Objects, faces, and spaces: Organizational principles of visual object perception as evidenced by individual differences in behavior

Heida Maria Sigurdardottir[1]* and Inga María Ólafsdóttir[2]

[1] Department of Psychology, University of Iceland, Reykjavik, Iceland
[2] Department of Psychology, Reykjavik University, Reykjavik, Iceland

* Corresponding author: Heida Maria Sigurdardottir (heidasi@hi.is)

This manuscript is a preprint (version 1). It has at the time of deposit not been sent out for peer review to an academic journal. Questions and comments, includings suggestions for references to additional literature, are welcome and should be directed to the corresponding author.

Abstract

What are the diagnostic dimensions on which objects truly differ visually, and do these dimensions provide the foundation for the organizational principles of visual object perception as evidenced by individual differences in behavior? Using data-driven methods, we utilized a deep convolutional neural network to construct a visual object space extracted from deep layer activations for a reference set of objects. The first two dimensions of object space capture attributes that distinguish one object from another and may map onto the anatomical organization of the ventral visual stream of non-human primates, but its existence in humans is unclear. If object space underlies human visual cognition, this should have a measurable effect on people´s object discrimination. We focused on the so-called face quadrant of object space, as there are considerable individual differences in face discrimination abilities, and a lively debate has revolved around whether such abilities are specific to faces. In this preregistered study, we collected data from a large (N=511) heterogeneous sample to amplify individual differences in visual discrimination. We primarily targeted people with self-declared face recognition abilities on opposite sides of the spectrum, ranging from poor to excellent face recognizers. We then administered a visual foraging task where faces and other objects had to be discriminated. Stimuli were carefully picked out of tens of thousands of images to either match or not match with the coordinates of faces in object space. Results showed that face discrimination did not rely on completely domain-specific abilities but also tapped into mechanisms that supported other object discrimination. Specifically, people who struggled with telling apart faces also had difficulties with discriminating other objects from the face quadrant of object space. To the best of our knowledge, the current study provides the first behavioral evidence for the existence of an object space in human visual cognition.

Keywords: object perception, face perception, high-level vision, foraging, neural networks, domain specificity, domain generality

Introduction

For decades, people have debated how the brain represents visual objects [1,2]. In anterior regions of the ventral visual pathway, which supports objects discrimination and recognition [1,3-5], neurons can show apparently idiosyncratic activation patterns, responding vigorously to only some objects [6], but what these preferred objects have in common may not be apparent (but see e.g. [7]). Some large-scale organization of the primate ventral stream is still known, where patches of cortex respond greatly to particular categories such as faces [8-10]. Consistent with this topographical organization, recognition deficits for faces as well as unusually good facial recognition ability have been documented, but domain-specificity vs. domain-generality is highly debated [11-28]; individual differences in object perception might not fall into neat categories, and regions that are supposedly category-selective also respond to other categories, albeit to a graded degree.

In the end, the brain might not "care" about our easily labelled or verbalized object categories. A recent paper provides strong evidence for the idea that the high-level ventral stream in non-human primates is anatomically organized not around categories, but according to the first two axes of *object space* – whose visual feature dimensions can be extracted from a deep neural network trained on visual image classification [9]. The two dimensions are hard to label but may be roughly approximated by how animate or inanimate objects appear to be, and how stubby or spiky they are. According to this line of thinking, categories such as faces evoke activity in specific anatomical regions of the ventral visual pathway due to the *visual* qualities that they happen to have, as almost all human faces fall within the stubby animate-looking quadrant of object space, also known as the face quadrant. Other objects that share these qualities also evoke activity in the same brain regions [9]. Evidence for a visual object space in the human visual system is however scarce. Two functional neuroimaging studies that specifically looked for its existence produced contradictory results. While one found some evidence for an anatomical organization according to object space [29], the other rather supported a categorical organization [30].

In the current preregistered study (https://osf.io/q5ne8)[1], we take a completely different approach. We reasoned that if object space underlies human visual cognition, this should have a measurable effect on people´s object discrimination abilities. We focused on the stubby animate-looking face quadrant of object space, as there are considerable individual differences in the ability to discriminate faces, and a lively debate has revolved around whether abilities can be specific to faces [10-13,15-17,19-21,23-25,28,31-51]. We hypothesized that face discrimination ability would be particularly associated with the ability to discriminate novel stubby animate-looking objects (preregistered hypothesis 1) and familiar stubby animate-looking objects (preregistered hypothesis 2). We separate novel and familiar objects as their association with faces may differ; while novel objects primarily have visual qualities, real objects differ in several ways, including appearance, semantics, and people's experience with them, and these are often hard to tease apart.

We took images of various things and ran them through a convolutional neural network pretrained on image classification[52-55] and extracted layer activations of a relatively deep layer. These layers may extract complex stimulus features similar to those represented in more anterior regions of the ventral visual pathway [9,56]. We then performed principal components analysis on these layer activations and extracted the first two principal components that explained the greatest amount of variance in the appearance of these

[1] This paper is on object space. We also preregistered hypotheses on so-called face space (https://osf.io/q5ne8), but these depended on there being measurable stable individual differences in the usage of face space. We found no such reliable individual differences, making our hypotheses regarding face space untestable.

reference objects. This allowed us to construct a two-dimensional object space that captures the two main dimensions on which objects visually differ from one another (figure 1). We then took tens of thousands of other object images and projected them onto our object space. Knowing the location of each image in object space, we could pick objects, both novel and familiar, from all quadrants of object space (stubby animate-looking, spiky animate-looking, stubby inanimate-looking, and spiky inanimate-looking). Our stimulus set included novel objects from each quadrant, previously used to anatomically map object space in non-human primates [9]. For each novel stubby animate-looking object, we found images of faces with closely matched object space coordinates. We also found matching stubby animate-looking birds, hands, seashells, and bowties. We did the same for the other three quadrants. Therefore, our stimulus set also included spiky animate-looking birds and hands, stubby inanimate-looking mattresses and buckets, and spiky inanimate-looking hangers and exercise equipment (figures 2 and 3).

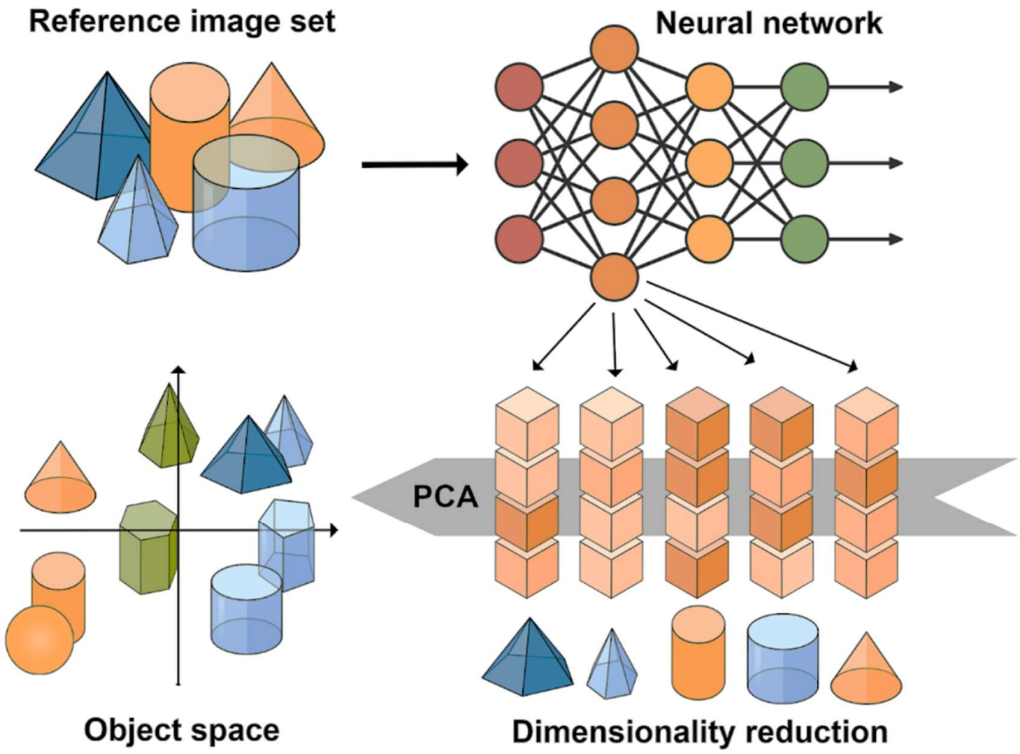

Figure 1. Defining object space. The image demonstrates our overall approach to defining a visual object space. The depicted objects, network structure, activation patterns, and object space are purely illustrative. A convolutional neural network (VGG-16bn; [52], see also [53]) pretrained on ecologically relevant image categories (ecoset; [54,55]) was used to extract deep layer activations for a reference image set (large and small animate and inanimate objects; [57]). Principal components analysis (PCA) on these activations was used to define an object space that should capture the main two dimensions on which the reference objects vary in appearance. Network activation patterns of other objects were then projected to this object space to extract their coordinates.

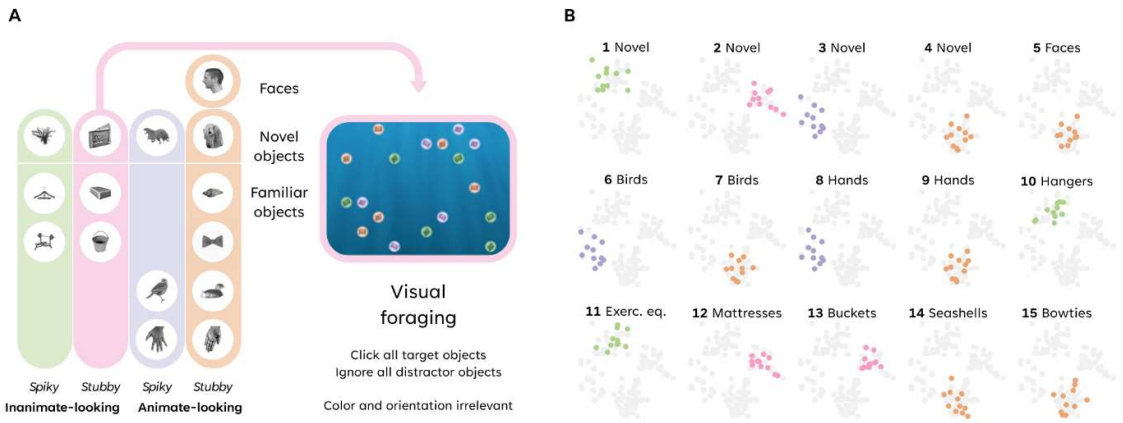

Figure 2. Visual foraging stimuli and their object space coordinates. A. Stimuli belonged to one of four quadrants of object space. They were used in a visual foraging task where people searched for targets among distractors. **B.** Object space coordinates of stimuli (11 exemplars of each type). Quadrants of object space are represented by different colors (green: spiky inanimate-looking; pink: stubby inanimate-looking; purple: spiky animate-looking; orange: stubby animate-looking; gray: all objects). Each dot represents object space coordinates of a particular exemplar object image.

These stimuli were then used in a visual foraging task (figure 2) where a large sample of participants (N = 511) with greatly varying face processing abilities had to search for targets among distractors. On each experimental trial, a sample stimulus was briefly shown at screen center, followed by the appearance of multiple targets identical to the sample, and multiple distractors that depicted another exemplar of the same object type. Participants were asked to click on all targets as fast as possible and avoid distractors. Color overlay and orientation of targets and distractors varied to discourage low-level visual matching strategies. The experiment also included control trials without objects where people searched for black targets among white distractors, or white targets among black distractors. Discrimination ability was then defined for each trial type. A person's face discrimination ability was therefore defined as their foraging speed on face trials, bucket discrimination ability was foraging speed on bucket trials, and so on. We then assessed whether individual differences in face discrimination depended on their location in visual object space.

Results

Almost all faces belong to the stubby animate-looking quadrant of object space (face quadrant). We hypothesized (preregistration: https://osf.io/q5ne8) that face foraging would be positively associated with foraging for novel stubby animate-looking objects (hypothesis 1) and with familiar stubby animate-looking objects (hypothesis 2).

We assessed these hypotheses by performing two preregistered hierarchical regressions with face discrimination ability as the dependent variable. In both cases, demographics (age, gender, and educational level) were entered at step 1, and control foraging performance was added at step 2. At stage 1 of each model, age, gender, and education accounted for around 37% of the variance in face discrimination. Higher educational level was associated with better face discrimination ability, while advanced age was associated with poorer face discrimination ability (supplementary figures s1 and s2). Adding the control trials at stage 2 greatly improved the model which now explained around 66% of the variance in the dependent variable. Control foraging should capture various factors of little interest that do not directly involve object discrimination ability, and likely provides an upper limit for a person's foraging speed.

On step 3 of the first hierarchical regression, we added three factors: foraging performance for novel spiky inanimate-looking objects, novel stubby inanimate-looking

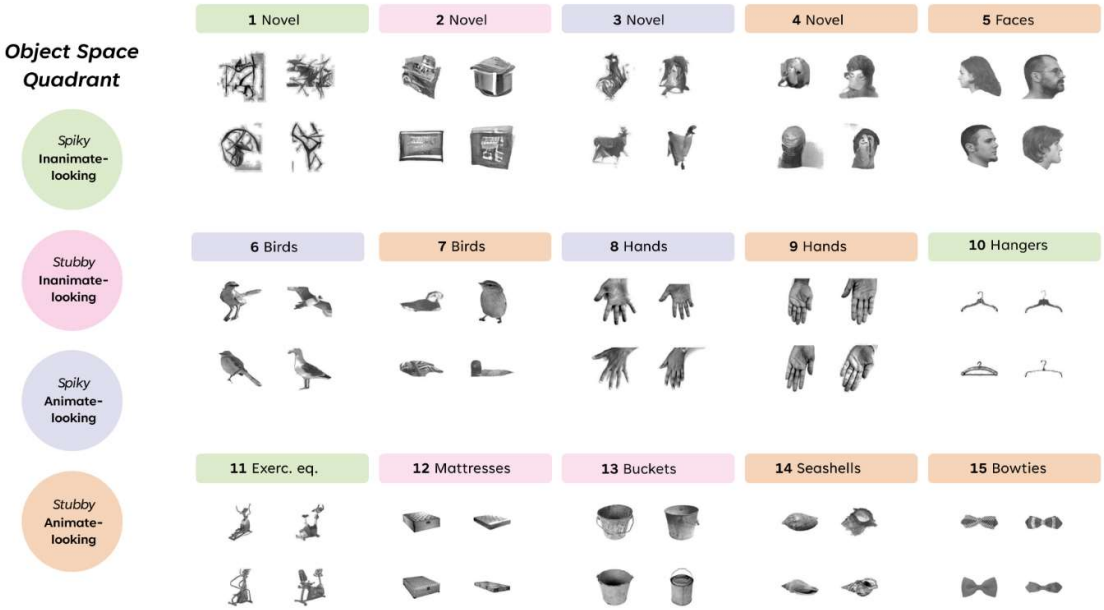

Figure 3. Example objects. Example stimuli of 15 object types are shown. Quadrants of object space are represented by different colors (orange: stubby animate-looking; purple: spiky animate-looking; pink: stubby inanimate-looking; green: spiky inanimate-looking; gray: all objects).

objects, and novel spiky animate-looking objects. The included variables now accounted for 82% of the variance in face discrimination ability. The greatly improved model likely reflects individual differences in a general object perception factor [58,59] that contributes to people's ability to discriminate all sorts of objects, including faces. Finally, and importantly for hypothesis 1, adding foraging performance for novel stubby animate-looking objects at stage 4 significantly (p = 4.95x10^{-8}, two-sided) improved the model, with 83% explained variance (table 1).

On step 3 of the second hierarchical regression, we added six factors: foraging performance for spiky animate-looking birds and hands, spiky inanimate-looking hangers and exercise equipment, and stubby inanimate-looking buckets and mattresses. In step 4 of this second regression, we added four factors: foraging performance for stubby animate-looking birds, hands, bowties, and seashells. This again greatly improved the model from that of stage 2, where the included factors explained a total of 83% of the variance in face foraging. Importantly for hypothesis 2, adding foraging performance for familiar stubby animate-looking objects (birds, hands, bowties, seashells) further improved the model (p = 7.28x10^{-4}, two-sided; table 1) where all included factors now explained a total of 84% of variance in face discrimination ability (table 1). Both preregistered hypotheses 1 and 2 were confirmed.

The variance in face foraging uniquely explained by foraging for stubby animate-looking objects was small. This is not unexpected given that a person's ability to forage for faces is already extremely well predicted from demographics and performance on other types of foraging trials, leaving little non-noise variance left to explain. We nonetheless chose to take a closer look at our results. The regression coefficient for novel stubby animate-looking object foraging was larger than all other types of foraging included as independent predictors in the final model, as predicted by hypothesis 1 (table 1). The regression coefficient for novel spiky inanimate-looking objects was the smallest of all novel object foraging types, in accordance with these objects being the furthest in object space from faces. Regression coefficients for familiar object foraging showed a less clear pattern (table 1). This can be deceptive as regression coefficients reflect the *unique*

contribution of a particular condition. When multiple conditions with a common factor (e.g., several different stubby animate-looking things in stage 4) are added, this common attribute of interest (e.g., "stubby-animateness") is essentially factored out. This less clear pattern could however also be due to the difficulty of fully separating semantics and previous experience with object categories from their visual qualities.

We can however do this for two categories as we had access to over ten thousand images of birds and hands, allowing us to include both stubby animate-looking and spiky animate-looking items. We therefore followed up with exploratory analyses where we focused on unique variance in face discrimination ability explained by stubby vs. spiky familiar objects of the same category (figure 4). This showed that face discrimination was more associated with average performance for stubby animate-looking birds and hands, partialling out average performance for spiky animate-looking birds and hands ($r_{partial}$ = 0.425), than with average performance for spiky animate-looking birds and hands, partialling out average performance for stubby animate-looking birds and hands ($r_{partial}$ = 0.299). We therefore are more confident in our assessment that individual differences in face discrimination are partially individual differences in the processing of stubby animate-looking objects.

To visualize our results for hypotheses 1 and 2, we performed multidimensional scaling on individual differences in foraging. As can be seen in figure 5, faces formed a separate cluster, indicating that people's performance on a given face foraging trial can generally best

Table 1. Association between face discrimination and discrimination of stubby animate-looking objects. The table shows the final hierarchical regression models for preregistered hypotheses 1 and 2 after completing all four steps. Dependent variable is face discrimination ability. Model 1: Residual standard error = 0.072 on 502 degrees of freedom, R^2 = 0.831, F-statistic = 309.1 on 8 and 502 degrees of freedom, $p < 2.2x10^{-16}$. Model 2: Residual standard error = 0.071 on 496 degrees of freedom, R^2 = 0.838, F-statistic = 183.5 on 14 and 496 degrees of freedom, $p < 2.2x10^{-16}$. B-values are unstandardized regression coefficients.

	Hypothesis 1: Novel		*Hypothesis 2: Familiar*		
Factor	b-value	Std. Error	b-value	Std. Error	Added on
Intercept	0.118	0.035	0.072	0.037	step 1
Age	-0.012	0.003	-0.009	0.003	
Gender	0.007	0.007	0.012	0.007	
Education	0.009	0.004	0.006	0.004	
Control foraging	0.073	0.025	0.082	0.027	step 2
Novel spiky-inanimate	0.126	0.041	-	-	step 3
Hanger spiky-inanimate	-	-	0.128	0.055	
Exer. eq. spiky-inanimate	-	-	0.005	0.041	
Novel stubby-inanimate	0.137	0.047	-	-	
Bucket stubby-inanimate	-	-	0.146	0.048	
Mattr. stubby-inanimate	-	-	0.032	0.037	
Novel spiky-animate	0.220	0.047	-	-	
Bird spiky-animate	-	-	0.160	0.048	
Hand spiky-animate	-	-	0.017	0.036	
Novel stubby-animate	0.296	0.054	-	-	step 4
Bird stubby-animate	-	-	0.145	0.054	
Hand stubby-animate	-	-	0.083	0.083	
Bowtie stubby-animate	-	-	0.081	0.046	
Seashell stubby-animate	-	-	0.043	0.053	

predict their performance on another trial of the same type. Of note, the face cluster is closer to the stubby animate-looking cluster than it is to the cluster for other objects, both for novel objects (figure 5a, hypothesis 1) and familiar objects (figure 5b, hypothesis 2), indicating that foraging performance for stubby animate-looking objects more closely predicts face foraging.

The preregistered analysis pipeline for hypotheses 1 and 2 concerned predicting foraging performance for faces matched for object space coordinates to other stubby animate-looking objects (see figure 2b). However, we assessed face processing in additional ways. We also included difficult faces foraging trials where cues unrelated to facial identity (e.g., viewpoint) were minimized. We did this by hand-picking faces of 11 male individuals looking slightly to the left with a similar expression, similar hairstyle, and no beard or glasses. These 11 faces formed a tight cluster in the stubby animate-looking quadrant of object space (face quadrant), consistent with their visual similarity. We refer to performance on these trials as self-matched face discrimination ability. Participants also filled out a questionnaire on their face recognition abilities in everyday life (20-item Prosopagnosia Index, PI20 [60]), and tried to recognize previously memorized faces (Cambridge Face Memory Test, CFMT [61]). As defined in our preregistration, we combined PI20 and CFMT into a single face recognition index.

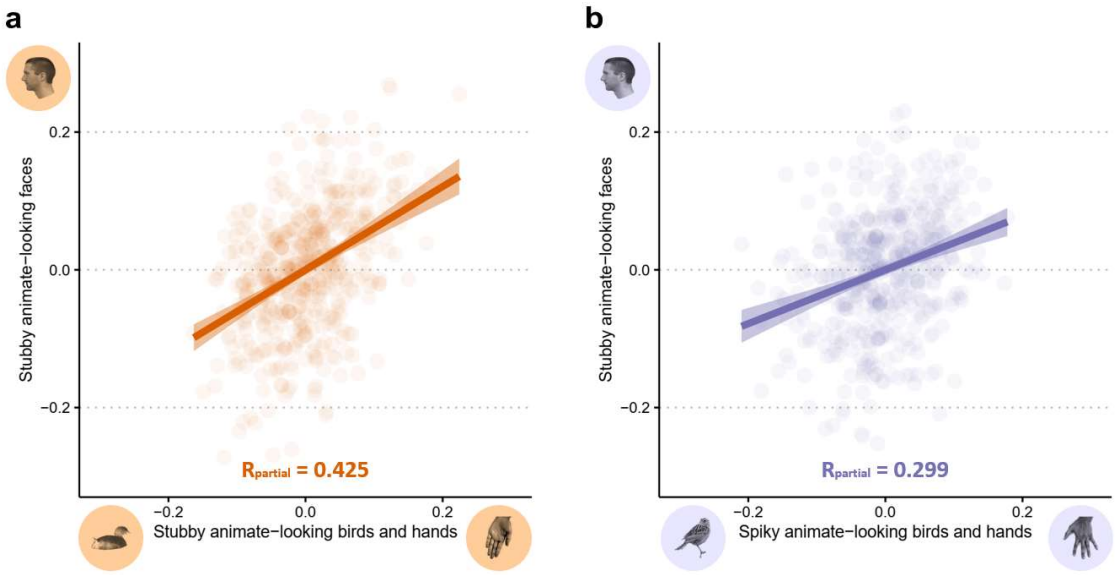

Figure 4. Association between face discrimination ability and stubby vs. spiky animate-looking birds and hands. A. Partial correlation plot between face discrimination ability and average discrimination ability for stubby animate-looking birds and hands; axes show residuals after partialling out average discrimination ability for spiky animate-looking birds and hands. **B.** Partial correlation plot between face discrimination ability and average discrimination ability for spiky animate-looking birds and hands; axes show residuals after partialling out average discrimination ability for stubby animate-looking birds and hands (targets per second).

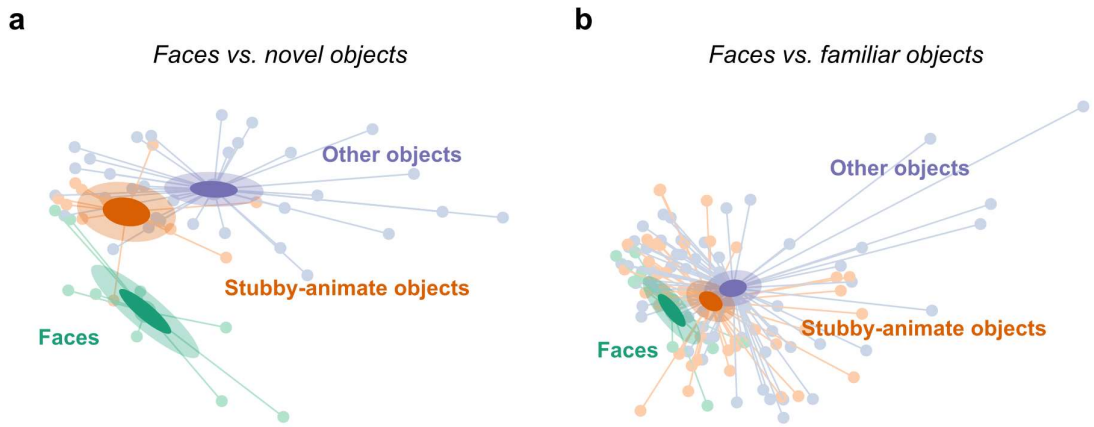

Figure 5. Clustering of individual differences in visual foraging. Multidimensional scaling was performed on the correlation matrix for performance (targets per second) on individual trials. This was done separately for a) trials with novel objects and faces matched to objects, and b) trials with familiar objects and faces matched to objects. Each dot depicts a single trial and the distance between two points is approximately equal to their dissimilarity. Confidence ellipses for barycenters of three clusters are shown, both for the 95% confidence level (light ellipses) and the 50% confidence level (dark ellipses).

We wanted to see if discrimination ability for stubby animate-looking objects could predict face processing more generally. We therefore took a better look at our additional measures of face processing ability. As our preregistered hypotheses held for both novel and familiar objects, we collapsed across them to make two indices (based on foraging excluding all face and control trials): discrimination ability for stubby animate-looking objects and for other objects. In an exploratory analysis, we asked: Is discrimination ability for stubby animate-looking objects associated with the ability to a) discriminate self-matched faces, and b) recognize faces? We performed two multiple regressions with demographics, control foraging, discrimination ability for stubby animate-looking objects, and discrimination ability for other objects as predictors. In both cases, discrimination ability for stubby animate-looking objects was an independent predictor of an individual's face processing ability (table 2). This further supports that individual differences in face processing are in part individual differences in the processing of stubby animate-looking objects.

Table 2. Predicting self-matched face discrimination and face recognition. The table shows two regression models, M1 with self-matched face discrimination ability and M2 for face recognition as dependent variables. M1: Residual standard error = 0.100 on 504 degrees of freedom, $R^2 = 0.681$, F-statistic = 179 on 6 and 504 degrees of freedom, $p < 2.2\text{x}10^{-16}$. M2: Residual standard error = 0.826 on 504 degrees of freedom, $R^2 = 0.11$, F-statistic = 10.33 on 6 and 504 degrees of freedom, $p = 8.34\text{x}10^{-11}$. B-values are unstandardized regression coefficients.

	M1: Self-matched faces			*M2: Face recognition*		
Factor	b-value	Std. Error	p-value	b-value	Std. Error	p-value
Intercept	0.038	0.049	0.447	-1.974	0.407	$1.64\text{x}10^{-6}$
Age	-0.015	0.005	0.002	0.080	0.039	0.042
Gender	0.016	0.010	0.123	0.025	0.083	0.765
Education	0.011	0.005	0.021	-0.021	0.041	0.610
Control foraging	-0.098	0.033	0.003	-0.763	0.272	0.005
Other objects	0.031	0.013	0.023	0.005	0.110	0.961
Stubby animate-looking	0.128	0.025	$4.47\text{x}10^{-7}$	0.544	0.207	0.009

Discussion

What are the diagnostic dimensions on which objects truly differ visually, and do these dimensions provide the foundation for the organizational principles of visual object perception as evidenced by individual differences in behavior? Using data-driven methods, we utilized a deep convolutional neural network [52,53] pretrained on ecologically relevant image categories [54,55] to construct a visual object space extracted from deep layer activations for a reference set of objects. The first two dimensions of object space capture attributes that distinguish one object from another and may map onto the anatomical organization of the ventral visual stream of non-human primates [9], but its existence in humans is unclear [29,30].

We focused on the face quadrant of object space, occupied by stubby animate-looking objects such as faces, as faces evoke activity in specific ventral visual regions [8-10,20], brain damage can lead to severe face recognition deficits[62], and abilities in face processing vary greatly in the general population [26,27,33,63]. If the human visual system is organized around an object space, the processing of faces should partially depend on neural structures that also process other stubby animate-looking objects. Supporting this, we found association-dissociation patterns in individual differences in visual cognition where face processing was specifically associated with the ability to tell apart other stubby animate-looking objects.

So far, two neuroimaging studies on a possible human object space have provided somewhat conflicting findings [29,30]. Using stimuli originally used to map object space in macaques [9], Coggan and Tong [29] found extensive activations in the human ventral stream related to aspect ratio (stubby-spiky) and animacy. Yargholi and Op de Beeck [30] on the other hand found little evidence for the human ventral stream being organized according to aspect ratio while anatomical organization was partially explained by animacy. While the former study [29] may better fit with the current results, activations in the human brain were not arranged in a clear sequence of complete maps of aspect-ratio/animacy object space as reported in non-human primates. At a first glance, the latter study [30] then appears to contradict the current study.

We believe that a likely explanation for the apparently diverging results has to do with differences in the operational definitions of variables. We defined our object space by data-driven techniques using deep neural networks that capture image statistics of objects. Yargholi and Op de Beeck [30] defined aspect ratio as a function of perimeter and area. As methodology was different, our results may not contradict those of Yargholi and Op de Beeck [30] after all. Their stimuli were not linearly separable by aspect ratio in a 2D object space derived from the latent space of an artificial neural network, so is still an open question whether the human ventral stream is anatomically organized around the dimensions of object space as defined in such a way. These dimensions may or may not map directly onto *any* human labels, but covary with several of them, including animate-inanimate, stubby-spiky, and big-small real-word size [57].

Another possible reason for discrepancies between studies is familiarity. In both Coggan and Tong [29] and Yargholi and Op de Beeck [30], stimuli were primarily or exclusively familiar objects, making it hard to separate appearance from meaning and prior experience. Of note, neither neuroimaging study separated *being* animate/inanimate from *looking* animate/inanimate, and in Coggan and Tong [29] all stubby animate-looking objects were human faces. Familiarity with objects may play a greater role when mapping out visual object space in humans compared to monkeys. Even though Bao et al. [9] and Coggan and Tong [29] largely used the same stimuli, many of the things that have meaning to a human – such as a trumpet, camera, or a duck – may as well be "novel objects" for monkeys as they have no actual prior experience with them. Visual and semantic properties may both be represented in anterior regions of the human visual system [64]. Indeed, our results were particularly clear for novel objects as opposed to familiar objects. As experience with objects can shape the visual system [65], the functional mapping of the anatomical organization of human object

space may be clearer if based on objects with which people have no prior experience.

We suggest that fundamental visual properties – here approximated by two primary dimensions of an object space captured by CNNs – are automatically extracted by the human visual system during an initial bottom-up sweep. These properties may guide the selection of appropriate methods of further object processing, as people are sensitive to various conceptual and perceptual properties of objects [66]. For example, the processing of animate objects may be skewed toward global, configural, or holistic processing [67-69] while inanimate objects could be preferentially processed in a piecemeal or feature-based manner [70], and the processing of animate objects could additionally need greater integration across time due to their dynamic nature [71] (see also [72]). Spiky objects could lend themselves to axis-based shape processing (shape skeletons) [73-80] or structural representations (parts and their relations) [2,81,82] while image-based or view-based representations [2,83-85] may be more appropriate for stubby objects.

To conclude, to the best of our knowledge, the current study provides the first behavioral evidence for the existence of an object space in human visual cognition.

Method

The study followed a protocol (SHV2021-001) reviewed by the Science Ethics Committee of Icelandic Universities. Study design, data collection, and data analysis were preregistered (https://osf.io/q5ne8). The study was run online through Pavlovia (https://pavlovia.org/) using PsychoPy [86,87].

Participants

Participants were unpaid volunteers recruited by word of mouth, email listservs, social media, and other media. We primarily targeted people with self-declared face recognition abilities on opposite sides of the spectrum, ranging from poor to excellent face recognizers, but all people were nonetheless encouraged to participate. A total of 514 people completed all tasks. Thereof, three participants had mean targets per second at least three standard deviations away from the sample mean in visual foraging. These data were excluded from further analysis, leaving us with a final sample of 511 participants: 357 female, 144 male, and 10 non-binary, other gender, or chose not to answer. 25 people were between the ages of 18-24, 86 were 25-34 years, 134 were 35-44 years, 132 were 45-54 years, and 134 were 55 years of age or older. All participants reported normal or corrected-to-normal vision and gave informed consent.

Procedure

Participants first completed a short speed click task where they clicked on several consecutive targets. Those who did this slowly were reminded to use a computer mouse instead of a touchpad and had to repeat the task before being allowed to continue. Participants then answered background questions on their age bracket, gender, and educational level. Next, people completed the 20-item prosopagnosia index (PI20 [60], Icelandic version), a self-report on problems with faces, and the Cambridge Face Memory Test (CFMT [61]) which measures people's memory for unfamiliar faces. Finally, participants completed a visual foraging task where they searched for target objects among distractor objects. Participants were given feedback on their general performance upon study completion.

Visual foraging

In this main experimental task, people searched for faces as well as other familiar and novel objects. These objects either shared or did not share visual qualities with faces as objectively defined by their coordinates in a visual object space derived from activation patterns in a deep layer of a neural network trained on object classification. We will refer to the dimensions of object space as animate-inanimate-looking and stubby-spiky-looking, following Bao et al. [9], but note that these are only approximate labels for these main dimensions on which objects differ visually.

The foraging task included 11 trials from each of the following 17 conditions (187 trials per person): seven stubby animate-looking object types (novel objects, hands, birds, bowties, seashells, faces matched to novel

objects, faces matched to each other), three spiky animate-looking object types (novel objects, hands, birds), three spiky inanimate-looking object types (novel objects, exercise equipment, hangers), three stubby inanimate-looking object types (novel objects, buckets, mattresses), and a feature foraging control condition. Each object image appeared on two trials in the experiment, once as a target and once as a distractor. Each target-distractor pair only appeared once, so if image A was a target and B a distractor in one trial, then B was a target and not-A a distractor in another trial. Each trial in the visual foraging task consisted of a sample stimulus shown at screen center for 1 second followed by 9 targets that were identical to the sample and 9 distractors that showed another object not identical to the sample (figure 2). Targets and distractors within a trial always belonged to the same condition. In all but control trials, stimulus orientation and color of foraging items was varied to discourage the use of low-level visual strategies. Participants foraged for targets as fast as they could by clicking each of them with the mouse. Targets disappeared as they were clicked; distractors jiggled and stayed on-screen. The trial ended when all targets had been clicked or when 10 seconds had passed from the onset of the foraging items, whichever came first. A trial was followed by feedback where 0 to 9 stars were shown depending on the number of targets clicked.

Image Sets

We defined object space using a reference set of 120 animals and 120 inanimate objects with large and small real-world size (image courtesy Konkle Lab: http://konklab.fas.harvard.edu/ [57]). Images for the foraging study were chosen from five additional image sets that contained over 30,000 images of various objects.

A set of 80 novel objects (images provided by the Tsao Lab: https://tsaolab.berkeley.edu/), were projected onto the object space. The images originally formed four groups of 20 items that mapped onto the four quadrants of object space (stubby-animate, stubby-inanimate, spiky-animate, and spiky-inanimate) in Bao et al. ([9] their extended data fig. 6) and evoked preferential activity in different networks of the macaque inferotemporal cortex. A set of 4760 images of various familiar objects (image credit: the University of California San Diego (UCSD) Vision and Memory Lab: https://bradylab.ucsd.edu/stimuli.html [88]) were also projected onto our object space. An image set of 3362 Caucasian faces of 106 individuals (image credit: Michael J. Tarr, Carnegie Mellon University, http://www.tarrlab.org/ – funding provided by NSF award 0339122 [89]) was projected onto our object space. Additionally, we used images from 11k Hands (image courtesy Mahmoud Afifi: https://sites.google.com/view/11khands [90]), and Caltech-UCSD Birds-200-2011 (image courtesy Caltech Vision / UCSD Vision: http://www.vision.caltech.edu/visipedia/CUB-200-2011.html [91]).

All objects were stripped of their background (if they had one) so that only a single object was shown on a blank background. To standardize size and borders, images were processed in IrfanView (https://www.irfanview.com). All images were cropped and then resized so that the longer side of each object (vertical or horizontal) was the same size as the longer side of all other objects (bounding box longer side 680 pixels). We then added a square, white background to all images, before we resized them again to 224 x 224 pixels. All images were therefore square, equally sized, with an equally wide, white border around them. Images were changed to grayscale and run through the SHINE Toolbox [92] to equate mean brightness and contrast.

Neural Network Activation

All images were run through the convolutional neural network VGG-16bn [52] trained on the ecoset image set [55] using batch normalization (code implemented through the Python toolbox THINGSvision [53] https://github.com/ViCCo-Group/THINGSvision). The activation from the second fully connected (fc) layer corresponding to each image was extracted – a vector of length 4096, with each vector item corresponding to the activation of a particular node.

Defining an Object Space

The layer activations corresponding to the reference set were run though Principal Component Analysis (PCA). The first two principal components were extracted; they explained a total of 20.1% of variance in the original 4096 variables.

Finding and Matching Stimuli

The four types of novel objects (that looked stubby-animate, stubby-inanimate, spiky-animate, and spiky-inanimate [9]) were projected onto our object space. 11 images from each type were selected to minimize overlap in object space location between images from different types. Images from four other image sets were then projected onto object space and for each novel object image, an image was found that projected onto the same or very similar location in object space. These images were therefore matched in visual qualities on dimensions of object space. Objects from six categories of familiar inanimate objects were matched with the novel objects: bowties and seashells with stubby animate-looking novel objects, mattresses and buckets with stubby inanimate-looking novel objects, and hangers and exercise equipment with spiky inanimate-looking novel objects. Few familiar inanimate objects projected to the spiky animate-looking quadrant of object space, and we therefore could not match them to corresponding novel objects. Hands and birds clustered in the animate-looking part of object space and were matched with animate-looking stubby and spiky novel objects; no good matches for hands or birds were found for inanimate-looking novel objects. Exclusion criteria for hand images were jewelry or salient sleeves. If two images that matched with novel objects were clearly of the same hand, the image with an object space distance further from its assigned novel object was excluded and a new match found. Exclusion criteria for bird images were more than one bird in each image. As expected, most faces projected to the stubby-animate quadrant (face quadrant) of object space and were therefore only matched with stubby animate-looking novel objects. Exclusion criteria for faces were as follows: sunglasses or obvious wigs. If two images of the same person were matched with different novel objects, the image with a distance in object space further from its assigned novel object was excluded and a new match found. In addition to these 15 types of target-distractor pairs, we included difficult face foraging trials where we handpicked 11 images of men that had similar facial features, hairstyles, and expressions, and were all looking slightly to the left. Lastly, we included 11 feature foraging control trials where participants searched for black discs among white discs and vice versa. 17 types of target-distractor pairs were therefore shown in the experiment: four types of novel objects, six types of familiar inanimate objects, two types of birds, two types of hands, two types of faces, and feature foraging control stimuli.

Target-Distractor Pairs

The novel objects were randomly paired together within each object space quadrant (stubby-animate, stubby-inanimate, spiky-animate, spiky-inanimate) as targets and distractors. Mean target-distractor distance was not significantly different between the four quadrants. Control trials consisted of black targets and white distractors or white targets and black distractors. For each of the other conditions, the target-distractor pairs were matched with the novel object pairs. For instance, if hand A was matched with novel object A and hand B with novel object B, then hand A and hand B were selected as a target-distractor pair, because novel objects A and B were paired together. Across the experiment, there were 11 unique target-distractor pairs per condition. The target and the distractor were separated in object space by a known distance varying from small (similar object space coordinates) to large (dissimilar object space coordinates). This distance should vary in a similar way for each experimental condition. Faces from the face quadrant of object space matched to each other are an exception; they were not chosen to match any other objects but were hand-picked by the researchers. These 11 faces formed a tight cluster in object space, consistent with their visual similarity. We include these faces even

though they are not matched to the other objects to include particularly challenging face foraging trials where cues unrelated to facial identity (e.g., viewpoint) are minimized.

Data Analysis

Target-distractor visual dissimilarity in the foraging task was defined as Euclidean distance in object space. Object discrimination performance was defined as the number of targets clicked per second. For each condition, separately for each participant, we calculated the median targets per second.

Two preregistered hierarchical regressions were conducted to assess whether face foraging was associated with a) novel stubby animate-looking objects, and b) familiar stubby animate-looking objects. Demographic variables were entered in the first step and were dummy coded as follows: Gender was coded as 1 = male, 0 = female/non-binary/other/choose not to answer. Age was coded as 1 = 18-24 years, 2 = 25-34 years, 3 = 35-44 years, 4 = 45-54 years, and 5 = 55 years and older. Educational level was coded as 1 = completed primary school, 2 = completed secondary school/gymnasium or compatible degrees, 3 = completed an undergraduate degree such as BSc, BA, BEd, and 4 = completed a graduate degree such as MA, MSc, PhD. Control foraging was added in step two of the hierarchical regressions, foraging for objects other than stubby animate-looking was added in step 3, and lastly foraging for stubby animate-looking objects was added in step 4.

For the purpose of a follow-up analysis on stubby and spiky-looking hands and birds, new indices of stubby-looking birds and hands and spiky-looking birds and hands were made. For each participant, the mean of a) the median targets per second for stubby-looking hands and stubby-looking birds, and b) spiky-looking birds and spiky-looking hands was calculated.

Because the correlation between scores on the PI20 and the CFMT reached a preregistered threshold (-0.5), these measures were combined into a single face recognition measure. For each participant, scores from PI20 and CFMT were z-scored, the z-score of PI20 was negated (so e.g., 1.5 became -1.5 and vice versa) and the mean of these were taken to indicate a person's face recognition ability.

For exploratory multiple regressions, new indices of mean stubby animate-looking objects and mean non-stubby animate-looking objects were calculated. Mean stubby-animate: For each participant, the mean of the median targets per second for novel stubby animate-looking objects, stubby animate-looking birds, stubby animate-looking hands, stubby animate-looking bowties, and stubby animate-looking seashells was calculated. Mean non-stubby-animate: For each participant, the mean of the median targets per second for novel spiky inanimate-looking objects, novel stubby-inanimate looking objects, novel spiky animate-looking objects, spiky animate-looking birds, spiky animate-looking hands, spiky inanimate-looking hangers, spiky inanimate-looking exercise equipment, stubby inanimate-looking buckets, and stubby inanimate-looking mattresses was calculated.

Open Science Practices

This study was preregistered (https://osf.io/q5ne8). Data and analysis code are publicly available (https://osf.io/2jn6c/).

Acknowledgements

We sincerely thank Sunneva Líf Albertsdóttir, Hannes Hansen, Hildur Franziska Hávarðardóttir, Martin Hebart, Unnur Andrea Ásgeirsdóttir, Ísak Örn Ívarsson, Bahareh Jozranjbar, Tim C. Kietzmann, Lukas Muttenthaler, Sarah Rutherford, Randi Starrfelt, Guðrún Þorláksdóttir, and all the people who shared images used in this study. This research was supported by the Icelandic Research Fund (grants no. 228916 and 218092) and the University of Iceland Research Fund.

References

1 Logothetis, N. K. & Sheinberg, D. L. Visual object recognition. *Annual review of neuroscience* **19**, 577-621 (1996).

2 Peissig, J. J. & Tarr, M. J. Visual object recognition: Do we know more now than we did 20 years ago? *Annu. Rev. Psychol.* **58**, 75-96 (2007).

3 Ungerleider, L., Mishkin, M., Ingle, D., Goodale, M. & Mansfield, R. Analysis of visual behavior. (1982).

4 Ungerleider, L. G. & Haxby, J. V. 'What'and 'where'in the human brain. *Current opinion in neurobiology* **4**, 157-165 (1994).

5 Milner, D. & Goodale, M. *The visual brain in action*. Vol. 27 (OUP Oxford, 2006).

6 Woloszyn, L. & Sheinberg, D. L. Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron* **74**, 193-205 (2012).

7 Tanaka, K. Inferotemporal cortex and object vision. *Annual review of neuroscience* **19**, 109-139 (1996).

8 Grill-Spector, K. & Weiner, K. S. The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience* **15**, 536-548 (2014).

9 Bao, P., She, L., McGill, M. & Tsao, D. Y. A map of object space in primate inferotemporal cortex. *Nature*, 1-6 (2020).

10 Kanwisher, N., McDermott, J. & Chun, M. M. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience* **17**, 4302-4311 (1997).

11 Farah, M. J. *Visual agnosia*. (MIT press, 2004).

12 Farah, M. J., Wilson, K. D., Drain, H. M. & Tanaka, J. R. The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision research* **35**, 2089-2093 (1995).

13 Robotham, R. J. & Starrfelt, R. Face and word recognition can be selectively affected by brain injury or developmental disorders. *Frontiers in psychology* **8**, 1547 (2017).

14 Behrmann, M., Nelson, J. & Sekuler, E. Visual complexity in letter-by-letter reading: Pure alexia is not pure. *Neuropsychologia* **36**, 1115-1132 (1998).

15 Behrmann, M. & Plaut, D. C. Distributed circuits, not circumscribed centers, mediate visual recognition. *Trends in cognitive sciences* **17**, 210-219 (2013).

16 Bukach, C. M., Gauthier, I. & Tarr, M. J. Beyond faces and modularity: the power of an expertise framework. *Trends in cognitive sciences* **10**, 159-166 (2006).

17 Burns, E. J., Arnold, T. & Bukach, C. M. P-curving the fusiform face area: Meta-analyses support the expertise hypothesis. *Neuroscience & Biobehavioral Reviews* **104**, 209-221 (2019).

18 Burns, E. J. *et al.* Intact word processing in developmental prosopagnosia. *Scientific Reports* **7**, 1-12 (2017).

19 Hills, C. S., Pancaroglu, R., Duchaine, B. & Barton, J. J. Word and text processing in acquired prosopagnosia. *Annals of neurology* **78**, 258-271 (2015).

20 Kanwisher, N. Domain specificity in face perception. *Nature neuroscience* **3**, 759-763 (2000).

21 Plaut, D. C. & Behrmann, M. Response to Susilo and Duchaine: beyond neuropsychological dissociations in understanding face and word representations. *Trends*

in cognitive sciences **17**, 546 (2013).

22 Roberts, D. J. *et al.* Processing deficits for familiar and novel faces in patients with left posterior fusiform lesions. *Cortex* **72**, 79-96 (2015).

23 Rubino, C., Corrow, S. L., Corrow, J. C., Duchaine, B. & Barton, J. J. Word and text processing in developmental prosopagnosia. *Cognitive Neuropsychology* **33**, 315-328 (2016).

24 Susilo, T. & Duchaine, B. Dissociations between faces and words: comment on Behrmann and Plaut. *Trends in cognitive sciences* **17**, 545 (2013).

25 Susilo, T., Wright, V., Tree, J. J. & Duchaine, B. Acquired prosopagnosia without word recognition deficits. *Cognitive Neuropsychology* **32**, 321-339 (2015).

26 Russell, R., Duchaine, B. & Nakayama, K. Super-recognizers: People with extraordinary face recognition ability. *Psychonomic bulletin & review* **16**, 252-257 (2009).

27 Ramon, M., Bobak, A. K. & White, D. Super-recognizers: From the lab to the world and back again. *British journal of psychology* **110**, 461-479 (2019).

28 Tarr, M. J. & Gauthier, I. FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature neuroscience* **3**, 764-769 (2000).

29 Coggan, D. D. & Tong, F. Spikiness and animacy as potential organizing principles of human ventral visual cortex. *Cerebral Cortex*, bhad108 (2023).

30 Yargholi, E. & Op de Beeck, H. Category trumps shape as an organizational principle of object space in the human occipitotemporal cortex. *J Neurosci* (2023).

31 Gauthier, I., Behrmann, M. & Tarr, M. J. Can face recognition really be dissociated from object recognition? *Journal of cognitive neuroscience* **11**, 349-370 (1999).

32 Gerlach, C. & Starrfelt, R. Patterns of perceptual performance in developmental prosopagnosia: An in-depth case series. *Cognitive Neuropsychology* **38**, 27-49 (2021).

33 Geskin, J. & Behrmann, M. Congenital prosopagnosia without object agnosia? A literature review. *The Face Specificity of Lifelong Prosopagnosia*, 4-54 (2020).

34 Barton, J. J., Albonico, A., Susilo, T., Duchaine, B. & Corrow, S. L. Object recognition in acquired and developmental prosopagnosia. *Cognitive Neuropsychology* **36**, 54-84 (2019).

35 Garrido, L., Duchaine, B. & DeGutis, J. Association vs dissociation and setting appropriate criteria for object agnosia. *Cognitive neuropsychology* **35**, 55-58 (2018).

36 Gray, K. L. & Cook, R. Should developmental prosopagnosia, developmental body agnosia, and developmental object agnosia be considered independent neurodevelopmental conditions? *Cognitive Neuropsychology* **35**, 59-62 (2018).

37 Campbell, A. & Tanaka, J. W. Decoupling category level and perceptual similarity in congenital prosopagnosia. *Cognitive Neuropsychology* **35**, 63-65 (2018).

38 Gerlach, C., Lissau, C. H. & Hildebrandt, N. K. On defining and interpreting dissociations. *Cognitive Neuropsychology* **35**, 66-69 (2018).

39 Towler, K. & Tree, J. The implications of highly associated face and object recognition

impairments for the cognitive architecture. *Cognitive Neuropsychology* (2018).
40 Rosenthal, G. & Avidan, G. A possible neuronal account for the behavioural heterogeneity in congenital prosopagnosia. *Cognitive Neuropsychology* **35**, 74-77 (2018).
41 Eimer, M. What do associations and dissociations between face and object recognition abilities tell us about the domain-generality of face processing? *Cognitive Neuropsychology* **35**, 80-82 (2018).
42 Susilo, T. Vol. 35 1-3 (Taylor & Francis, 2018).
43 Nestor, A. Congenital prosopagnosia: Deficit diagnosis and beyond. *Cognitive neuropsychology* **35**, 78-79 (2018).
44 Ramon, M. The power of how—lessons learned from neuropsychology and face processing. *Cognitive Neuropsychology* **35**, 83-86 (2018).
45 de Gelder, B. & Van den Stock, J. Face specificity of developmental prosopagnosia, moving beyond the debate on face specificity. *Cognitive Neuropsychology* **35**, 87-89 (2018).
46 Barton, J. J. in *The Face Specificity of Lifelong Prosopagnosia* 90-93 (Routledge, 2020).
47 Starrfelt, R. & Robotham, R. J. On the use of cognitive neuropsychological methods in developmental disorders. *Cognitive neuropsychology* **35**, 94-97 (2018).
48 Rossion, B. Prosopdysgnosia? What could it tell us about the neural organization of face and object recognition? *Cognitive Neuropsychology* **35**, 98-101 (2018).
49 Burns, E. J. & Bukach, C. M. Face processing still predicts reading ability: Evidence from developmental prosopagnosia. A reply to Gerlach and Starrfelt (2022). *Cortex* **154**, 340-347 (2022).
50 Gerlach, C. & Starrfelt, R. Face processing does not predict reading ability in developmental prosopagnosia: A commentary on Burns & Bukach (2021). *Cortex* **154**, 421-426 (2022).
51 Burns, E. J. & Bukach, C. M. Face processing predicts reading ability: Evidence from prosopagnosia. *Cortex* **145**, 67-78 (2021).
52 Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
53 Muttenthaler, L. & Hebart, M. N. THINGSvision: a Python toolbox for streamlining the extraction of activations from deep neural networks. *Frontiers in Neuroinformatics* **15**, 679838 (2021).
54 Mehrer, J., Kietzmann, T. C. & Kriegeskorte, N. in *Conference on Cognitive Computational Neuroscience. New York, NY, USA.*
55 Mehrer, J., Spoerer, C. J., Jones, E. C., Kriegeskorte, N. & Kietzmann, T. C. An ecologically motivated image dataset for deep learning yields better models of human vision. *Proceedings of the National Academy of Sciences* **118**, e2011417118 (2021).
56 Güçlü, U. & van Gerven, M. A. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience* **35**, 10005-10014 (2015).
57 Konkle, T. & Caramazza, A. Tripartite organization of the ventral stream by animacy and object size. *Journal of Neuroscience* **33**, 10235-10242 (2013).

58 Hendel, R. K., Starrfelt, R. & Gerlach, C. The good, the bad, and the average: Characterizing the relationship between face and object processing across the face recognition spectrum. *Neuropsychologia* **124**, 274-284 (2019).
59 Richler, J. J. *et al.* Individual differences in object recognition. *Psychological Review* **126**, 226 (2019).
60 Shah, P., Gaule, A., Sowden, S., Bird, G. & Cook, R. The 20-item prosopagnosia index (PI20): a self-report instrument for identifying developmental prosopagnosia. *Royal Society open science* **2**, 140343 (2015).
61 Duchaine, B. & Nakayama, K. The Cambridge Face Memory Test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants. *Neuropsychologia* **44**, 576-585 (2006).
62 Davies-Thompson, J., Pancaroglu, R. & Barton, J. Acquired prosopagnosia: structural basis and processing impairments. *Frontiers in Bioscience-Elite* **6**, 159-174 (2014).
63 Wilmer, J. B. Individual differences in face recognition: A decade of discovery. *Current Directions in Psychological Science* **26**, 225-230 (2017).
64 Proklova, D., Kaiser, D. & Peelen, M. V. Disentangling representations of object shape and object category in human visual cortex: The animate–inanimate distinction. *Journal of cognitive neuroscience* **28**, 680-692 (2016).
65 Sigurdardottir, H. M. & Gauthier, I. Expertise and object recognition. *Brain mapping: An encyclopedic reference* **2**, 523-527 (2015).
66 Hebart, M. N., Zheng, C. Y., Pereira, F. & Baker, C. I. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature human behaviour* **4**, 1173-1185 (2020).
67 Maurer, D., Le Grand, R. & Mondloch, C. J. The many faces of configural processing. *Trends in cognitive sciences* **6**, 255-260 (2002).
68 Robbins, R. A. & Coltheart, M. The effects of inversion and familiarity on face versus body cues to person recognition. *Journal of Experimental Psychology: Human Perception and Performance* **38**, 1098 (2012).
69 Robbins, R. A. & Coltheart, M. Left–right holistic integration of human bodies. *Quarterly Journal of Experimental Psychology* **65**, 1962-1974 (2012).
70 Peterson, M. A. & Rhodes, G. *Perception of faces, objects, and scenes: Analytic and holistic processes*. (Oxford University Press, 2003).
71 Wallis, G. & Bülthoff, H. H. Effects of temporal association on recognition memory. *Proceedings of the National Academy of Sciences* **98**, 4800-4804 (2001).
72 Miyashita, Y. Inferior temporal cortex: where visual perception meets memory. *Annual review of neuroscience* **16**, 245-263 (1993).
73 Sigurdardottir, H. M., Michalak, S. M. & Sheinberg, D. L. Shape beyond recognition: form-derived directionality and its effects on visual attention and motion perception. *Journal of Experimental Psychology: General* **143**, 434 (2014).
74 Blum, H. A transformation for extracting new descriptions of shape. *Models for the perception of*

speech and visual form, 362-380 (1967).
75 Kimia, B. B. On the role of medial geometry in human vision. *Journal of Physiology-Paris* **97**, 155-190 (2003).
76 Lin, J.-C. The family of universal axes. *Pattern Recognition* **29**, 477-485 (1996).
77 Hung, C.-C., Carlson, E. T. & Connor, C. E. Medial axis shape coding in macaque inferotemporal cortex. *Neuron* **74**, 1099-1113 (2012).
78 Kovács, I., Fehér, A. & Julesz, B. Medial-point description of shape: A representation for action coding and its psychophysical correlates. *Vision research* **38**, 2323-2333 (1998).
79 Kovács, I. & Julesz, B. Perceptual sensitivity maps within globally defined visual shapes. *Nature* **370**, 644-646 (1994).
80 Firestone, C. & Scholl, B. J. "Please tap the shape, anywhere you like" shape skeletons in human vision revealed by an exceedingly simple measure. *Psychological science* **25**, 377-386 (2014).
81 Marr, D. & Nishihara, H. K. Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences* **200**, 269-294 (1978).
82 Biederman, I. Human image understanding: Recent research and a theory. *Computer vision, graphics, and image processing* **32**, 29-73 (1985).
83 Riesenhuber, M. & Poggio, T. Models of object recognition. *Nature neuroscience* **3**, 1199-1204 (2000).
84 Tarr, M. J. & Pinker, S. Mental rotation and orientation-dependence in shape recognition. *Cognitive psychology* **21**, 233-282 (1989).
85 Tarr, M. J. & Bülthoff, H. H. Image-based object recognition in man, monkey and machine. *Cognition* **67**, 1-20 (1998).
86 Peirce, J. W. PsychoPy—psychophysics software in Python. *Journal of neuroscience methods* **162**, 8-13 (2007).
87 Peirce, J. *et al.* PsychoPy2: Experiments in behavior made easy. *Behavior research methods* **51**, 195-203 (2019).
88 Konkle, T., Brady, T. F., Alvarez, G. A. & Oliva, A. Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *Journal of experimental Psychology: general* **139**, 558 (2010).
89 Righi, G., Peissig, J. J. & Tarr, M. J. Recognizing disguised faces. *Visual Cognition* **20**, 143-169 (2012).
90 Afifi, M. 11K Hands: Gender recognition and biometric identification using a large dataset of hand images. *Multimedia Tools and Applications* **78**, 20835-20854 (2019).
91 Wah, C., Branson, S., Welinder, P., Perona, P. & Belongie, S. The caltech-ucsd birds-200-2011 dataset. (2011).
92 Willenbockel, V. *et al.* Controlling low-level image properties: the SHINE toolbox. *Behavior research methods* **42**, 671-684 (2010).

Supplementary Information

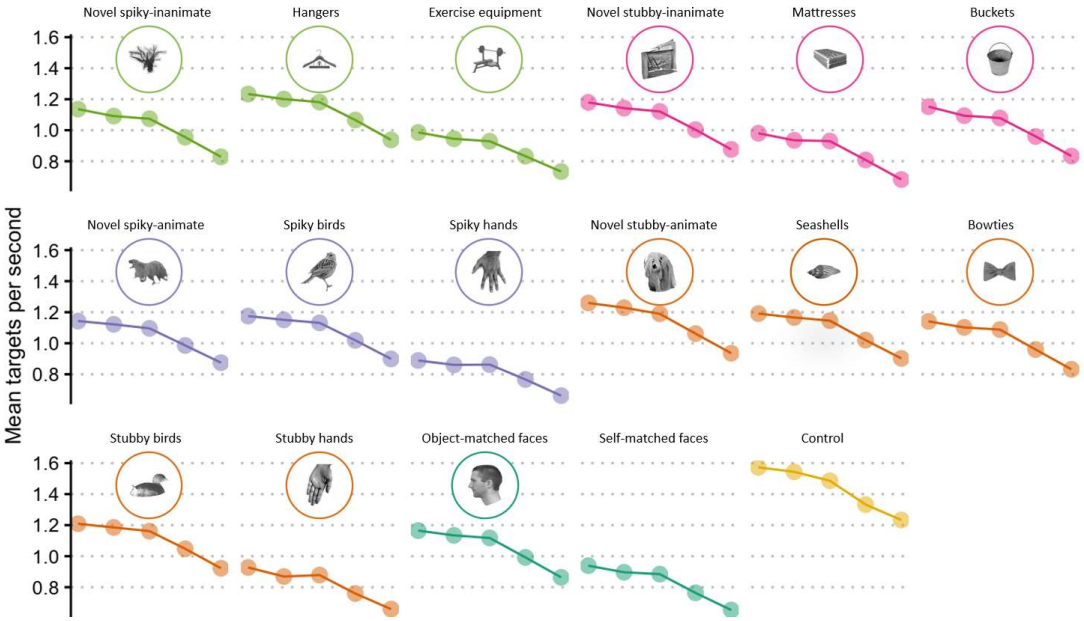

Supplementary figure s1. Foraging by age. The five dots in each subplot represent the five age groups. Foraging performance declined with age for all trial types.

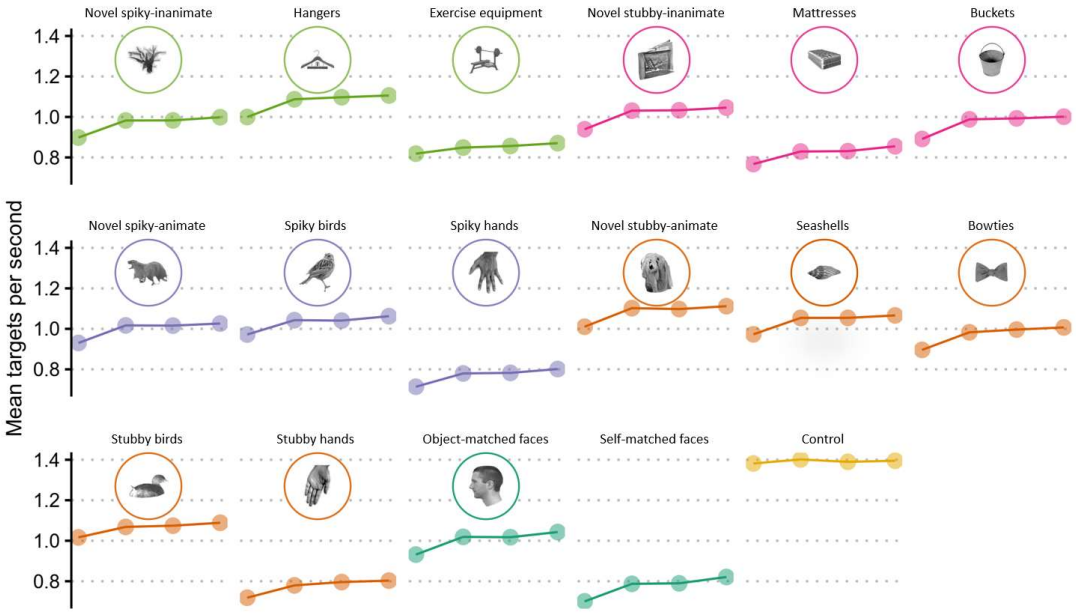

Supplementary figure s2. Foraging by education. The four dots in each subplot represent the five educational levels. Foraging performance increased with educational level, with the notable exception of control foraging.