

Mapping visual features onto numbers

Erik Brockbank (ebrocbank@ucsd.edu)

Department of Psychology, 9500 Gilman Dr.
La Jolla, CA 92093-109 USA

Edward Vul (evul@ucsd.edu)

Department of Psychology, 9500 Gilman Dr. #109
La Jolla, CA 92093-109 USA

Abstract

Modern society frequently requires that we express our subjective senses in objective, shared formal systems; this entails mapping multiple internal variables onto a common scale. Here we ask whether we accomplish this feat in the case of estimating number by learning a single mapping between explicit numbers and one integrated subjective estimate of numerosity, or if we separately map different perceptual features onto numbers. We present people with arrays of dots and ask them to report how many dots there are; we rely on the systematic under/overestimation of number at higher quantities to estimate error in the mapping function. By comparing how this error changes over time, as the mapping fluctuates for different visual cues to numerosity, we can evaluate whether these cues share a single mapping, or are mapped onto number individually. We find that area, size, and density all share a common mapping, indicating that people obtain a unified subjective estimate of numerosity before mapping it onto the formal number line.

Keywords: numerosity; number; estimation

Introduction


Human reasoning and planning frequently involves mapping internal estimates onto formal systems: we can compare the weights of two rocks using our subjective sense of weight, but to provide an estimate of one rock's weight in kilograms requires mapping that subjective sense of weight onto a formal metric system. This task of expressing our internal subjective senses in objective, standard systems is commonplace, from making time estimates to evaluating prices. To accomplish this we somehow learn to map from perceptual and internal states onto formal systems like weight in kilograms. The task is often complicated by the fact that we might have many subjective variables that must map on to the same formal system: weight might be estimated by the pressure a rock exerts on our hand, or by its inertia as we try to move it. How do we deal with multiple subjective cues: do we map each one individually onto the formal system, or do we combine them to come up with a single subjective estimate, and then map that estimate onto the formal system? In this paper, we approach this question for people's ability to estimate numerosity.

Based on a quick glance at a display of many objects, people can estimate the number of objects present in the display. Even if there is insufficient time to explicitly count the objects, there are enough visual features that correlate with number, that a number estimate may be obtained just based on these internal, analog signals which together give us a sense


of "Approximate Magnitude". Imagine for example stepping into a room full of people: as you look around, you can get a rough sense of how many are present just based on the density of the crowd and the size of the room faster than you would be able to count each person individually. In general, displays with higher numerosity tend to have objects distributed over a larger portion of the visual field (area) and the number of objects in a constant area tends to be higher, either because the objects themselves tend to be smaller (size), or because the inter-object distances are smaller (spacing/density). These separate cues to numerosity may be treated in different ways by the visual system. They may be combined into one internal representation of numerosity which forms the basis of estimation. Or, because visual cues to numerosity all tend to correlate together in the real world, it may not be necessary to undertake the extra calculation of integrating them to form an internal sense of number. Instead, these features may be mapped onto numbers directly. Both explanations posit internal representations which must be mapped onto formal numbers when making estimates, but differ as to how this mapping occurs.

A large body of research has examined the representations that support our internal sense of number but comparatively little work examines how we might map from that internal sense to number estimates. The degree to which we directly perceive and represent number is an area of active debate (for a recent review see Leibovich, Katzin, Harel, and Henik (2017)). Researchers have proposed that an internal sense of number, the "Approximate Number System" or ANS, exists in many animals and is developed by infants at a young age (Feigenson, Dehaene, & Spelke, 2004). However, competing accounts emphasize that perceptual features of a quantity such as size, area, and density are highly correlated with number: this has led some to argue that our ability to estimate numerical quantities can be served directly by these continuous magnitudes without any internal number sense (Gebuis & Reynvoet, 2012) or that insofar as we have an internal representation of number, it is assembled directly from our sense of continuous magnitudes (Leibovich et al., 2017).

The present experiment is agnostic about the precise mechanisms for visual processing and internal representation of numerosity. We are interested in understanding how people map from various internal representations to the formal number line during numerical estimation. One hypothesis is that



(a) A relationship between visual properties and number estimates that relies on multiple independent mappings (numbered 1, 2, and 3) from distinct visual features to an estimate “n”



(b) A relationship between visual properties and number that specifies a *single mapping* (numbered 1) between some internal quantity representation (the red incline) and an estimate “n”

Figure 1: Two different ways of thinking about how visual cues to magnitude map onto numbers

people have multiple mappings which take as their inputs features associated with numerosity such as size, density, and area. Another hypothesis is that people instead have a single mapping from some internal representation—whether a number sense or a broader integrated magnitude—to an estimated quantity. Both mapping hypotheses are plausible a priori and might inform the broader debate about how people perceive or represent number. In what follows, we discuss in greater detail the research on number representation, which supports the availability of various possible *inputs* to this mapping function.

ANS and Continuous Magnitudes

The predominant theory in number processing holds that people have an internal approximate number system which they map onto the formal number system for purposes of estimation and other related tasks (Izard & Dehaene, 2008). Work in this space has sought to model the characteristics of this number system, including how it is represented internally (Izard & Dehaene, 2008) and how it develops in infants and young children (Carey, 2009). Research on development of the approximate number system has found that ability to distinguish between distinct numbers—the *acuity* of the approximate number system—develops independently of ability to discriminate area, density, length, and time (Odic, 2018) and that acuity of number sense in children is correlated with mathematical ability later in life (Halberda, Mazocco, &

Feigenson, 2008). In line with the idea that numerosity is a core part of how we represent the world around us, some have argued that numerosity is even available as a primary feature of perception and not reducible to related properties like texture density (Burr & Ross, 2008). In support of this, it has been shown that numerosity estimates are subject to visual adaptation effects, much like other visual properties such as color and motion (Burr & Ross, 2008).

In contrast to proposals that humans have an internal approximate number system, some have argued that number estimation is inferred directly from visual properties that correlate with number (Dakin, Tibber, Greenwood, & Morgan, 2011). Evidence that people’s ability to estimate quantities stems directly from their processing of visual cues comes primarily from work showing that people struggle to infer numerosity independently of the information they receive from visual cues (Leibovich et al., 2017). For example, Gebuis and Reynvoet (2012) presented participants with a series of dot arrays which manipulated the convex hull, aggregate surface, and density of the dots such that none of these visual properties correlated with true quantity across all the trials. They found that participants’ estimates of the number of dots in the arrays were largely explained by each of these features even though these features provided no information about the true number of dots. They argue that people are therefore unable to estimate numerosity independently of the visual cues which tend to provide certain signals about numerosity. More recent work has argued that the basis for our sense of number is our ability to process *continuous magnitudes* (density, area, size, etc.): Leibovich et al. (2017) challenge the degree to which research on the approximate number system is able to isolate numerosity from visual cues and argue for a more general magnitude system from which number is inferred.

Regardless of whether people have an internal sense devoted specifically to numerosity or assemble their sense of quantity from continuous magnitudes that correlate with number, it’s clear from the existing research that a.) numerosity and visual features such as area, size, and density are closely tied and b.) that a mapping from internal quantity estimates onto the formal number system could plausibly take as its inputs any combination of visual features and numerical representation. In light of this, we propose two hypotheses about how such a mapping might work. One holds that we have multiple mappings from size, area, and density features to number estimates. These mappings could independently serve our estimation needs. The other holds that we have a single mapping from some internal quantity representation onto the number line. This internal representation could be our approximate number sense or a quantity estimate assembled by combining information from size, area, and density. In what follows, we summarize research which has examined people’s performance on number estimation tasks and describe a novel method of investigating mappings from internal number to formal number.

Individual differences and drift in mental number-line calibration

Work investigating the approximate number system has sought to understand how we map from our internal sense of number to the verbal number system. Several key findings have informed this line of inquiry. First, people's mapping from internal quantity representations to formal numbers is often miscalibrated (Izard & Dehaene, 2008). Specifically, individuals asked to estimate quantities outside the *subitizing* range tend to systematically over- or underestimate those quantities. This relationship follows a power law: the higher the true quantity, the more people reliably over- or underestimate (Izard & Dehaene, 2008). Second, the amount that people are miscalibrated in their estimations varies considerably across individuals (Vul, Barner, & Sullivan, 2013). Some people reliably overestimate while others reliably underestimate, suggesting that whatever mapping we use onto formal number varies from person to person. Finally, the amount that people are miscalibrated in their estimations varies *within* individuals. In other words, the degree to which people over- or underestimate has been shown to *drift* over many successive estimations (Vul et al., 2013).

In this study, we use the slow drift of individuals' number estimates to investigate the mapping between internal number and quantity estimates across various visual conditions. If people rely on multiple mappings from independent visual features to numerosity estimates, then we would expect these mappings to drift independently over many estimates based on different visual features. However, if people utilize a singular mapping function from some internal quantity estimate onto number, then we would expect the drift in their estimates to be invariant to changes in the visual cues used to form each estimate. This difference is illustrated in Figure 1.

Experiment

We tested the degree to which the numerosity of a display is estimated through independent mappings from correlated visual features (e.g., size, area, and density) to number, or if these features map onto a single internal numerosity estimate, which is then mapped onto a symbolic number. We presented participants with an estimation task in which the stimuli varied along one dimension (size, area, or density) as magnitude changed, while holding the other two dimensions constant. We compared the drift in participants' magnitude estimations across the three conditions to assess whether number estimates may be independently obtained from size, area, and density cues.

Participants

Participants were 57 undergraduates at the University of California, San Diego who received course credit for their participation.

Methods

Participants were shown a series of dot arrays on a white background like those in Figure 2. The dots appeared on

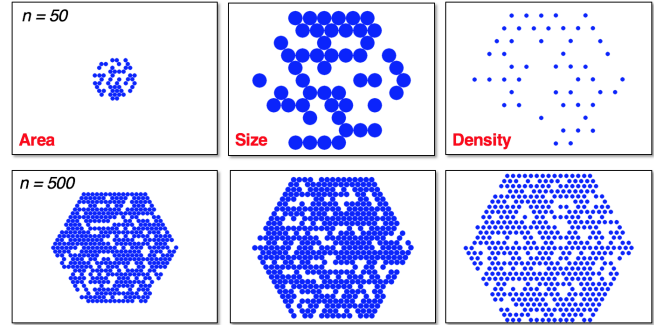


Figure 2: Sample stimuli for $n = 50$ and $n = 500$ across *area*, *size*, and *density* conditions.

the screen for 500ms and then disappeared. Participants were then prompted to guess the number of dots on the screen. We did not use a mask between trials, as any aid that participants received in estimating due to an after image would have been consistent across all trials. For the first 25 trials, participants were given feedback after each round about the true number of dots they had just seen. Participants were awarded points after each round on a logarithmic scale based on the difference between their guess and the true number of dots. Participants performed 1,000 estimation trials or 50 minutes on the task, whichever came first.

Stimuli

The number of dots on each trial was selected by sampling a number between 10 and 750 from an exponential distribution with a mean of 100. Each trial of the experiment was randomly selected to vary one of the *area*, *density*, or *size* of the dots while keeping the other two constant.

Trials that varied the *area* of the dots used a predetermined constant for spacing between dots and dot size so that the number of dots on the screen was indicated by how much area the dots occupied. For trials that varied the *density* of dots, dots were populated in a constant area on the screen with a constant size: when there were more dots, the spacing between them was lower and when there were fewer dots, there was greater spacing between them. Finally, trials that varied the *size* of the dots used a consistent area of the screen and a consistent spacing between dots, generating larger dots when there were fewer in a given trial and smaller dots when the magnitude was greater. See Figure 2 for examples of dot arrays that varied each visual feature. In each trial, a random selection determined whether the dot display would vary size, area, or density so that over the course of the experiment, all three features would present cues to numerosity but on any given trial, only one would be informative.

Results

To understand how well participants estimate visual quantities based on size, area, and density inputs, we can compare their estimates for trials in which each of these features were informative to the actual numbers presented. For




Figure 3: Individual data for *Number presented* and *number reported* from three sample participants. The degree to which each participant underestimates, overestimates, or is accurate reflects individual differences in this task

perfect estimators, each estimate plotted this way would lie along the identity line. Estimates show a high degree of variance across individuals: Figure 3 shows data from three sample participants which illustrate this. Combining this data across all participants, Figure 4 shows participants' accuracy by plotting their estimates alongside the true number presented for trials varying size, area, and density. Consistent with earlier findings, people are accurate up to numerosities of about 20-30, but they reliably underestimate larger numbers on average (even setting aside the individual variance: see Vul et al., 2013). The underestimation pattern in a given set of trials can be described as a bilinear function which follows the identity line up to a threshold, and deviates from the identity line with some slope thereafter. This slope amounts to the “calibration” of the mental number line (Izard & Dehaene, 2008), and was precisely shown to (a) vary across subjects, and (b) within subjects, slowly drift over the course of an experiment (Vul et al., 2013). Figure 4 shows that the calibration of the mapping to the formal number line is similar regardless of whether area, size, or density is the numerosity-informative variable.

Previous research has shown that perception of structure and groupings can lead to systematic underestimation. For example, objects connected by lines are underestimated relative to disconnected objects (Franconeri, Bemis, & Alvarez, 2009). When dot arrays are seen as grouped, the degree to which they're clustered increases underestimation (Im, Zhong, & Halberda, 2016). Even sub-conscious processes like statistical learning of co-occurrence in colored dot arrays can lead to underestimation (Zhao & Yu, 2016). In the stimuli presented here, it is possible that perception of grouping among dots on various trials led to underestimation. However, given the similarity in patterns of underestimation in




Figure 4: *Number presented* and *number reported* across estimate conditions. The red line is the median response for each number presented. Participants show similar underestimations across estimate conditions.

Figure 4 across modalities, it's unlikely that underestimation due to perceptual grouping affected the size, area, or density informative trials more than any other.

Consistent individual differences across modalities

Dividing each participant's estimates into blocks of size 50 (after the initial 25 “calibration” trials), we can evaluate the best fitting slope estimates for each block in each condition. For example, block 11 for each participant will contain trials 476–525. Of these, some number will be area trials, some will be size trials, and some will be density trials. We can extract the trials belonging to each estimate condition (size, area, density) for a given block of trials for a given participant, and compute a best fitting slope for that subset of number estimates for that participant. This gives us a slope for each estimate condition, for each participant, in each block of 50 trials over the course of the experiment.

To the degree that processes of estimating quantity based on changes in size, area, and density are independently calibrated to the data participants have seen, individual differences in slopes should not be consistent across these different modalities. However, if different variables are mapped onto a subjective quantity estimate, and the uncertain, idiosyncratic mapping lies between approximate number and reported number, then an individual's slope for area-determined numerosity will be consistent with their density- and size-determined numerosity as well.

Figure 5 shows each of the possible correlations between best fitting slopes for size, area, and density estimates across participants in each block: size-area slope correlations, size-density slope correlations, and density-area slope correlations. Block zero (trials 1-25) includes the trials in which participants received feedback after each guess. In these tri-




Figure 5: Slope correlations of size-area, size-density, and density-area for the first 11 blocks of the experiment. Participants received feedback in block 0. After that, correlations across estimate conditions are relatively high and are tightly coupled in remaining blocks (error bars indicate 95% confidence intervals on the correlation coefficient).

ally, participants were fairly accurate in their estimates. Correlation between slope estimates in block zero is therefore low due to low variance in slope estimates across all conditions as a result of the feedback. However, from block one onwards, correlation of slope estimates between size and density trials, size and area trials, and density and area trials increases to 0.6–0.8. Critically, each of these three slope estimate correlations (size-area, size-density, and density-area) remain high and tightly in tandem from block one onwards. Such closely aligned correlations would be unusual if variations in stimulus size, density, and area across trials each directly and independently enabled an estimate of quantity.

Consistent within-individual drift across modalities

The correlation between different slope estimates for each block, described above, indicates how similar size, area, and density estimates were to each other across participants in each set of 50 trials throughout the experiment. In other words, this shows whether individual differences in numerosity estimations are consistent across modalities. Another feature of numerosity estimates is their drift in calibration over time *within individuals*. Here, we examine the data from estimation across modalities in light of this pattern: if size, area, and density each independently map to a formal number estimate, the drift in calibration for each of these modalities should be independent. However, if each modality maps to an internal estimate and the drift reflects changes in the mapping of internal estimates to formal number, then we will not detect any difference in drift across modalities. Figure 6 shows the correlation between slopes in blocks 1–11 for each possible comparison of estimate conditions: autocorrelation of e.g. density slopes across blocks and correlations across modali-

ties of e.g. size to area slopes between each block. The overall pattern of correlations across blocks looks very similar for each of these comparisons, further reinforcing the idea that these features do not map separately onto number estimates. Individual drift in over- and underestimations can be seen in the lower correlation between blocks that are farther apart: this pattern is also similar across comparisons.

To better compare the slope correlations within and across modalities, we group pairs of blocks based on their temporal separation: their *trial distance*. For example, the correlation between blocks 1 and 4, and 2 and 5, and 3 and 6, all have a trial distance of 3 blocks (150 trials). The decline in correlations over longer trial distances indicates the drift of mapping over time. We can thus compare these cross-correlation functions for different modality-modality comparisons. Comparing slope estimates for each block of a given condition to those that are all an equal distance away in the same or alternate estimate conditions gives us a correlation between trials across a range of trial distances. Figure 7 shows these correlations by distance for the same combinations of estimate conditions shown in Figure 6. Across all comparisons, the correlations decrease as distance between trials increases. This drift in estimate slopes—the slopes of trial blocks farther from each other are less similar to the slopes of closer blocks—reflects the drift in calibration of the mapping function onto precise quantity over time (Vul et al., 2013).

To ensure that the “drift” shown in Figure 7 is not attributable to differential distributions of each trial type across blocks of increasing distances, we shuffled trial order for each participant and re-calculated the correlation of slopes by trial distance. For the shuffled data, the correlations within and between modalities were very stable across *all* trial distances: in other words, there was no sign of systematic drift. We fit linear models to the mean correlations at each trial distance for each modality comparison to ensure that indeed there was no drift in the shuffled data: for the six comparisons shown in Figure 7 (with shuffled trial order), none had a slope significantly different from 0¹.

Figure 7 illustrates that the drift in calibration occurs not only within each estimate condition but also across them: size-density, size-area, and density-area slope comparisons show similar decreasing correlations at greater trial distances. Most importantly, these correlations over trial distances are indistinguishable whether we consider within-modality correlations (e.g., area-area) or across-modality correlations (e.g., area-size). The correlation of slopes over varying trial distances is indistinguishable within and across modalities. If visual cues to density, size, and area each map to a subjective numerosity which *then* maps to precise quantity estimates, the similarity of slope correlations within and across esti-

¹Shuffled trial order correlation slopes (per block): size-size 95% CI = [-0.013, 0.001] $p = 0.07$, density-density 95% CI = [-0.010, 0.002] $p = 0.16$, area-area 95% CI = [-0.001, 0.015] $p = 0.08$, size-density 95% CI = [-0.019, 0.003] $p = 0.12$, size-area 95% CI = [-0.014, 0.003] $p = 0.19$, density-area 95% CI = [-0.006, 0.013] $p = 0.39$

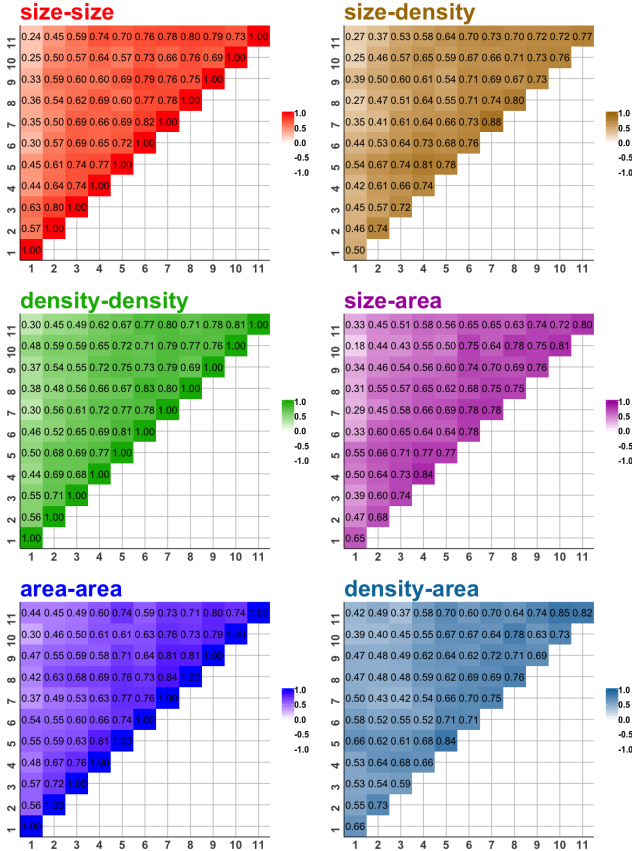


Figure 6: Slope comparisons between each block across all estimate conditions. The correlation pattern is similar across all these comparisons, a pattern seen when plotting correlation by trial distance as well.

mate conditions could be accounted for by separate mapping functions from visual properties to approximate number being similarly calibrated. The drift would then be attributed to the mapping from this internal numerosity to the number reported. If the estimation of reported quantities was accomplished by separate mapping functions from each of the visual modalities, it would be improbable for these mapping functions to change in lock-step, thus rendering the within- and across-modality correlations identical.

Conclusion

We asked how people map their analog, perceptual features onto explicit numbers. Specifically, we investigated whether (a) people have one mapping from a cohesive, internal estimate of numerosity/magnitude onto the number line, or if (b) people have multiple mappings onto number from different visual features that tend to correlate with number. In this experiment, we asked participants to estimate the number of dots present in a display, while we varied which visual feature varied with number. The numerosity of the dot arrays in a given trial could only ever be inferred from a single perceptual modality (one of size, area, or density) of the dots.

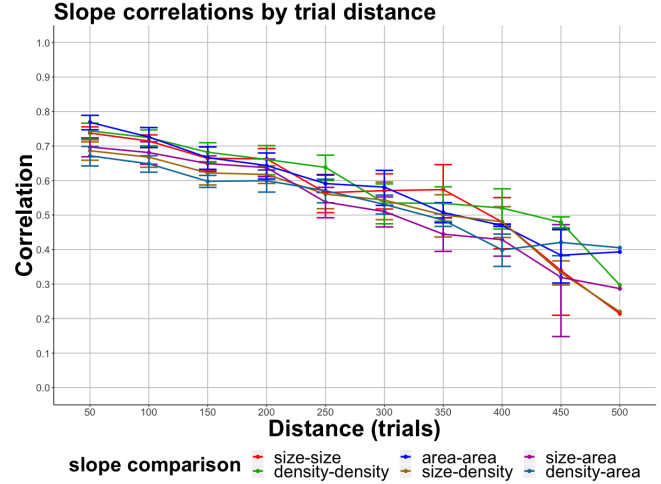


Figure 7: Slope comparisons *within* and *across* estimate conditions by trial distance. The correlation of slopes over varying trial distances is indistinguishable within and across estimate conditions (error bars represent standard error across measurements of each distance).

Using results from prior research indicating that people drift in their mapping of numerosity to number over the course of many trials (Vul et al., 2013), we analyzed whether such drift occurs independently across estimates based on size, density, and area. We find that the (mis)calibration of numerosity onto number is indistinguishable between size, area, and density trials. Moreover, as this mapping drifts over the course of many trials, it changes in lock-step for all the modalities, indicating that there is only one mapping function that drifts, which is shared across all modalities. From this we conclude that size, area, and density all share a common mapping onto formal number. In other words, perceptual features that are cues to numerosity must be combined into an internal representation of numerosity which is then mapped onto the formal number line when reporting an exact number.

It would be rash to generalize these results to all cases in which we might map subjective senses onto objective, external standards (e.g, estimating weight in kilograms, or our willingness-to-pay in dollars). It seems likely that some formal systems do not have a corresponding unified internal representation, and instead have an assortment of independent mapping functions which may be inconsistent and incommensurate. However, these results are encouraging that for at least some formal systems, we have unified, coherent internal representations that serve as their substrate. We postulate that the methods we develop here—of relying on fluctuations in the mapping of subjective states onto formal, objective systems—might be used to identify whether notions like “subjective value” are indeed unified monolithic entities, or if they are an ensemble of related, but independent, internal senses.

Acknowledgments

We thank our reviewers for their thoughtful suggestions.

References

- Burr, D., & Ross, J. (2008). A visual sense of number. *Current biology*, *18*, 425–428.
- Carey, S. (2009). Where our number concepts come from. *The Journal of philosophy*, *106*.
- Dakin, S., Tibber, M., Greenwood, J., & Morgan, M. (2011). A common visual metric for approximate number and density. *Proceedings of the National Academy of Science*, *108*, 19552–19557.
- Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in cognitive sciences*, *8*, 307–314.
- Franconeri, S., Bemis, D., & Alvarez, G. (2009). Number estimation relies on a set of segmented objects. *Cognition*, *113*(1), 1–13.
- Gebuis, T., & Reynvoet, B. (2012). The role of visual information in numerosity estimation. *PLoS One*, *7*.
- Halberda, J., Mazocco, M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature*, *455*.
- Im, H., Zhong, S., & Halberda, J. (2016). Grouping by proximity and the visual impression of approximate number in random dot arrays. *Vision research*, *126*, 291–307.
- Izard, V., & Dehaene, S. (2008). Calibrating the mental number line. *Cognition*, *106*, 1221–1247.
- Leibovich, T., Katzin, N., Harel, M., & Henik, A. (2017). From sense of number to sense of magnitude: The role of continuous magnitudes in numerical cognition. *Behavioral and Brain Sciences*, *40*.
- Odic, D. (2018). Children's intuitive sense of number develops independently of their perception of area, density, length, and time. *Developmental Science*, *21*.
- Vul, E., Barner, D., & Sullivan, J. (2013). Slow drift of individuals' magnitude-to-number mapping. In *Proceedings of the 35th annual meeting of the cognitive science society*.
- Zhao, J., & Yu, R. (2016). Statistical regularities reduce perceived numerosity. *Cognition*, *146*, 217–222.