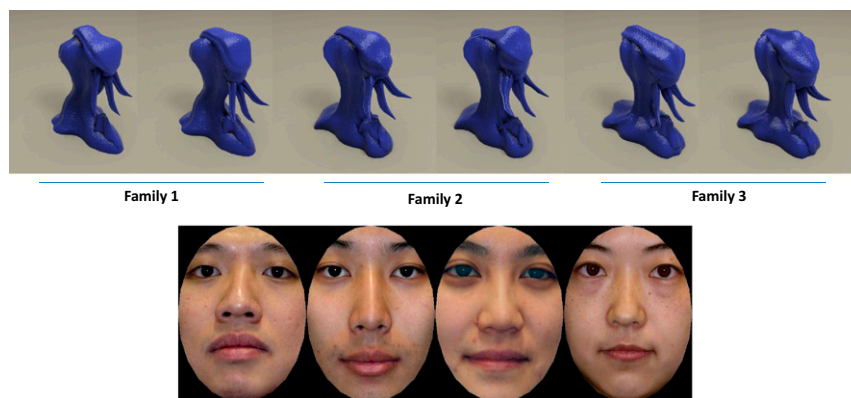
The elucidation of representational change has gained in popularity recently, especially in neuroscience studies that explore neural patterns evoked in response to specific visual stimuli. Representational similarity analyses (RSA), which quantify representational distances between any number of categories, conditions, or objects, have also become increasingly

popular (15, 16). RSA alone, and to an even greater degree when paired with multivariate decoding techniques, provide insights into the representational bases of stimuli and permit powerful representational comparisons across both visual categories and imaging techniques (17, 18).

Thus far, most studies that explore learning-induced representational or neural changes have compared changes between categories or conditions (8), rather than changes of individual exemplars across time. The few perceptual learning studies that have measured changes between individual stimuli typically train individuals on an entire set of stimuli and then report the results for the entire group of stimuli (19), averaging over individual stimuli, rather than comparing dynamic changes in representations of individual instances. Although these approaches are valuable, they do not necessarily reflect the type of perceptual learning in which observers encounter only a single new exemplar or two new exemplars at a time, as is the case for observers under more naturalistic conditions. As such, these existing approaches limit the extent to which one can detect changes in object-level representations in response to single, novel exemplars.

### The Present Study

In the present study, we adopt a behavioral representational approach in the context of a perceptual training paradigm in which individuals encounter and learn only a few exemplars at a given time, and in which the exemplars are drawn either from a category with which individuals have substantial experience (faces) or from a novel category (novel computer-generated objects [YUFOs]). We quantified the representational space of the trained and nontrained exemplars pre-, mid-, and post-learning. We hypothesized that exemplar learning has differential effects on representational space depending on the previous experience of the learner. Specifically, we predicted that exemplar learning within a familiar category will result in representational changes specific to the trained objects; because the representational space for an existing category is relatively stable and the dimensions well established, any changes are expected to be small and specific to the novel exemplars. In contrast, we predicted that learning exemplars within a novel category will drive more generalized changes in representational space in addition to changes specific to the object being learned. Finally, we hypothesized that, whereas exemplar learning would not substantially change the dimensionality of the preexisting representational space for faces, we predicted changes in the dimensionality of representational space for the YUFOs as the relevant dimensions along which exemplars vary have not yet been extracted.



**Fig. 1.** Stimuli for novel object (*Upper*) and ORF (*Lower*) experiments. The novel face image set consists of 30 young faces of East Asian descent. The novel object set (YUFOs) also contains 30 objects. In this figure, novel objects are not to scale with face stimuli.

Of note, for the familiar category, we implemented training with other-race faces (ORFs), rather than with own-race faces, to ensure that performance was not at ceiling (20) and that participants could still show improvements on the training task over subsequent sessions. Although ORFs are not necessarily a category of visual expertise per se, participants likely bring to bear their previous experience with own-race faces when confronted with novel ORFs. Furthermore, inversion effects have been documented during ORF perception, suggesting similarity in the visual processes of own-race faces (21, 22).

## Materials and Methods

**Participants.** This study was approved by the Institutional Review Board of Carnegie Mellon University and informed consent was obtained from all participants. Thirty-seven participants completed the study (female = 24; average age = 27, SD = 2.9) and were paid for their participation. Participants who completed the version of the study with ORFs were Caucasian, ensuring that individuals would not be making own-race face judgements (23). All participants reported that they had normal or corrected-to-normal vision.


### Stimuli.

**YUFOs.** We used computer-generated YUFOs (24) as they are well controlled for color, lighting, size, variability in shape, and alignment at multiple viewing angles and do not obviously resemble faces (Fig. 1). These novel objects have been used successfully in previous studies to advance our understanding of how humans learn to recognize novel objects (25–27), although none of these studies explored the questions about representation being addressed here. The YUFO stimulus set is divided into subgroups or “families,” each of which contains 12 unique objects. We used all of two family sets and half of an additional family set, totaling 30 unique objects belonging to one of three families. A family simply refers to a group of stimuli that are specifically designed to be similar to each other compared to two objects from different families. Different families are intended to represent more basic-level category differences compared to within-family differences, which are more consistent with subordinate-level differences. The family structure of this YUFO set permits experimental control over the extent to which individuals must perform more or less fine-grained visual analysis. Of note, these levels of differentiation do not necessarily correlate with traditional taxonomic concepts of basic and subordinate, but rather refer to overall difficulty of the visual discrimination. Importantly, as detailed below, stimuli used for training were pseudorandomly selected and counterbalanced such that each stimulus was presented roughly the same number of times.

**Other-race faces.** For the ORF image set, we selected face images from the Multi-Pie face database (28). The stimuli consisted of 30 face identities, consisting of an equal number of men and women of Asian descent with no obvious facial hair or accessories (e.g., glasses, piercings, and so forth). The images were aligned at the eyes, cropped to exclude external facial features, and manipulated such that all faces had the same mean color values in  $L^*a^*b$  color space. There were two expressions (happy/neutral) and nine viewing angles per identity. The inclusion of multiple expressions in the ORF set, of which there is no equivalent in the YUFO stimulus set, was designed to induce additional difficulty for participants when making comparisons of faces. This ensured that participants would not be at ceiling on the training paradigm at the outset of the experiment.

Although the structure of the stimulus sets for the ORFs and YUFOs is not identical and this is extremely difficult to achieve, given their obvious constraints, we would predominantly be exploring representational changes within each set over the course of training rather than focusing on direct comparisons between the effects of training of the different sets. Furthermore, and critical to our design, as we note below, behavioral performance for ORF and YUFO discriminations are comparable prior to the implementation of training, ensuring perceptual equivalence and titration of the two sets a priori.

**Learning Paradigm.** Independent of which stimulus set was trained, to evaluate learning-induced representational changes, we conducted an extended learning paradigm consisting of 26 sessions (Fig. 2). There were 20 total days of perceptual training and 6 additional days in which participants made similarity ratings (2 d at each of the beginning, midpoint, and end of training). This design enabled us to quantify representational space separately for the first and second half of training. By including two sets of exemplars for learning and for obtaining separate ratings, we could uncover



**Fig. 2.** Learning paradigm flow diagram. Representational space was quantified at the beginning, middle, and end of training, using 2 d of pairwise similarity ratings at each point. Perceptual training involved a visual search task completed over 20 d, with a break after day 10 to quantify the representational space with similarity ratings. Participants were instructed to complete one session per day, consecutively, for 26 d.

the extent to which there was generalization across the sets as well as any benefits of additional training. Second, several other studies have employed training paradigms of about 10 d and have provided evidence of success with such a timeline (7, 29), but because we required considerable data to obtain reliable estimates per identity, we used this extended design.

**Similarity Ratings.** As noted above, participants completed three sets of similarity ratings. Each set of ratings consisted of pairwise similarity judgements for all possible combinations of 60 total images of 30 total objects (two images per object) using an ordinal scale with 1 (clearly different objects) to 7 (clearly the same object). For the YUFO experiment, the two images of each object differed in viewing angle (always +30° or –45° from center facing). YUFO images subtended an approximate visual angle of 8° horizontally and vertically. For the ORF version of the experiment, the two images of each face were forward facing, but differed in expressions (neutral/happy). Face images subtended an approximate visual angle of 6° horizontally and 8° vertically.

Each trial began with a fixation screen for 100 ms. This was followed by the first image for 250 ms and then immediately by the second image for 250 ms. Finally, a blank screen appeared with a central “?” directing the subject to respond using the keyboard (1–7 buttons). The “?” was left on the screen until the subject responded. Each object image was spatially jittered across the central horizontal axis of the screen to reduce feature-by-feature judgements: Horizontal image positions were randomly drawn from a Gaussian distribution centered at the vertical axis and with a SD of 3°. Each set of ratings included 1,830 trials (each image was also rated against itself), split evenly over two sessions on consecutive days, with 915 trials of similarity ratings per session. At the beginning of each rating session, 10 practice trials of randomly selected stimuli allowed participants to acclimate to the task.


Although participants may have learned to differentiate or individuate unique exemplars over the course of rating (and not just training) sessions, because the number of trials, number of unique objects, and number of stimuli were identical for ORFs and YUFOs, any changes in similarity ratings observed between the two classes cannot be obviously explained by the rating procedure itself.

**Perceptual Training.** Participants completed 20 total sessions of the perceptual training task, divided into two periods of 10 d each (separated by ratings session). During each session, held once per day, participants completed 240 trials of a visual search task. The task explicitly trained individuals to recognize 4 of the possible 30 exemplars. Participants were instructed that the task was specifically designed to help them recognize four specific exemplars. An additional 8 exemplars served as distractors (taken from the set of 30 used during similarity ratings).

For those completing the task with YUFOs, on each trial (Fig. 3), one of the four possible training objects was centrally presented for 200 ms at a horizontal and vertical visual angle of 8°, followed by a 400-ms fixation. Thereafter, four object images, each with a horizontal and vertical visual angle of 5°, were displayed in a circle, each equidistant from the center of the screen, and remained on the screen until the participant selected, with one of the four arrow keys, the image of the object initially presented. To encourage engagement in the task, participants were awarded points based on accuracy and reaction time. Participants received at least 50 points per correct trial in addition to points based on a reaction time exponential decay function. Negative points were awarded for incorrect trials corresponding to the trial number (–17 points for incorrect response on trial 17). This made incorrect responses increasingly punitive over the course of a training session. This design encouraged participants to complete the task as quickly and accurately as possible.

For the YUFO version of the experiment, individuals followed one of two training schemes, summarized in Table 1. The choice to implement two





**Fig. 3.** Participants completed 240 match-to-sample trials per training session to facilitate learning of four specific novel objects. Viewing angle in the first image always differs from all viewing angles of all objects in the visual search images. In the face version of this experiment, faces differed in viewing angle in the same manner, but also differed in expression.

versions of the YUFO experiment stems from previous research indicating that subordinate-level processing is the entry point to visual expertise [for example, in wading birds (30)]. However, any results we obtain might simply reflect exposure to a novel category more generally, rather than the specific effect of subordinate-level discrimination. To test this possible explanation, we implemented two versions of the YUFO experiment that differed in the number of subordinate-level discriminations. For participants completing YUFO-A, there was a mix of subordinate-level and family-level discriminations, whereas for those completing YUFO-B, perceptual training included only subordinate-level discriminations. In the analyses below, we explore the extent to which this manipulation differentially impacted changes in representational space.

In the first case (YUFO-A), participants ( $n = 11$ ) first learned two members each of families 2 and 3, and then learned four members of family 1 in the second training block. In this case, distractors for the first training block included a random four objects from family 2 and four from family 3; for the second training block, distractors were reused from the first training block. In the second case (YUFO-B), individuals ( $n = 13$ ) learned either four members of family 2 or four members of family 3 during the first training block, and then the other family exemplars during the second learning block, counterbalanced across subjects. In this case, the distractors for each block were simply the remaining eight objects from the same family as the items being learned.

In both YUFO experiments, the sample image and subsequent test images always differed in viewing angle, but were otherwise selected randomly. There were nine possible viewing angles for each novel object, including forward facing ( $0^\circ$ ), and positive and negative  $15^\circ$ ,  $30^\circ$ ,  $45^\circ$ , and  $60^\circ$  from center.

For the ORF version of the experiment, individuals ( $n = 13$ ) learned a randomly selected four male or four female identities. An 8 additional gender-matched distractors were randomly selected from the larger set of 15 same-gender identities (from the 30 identities used during similarity ratings). During the second training block, participants then learned a new set of randomly selected four identities from the gender that differed from that used in first training block. Training order was counterbalanced across subjects. Just as with YUFOs, the sample image ( $8 \times 6^\circ$  in visual angle) and subsequent test images ( $6 \times 5^\circ$  in visual angle) always differed in viewing angle, but were otherwise selected randomly. In addition, the initial sample image always differed in expression (happy/neutral) from the subsequent test images. Just as in the YUFO experiment, there were nine viewing angles for each face identity, including forward facing ( $0^\circ$ ), and positive and negative  $15^\circ$ ,  $30^\circ$ ,  $45^\circ$ , and  $60^\circ$  from center. Again, the rating and training procedures for YUFOs and ORFs were exactly the same, with the exception of the stimuli used.

In balancing these training regimes, we prioritized equating initial training difficulty across YUFOs and ORFs rather than precisely balancing the number of stimuli or their orientation across the stimulus types. Given that our central hypothesis focuses on the effect of preexisting category experience on subsequent exemplar learning, we made every effort to ensure that, at the outset, the training regimens yielded equivalent performance and, as shown

below, there was no initial difference across the two classes (Fig. 4). In confirming a level playing field at the outset, any results between the two stimulus sets likely reflect differences in preexisting category experience rather than differences evoked by unequal difficulties of the training task.

## Analysis and Results


**Perceptual Training Results.** In the analyses of the training data, we utilized as the dependent measure inverse efficiency (31), calculated as reaction time divided by accuracy. When participants are instructed to complete the task as quickly and as accurately as possible, some preferentially respond to minimize one measure or the other. In addition, participants were awarded points during the training sessions based on their accuracy and reaction time (SI Appendix, Fig. S1 summarizes group-level performance with respect to accuracy and reaction time, separately). Inverse efficiency has been shown to incorporate varying strategies across participants effectively (31–33). A summary of group-level responses is shown in Fig. 4, which plots inverse efficiency over training sessions. In the analyses below, we use the term “training section” to denote a within-subjects factor with two levels (training period 1/training period 2; T1/T2).

Using inverse efficiency as the dependent variable, a repeated-measures ANOVA with training section (T1/T2) and session (sessions 1 to 10) as within-subjects factors, and experiment (YUFO-A/YUFO-B/ORFs) as a between-subjects factor revealed no significant three-way interaction,  $F(18, 306) = 0.74$ ;  $P = 0.76$ ;  $\eta^2 = 0.03$ . There was a significant two-way interaction of training section  $\times$  session [ $F(9, 306) = 7.37$ ;  $P < 0.001$ ;  $\eta^2 = 0.17$ ], reflecting the better performance in later sessions in T1 than in T2, but no significant interaction of training section  $\times$  experiment [ $F(2, 34) = 0.15$ ;  $P = 0.86$ ;  $\eta^2 = 0.005$ ]. There were also main effects of training section [ $F(1, 34) = 23.1$ ;  $P < 0.001$ ;  $\eta^2 = 0.40$ ] and session [ $F(9, 306) = 74.4$ ;  $P < 0.001$ ;  $\eta^2 = 0.65$ ], which is unsurprising given the interaction of these factors. These main effects indicate that participants clearly improved across sessions and that performance was better in the second (sessions 11 to 20) than first (sessions 1 to 10) half of training in all three experiments to an equivalent degree.

Because distinctions between learning ORFs and YUFOs is of primary interest here, we were especially interested in the main effect of experiment [ $F(2, 34) = 5.84$ ;  $P = 0.007$ ;  $\eta^2 = 0.26$ ], which was qualified by a two-way interaction of experiment  $\times$  session [ $F(18, 306) = 2.59$ ;  $P < 0.001$ ;  $\eta^2 = 0.05$ ]. We therefore conducted additional post hoc tests to understand further the effect of experiment on inverse efficiency performance. Close inspection of Fig. 4 shows that the differential effect of experiment on inverse efficiency across session results from a greater improvement in the first few sessions in both T1 and T2. To evaluate this, we conducted a post hoc analysis between experiments for sessions 1 to 3 and 11 to 13 by calculating the difference in inverse efficiency between session 1 and session 3 for T1 and between session 11 and session 13 for T2, and compared these between experiments. In T1, a one-way ANOVA was significant [ $F(2, 34) = 4.02$ ;  $P = 0.02$ ;  $\eta^2 = 0.19$ ], and post hoc Bonferroni-corrected pairwise comparisons revealed more rapid learning (lower inverse efficiency) for those in the face experiment than

**Table 1.** Individuals completed one of three possible experiments during the 26-d learning paradigm

	Version YUFO-A	Version YUFO-B	Version ORF
Ratings 1	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	15 Male, 15 Female
Training 1	2 F <sub>2</sub> and 2 F <sub>3</sub>	4 F <sub>2</sub> or 4 F <sub>3</sub>	4 Male or 4 Female
Ratings 2	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	15 Male, 15 Female
Training 2	4 F <sub>1</sub>	4 F <sub>2</sub> or 4 F <sub>3</sub>	4 Female or 4 Male
Ratings 3	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	12 F <sub>2</sub> , 12 F <sub>3</sub> , 6 F <sub>1</sub>	15 Male, 15 Female

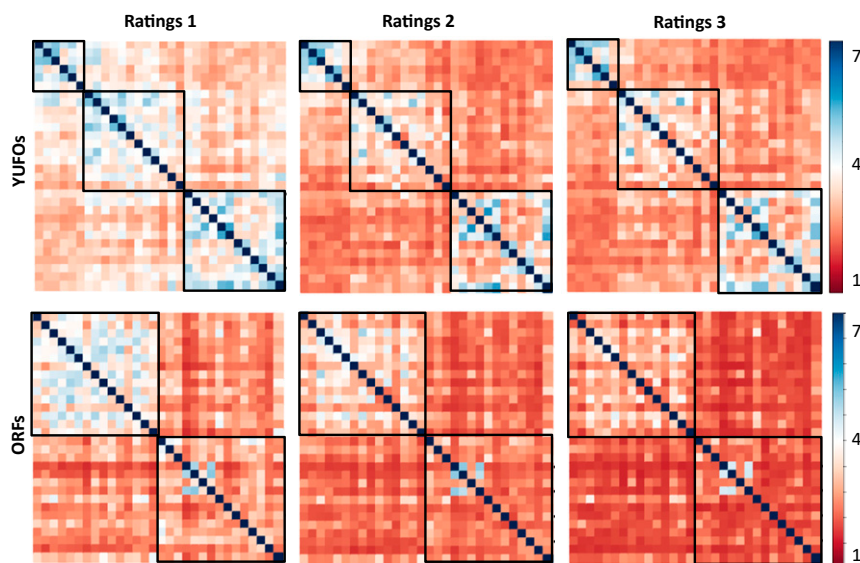


**Fig. 4.** Group-level summary of performance (inverse efficiency) during visual search paradigm across training sessions. Lower scores represent better performance. Error bars represent  $\pm 1$  SEM. Note that individuals completed a second group of similarity ratings between sessions 10 and 11. Participants also switched to a new random set of four objects to learn for the second half of training, starting at session 11. Here, lines connecting sessions 10 and 11 show the cost associated with changes in training stimuli.

those in the YUFO-B experiment [ $t(25) = 2.68$ ;  $P = 0.03$ ]. However, there were no differences between the face and YUFO-A experiments [ $t(23) = 2.10$ ;  $P = 0.13$ ] or between the two YUFO experiments themselves [ $t(23) = 0.47$ ;  $P = 1.00$ ]. The same pattern of results was observed in the second section of training (sessions 11 to 13). The one-way ANOVA was significant [ $F(2, 34) = 5.65$ ;  $P = 0.008$ ;  $\eta^2 = 0.25$ ], and post hoc Bonferroni-corrected  $t$  tests revealed more rapid learning in the face experiment than in the YUFO-B experiment [ $t(25) = 3.35$ ;  $P = 0.006$ ]. Again, there were no differences between ORFs and YUFO-A experiments [ $t(23) = 1.84$ ;  $P = 0.22$ ] or between the YUFO experiments themselves [ $t(23) = 1.36$ ;  $P = 0.54$ ], suggesting that YUFO-A fell intermediate between the other two experiments.

Analysis of the main effect of experiment by one-way ANOVA revealed no difference in average inverse efficiency in training section 1 (T1) across experiments [ $F(2, 34) = 2.24$ ;  $P = 0.12$ ;  $\eta^2 = 0.15$ ]. There was, however, a significant difference in average performance across training section 2 (T2) [ $F(2, 34) = 17.7$ ;  $P < 0.001$ ;  $\eta^2 = 0.51$ ]. Post hoc Bonferroni-corrected pairwise comparisons revealed that those completing the YUFO-B experiment performed more poorly than both those in YUFO-A [ $t(23) = 4.21$ ;  $P < 0.001$ ] and in ORF experiments [ $t(25) = 5.70$ ;  $P < 0.001$ ], which did not differ from each other [ $t(23) = 1.25$ ;  $P = 0.65$ ].

Together, the analyses of the training data reveal that participants improved over successive sessions, and that performance was better overall in the second relative to the first half of training. There were some modulating effects of particular



**Fig. 5.** Group level similarity matrices for the average of both YUFOs (Upper) and ORFs (Lower) derived from pairwise similarity ratings (Euclidean distance in similarity space) across three sessions. Red (1, very different) and blue (7, very similar) correspond to different ends of the rating spectrum. A general increase in red (more negative) corresponds to objects being rated less similar, hence, moving farther apart. Black boxes correspond to the structure of stimulus sets. For YUFOs, stimuli 1 to 6 are from family 1; 7 to 18 are from family 2; 19 to 30 are from family 3. For ORF stimuli 1 to 15 are male faces and 16 to 30 are female faces.

experiments; for example, when the number of subordinate-level discriminations was matched (faces vs. YUFO-B), those in the face training experiment performed better overall in T2 than those in YUFO-B experiment and, in both training sections (T1/T2), they also improved disproportionately across the first few sessions of training. There were no obvious differences between YUFO-A and faces or YUFO-B and YUFO-A as a function of session or training section. These findings suggest that the two YUFO experiments employed here overlap substantially, and therefore do not permit further inferences regarding subordinate and basic levels of processing and their role in the development of visual expertise. Importantly, these results are not inconsistent with findings that suggest that subordinate-level discrimination is the “entry point” to visual expertise. Many possible explanations can be offered for this inconsistency, including the fact that both YUFO tasks employed substantial subordinate-level discriminations (just more in YUFO-B than in YUFO-A) and the fact that a subordinate-level advantage may depend on some familiarity with category and YUFOs are entirely novel for the observers. Other assays of posttraining performance might uncover differences between the YUFO experiments (e.g., differential inversion effects), but under these testing conditions, no differences were evident. Most importantly, these results show that initial training difficulty was well matched between conditions and that the pattern of observed training results was relatively similar between conditions, particularly during T1.

Having evaluated the effects of acquisition of novel exemplars, we next carried out more detailed representational analyses to assess the fate of individual exemplars across the similarity rating sessions.

### Similarity Ratings.


**Magnitude of changes in representational space.** To compare similarity ratings across the three training scenarios (faces, YUFO-A, and YUFO-B), we converted all pairwise ratings to a  $30 \times 30$  similarity matrix for each participant, where each cell reflects the average of similarity judgments between the two unique object identities. There were two images of each unique object, yielding four pairwise judgements between each pair of unique objects. The raw group-level similarity matrices are shown in Fig. 5. Histograms of raw response frequency at the group level are included in *SI Appendix, Fig. S2*. To obtain the difference between ratings sessions, we simply created a difference matrix by subtracting the matrix of the first rating session from the second rating session (and separately, subtracting the matrix of the second rating session from that of the third, yielding two difference matrices per participant, corresponding to T1 and T2). In the analyses below, we use the term “training section” to denote a within-subjects factor with two levels (T1/T2). Note that a negative change in similarity corresponds to objects moving farther apart in representational space and this is shown in red colors in Fig. 5.

In the following analyses, we tested three specific hypotheses. First, we hypothesized that there should be greater differences in the magnitude of distance change relative to all other objects for trained versus nontrained objects, corresponding to better differentiation of trained objects generally. For each object (each row of  $30 \times 30$  difference matrix), we extracted the mean value of distance change between that object and the 29 other objects. We then separated these values based on whether the object was one of the four objects trained or not. Second, we hypothesized that the observable distance changes would not be uniformly distributed across the representational space (difference matrix). We predicted that there would be greater distance changes for any two objects located more closely in representational space prior to the start of training, relative to two objects that were more distant pretraining. Said another way, we predicted a

significant negative correlation between the amount of post-training distance change and the initial distance between two objects, with objects nearer to each other exhibiting greater change in representational distance. In perception, this corresponds to a greater improvement in the discriminations between more similar items relative to those that are less similar before training. To test this prediction, we obtained correlations between object similarity, defined by subject-level average of pre-training and posttraining similarity ratings, and the magnitude of distance changes. Again, we obtained average correlations for trained objects and all other, nontrained objects with the distance change observed for those objects. Finally, we expected that results observed with respect to the above hypotheses would differ for those trained on ORFs compared to those trained on YUFOs, because of differences in preexisting category experience.

Before contrasting data from learning novel objects against data from learning ORFs, we first compared the data between YUFO experiments (A and B). Although these experiments differed during the training portion of the paradigm, these participants completed the exact same ratings sessions, and given that both versions of training encompassed at least some subordinate-level discrimination, the changes in similarity ratings may not have differed. A repeated-measures ANOVA with training section (T1/T2) and learning status (Learned/Not Learned) as within-subjects factors, and experiment version (YUFO-A, YUFO-B) as a between-subjects factor, revealed no three-way interaction of training section  $\times$  learning status  $\times$  experiment,  $F(1, 22) = 0.33$ ;  $P = 0.86$ ;  $\eta^2 = 0.001$ . There were also no two-way interactions with experiment: Training section  $\times$  experiment [ $F(1, 22) < 0.001$ ;  $P = 0.99$ ;  $\eta^2 < 0.001$ ] or learning status  $\times$  experiment [ $F(1, 22) = 0.003$ ;  $P = 0.95$ ;  $\eta^2 < 0.001$ ]. Finally, there was no main effect of experiment [ $F(1, 22) = 2.77$ ;  $P = 0.11$ ;  $\eta^2 = 0.11$ ]. However, there were main effects of training section [ $F(1, 22) = 20.8$ ;  $P < 0.001$ ;  $\eta^2 = 0.48$ ], showing greater distance changes during T1 compared to T2, and of learning status [ $F(1, 22) = 8.86$ ;  $P = 0.007$ ;  $\eta^2 = 0.29$ ] with greater distance changes for trained relative to nontrained YUFOs. Given the lack of differences between the two YUFO experiments, we combined these data points for comparison with the ORF data.

**Selectivity of changes in representational space.** The mean distance changes for ORF and YUFO training are plotted in Fig. 6. A repeated-measures ANOVA with training section (T1/T2) and learning status of exemplar (Trained/Not Trained) as within-subjects factors and experiment (faces/YUFOs) as the between-subjects factor



**Fig. 6.** Mean distance changes in representational space. Distance changes are separated by experiment version (ORFs/YUFOs) and by training section (T1/T2). Negative distance changes correspond to moving farther apart in representational space. Error bars represent  $\pm 1$  SEM.




revealed a reliable three-way interaction,  $F(1, 35) = 5.94$ ;  $P = 0.02$ ;  $\eta^2 = 0.12$ . There was also a significant two-way interaction of training section  $\times$  learning status [ $F(1, 35) = 4.87$ ;  $P = 0.03$ ;  $\eta^2 = 0.11$ ], and marginally significant interactions of learning status  $\times$  experiment [ $F(1, 35) = 3.64$ ;  $P = 0.06$ ;  $\eta^2 = 0.05$ ] and training section  $\times$  experiment [ $F(1, 35) = 4.10$ ;  $P = 0.05$ ;  $\eta^2 = 0.07$ ]. The main effects of learning status [ $F(1, 35) = 26.7$ ;  $P < 0.001$ ;  $\eta^2 = 0.41$ ] and training section were also significant [ $F(1, 35) = 18.5$ ;  $P < 0.001$ ;  $\eta^2 = 0.32$ ], indicating that learned objects exhibited greater mean distance changes relative to objects not learned and that the mean distance changes were greater in T1 compared to T2. Importantly, given the difference in the numbers of participants trained with faces ( $n = 13$ ) and YUFOs ( $n = 24$ ), we tested for equality of variances between experiments, using subject-level mean distance changes. No differences in variances were observed in any group of distance changes using Levene's test of equality of variances: T1 Learned [ $F(1, 35) = 0.40$ ;  $P = 0.53$ ]; T1 Not Learned [ $F(1, 35) = 4.31$ ;  $P = 0.05$ ]; T2 Learned [ $F(1, 35) = 0.58$ ;  $P = 0.45$ ]; T2 Not Learned [ $F(1, 35) = 0.38$ ;  $P = 0.54$ ].

Having set aside potential differences in equality of variance and given the significant three-way and two-way interactions, we completed a series of additional pairwise comparisons that address our a priori hypotheses more directly. In the ORF group, across training sections (T1 vs. T2), distances changes were larger for trained ORFs,  $t(12) = 2.70$ ;  $P = 0.01$ ;  $d = 0.75$ , than for nontrained ORFs,  $t(12) = -1.17$ ;  $P = 0.26$ ;  $d = 0.32$ . The mean distance between the trained ORFs and all other ORFs increased to a greater degree than was the case for nontrained ORFs [ $t(12) = 3.68$ ;  $P = 0.003$ ;  $d = 1.02$ ] during T1 and during T2 [ $t(12) = 2.69$ ;  $P = 0.02$ ;  $d = 0.75$ ]. In the YUFO learning experiment, across training sections (T1/T2) distances between trained or nontrained YUFOs versus all other YUFOs differed [ $t(23) = 2.32$ ;  $P = 0.03$ ;  $d = 0.47$ ], with greater distance changes for trained over nontrained YUFOs. Across training sections (T1 vs. T2) there were differences in distance change for both trained YUFOs [ $t(23) = 4.58$ ;  $P < 0.001$ ;  $d = 0.93$ ] and nontrained YUFOs [ $t(23) = 4.76$ ;  $P < 0.001$ ;  $d = 0.97$ ]. This is not surprising given that the distance changes during T2, for both trained [ $t(23) = 0.91$ ;  $P = 0.37$ ] and nontrained YUFOs [ $t(23) = 0.39$ ;  $P = 0.70$ ] did not differ from zero.

Together, our analyses of distance changes revealed several clear outcomes. First, the evidence indicates that distance changes for trained objects were greater than for nontrained objects both in those trained on ORFs and in those trained on YUFOs. However, given the reliable interaction with experiment and an effect size that is over twice as large for the ORF experiment relative to the YUFO experiment, the difference between trained and nontrained objects was much larger for ORFs than YUFOs in T1. This suggests that, as one gains experience with a visual category, subsequent representational changes that occur with exposure to a new exemplar become increasingly specific to the novel object. Second, in both experiments the distance changes were greater during T1 compared to T2. Interestingly, those learning ORFs exhibited continued nonzero distance change in T2 for trained over nontrained ORFs, suggesting a much richer and more complex representational space. In contrast, those learning YUFOs showed no additional distance changes for trained over nontrained YUFOs at the group level during T2.

**Relativity of changes in representational space.** Next, we tested our second hypothesis that the distance change between objects would occur preferentially for more similar pairs of objects. That is, we predicted that objects initially perceived to be closer together would experience a greater shift in representational distance relative to two objects that were initially far apart. To test this prediction, for each object learned by each participant, we correlated the distance in space with the distance between the



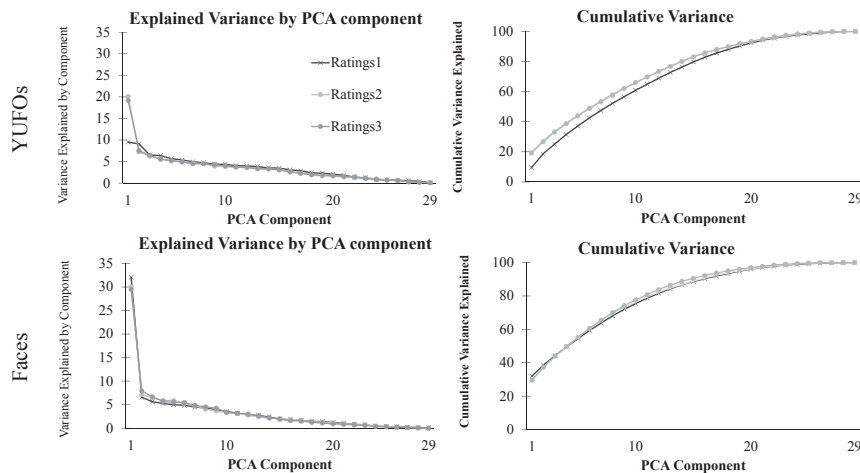
**Fig. 7.** Correlations between pretraining similarity and representational distance change. Both ORFs and YUFOs are included across two consecutive training sections (T1 and T2). Negative correlations correspond to greater separation in representational space (moving apart) for objects that were closer together in space initially. Error bars correspond to  $\pm 1$  SEM.

object and all 29 other objects.<sup>†</sup> We then separated these pairwise correlations based on training status (Trained/Not Trained) and tested for differences at the group level using the approach used for absolute magnitude differences above. See Fig. 7 for a summary of group-level results.

A repeated-measures ANOVA with training section (T1/T2) and learning status (Trained/Not Trained) as within-subjects factors, and experiment (ORFs/YUFOs) as a between-subjects factor did not reveal a significant three-way interaction of training section  $\times$  learning status  $\times$  experiment,  $F(1, 35) = 0.51$ ;  $P = 0.47$ ;  $\eta^2 = 0.01$ . There were no significant two-way interactions either: Training section  $\times$  experiment [ $F(1, 35) = 2.82$ ;  $P = 0.10$ ;  $\eta^2 = 0.07$ ] and learning status  $\times$  experiment [ $F(1, 35) = 1.89$ ;  $P = 0.17$ ;  $\eta^2 = 0.04$ ], or learning status  $\times$  training section [ $F(1, 35) = 0.73$ ;  $P = 0.39$ ;  $\eta^2 = 0.02$ ]. There were, however, significant main effects of learning status [ $F(1, 35) = 7.30$ ;  $P = 0.01$ ;  $\eta^2 = 0.17$ ] and experiment [ $F(1, 35) = 22.5$ ;  $P < 0.001$ ;  $\eta^2 = 0.39$ ], but not training section [ $F(1, 35) = 0.94$ ;  $P = 0.34$ ;  $\eta^2 = 0.02$ ]. These main effects reveal larger negative correlations overall for trained objects relative to nontrained objects and larger negative correlations for ORFs relative to YUFOs. Because the interactions with experiment were not significant, we did not conduct follow-up analyses.

Together, the findings from the analyses of correlations between similarity and distance change revealed one main difference between ORFs and YUFOs in representational changes induced by exemplar learning: Overall correlations were more negative for ORFs compared to YUFOs. This means that changes in representational distance occurred more locally for ORFs. Said another way, larger magnitude distance changes tended to occur in closer proximity to the objects being learned, rather than farther away, in representational space. Importantly, the correlations for YUFOs did not differ from zero, suggesting there was no differential localization of distance change overall; rather, distance changes occurred more uniformly across the representational space. Note, this pattern of results was observed despite similar variability in pairwise representational distance between YUFOs at the outset compared to between ORFs (Fig. 5 and *SI Appendix*, Fig. S2), potentially allowing for similar correlations with pretraining similarity for YUFOs compared to

<sup>†</sup>More specifically, we quantified similarity distance at the subject level. We averaged the pretraining and posttraining similarity spaces for a given subject and then correlated the average similarity distances with the changes in distance before and after training at the individual subject level.



**Fig. 8.** Principle component analysis (PCA) of the group-level similarity matrices from each of three ratings sessions. Components are plotted, on the *Left*, by variance explained, in decreasing order. On the *Right*, cumulative variance explained by adding each additional component is plotted.

ORFs. However, the opposite pattern was observed here, with greater correlations uncovered in the representational space for ORFs. A second main effect of learning status revealed that trained objects exhibited more local distance changes relative to nontrained objects. Close inspection of Fig. 7 suggests this effect was driven more by the difference between trained and non-trained ORFs than YUFOs, although the interaction with experiment did not quite reach significance.

Finally, to further bolster our earlier claim that our training paradigm was well balanced with respect to initial difficulty, and to support the claim that the pattern of observed results was a product of previous category experience rather than substantially different training experiences, we present two supplementary figures (*SI Appendix, Figs. S3 and S4*). These figures represent a visual summary of the group-level results when every participant ( $n = 13$ ) trained on faces is closely matched with a single participant from the YUFO experiment. Note that these results are nearly identical to and correspond with the results shown in those shown in Figs. 4, 6, and 7, and further confirm that the results reflect preexisting category experience rather than differences in the (matched) training.

**Dimensionality of Representational Space.** We conducted a final analysis to explore possible changes in the underlying dimensions of representational space as a function of training and visual category. We performed principal component analyses on each of the group-level  $30 \times 30$  similarity matrices (shown in Fig. 5) and plotted the percentage of explained variance by component, as well as the cumulative variance across components (Fig. 8).

This analysis revealed that the overall dimensionality of the representational space for ORFs was relatively constant across ratings sessions.<sup>‡</sup> The number of orthogonal components to explain any amount of variance did not change dramatically across the experiment. However, this does not mean that individuals could not change the relative weights of multiple dimensions concurrently, facilitating better separation of the stimuli, as was observed in the analyses above. The results suggest that individuals have likely previously extracted the dimensions along which faces vary, and perhaps, alter weights of different dimensions to meet

the demands of the perceptual task rather than massively reorganize the dimensions.

In contrast, there was a change in the dimensionality of the YUFO representational space between the first and second ratings sessions (across T1). The analysis revealed that the 1st component accounted for only half of the variance relative to subsequent ratings sessions, and a larger amount of variance was explained by the 2nd through 29th components. This indicates that additional components from the first ratings section were required to explain the same amount of variance as the smaller number of components in the second and third ratings sessions. While we make no inference about the specific meaning of individual components, we conclude that, overall, the dimensionality of the YUFO representational space appears to have decreased with experience. Finally, this analysis revealed that in all three ratings sessions, the representational space for ORFs appears to have a different number of dimensions relative to the representational space for YUFOs, perhaps indicating differences in the complexity of the stimulus or in the complexity of the visual processes involved for each stimulus set.

## Discussion

The aim of the present study was to elucidate how novel exemplars from categories of visual objects are incorporated into representational space and the extent to which this process is governed by preexisting category experience. To this end, we implemented a 26-d visual search training paradigm to train individuals to recognize four exemplar objects at a time. We quantified representational space before, midway, and after training using pairwise similarity ratings between all 30 objects. Three groups of participants completed the study with one group studying and rating ORFs and two groups studying and rating three-dimensional-generated novel objects, YUFOs, with each group having a slightly different training regimen (although, because there were no significant differences in the representational analyses, we combined the data from these two groups in the analyses). We then measured representational distance changes at the individual stimulus level and examined to what extent the changes observed were governed by previous category experience. In the analyses of our data, we tested three distinct hypotheses in the setting of two unique stimulus categories.

In our analyses of the visual search paradigm, we found that performance generally improved over the course of training, and that this improvement did not generally differ as a function of visual category, particularly during T1. Individuals training on

<sup>‡</sup>To aid understanding of the data, *SI Appendix, Figs. S5 and S6* show two dimensional multidimensional scaling solutions for group-level representational space at pre-, mid-, and posttraining. Note that we make no inferences from these figures, as each individual participant was randomly assigned to train on a different subset of stimuli.



ORFs and those training on YUFOs improved roughly to the same extent in both training sections (T1/T2). However, when the amount of subordinate-level discrimination was matched (ORFs vs. YUFO-B), there was a small category difference improvement over the first few sessions of the task: Individuals learning ORFs were significantly better than those learning YUFOs between the first and third sessions in both T1 and T2. Importantly, the difference occurred only between experiments where participants completed trials that contained exclusively subordinate-level discriminations (matching for difficulty). The lack of a greater difference in performance for the different categories in our visual search task is not entirely surprising, perhaps reflecting the significant overlap of our training tasks. That is, both groups completed a substantial number of trials that tapped subordinate-level discrimination. More generally, advantages of expertise are commonly found in visual search tasks, but these tasks are typically more complex or involve a greater number of stimuli (34, 35) than the tasks adopted here. More importantly, the lack of a stimulus category difference in improvement, during the T1 training section in particular, means that any observed differences in subsequent RSA of T1 are more likely to have been driven by previous category experience and not by an interaction with the training paradigm.

The analyses of changes in representational space revealed two dissociations between face and novel-object learning. We use the terms “what” and “where” to facilitate their distinction for the reader. First, there was category specificity in the magnitude of distance change, revealing what changed. For those trained on ORFs, the distance between trained ORFs and all other ORFs increased about twice as much as the distance between non-trained ORFs and all other ORFs. For novel objects, the same trend was found, but the difference between trained and non-trained YUFOs was much smaller. These findings suggest that, as one gains experience with a visual category, the separation observed in representational space becomes more specific to the newly learned exemplars, or less generalized to the space as a whole (i.e., the new individual instances become increasingly differentiated relative to existing representations). The second category-specific dissociation revealed differences in where the representational changes happen for faces relative to novel objects. Overall, the correlation between pretraining similarity and posttraining distance change was greater for ORFs relative to YUFOs. Additionally, close inspection of Fig. 7 reveals that this correlation appears even more negative for trained relative to nontrained objects. This indicates that as one gains experience with a visual category, the change observed in representational space occurs increasingly locally. Together, these dissociations in what change happens and where this change occurs during exemplar learning help to characterize the representational origin of the learning in highly skilled visual perception.

Although clear trends emerged during T1, there were also differences at T2. For ORFs, the specificity for trained relative to nontrained faces persisted, but was smaller in magnitude. Interestingly, there appeared to be no reliable change at all during T2 at the group level for those trained on YUFOs, despite clear change during T1. This is somewhat counterintuitive given that faces at the start of T2 were already farther apart on average (see more red, overall, for faces compared to YUFOs in middle column of Fig. 5) compared to YUFOs. Said another way, YUFOs had more room to move apart in space (on the similarity scale) relative to ORFs, yet they did not appear to undergo a net distance change at all on average during T2. This does not mean that there was no reorganization of space, just that, on average, the net distance between objects did not change much over T2. There are a few possible explanations for this. First, ORFs may have continued to move apart in space, while YUFOs did not, because individuals have extensive face experience that they can exploit in the task and continue to make fine-grained differentiations.

For those learning YUFOs, there is no obviously related visual category with which they have experience and, therefore, they do not have other representations or dimensions that can be easily leveraged in further differentiating YUFOs.

Alternatively, there may be differences in the way face and YUFO stimuli are processed by the learner. While faces are processed configurally as is often assumed to be true for domains of expertise (36, 37), YUFOs may be processed, in our task, more componentially. That is, individuals trained on YUFOs may have simply extracted a single feature or two and then differentiated objects along only these dimensions (for example, the top or “hat” part of the YUFO; see Fig. 1). After T1, YUFOs may have been sufficiently differentiated, in the similarity ratings portion of the experiment, along those dimensions that the net distance change was zero in T2. These two explanations are not mutually exclusive, and may both be contributing to the observed pattern of results.

In addition to the observed distance changes, we also found that the dimensionality of representational space decreased with training for YUFOs, but remained constant for ORFs. This suggests that individuals trained on faces may have already extracted the relevant dimensions along which faces vary, including ORFs, and are thus stable over time. Individuals may simply change the relative weightings of these existing dimensions to facilitate better perceptual discrimination of ORFs. In contrast, the individuals trained on YUFOs had no prior experience with these computer-generated objects, and so they could not have known along which dimensions the YUFOs would vary. In our training paradigm, individuals improved their ability to discriminate highly similar YUFOs, implying that individuals must have extracted at least some relevant features. This change could cooccur with “discarding” nondiagnostic features, leading to a lower dimensional representational space. The difference, then, between experienced and expert perception may lie primarily in additional fine-tuning of established dimensions.

A plausible neural implementation of this fine-tuning across exemplars has been clearly demonstrated in the motor domain in the context of brain computer interfaces (13, 14). By reducing a high-dimensional neural space to a smaller set of orthogonal dimensions, the intrinsic manifold, Sadtler et al. (14) could successfully predict the extent of generalization on a brain computer interfaces motor task. If a novel task could be explained using the same low-dimensional space—that is, the new task exists within the intrinsic manifold of the old task, albeit with different dimensional weights—generalization to the new task is possible. In the present study, the net dimensionality of the representational space for faces did not change with training, perhaps suggesting that ORFs lie within the existing intrinsic manifold of faces more generally. Participants of the present study, then, were likely reweighting existing dimensions of face perception to permit the pattern of observed results. In contrast, the intrinsic manifold observed for YUFOs appears to emerge and change with training, suggesting that participants have to extract *de novo* dimensions along which features vary. The difference then, between expert and novice visual perception, may lie primarily in the degree to which an intrinsic manifold, extracted from a high-dimensional neural space, permits efficient generalization to novel exemplars.

In the present study, we implemented a face-training paradigm with other-race faces instead of own-race faces. We made this decision for two reasons. Pragmatically, our method of quantifying representational space using similarity ratings has a somewhat artificial ceiling, such that objects already far apart in space might still exhibit distance changes with learning, but such changes would likely not be detectable with our scale. The use of ORFs in the present study ensured that we would be able to detect both small and large distance changes, such as those observed in Fig. 7. The second reason we used ORFs is to emphasize

the spectrum of visual expertise, rather than postulate a binary distinction. While our participants may not have been experts in ORFs in the strictest sense (38), they almost certainly had some ORF experience, as well as a lifetime of face experience generally. Thus, the differences observed here are likely just the differences between two points on the spectrum of experience. One can even extrapolate out to an extreme case of expertise, own-race recognition, and speculate what one would observe in the same experiment theoretically. In this case, we predict an even more extreme pattern of what and where findings compared to those observed in our ORF condition (provided that the tasks are sufficiently difficult to keep the observer off the ceiling). That is, distance changes might only occur for newly encountered objects and these changes might only occur locally in representational space. This sliding scale of the specificity of findings that would be expected to occur with varying experience further emphasizes the need to study expertise along a continuum; there does not seem to be any reasonable point at which to impose a binary distinction of expert and nonexpert processes in the context of the findings observed here.

The present experiment demonstrates the utility of a generalized representational approach in comparing behavior across stimulus categories. This approach has already substantially advanced neuroscientific investigations (39), and stands to benefit behavioral investigations as well. In the domain of faces, representational approaches have already led to a unifying theory of face perception, Face Space (40), whose predictions have held true across a large body of research (12). In the present study, we found that a more generalized representational approach allows specific predictions that might separate different types of representational spaces (i.e., expert from novice). The motivation for our approach is bolstered by the field, which has identified several behavioral assays of visual expertise, including the inversion effect and subordinate-level processing. Many of these assays make assumptions about the underlying representations, without explicitly quantifying the representational space. Our approach offers a theoretical framework under which to bring together disparate findings of the visual expertise field, and explicitly test predictions about “expert” representations. More generally, comparing changes across representational spaces also offers the advantage of directly examining domain general processes across disparate visual object categories: For example, during exemplar learning.

Although we specifically investigate the effect of previous experience on representational changes during exemplar learning, there

are a countless number of future directions worth exploring to understand better the principles that contribute to the development of representational spaces that subserve highly skilled visual perception. In future studies, it will be important to elucidate the contributions of other aspects of learning, such as the coverage of the space during learning, the relative emphasis or effort on different objects within a category, the specific training task employed, as well as contributions of nonperceptual features, such as semantics, to the development of visual expertise. As an additional future direction more specific to the present study, one might investigate the effects of sampling the representational space itself on learning, perhaps by comparing the ratings of a group of participants that completed training to a second “wait-list” group that had no training. Understanding the contributions of these features of learning and confirming them with other experimental approaches will elucidate further the stability of the findings from the present study.

In summary, we conducted a multiday exemplar learning paradigm in which we quantified representational changes at the individual stimulus level before, in the middle, and after training. Participants completed this learning paradigm with faces or computer-generated novel objects, YUFOs. When individuals had substantial previous category experience (faces), changes in representational space were more specific to the trained exemplars in both magnitude and in representational locality. When individuals learned novel objects (YUFOs), representational changes were more generalized across the space, both in magnitude and representational locality. Finally, learning a novel visual category was associated with a reduction in the dimensionality of the representational space. Together, these findings offer a representational mechanism by which highly skilled visual processing emerges over the course of experience, as well as a theoretical representational framework in which to explicitly test representational assumptions that have arisen from the field of visual expertise.

**Data and Experimental Materials.** Data and experimental materials are publicly available at DOI 10.1184/R1/11869524.

**ACKNOWLEDGMENTS.** We thank Carl Olson, Michael Tarr, and David Plaut for their constructive comments in the design and execution of this research and helpful feedback on this manuscript; and Byron Yu for his feedback on this manuscript. Funding for this work was provided by a Richard King Mellon Foundation Presidential Fellowship in Life Sciences and by predoctoral Fellowship NIH 5T32GM081760-09 (to E.C.).

1. F. Martens, J. Bulthé, C. van Vliet, H. Op de Beek, Domain-general and domain-specific neural changes underlying visual expertise. *Neuroimage* **169**, 80–93 (2018).
2. J. Shen, M. L. Mack, T. J. Palmeri, Studying real-world perceptual expertise. *Front. Psychol.* **5**, 857 (2014).
3. J. Tanaka, M. Taylor, Object categories and expertise: Is the basic level in the eye of the beholder? *Cognit. Psychol.* **23**, 457–482 (1991).
4. R. Diamond, S. Carey, Why faces are and are not special: An effect of expertise. *J. Exp. Psychol. Gen.* **115**, 107–117 (1986).
5. A. Harel, D. Kravitz, C. I. Baker, Beyond perceptual expertise: Revisiting the neural substrates of expert object recognition. *Front. Hum. Neurosci.* **7**, 885 (2013).
6. M. Bilalic, T. Grottenhaler, T. Nägele, T. Lindig, The faces in radiological images: Fusiform face area supports radiological expertise. *Cereb. Cortex* **26**, 1004–1014 (2016).
7. I. Gauthier, M. J. Tarr, Becoming a “Greeble” expert: Exploring mechanisms for face recognition. *Vision Res.* **37**, 1673–1682 (1997).
8. M. J. Tarr, I. Gauthier, FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nat. Neurosci.* **3**, 764–769 (2000).
9. F. A. Soto, F. G. Ashby, Categorization training increases the perceptual separability of novel dimensions. *Cognition* **139**, 105–129 (2015).
10. V. Bruce, A. Young, Changing faces: Visual and non-visual coding processes in face recognition. *Br. J. Psychol.* **3**, 105–116 (1986).
11. A. Nestor, D. C. Plaut, M. Behrmann, Feature-based face representations and image reconstruction from behavioral and neural data. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 416–421 (2016).
12. T. Valentine, M. B. Lewis, P. J. Hills, Face-space: A unifying concept in face recognition research. *Q. J. Exp. Psychol. (Hove)* **69**, 1996–2019 (2016).
13. M. D. Golub *et al.*, Learning by neural reassociation. *Nat. Neurosci.* **21**, 607–616 (2018). Correction in: *Nat. Neurosci.* **21**, 1138 (2018).
14. P. T. Sadler *et al.*, Neural constraints on learning. *Nature* **512**, 423–426 (2014).
15. I. Charest, N. Kriegeskorte, The brain of the beholder: Honouring individual representational idiosyncrasies. *Lang. Cogn. Neurosci.* **30**, 37–41 (2015).
16. N. Kriegeskorte *et al.*, Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* **60**, 1126–1141 (2008).
17. E. Freud, J. C. Culham, D. C. Plaut, M. Behrmann, The large-scale organization of shape processing in the ventral and dorsal pathways. *eLife* **6**, 1–26 (2017).
18. A. Nestor, D. C. Plaut, M. Behrmann, Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9998–10003 (2011).
19. J. Tanaka, B. Heptonstall, S. Hagen, Perceptual expertise and the plasticity of other-race face recognition. *Vis. Cogn.* **21**, 1–19 (2013).
20. A. W. Young, V. Bruce, Understanding person perception. *Br. J. Psychol.* **102**, 959–974 (2011).
21. S. Caharel *et al.*, Other-race and inversion effects during the structural encoding stage of face processing in a race categorization task: An event-related brain potential study. *Int. J. Psychophysiol.* **79**, 266–271 (2011).
22. K. Crookes, S. Favelle, W. G. Hayward, Holistic processing for other-race faces in Chinese participants occurs for upright but not inverted faces. *Front. Psychol.* **4**, 29 (2013).
23. P. Chiroro, T. Valentine, An investigation of the contact hypothesis of the own-race bias in face recognition. *Q. J. Exp. Psychol.* **48**, 879–894 (1995).
24. I. Gauthier, M. J. Tarr, A. W. Anderson, P. Skudlarski, J. C. Gore, Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nat. Neurosci.* **2**, 568–573 (1999).

25. J. A. Collins, K. M. Curby, Conceptual knowledge attenuates viewpoint dependency in visual object recognition. *Vis. Cogn.* **21**, 945–960 (2013).
26. I. Gauthier, T. W. James, K. M. Curby, M. J. Tarr, The influence of conceptual knowledge on visual discrimination. *Cogn. Neuropsychol.* **20**, 507–523 (2003).
27. B. Rossion, C. Kung, M. J. Tarr, Visual expertise with nonface objects leads to competition with the early perceptual processing of faces in the human occipitotemporal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 14521–14526 (2004).
28. R. Gross, I. Matthews, J. Cohn, T. Kanade, S. Baker, Multi-PIE. *Proc. Int. Conf. Autom. Face Gesture Recognit.* **28**, 807–813 (2010).
29. B. Rossion, C. C. Kung, M. J. Tarr, Visual expertise with nonface objects leads to competition with the early perceptual processing of faces in the human occipitotemporal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 14521–14526 (2004).
30. J. W. Tanaka, The entry point of face recognition: Evidence for face expertise. *J. Exp. Psychol. Gen.* **130**, 534–543 (2001).
31. J. T. Townsend, F. G. Ashby, “Methods of modeling capacity in simple processing systems” in *Cognitive Theory*, J. Castellan, F. Restle, Eds. (Erlbaum, Hillsdale, New-Jersey, 1978), Vol. 3, pp. 200–239.
32. E. Collins, J. Park, M. Behrmann, Numerosity representation is encoded in human subcortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E2806–E2815 (2017).
33. E. Freud, M. Behrmann, The life-span trajectory of visual perception of 3D objects. *Sci. Rep.* **7**, 11034 (2017).
34. C. Goulet, C. Bard, M. Fleury, Expertise differences in preparing to return a tennis serve: A visual information processing approach. *J. Sport Exerc. Psychol.* **11**, 382–398 (1989).
35. A. M. Williams, K. Davids, Visual search strategy, selective attention, and expertise in soccer. *Res. Q. Exerc. Sport* **69**, 111–128 (1998).
36. I. Gauthier, C. Bukach, Should we reject the expertise hypothesis? *Cognition* **103**, 322–330 (2007).
37. B. Rossion, I. Gauthier, V. Goffaux, M. J. Tarr, M. Crommelinck, Expertise training with novel objects leads to left-lateralized facelike electrophysiological responses. *Psychol. Sci.* **13**, 250–257 (2002).
38. A. W. Young, A. M. Burton, Are we face experts? *Trends Cogn. Sci.* **22**, 100–110 (2018).
39. N. Kriegeskorte, M. Mur, P. Bandettini, Representational similarity analysis —Connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
40. T. Valentine, A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Q. J. Exp. Psychol. A* **43**, 161–204 (1991).