Limitations on animacy categorization in ensemble perception

Vladislav A. Khvostov¹, Yuri A. Markov¹, Timothy F. Brady², & Igor S. Utochkin¹

¹HSE University, Russian Federation

²Department of Psychology, University of California, San Diego

Correspondence should be sent: Vladislav A. Khvostov
Psychology Dept, HSE University
101000, Armyansky per., 4, room 419
Moscow, Russian Federation
Corresponding author e-mail: khvostov.vladislav@gmail.com

Abstract

People can rapidly and efficiently categorize the animacy of individual objects and scenes, even with few visual features available. Does this necessarily mean that the visual system has an unlimited capacity to process animacy across the entire visual field? We tested this in an ensemble task requiring observers to judge the relative numerosity of animate vs. inanimate items in briefly presented sets of objects. We generated a set of morphed “animacy continua” between pairs of animal and inanimate object silhouettes and tested them in both individual object categorization and ensemble enumeration. For the ensemble task, we manipulated the ratio between animate and inanimate items present in the display and used both “segmentable” (including only definitely animate and inanimate items) and “non-segmentable” (middle-value, ambiguous morphs pictures were shown along with the extremes) distributions of items. We found that observers failed to integrate animacy information from multiple items, as they showed very poor performance in the ensemble task and were not sensitive to the distribution type despite their categorization rate for individual objects being near 100%. A control condition using the same design with color as a category-defining dimension elicited both good individual object and ensemble categorization performance and a strong effect of the segmentability type. We conclude that good individual categorization does not necessarily allow people to build ensemble animacy representations, thus showing the limited capacity of animacy perception.

Keywords: ensemble statistics, animacy, rapid categorization, summary statistics

1. Introduction

Despite severe attentional and working memory limits, the visual system is capable of processing all input to some degree. This is one of the central claims made by the broad literature on ensemble perception (Alvarez, 2011; Whitney & Yamanashi Leib, 2018). Numerous studies described in this literature have shown that despite limits on processing a large number of individual items as individuals, the visual system rather precisely represents the approximate number of objects in a scene (Burr & Ross, 2008; Chong & Evans, 2011; Halberda et al., 2008), as well as the average (Alvarez & Oliva, 2009; Alvarez & Oliva, 2008; Ariely, 2001; Chong & Treisman, 2003, 2005; Parkes et al., 2001; Watamaniuk & Duchon, 1992), variance/range (Dakin & Watt, 1997; Morgan et al., 2008; Norman et al., 2015; Solomon et al., 2011) of their visual features, and even information about the whole distribution of values present in a feature (Chetverikov et al., 2016, 2017a, 2017b, 2017c; Oriet & Hozempa, 2016). Easy and fast segmentation of objects based on their color into distinct sets of objects with distinct summaries (Chong & Treisman, 2005; Halberda et al., 2006; Im & Chong, 2014) provides additional evidence for the availability of feature information about many objects and even many sets at one time. However, most of these facts about ensemble perception are demonstrated on low-level, simple features such as color, size, or orientation (although facial features are an exception, Haberman et al., 2009, 2015; Haberman & Whitney, 2007, 2009, 2010, 2011). Can the visual system represent the information about less perceptual and more abstract features — such as animacy — in a similar manner?

Intuitively, it seems animacy perception should require fine discrimination and recognition associated with the slow, serial deployment of focused attention (Wolfe et al., 2011) making it impossible to get the “gist” of animacy from the entire scene simultaneously. However, multiple studies cast doubt on this intuition. For example, people can efficiently search for an animal among artifacts and vice versa (Levin et al., 2001); animals can be detected

very quickly (Evans & Treisman, 2005; Thorpe et al., 1996) on the periphery in the near absence of attention (Li et al., 2002) and in parallel (Rousselet et al., 2002), showing early neural correlates for such detection (Fabre-Thorpe et al., 1998; Long et al., 2017; Thorpe et al., 1996; VanRullen & Thorpe, 2001). Recently, Leib, Kosovicheva, and Whitney (2016) even showed that people can compute the average animacy of multiple objects.

Based on these findings, some researchers argue that animacy is similar to low-level features, in that it can be computed for each object in an ensemble (e.g. Leib et al., 2016). In other words, this framework suggests that animacy works like a “preattentive” feature, which means that the visual system can represent the animacy of each item in the visual field in parallel (Treisman & Gelade, 1980), even with limited or no conscious access to every single item -- at least in a way sufficient to compute ensemble statistics about this feature (e.g. Leib et al., 2016).

However, there is an alternative view: In particular, it may be that differences in more basic features between animate and inanimate objects are responsible for the rapid animacy perception and computing ensemble statistics of ‘animacy’ (Banno & Saiki, 2015; Evans & Treisman, 2005; Levin et al., 2001; Rosenholtz et al., 2012; Treisman, 2006). Among the main candidates are the statistics of edge co-occurrences (Perrinet & Bednar, 2015), some mid-level characteristics like curvilinearity (Levin et al., 2001; Long et al., 2017; Long et al., 2018; Zachariou et al., 2018), shape (Schmidt et al., 2017), and higher-order texture statistics (Banno & Saiki, 2015; Rosenholtz et al., 2012), which differ between animate and non-animate images on average. According to this view, animacy cannot be called an independent preattentive feature, like color, orientation, or size (Wolfe & Horowitz, 2017; Wolfe & Utochkin, 2019); instead, animacy, in some stimulus sets, simply happens to be correlated with basic features — allowing people who are sensitive to this correlation to compute ‘animacy’ quickly and easily

in such stimulus sets but without actual preattentive feature detectors for animacy (Treisman, 2006).

The main goal of the present work is to test whether the good ability to categorize individual objects or scenes as animate or inanimate necessarily leads to high efficiency in simultaneously categorizing multiple objects in an ensemble. Given the debate about the low- vs. high-level nature of animacy features, we used specific stimuli with reduced information about low-level features, but which still made animacy perception of an individual item possible: silhouettes of animals and artifacts. To estimate categorization and category-based ensemble perception, we used a subset enumeration task (Halberda et al., 2006, 2008). Observers were shown various numbers of animal and artifact silhouettes and had to report which category was more numerous in a display. If participants can genuinely categorize the animacy of objects in an ensemble simultaneously they should show enumeration accuracy comparable with that in subsets defined by distinct basic features such as color (Halberda et al., 2006). Poor performance would indicate the inability to simultaneously categorize multiple objects by animacy, and might mean that previous work is more dependent on non-animacy specific visual features (e.g. Long et al., 2017; Zachariou et al., 2018).

Previous studies have shown that animacy is not a binary feature (either animate or inanimate). Rather, it is a continuum of values with a lot of intermediate points (Leib et al., 2016; Schmidt et al., 2017; Sha et al., 2015). Importantly, for low-level features (color, length, or orientation), manipulations of the distribution of features have been shown to strongly affect performance in tasks involving categorization of multiple items (Im et al., 2020; Utochkin et al., 2018; Utochkin & Yurevich, 2016). In particular, if the distribution has several sharp peaks with a gap between them (a *segmentable* distribution; e.g., steep and shallow orientations, without intermediate values), categorization performance is better than if the stimuli form one smooth peak (a *non-segmentable* distribution; many intermediate values of orientation along

with steep and shallow ones). Based on this finding, we manipulated the shape of the distribution as a diagnostic factor in our study, which could provide additional insight on the question of whether participants are performing the task with animacy in the same parallel way they perform it with low-level features. In particular, if participants can perceive animacy across multiple objects, their ensemble enumeration performance should benefit from sharply distributed segmentable sets of animate and inanimate exemplars (i.e., all objects are either animate or inanimate, with no in-between) compared to smoothly distributed, non-segmentable sets (intermediate values of animacy along with animate and inanimate objects) as it does for low-level features (Utochkin, 2015).

In all our tasks, we also had participants perform a control condition, where we manipulated color as a category-defining feature. Color is a prototypical low-level feature that can be processed in parallel and has previously been shown to be good for subset segmentation in the subset enumeration task (Halberda et al., 2006) and, as a basic low-level feature (Treisman & Gelade, 1980) is the kind of feature prone to segmentability effects (Utochkin, 2015). This allows us to compare performance in our animacy task with a feature that can be easily categorized in an ensemble-based manner.

2. Experiment 1a

2.1 Method

2.1.1 Participants

Thirty undergraduate students (see the justification of a sample size in the “design, power, and data analysis” section) at the University of California, San Diego were recruited for participation in Experiment 1 (22 female, mean age = 20.5 years, SD = 2.32). All participants reported normal or corrected-to-normal vision without color-vision deficiency and provided written informed consent.

2.1.2 Stimuli

We specifically designed a stimulus set that limited the low-level and mid-level correlates that are associated with animacy in real images, while still allowing clear animacy perception.. In particular, we created morphed silhouettes of various animals and artifacts using the “Neural Best-Buddies” neural network (Aberman et al., 2018) and a morphing program (Liao et al., 2014). We first morphed two silhouettes of animate creatures with similar shape (e.g., tyrannosaurus and kangaroo, Figure S1) to create a novel "animate" stimulus and two silhouettes of inanimate objects (e.g., table and telescope, Figure S1) to create a novel "inanimate" stimulus. 10 animate and 10 inanimate stimulus endpoints were created. These stimuli were used to decrease individual recognizability of real-world exemplars and to highlight the categorical features of animate/inanimate. We then took one animate and one inanimate stimuli with similar shape and created 11 morphs between them (with a 10% morphing step) to obtain an animacy continuum. We created 10 such continua from different pairs of animate and inanimate stimuli and tested these continua in an online experiment aimed at stimulus validation (Experiment A, see Appendix A). Observers ($N=100$) were shown each of 110 unique images (10 continua, 11 morphs each) for 500 ms and answered whether a shown object is animate or inanimate using “right” and “left” buttons (for details see Appendix A). The probability of each image being categorized as animate or inanimate was calculated from response frequencies. We excluded 3 continua based on poor performance (categorization probabilities were always far from 0% or 100% animate, even at the endpoints of the continuum). Within each remaining continuum, we chose both the extremes and the midpoint as stimuli: 0% and 100% animate and one middle image that provided a categorization rate closest to 0.5. Therefore, a 3 (“definitely animate”, “half-animate”, “definitely inanimate”) $\times$ 7 (different continua) set of pictures were used in Experiment 1 (Figure 1).

Inanimate object (0% animate)
Half- animate object
Animate object (100% animate)

Figure 1. “Animate”, “half-animate”, and “inanimate” objects for all 7 continua used in Experiment 1.

Stimuli in Experiment 1 was presented using PsychoPy (Peirce et al., 2019). A $19.5^\circ \times 19.5^\circ$ square at the center of the screen was used as a “working” field for presenting stimuli; the rest of the screen remained empty. The working field was divided into $5 \times 5 = 25$ cells by an imaginary grid (cell side was 3.94°). Each cell could be used for positioning a single item of a stimulus set. Within the cell, an item could be randomly jittered within $\pm 0.81^\circ$ along both horizontal and vertical directions. Also, to prevent image grouping by low-level similarity the images could differ in orientation, size, and a mirror-flip manipulation.

A display for the ensemble (subset) enumeration task consisted of variable numbers of animate and inanimate objects. The numeric ratio between categories could be 2:1 or 10:3 (each category prevailed in half of all trials). We also manipulated the shape of the animacy distribution: it could be either *segmentable* (the display contained only “definite” objects from the animacy continuum, in the ratio of 2:1 or 10:3) or *non-segmentable* (“half-animate” objects were added in the same number as the number of a minor category of “definite” objects so that proportions were 2:1:1 or 10:3:3). Note that the correct answer for trials with a non-segmentable distribution was the same as with segmentable one since the additional objects do not change the average animacy value due to their “half-animacy”. As the *non-segmentable* and *segmentable* conditions differ in terms of a total set size (16 vs. 12 and 16 vs. 13 for 2:1

and 10:3, correspondingly), we added one more segmentable condition (*segmentable-increased set size*), where the set size was equated with the corresponding non-segmentable set sizes but only “definite” morphs were presented (the ratios were approximately preserved provided rounding to nearest integers). The segmentable displays aimed to provide a two-peak animacy distribution in which each peak corresponds to a single category (animate and inanimate). This should make categorization relatively easy (Utochkin, 2015; Utochkin et al., 2018) if animacy is identified simultaneously. The non-segmentable trials provided a smooth, wide distribution where an ensemble of clear categories is filled by multiple items sharing animate and inanimate features and, thus, complicating categorization.

We had 3 conditions that formed the basis of which stimulus set was used, and what ensemble enumeration task needed to be performed (we will refer to this factor as *feature space*): *animacy*, *mixed animacy*, and *color*. In the *animacy* condition, displays consisted of objects drawn from a single morph continuum (e.g., 8 identical animate and 4 identical inanimate silhouettes). To test whether people can perceive the animacy of multiple objects as an abstract feature and rule out the strategy of just learning the difference between animate and inanimate morphs in each continuum, we introduced the *mixed animacy* condition where each display could contain objects from all 7 continua preserving the animate/inanimate ratios. The *color* condition served as a baseline to measure how well observers perform numerosity judgments with a basic low-level feature known to support good segmentation (Halberda et al., 2006). Here, we used displays from the *animacy* condition with an addition that all silhouettes were colored independently from their animacy. The colors were drawn from a hue distribution along the approximately perceptually uniform CIELAB color wheel (Suchow et al., 2013; Schurgin et al., 2020). The hue distribution ranged from 270° (violet) to 300° (pink) with a 15° step between nearest colors (three hues overall), designed to match the difficulty of single-item

categorization to that of animacy. The ratio and segmentability manipulations with color were identical to those with animacy as described above (Figure 2).

Figure 2. The example of stimuli displays used in ensemble enumeration (top) and object categorization tasks (bottom). **Top:** In the ensemble enumeration task, participants said whether the display as a whole was more animate or more inanimate; or more pink or more purple. The rows represent the 3 conditions of segmentability within the ensemble task (where segmentable means only extreme values are shown; whereas non-segmentable means intermediate values are also present). The columns represent the 3 feature spaces (Animacy: only 1 kind of animate and inanimate stimulus is present; Mixed animacy: a mix of different animate and inanimate items are used). The ratio of animate/inanimate (or violet/pink) objects

is 2:1 (2:1:1 for non-segmentable case) for all 9 displays of ensemble enumeration task. **Bottom:** The object categorization task served as our baseline for how individually distinguishable the items were. Participants choose which of the two options was more animate or more pink.

2.1.3 Procedure

Participants were seated approximately 50 cm from a monitor. Each trial of the *ensemble* enumeration task was started by pressing the spacebar. Following a 200-ms blank screen, a stimulus set was presented for 350 ms. The participants had to report which category (animate/inanimate in the *animacy/mixed animacy* conditions or pink/violet in the *color* condition) was more numerous by pressing “L” or “A” buttons on a keyboard, respectively.

To test that our images can be accurately categorized by animacy/color as individual objects, we ran a set of *object* categorization tasks before the main ensemble tasks. In each trial, two silhouettes (0% animate and 100% animate - for *animacy* and *mixed animacy* condition, or pink and violet - for *color* condition) were shown on opposite sides from the center of a screen for 350 ms. Participants reported whether a more animate or a pink item is on the left or right by pressing an “A” or “L” button.

In both *object* categorization and *ensemble* enumeration tasks, participants were asked to answer as fast as possible. Feedback about response correctness was provided after each trial for 300 ms. Overall, each participant performed 6 blocks: three *object* categorization blocks (color, animacy, and mixed animacy) with 140 trials in each, then three *ensemble* enumeration blocks with 420 trials in each (70 trials for each combination of factors). The order of each triplet of blocks was randomized across participants.

2.1.4 Design, power, and data analysis

In Experiment 1a, we used a 2 (Task: object vs. ensemble) $\times$ 3 (Feature space: color, animacy, mixed animacy) $\times$ 3 (Segmentability: non-segmentable, segmentable, segmentable-increased set size) $\times$ 2 (Ratio: 2:1 vs 10:3) within-subject design. It was not a full-factor design because the last two factors can be manipulated only within the ensemble enumeration task

(not the object categorization task). We analyzed the percentage of correct responses in both tasks.

We determined a sample size for this experiment based on the pilot study (N=6) with a similar experimental design (the data can be found at <https://osf.io/92ez7/>). We evaluated the expected mean differences, SD, and other parameters for the smallest effect of interest in our within-subject design (the interaction between ratio and segmentability factors). With these numbers and the desired power at .8 (with $\alpha = .05$), our power analysis (GLIMMPSE program: <https://glimmpse.samplesizeshop.org/>) indicated that we needed around 23 participants. Concerning possible technical problems and some variations between the pilot and the main experiment, we recruit a sample of 30 observers.

2.2 Results

Overall, people were much better at the object categorization task (*which is animate? or which is pink?*) than the ensemble enumeration task. Thus, the Task $\times$ Feature space ANOVA (Figure 3) showed a strong main effect of the task ($F(1,29) = 287.96$, $p < .001$, $\eta^2_p = .91$, $BF_{10} = 1.2 \times 10^{37}$), such that performance in the *ensemble* tasks was worse than in the *object* tasks.

Figure 3. The proportion of correct responses as a function of feature space and task in Experiment 1. Error bars denote ± 1 standard error of the mean. Participants were excellent at all object categorization tasks but were far better at the ensemble enumeration task for color than for animacy.

There was a large effect of feature space, such that people were much better at color ensembles than either kind of animacy ensembles. In particular, we found a feature space effect ($F(2,58) = 80.83$, $p < .001$, $\eta^2_p = .74$, $BF_{10} = 491.72$) that can be better illustrated using an interaction between the task and feature space factors ($F(2,58) = 60.38$, $p < .001$, $\eta^2_p = .68$, $BF_{12} = 1.88 \times 10^{11}$, where BF_{12} is a ratio of BF_{10} for a model taking the main effects and interaction into account to BF_{10} for a model taking only the main effects into account). Critically, performance in the color ensemble condition was higher than in the two animacy ensemble tasks ($ts(29) > 10.78$, $ps < .001$, Bonferroni corrected $\alpha = .008$, d_z 's > 1.97 , BF_{10} 's $> 7.87 \times 10^8$). The color object condition was slightly easier than the mixed animacy object condition ($t(29) = 3.09$, $p = .004$, Bonferroni corrected $\alpha = .008$, $d_z = 0.57$, $BF_{10} = 9.15$) whereas the difference between all other conditions either did not differ from each other or did not

survive from correction for multiple comparison ($ts(29) = 2.8$, $ps = .009$, Bonferroni corrected $\alpha = .008$, d_z 's = 0.51, BF_{10} 's = 4.96).

The main finding from this analysis is that despite almost ideal performance (> 94% of correct responses) in detecting which of two objects is animate, participants were poor at the ensemble animacy tasks (~64%). Substantially better performance in the color ensemble tasks shows that the task itself can be successfully performed using a basic feature for categorization. The poor performance in the animacy ensemble task is a surprising result given the relatively large numeric ratios we used: Our smallest ratio was the same as the largest ratio from Halberda et al. (2008), where it provided almost 100% correct performance for color-segmented stimuli. A simple calculation assuming that observers just select one random item in each ensemble trial, correctly identify its animacy value in 94% of cases (as measured in the object categorization animacy task) and respond based only on this item, predict almost the same overall performance in this task (~66%). Based on this finding, we have to conclude that despite an almost ideal categorization rate for individual objects our participants failed to categorize the animacy of multiple objects simultaneously which is at odds with previous work (Leib et al., 2016; Rousselet et al., 2002).

Another argument for the impossibility of integrating animacy information from multiple objects comes from a second set of analyses, considering the ensemble enumeration tasks only and how segmentability and ratio plays a role in these tasks (ANOVA's of Segmentability $\times$ Ratio, Figure 4). In the *color* feature space, we found the main effects of ratio ($F(1,29) = 96.85$, $p < .001$, $\eta^2_p = .77$, $BF_{10} = 2.98 \times 10^{12}$) and segmentability ($F(2,58) = 42.36$, $p < .001$, $\eta^2_p = .59$, $BF_{10} = 5.38 \times 10^4$). Importantly, performance was higher with the greater (10:3) vs smaller (2:1) ratio ($t(29) = 9.84$, $p < .001$, $d_z = 1.8$, $BF_{10} = 2.89 \times 10^{12}$) and in segmentable conditions vs non-segmentable one ($ts(29) > 5.5$, $ps < .001$, d_z 's > 1, BF_{10} 's > 6.36×10^3). We also found the two-way interaction ($F(2,58) = 12.06$, $p < .001$, $\eta^2_p = .29$, $BF_{12} = 1.81 \times 10^3$),

showing that the smaller ratio was more diagnostic ($ts(29) > 3.95$, $ps < .001$, Bonferroni corrected $\alpha = .008$, d_z 's > 0.72 , BF_{10} 's > 65.83), while the larger ratio does not provide any substantial difference between the segmentability conditions, probably because of the ceiling effect. In sum, these results show that color-based ensemble enumeration benefits substantially from color segmentability.

In contrast with color, for both *animacy* and *mixed animacy*, we found no evidence for a segmentability effect ($Fs(2,58) < 3.54$, $ps > .035$, η_p^2 's $< .11$, BF_{10} 's $< .5$) or for any interaction between segmentability and ratio ($Fs(2,58) < 1.92$, $ps > .16$, η_p^2 's $= .63$, BF_{12} 's $< .56$). We found only an effect of ratio ($Fs(1,29) > 13.81$, $ps < .001$, η_p^2 's $> .32$, BF_{10} 's > 272.24) showing better performance in the greater ratios condition compared to the smaller one, which would be expected to occur even if subsampling only a single item.

Figure 4. The proportion of correct responses as a function of ratio and segmentability for three feature spaces. Error bars denote ± 1 standard error of the mean.

Overall, the very low performance level in the animacy ensemble task is consistent with two possibilities. Either responses in the animacy ensemble task were based on the information sampled from a very small number of items (see Allik et al., 2013; Myczek & Simons, 2008), or multiple animacy representations put together were too noisy to build a reliable impression about overall display animacy. Also, there was no effect of distribution segmentability for animacy (unlike color) suggesting that participants' performance was not sensitive to changes

in the shape of the animacy distribution. Again, it can be explained either by sampling of a very limited number of items or by the high level of noise in global animacy representations, so that adding a few uncertain “half-animate” objects could not make global ensemble discrimination significantly worse. In other words, in terms of the global impression, animate, inanimate, and half-animate objects could look almost indistinguishable. In either way, we can conclude that our observers failed to successfully integrate animacy information from multiple items, i.e., they could not use the information from all or at least most of the objects to respond in the animacy ensemble task (cf. Leib et al., 2016).

3. Experiment 1b

Experiment 1b aimed to replicate the findings from Experiment 1a, and ensure generalizability to small differences in conditions (e.g., we used 3 ratios).

3.1 Method

3.1.1 Participants

Thirty-one undergraduate students at the University of California, San Diego participated in Experiment (20 female, mean age = 20.9 years, SD = 2.08). The sample size was based on the power analysis described in Experiment 1a. All participants reported normal or corrected-to-normal vision without color-vision deficiency and provided written informed consent.

3.1.2 Apparatus, stimuli, and procedure

Apparatus, stimulus set, and procedure were the same as in Experiment 1a, except that participants performed only the ensemble enumeration tasks, without the preceding object categorization task. Manipulations of segmentability were also identical to Experiment 1a, whereas we used three different ratios in this experiment: 2:1, 5:2, and 10:3 and only two feature spaces: *animacy* and *color*.

3.2 Results

Our results (Figure 5) replicated Experiment 1a, with color ensembles much better than animacy ensembles, and effects of segmentability larger in color than in animacy, although the small segmentability effect in animacy condition reached (uncorrected) significance in this experiment.

Figure 5. The proportion of correct responses as a function of ratio and segmentability for two feature spaces. Error bars denote ± 1 standard error of the mean.

In particular, the two-way ANOVA (Segmentability $\times$ Ratio) for *color* showed the main effects of ratio ($F(2,60) = 27.74$, $p < .001$, $\eta^2_p = .52$, $BF_{10} = 9.68 \times 10^9$) and segmentability ($F(2,60) = 27.74$, $p < .001$, $\eta^2_p = .48$, $BF_{10} = 3.15 \times 10^6$). Specifically, performance was higher in greater (10:3 and 5:2) vs. smaller (2:1) ratio ($ts(29) > 5.09$, $ps < .001$, Bonferroni corrected $\alpha = .017$, d_z 's > 0.91 , BF_{10} 's $> 1.36 \times 10^5$), in the *segmentable* and *segmentable-increased set size* conditions vs. the *non-segmentable* one ($ts(29) > 5.58$, $ps < .001$, Bonferroni corrected $\alpha = .017$, d_z 's > 1 , BF_{10} 's $> 2.66 \times 10^6$) and also in the *segmentable-increased set size* vs. the *segmentable* condition ($t(29) = 2.6$, $p = .014$, $d_z = .47$, $BF_{10} = 2.2$). Also, we found a two-way effect of segmentability and ratio ($F(4,120) = 6.64$, $p < .001$, $\eta^2_p = .17$, $BF_{12} = 54$). For the 2:1 and 5:2 ratios, performance in the *non-segmentable* condition was lower than in the two

segmentable conditions; and for the 10:3 ratio, performance in the *segmentable-increased set size* condition was higher than that in the *non-segmentable* condition ($ts(30) > 4.58$, $ps < .001$, d_z 's > 0.82 , BF_{10} 's > 333.44).

For *animacy*, we also found the strong effect of ratio ($F(2,60) > 30.24$, $p < .001$, $\eta^2_p > .5$, $BF_{10} = 1.1 \times 10^9$): performance within 2:1 ratio $< 5:2 < 10:3$ ($ts(30) > 3.69$, $ps < .002$, d_z 's > 0.66 , BF_{10} 's > 45.81). An effect of segmentability was also present ($F(2,58) = 8.3$, $p < .001$, $\eta^2_p = .22$, $BF_{10} = 8.11$). We also found a two-way interaction between these factors ($F(4,120) = 5.99$, $p < .001$, $\eta^2_p = .12$, $BF_{12} = 152$), but only two post-hoc comparison survived correction for multiple comparison: for the 2:1 ratio performance in *non-segmentable* condition was lower than in *segmentable-increased set size* condition and for the 10/3 ratio performance in the *segmentable-increased set size* condition was lower than in *segmentable* condition ($ts(30) > 3.24$, $ps < .001$, Bonferroni corrected $\alpha = .005$, d_z 's > 0.58 , BF_{10} 's > 12.92). The rest condition did not differ from each other ($ts(30) < 3.06$, $ps > .005$, d_z 's < 0.55 , BF_{10} 's < 8.57).

Thus, Experiment 1b replicates the main results of Experiment 1a: good performance (~89%) in the *color* ensemble task and around 65% in the *animacy/mixed* animacy ensemble task coupled with a large effect of segmentability for color and a tiny one for animacy. This again shows that our observers failed to successfully integrate animacy information from multiple items while performing well the same task for color.

4. Experiment 2

The main result of Experiment 1a and 1b is that, despite almost ideal object categorization performance for both animacy and color feature spaces, participants performed well only in the color ensemble task but not in the animacy ensemble task. We interpreted this as evidence that participants could integrate color information from the entire visual field but were not able to do it for animacy. Could these results instead be explained in terms of the

quantitative difference in feature discriminability in color vs. animacy, rather than in terms of the qualitative difference between color and animacy dimensions? Is it possible that we picked too distinct colors, but that in spite of our near-ceiling performance in the object categorization task for animacy, had less distinctive items for animacy. We address this issue in Experiment 2 where we parametrically manipulated feature discriminability in both the color and animacy dimensions. This allowed us to test whether our results of limited animacy performance in the ensemble enumeration task (as compared to color) would still hold if we used the most discriminable values from our animacy space and less discriminable values from color space.

4.1 Method

4.1.2 Participants

Thirty observers (16 female, mean age = 30.4 years, SD = 11.11) were recruited via Prolific (www.prolific.ac; Palan & Schitter, 2018; Peer et al., 2017) and run using Pavlovia (<https://pavlovia.org>). The sample size was based on the power analysis described in Experiment 1a. They gave informed consent and were paid £5.58 per hour (the experiment lasted approximately 35 minutes).

4.1.3 Stimuli

This experiment was similar to Experiment 1, with a few small differences. First, we did not manipulate segmentability and ratio in this experiment - only the ‘segmentable’ condition with a 2:1 ratio was used (there were always 12 objects on the screen). Second, we parametrically manipulated feature *discriminability between the to-be-enumerated subsets*. For color, this factor had two levels of pink-violet separation: easy (30° apart along the CIE Lab color wheel, as in Experiment 1) and hard (15° apart along that color wheel). For animacy, we also had two discriminability levels: easy (100% animate vs. 100% inanimate morphed silhouettes, as in Experiment 1) and hard (80% animate/20% inanimate vs. 20% animate/80% inanimate morphs, which are less discriminable). We also added an additional control condition

(animacy non-morphed), 100% original silhouettes of real objects (e.g., 100% cat vs. 100% tractor) having maximum separability physically possible in the animacy space we used (see Figure 6 for example displays). These new conditions allowed us to test whether we can make performance in the color ensemble task as bad as in the animacy ensemble task by just decreasing color discriminability or whether we can make performance in the animacy task as good as in the color task by increasing animacy discriminability.

Figure 6. The example of stimuli displays used in the ensemble enumeration tasks of Experiment 2. Rows represent the animacy and the color feature spaces, columns represent discriminability levels.

4.1.4 Procedure and design

The procedure of Experiment 2 was identical to Experiment 1. We had a 2 (Task: object vs. ensemble) $\times$ 5 (Feature space: color easy, color hard, animacy easy, animacy hard, animacy non-morphed). We analyzed the percentage of correct responses in both tasks.

4.2 Results

As in Experiment 1, observers were much better at the object categorization task (*which is animate?* or *which is pink?*) than the ensemble enumeration task (Figure 7). Thus, an ANOVA showed a strong main effect of the task ($F(1,29) = 311, p < .001, \eta^2_p = .915, BF_{10} = 5.08 \times 10^{57}$), such that performance in the ensemble tasks was worse than in the object tasks.

Figure 7. The proportion of correct responses as a function of feature space, discriminability, and task in Experiment 2. Error bars denote ± 1 standard error of the mean.

In addition, there was a large effect of feature space, such that people were much better at both color ensembles than either kind of animacy ensembles. In particular, we found a feature space effect ($F(4,116) = 47.896, p < .001, \eta^2_p = .623, BF_{10} = 6.48 \times 10^3$) that can be better illustrated using an interaction between the task and feature space factors ($F(4,116) = 36.329, p < .001, \eta^2_p = .556, BF_{12} = 1.05 \times 10^{15}$). Critically, performance in both color ensemble tasks was much higher than all animacy tasks ($ts(29) > 6.234, ps < .001$, Bonferroni corrected $\alpha = .003, d_z$'s $> 1.137, BF_{10}$'s $> 2.08 \times 10^4$). Also, performance in the “color easy” ensemble task was higher than in the “color hard”, and performance in the “animacy easy” ensemble task was higher than in the “animacy hard” ($ts(29) > 3.759, ps < .001$, Bonferroni corrected $\alpha = .008, d_z$'s $> 0.686, BF_{10}$'s > 42.083). Whereas all other conditions belonging to the same type of the task did not differ from each other ($ts(29) < 2.577, ps > .014$, Bonferroni corrected $\alpha = .008, d_z$'s $< 0.471, BF_{10}$'s < 3.138). Therefore, our results show that observers are constantly rather

good at the color ensemble task even at low color discriminability and that they are constantly bad at the animacy ensemble task even at the highest physically possible discriminability.

The results of this experiment strongly replicate our initial finding. Even when less discriminable colors are used, performance in the color ensemble task went down, but just slightly (from 85% to 79% correct) and still stayed much better than in the animacy ensemble task (61-67% correct), even when the stimuli were the original silhouettes of familiar, everyday objects. Therefore, it appears that our original results cannot be explained by greater discriminability of chosen colors compared to animacy. Rather, we show that there are important differences between the two feature spaces: Even if individual categorization is at the ceiling for both animacy and color, ensemble tasks are much more difficult for animacy than for color.

5. Experiment 3

The presentation of silhouettes in the object categorization task and the ensemble enumeration task of Experiment 1a and 2 was slightly different. In the ensemble enumeration task, items differed in orientation and size (see Figures 2 and 6), while in the object categorization task, they were always of the same size and upright. Such variation was necessary for the ensemble task to ensure that participants could not use straightforward alignments of basic visual features to get animacy information about multiple items at once. Experiment 3 aimed to make sure that this variation in visual features per se did not prevent the correct animacy perception and was not the cause of the difference in performance between object and ensemble tasks.

5.1 Method

5.1.1 Participants

Thirty observers (9 female, mean age = 25.97 years, SD = 8.56) were recruited via Prolific (www.prolific.ac; Palan & Schitter, 2018; Peer et al., 2017) and run using Pavlovía (<https://pavlovía.org>). They were paid £5.58 per hour (the experiment lasted approximately 12 minutes). The sample size was based on the power analysis described in Experiment 1a.

5.1.2 Apparatus, stimuli, and procedure

Apparatus, stimulus set, and procedure were similar to those in Experiment 2, with three exceptions. First, participants performed only the object categorization tasks, without any ensemble enumeration tasks. Second, we used only the animacy conditions (animacy easy, hard, and non-morphed) in this experiment since the manipulations with size and orientation did not change anything in terms of color. Third, we introduced a new manipulation, the *variation* which defines whether the size and orientation of two silhouettes in a trial are randomized (as in ensemble enumeration task in Experiment 1a and 2). The silhouette size was either always equal to 95 pixels within blocks without variation or was chosen out of three options (80, 95, or 110) in the blocks with variation; orientation was either vertical or randomly chosen from an interval $[-85^\circ; 85^\circ]$.

Thus, we a 3 (Discriminability: easy, hard, silhouettes) $\times$ 2 (Variation: with vs. without) within-subject design.

5.2 Results

The two-way ANOVA (Discrimination $\times$ Variation) showed a small effect of variation ($F(1,29) = 6.56$, $p = .016$, $\eta^2_p = .184$, $BF_{10} = 1.963$), no effect of discriminability ($F(2,58) = 0.983$, $p = .38$, $\eta^2_p = .033$, $BF_{10} = 0.47$) and the absence of interaction of these two factors ($F(2,60) = 0.568$, $p = .57$, $\eta^2_p = .019$, $BF_{12} = 0.13$). There was no reliable difference between corresponding pairs of conditions with and without variation after correction for multiple

comparisons ($ts(29) < 2.51$, $ps > .017$, Bonferroni corrected $\alpha = .016$, d_z 's < 0.459 , BF_{10} 's < 2.756).

While there was some effect of variation, observers showed high performance (92-96%) in all six blocks (Figure 8), without any large differences between blocks with and without variation. We can conclude that the variation in terms of size and orientation that was applied in the ensemble task was not itself sufficient to drive the very poor performance in the ensemble task. Thus, the dramatic difference between object and ensemble categorization in Experiments 1 and 2 was not due to the variation in size and orientation between silhouettes in the ensemble enumeration task.

Figure 8. The proportion of correct responses as a function of discrimination difficulty and variation. Error bars denote ± 1 standard error of the mean.

6. General Discussion

The present study shows that excellent categorization of individual items based on a high-level feature (animacy) does not always entail efficient calculation of ensemble statistics of this feature. Our observers had almost perfect performance when the animacy of only two items should be compared, but their performance dropped dramatically when the task required the parallel categorization of multiple objects. At the same time, a low-level feature space such as color yielded much better performance in the ensemble tasks, consistent with the parallel ability to categorize for colors but not for animacy. Therefore, observers could integrate more information about color ensembles than about animacy ensembles. We also found a stronger segmentability effect for colors than for animacy, consistent with greater sensitivity to the entire distribution of items in the case of color than animacy.

The poor performance in the animacy ensemble tasks found in our work is at odds with some previous reports (e.g., Leib et al., 2016). One of the crucial differences between these studies is the type of stimuli. Whereas most of the previous studies used real animate and inanimate images, our stimulus set consisted of unicolor silhouettes. The real pictures likely contain many correlates of animacy in terms of relatively low-level and mid-level features, which have previously been shown to be very important for distinguishing between animate and inanimate objects in tasks involving categorization of multiple objects at once. For example, animacy is related to the statistics of edge co-occurrences (Perrinet & Bednar, 2015), mid-level perceptual features, including curvature (Long et al., 2017; Long et al., 2018; Zachariou et al., 2018) or rectilinearity (Levin et al., 2001), as well as more complex shape characteristics (Schmidt et al., 2017), and higher-order texture statistics (Banno & Saiki, 2015; Rosenholtz et al., 2012). Therefore, some of these features, rather than animacy itself, might support performance in tasks requiring instantaneous animacy categorization of multiple objects. Our study showed that making most of these features unavailable by using silhouettes leads to an inability to categorize multiple features based on their animacy, even though the

remaining silhouette features were sufficient for very good individual categorization. Thus, our work is consistent with the intuitive view that genuine animacy perception should require fine discrimination, associated with the slow, serial deployment of focused attention. Our findings suggest that further experiments are necessary to figure out the nature and level of features underlying ensemble perception of complex object properties such as animacy. To speak to animacy *per se*, as opposed to low- and mid-level correlates of animacy, such experiments would require careful control for low- and mid-level visual differences between the items representing high-level categories (e.g. Long et al., 2017).

The results of this work also showed some evidence against the view of animacy as a basic “feature” that can be detected in parallel across the entire visual field, like color, orientation, size, or motion (Wolfe & Horowitz, 2017; Wolfe & Utochkin, 2019). Our observers systematically failed to ascribe and integrate animacy values of multiple objects in the visual field, which contradicts the idea of a basic, “preattentive” animacy feature. In fact, as shown in Experiment 1, our data were broadly consistent with the level of performance you would expect from sampling only a single item and using it to make the ensemble enumeration judgment. While our study alone is not sufficient to provide evidence that animacy is not a basic feature, we clearly show that animacy is at least a much weaker feature than more basic and low-level sensory dimensions such as color.

In sum, our results suggest that although the visual system can be rather efficient at rapid animacy categorization of individual objects based on the minimum of features, it can fail to get the animacy “gist” (ensemble summary statistics) of the same objects simultaneously. At the same time, low-level features such as color support both individual categorization and ensemble categorization. This conclusion is suggestive for thinking about experimental designs and findings aimed to investigate the roles of high-capacity, non-selective vs. limited-capacity,

selective processes (Wolfe et al., 2011) in the perception of high-level visual or conceptual properties.

Acknowledgments

Experiments 1a, 2, 3 were supported by the Russian Science Foundation (grant 18-18-00334) to I.S.U. Experiment 1b was supported by NSF CAREER grant BCS-1653457 to T.F.B.

Open science statement

The data from all experiments reported in this article can be accessed at <https://osf.io/92ez7/>.

References

- Aberman, K., Liao, J., Shi, M., Lischinski, D., Chen, B., & Cohen-Or, D. (2018). Neural Best-Buddies: Sparse Cross-Domain Correspondence. *ACM Transactions on Graphics*, 37(4), 1–14. <https://doi.org/10.1145/3197517.3201332>
- Allik, J., Toom, M., Raidvee, A., Averin, K., & Kreegipuu, K. (2013). An almost general theory of mean size perception. *Vision Research*, 83, 25–39. <https://doi.org/10.1016/j.visres.2013.02.018>
- Alvarez, G. A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences*, 106(18), 7345–7350. <https://doi.org/10.1073/pnas.0808981106>
- Alvarez, G.A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15(3), 122–131. <https://doi.org/10.1016/j.tics.2011.01.003>
- Alvarez, G. A., & Oliva, A. (2008). The Representation of Simple Ensemble Visual Features Outside the Focus of Attention. *Psychological Science*, 19(4), 392–398. <https://doi.org/10.1111/j.1467-9280.2008.02098.x>
- Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, 12(2), 157–162.
- Banno, H., & Saiki, J. (2015). The use of higher-order statistics in rapid object categorization in natural scenes. *Journal of Vision*, 15(2). <https://doi.org/10.1167/15.2.4>
- Burr, D., & Ross, J. (2008). A Visual Sense of Number. *Current Biology*, 18(6), 425–428. <https://doi.org/10.1016/j.cub.2008.02.052>
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition*, 153, 196–210. <https://doi.org/10.1016/j.cognition.2016.04.018>

- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017a). Chapter 5 - Learning features in a complex and changing environment: A distribution-based framework for visual attention and vision in general. In C. J. Howard (Ed.), *Progress in Brain Research* (Vol. 236, pp. 97–120). Elsevier. <https://doi.org/10.1016/bs.pbr.2017.07.001>
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017b). Representing Color Ensembles. *Psychological Science*, 28(10), 1510–1517. <https://doi.org/10.1177/0956797617713787>
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017c). Set size manipulations reveal the boundary conditions of perceptual ensemble learning. *Vision Research*, 140, 144–156. <https://doi.org/10.1016/j.visres.2017.08.003>
- Chong, S.C., & Evans, K. K. (2011). Distributed versus focused attention (count vs estimate): Distributed versus focused attention. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(6), 634–638. <https://doi.org/10.1002/wcs.136>
- Chong, S.C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 12.
- Chong, S.C., & Treisman, A. (2005). Statistical processing: Computing the average size in perceptual groups. *Vision Research*, 45(7), 891–900. <https://doi.org/10.1016/j.visres.2004.10.004>
- Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, 37(22), 3181–3192. [https://doi.org/10.1016/S0042-6989\(97\)00133-8](https://doi.org/10.1016/S0042-6989(97)00133-8)
- Evans, K. K., & Treisman, A. (2005). Perception of objects in natural scenes: Is it really attention free? *Journal of Experimental Psychology. Human Perception and Performance*, 31(6), 1476–1492. <https://doi.org/10.1037/0096-1523.31.6.1476>

- Fabre-Thorpe, M., Richard, G., & Thorpe, S. J. (1998). Rapid categorization of natural images by rhesus monkeys. *Neuroreport*, 9(2), 303–308.
- Haberman, J., Harp, T., & Whitney, D. (2009). Averaging facial expression over time. *Journal of Vision*, 9(11), 1.1-13. <https://doi.org/10.1167/9.11.1>
- Haberman, J., Lee, P., & Whitney, D. (2015). Mixed emotions: Sensitivity to facial variance in a crowd of faces. *Journal of Vision*, 15(4), 16. <https://doi.org/10.1167/15.4.16>
- Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), R751–R753.
<https://doi.org/10.1016/j.cub.2007.06.039>
- Haberman, J., & Whitney, D. (2009). Seeing the mean: Ensemble coding for sets of faces. *Journal of Experimental Psychology. Human Perception and Performance*, 35(3), 718–734. <https://doi.org/10.1037/a0013899>
- Haberman, J., & Whitney, D. (2010). The visual system discounts emotional deviants when extracting average expression. *Attention, Perception & Psychophysics*, 72(7), 1825–1838. <https://doi.org/10.3758/APP.72.7.1825>
- Haberman, J., & Whitney, D. (2011). Efficient summary statistical representation when change localization fails. *Psychonomic Bulletin & Review*, 18(5), 855–859.
<https://doi.org/10.3758/s13423-011-0125-6>
- Halberda, J., Mazocco, M. M. M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature*, 455(7213), 665–668.
<https://doi.org/10.1038/nature07246>
- Halberda, J., Sires, S. F., & Feigenson, L. (2006). Multiple Spatially Overlapping Sets Can Be Enumerated in Parallel. *Psychological Science*, 17(7), 572–576.
<https://doi.org/10.1111/j.1467-9280.2006.01746.x>

- Im, H. Y., & Chong, S. C. (2014). Mean Size as a Unit of Visual Working Memory. *Perception*, 43(7), 663–676. <https://doi.org/10.1068/p7719>
- Im, H. Y., Tiurina, N. A., & Utochkin, I. S. (2020). An explicit investigation of the roles that feature distributions play in rapid visual categorization. *Attention, Perception, & Psychophysics*. <https://doi.org/10.3758/s13414-020-02046-7>
- Leib, A. Y., Kosovicheva, A., & Whitney, D. (2016). Fast ensemble representations for abstract visual impressions. *Nature Communications*, 7(1). <https://doi.org/10.1038/ncomms13186>
- Levin, D. T., Takarae, Y., Miner, A. G., & Keil, F. (2001). Efficient visual search by category: Specifying the features that mark the difference between artifacts and animals in preattentive vision. *Perception & Psychophysics*, 63(4), 676–697.
- Li, F. F., VanRullen, R., Koch, C., & Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences*, 99(14), 9596–9601. <https://doi.org/10.1073/pnas.092277599>
- Liao, J., Lima, R. S., Nehab, D., Hoppe, H., Sander, P. V., & Yu, J. (2014). Automating Image Morphing Using Structural Similarity on a Halfway Domain. *ACM Transactions on Graphics*, 33(5), 1–12. <https://doi.org/10.1145/2629494>
- Long, B., Störmer, V. S., & Alvarez, G. A. (2017). Mid-level perceptual features contain early cues to animacy. *Journal of Vision*, 17(6), 20. <https://doi.org/10.1167/17.6.20>
- Long B., Yu C., & Konkle T. (2018). Mid-level visual features underlie the high-level categorical organization of the ventral stream. *Proceedings of the National Academy of Sciences*, 115 (38), 9015-9024. <https://doi.org/10.1073/pnas.1719616115>
- Morgan, M., Chubb, C., & Solomon, J. A. (2008). A “dipper” function for texture discrimination based on orientation variance. *Journal of Vision*, 8(11), 9–9. <https://doi.org/10.1167/8.11.9>

- Myczek, K., & Simons, D. J. (2008). Better than average: Alternatives to statistical summary representations for rapid judgments of average size. *Perception & Psychophysics*, 70(5), 772–788. <https://doi.org/10.3758/PP.70.5.772>
- Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2015). Direct encoding of orientation variance in the visual system. *Journal of Vision*, 15(4), 3. <https://doi.org/10.1167/15.4.3>
- Oriet, C., & Hozempa, K. (2016). Incidental statistical summary representation over time. *Journal of Vision*, 16(3), 3–3. <https://doi.org/10.1167/16.3.3>
- Palan, S., & Schitter, C. (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17, 22–27. <https://doi.org/10.1016/j.jbef.2017.12.004>
- Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, 4(7), 739–744. <https://doi.org/10.1038/89532>
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163. <https://doi.org/10.1016/j.jesp.2017.01.006>
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
- Perrinet, L. U., & Bednar, J. A. (2015). Edge cooccurrences can account for rapid categorization of natural versus animal images. *Scientific Reports*, 5, 11400. <https://doi.org/10.1038/srep11400>

- Rosenholtz, R., Huang, J., Raj, A., Balas, B. J., & Ilie, L. (2012). A summary statistic representation in peripheral vision explains visual search. *Journal of Vision*, 12(4), 14–14. <https://doi.org/10.1167/12.4.14>
- Rousselet, G. A., Fabre-Thorpe, M., & Thorpe, S. J. (2002). Parallel processing in high-level categorization of natural images. *Nature Neuroscience*, 5(7), 629–630. <https://doi.org/10.1038/nn866>
- Schmidt, F., Hegele, M., & Fleming, R. W. (2017). Perceiving animacy from shape. *Journal of Vision*, 17(11), 10. <https://doi.org/10.1167/17.11.10>
- Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical Scaling Reveals a Unified Theory of Visual Memory Strength. *Nature Human Behaviour*, 4, 1156–1172.
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, 27(4), 665–678. https://doi.org/10.1162/jocn_a_00733
- Solomon, J. A., Morgan, M., & Chubb, C. (2011). Efficiencies for the statistics of size discrimination. *Journal of Vision*, 11(12), 13–13. <https://doi.org/10.1167/11.12.13>
- Suchow, J. W., Brady, T. F., Fougner, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, 13(10). <https://doi.org/10.1167/13.10.9>
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 520–522. <https://doi.org/10.1038/381520a0>
- Treisman, A. (2006). How the deployment of attention determines what we see. *Visual Cognition*, 14(4–8), 411–443. <https://doi.org/10.1080/13506280500195250>
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. [https://doi.org/10.1016/0010-0285\(80\)90005-5](https://doi.org/10.1016/0010-0285(80)90005-5)

- Utochkin, I. S. (2015). Ensemble summary statistics as a basis for rapid visual categorization. *Journal of Vision*, 15(4), 8. <https://doi.org/10.1167/15.4.8>
- Utochkin, I. S., Khvostov, V. A., & Stakina, Y. M. (2018). Continuous to discrete: Ensemble-based segmentation in the perception of multiple feature conjunctions. *Cognition*, 179, 178–191. <https://doi.org/10.1016/j.cognition.2018.06.016>
- Utochkin, I. S., & Yurevich, M. A. (2016). Similarity and heterogeneity effects in visual search are mediated by “segmentability”. *Journal of Experimental Psychology: Human Perception and Performance*, 42(7), 995–1007. <https://doi.org/10.1037/xhp0000203>
- VanRullen, R., & Thorpe, S. J. (2001). Is it a Bird? Is it a Plane? Ultra-Rapid Visual Categorisation of Natural and Artifactual Objects. *Perception*, 30(6), 655–668. <https://doi.org/10.1068/p3029>
- Watamaniuk, S. N. J., & Duchon, A. (1992). The human visual system averages speed information. *Vision Research*, 32(5), 931–941. [https://doi.org/10.1016/0042-6989\(92\)90036-I](https://doi.org/10.1016/0042-6989(92)90036-I)
- Whitney, D., & Yamanashi Leib, A. (2018). Ensemble Perception. *Annual Review of Psychology*, 69(1), 105–129. <https://doi.org/10.1146/annurev-psych-010416-044232>
- Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, 1(3), 0058. <https://doi.org/10.1038/s41562-017-0058>
- Wolfe, J. M., Võ, M. L.-H., Evans, K. K., & Greene, M. R. (2011). Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences*, 15(2), 77–84. <https://doi.org/10.1016/j.tics.2010.12.001>
- Wolfe, J.M., & Utochkin, I. S. (2019). What is a preattentive feature? *Current Opinion in Psychology*, 29, 19–26. <https://doi.org/10.1016/j.copsyc.2018.11.005>

Zachariou, V., Del Giacco, A. C., Ungerleider, L. G., & Yue, X. (2018). Bottom-up processing of curvilinear visual features is sufficient for animate/inanimate object categorization. *Journal of Vision*, 18(12):3, 1–12. <https://doi.org/10.1167/18.12.3>

Appendix A

Experiment A

Experiment A was used to create and validate the stimulus set.

Method

Participants

We collected data online from 100 participants (38 female, mean age = 27.75 years, SD = 8.75 on Prolific (www.prolific.ac; Palan & Schitter, 2018; Peer et al., 2017). The participants received £0.42 for participation. They completed the task using a laptop or desktop computer. At the beginning of the experiment, they gave informed consent online.

Stimuli

Stimuli were presented via Pavlovia.org based on PsychoPy 3.1 (Peirce et al., 2019). In our experiments, we used morphed silhouettes (500 × 500 pixels) of various animals and inanimate objects. All images were unambiguous in terms of animacy (we did not use pictures of toys, robots, etc). We used the “Neural Best-Buddies” neural network (Aberman et al., 2018) and morphing program (Liao et al., 2014). To decrease the familiarity of each particular silhouette as a recognizable everyday object, we first morphed between pairs of animate creatures (e.g., cat and owl, 50% from each, Figure S1) and pairs of inanimate objects (e.g., tractor and lift truck, 50% of each, Figure S1). At the second stage of morphing, we morphed between these combined silhouettes from opposite animacy categories (e.g., the “cat-owl” silhouette was morphed with the “tractor-lift truck”, Figure S1). For each continuum, hence, four silhouettes were used as initial sources (two images of animals and two images of inanimate objects) which were rather close in shape to avoid poor morphing performance.

Overall, we created 10 morph continua. The morphing step was 10% so that we obtained 11 images in each morph continuum. We manually removed all “grey” regions in morphed silhouettes caused by the imperfection of the morphing algorithm.

Figure S1. An example of the morphing algorithm for one of the animacy continua.

Procedure

Each trial started with a blank screen for 500 ms. It was followed by a morph presentation for 500 ms, again followed by a blank screen until response. Participants had to press the “L” on a keyboard if the image looked more animate than inanimate, or the “A” if the image looked more inanimate than inanimate. Each participant saw each image once (110 trials in total). The experiment lasted about 5 minutes.

Results and data analysis

For each image, we estimated the probability of it being categorized as animate. Based on these probabilities we selected continua providing fairly unambiguous categorization at endpoints: below .25 probability of the “animate” response for the inanimate end in the

morphing space and above .75 probability for the animate end in that space. Thus, we removed from further experiments (Figure S2) pairs 1, 6, and 8. It is important to note that despite the physically smooth transition from inanimate to an animate endpoint, there was no such smoothness in categorization probability as a function of the morphing step. On the contrary, a quite general pattern (Figure S2) for all of the continua was that one-two images around the middle (with some slight shifts towards an animate or inanimate side varying across the continua) that yielded ambiguous categorization ($\sim .50$ “animate” responses), whereas the rest of the images were categorized as fairly animate and inanimate. This is consistent with categorical perception, which is found in most domains where binary categories are imposed on continuous stimuli (Harnad, 2003; Goldstone & Hendrickson, 2010; Calder, Young, Perrett, Etcoff, & Rowland, 1996).

Figure S2. Results of categorization in online stimulus validation experiment. The red background ones were subsequently excluded for not reaching the required threshold of discriminability (e.g. unambiguous classification of the most extreme stimuli).

References

Harnad, S. (2003) Categorical Perception. *In, Encyclopedia of Cognitive Science.*

Encyclopedia of Cognitive Science (01/12/03) Nature Publishing Group: Macmillan.

Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley*

Interdisciplinary Reviews: Cognitive Science, 1(1), 69-78.

Calder, A. J., Young, A. W., Perrett, D. I., Etcoff, N. L., & Rowland, D. (1996). Categorical

perception of morphed facial expressions. *Visual Cognition*, 3(2), 81-118.