Title

Arc Length as a Geometric Constraint for Psychological Spaces

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Ralston, Robert
Sloutsky, Vladimir

Publication Date

2021

Peer reviewed

Arc Length as a Geometric Constraint for Psychological Spaces

Robert Ralston (ralston.123@osu.edu)
Department of Psychology, 1835 Neil Avenue
Columbus, OH 43210 USA

Vladimir Sloutsky (sloutsky@psy.ohio-state.edu)
Department of Psychology, 1835 Neil Avenue
Columbus, OH 43210 USA

Abstract

Many cognitive models assume that stimuli can be represented as points in a latent psychological space. However, it has been difficult to provide these spaces with a geometric structure where the distance between items accurately reflects their subjective dissimilarity. In this paper, we propose a new method to give psychological spaces a geometric structure by equating the amount of change undergone by a stimulus with the arc length of a curve in psychological space. We then assess our method with a categorization experiment where participants classified continuously changing visual stimuli according to their rate of change. Our results indicate that individuals' judgements are well predicted by arc length, suggesting that it may be a promising geometric constraint for psychological spaces in other contexts.

Keywords: similarity; psychophysics; selective attention; psychological spaces

One of the fundamental principles of behavior is that sufficiently similar stimuli evoke similar responses. To embody this principle, many cognitive models assume that stimuli lie in a latent *psychological space*, where subjectively similar stimuli are assigned coordinates with similar values. For example, in the Generalized Context Model of classification (GCM; Nosofsky, 1986), items are represented as points in a metric space, where the similarity of items is a decreasing function of the distance between them. Models such as the GCM that employ a psychologically meaningful notion of distance are known as geometric models and have been applied in a variety of domains, such as categorization, recognition, and semantics (e.g., Love, Medin, & Gureckis, 2004; Gärdenfors, 2017).

Despite their successes, many researchers are skeptical that distance in a psychological space provides a good model of subjective similarity. A major reason for this skepticism is that explicit similarity judgements between two items may require a more complex comparison mechanism than is allowed by a metric (Tversky, 1977; Hahn, Chater, & Richardson, 2003).

In this paper, we propose a new method to give psychological spaces a geometric structure reflecting the subjective difference between stimuli. However, unlike previous approaches, our method does not make use of similarity judgements or the confusability of items. In everyday life, we encounter many stimuli that change through time. For example, the individuals we know may change in height, weight, age, hair color, spatial location, and in a myriad of other ways. Additionally, not all stimuli appear to change at the same rate; it may seem like a stimulus has changed a greater or lesser amount through a period of time.

In the sections that follow, we propose an account of psychological space that equates the amount of change an individual perceives a stimulus to undergo with the arc length of the curve that the changing stimulus follows in psychological space. This constraint paired with certain psychophysical hypotheses is sufficient to induce a geometric structure on psychological space and thus a way to calculate the distance between any two stimuli without relying on previously criticized tasks such as explicit similarity judgements. In the first section, we present a mathematical model of psychological space and derive an expression for the arc length of continuous curves, corresponding to stimuli that change continuously in time. Then, in section two, we test whether the predicted arc length of a curve is related to individuals' judgements of the amount of change undergone by a stimulus.

Arc Length in Psychological Space

Finding Psychological Space

Suppose that a collection of stimuli varies according to n continuous, measurable dimensions. We will place these stimuli into a psychological space by finding a function Φ that maps the objective coordinates of a stimulus – its values on each measurable dimension – onto a manifold $\mathcal{M}$ embedded in a typically higher-dimensional space. A high-level depiction of this approach can be seen in Figure 1.

Formally, we will assume that the measurable dimensions of a stimulus set form a coordinate system $\Phi: A \subset \mathbb{R}^n \to \mathbb{R}^m$ for a differentiable manifold $\mathcal{M}$, embedded in $\mathbb{R}^m$, corresponding to an individual's representation of those stimuli in psychological space. For example, Gabor patches are grating-like stimuli that vary according to two continuous dimensions, spatial frequency and orientation (Turner, 1986). For these stimuli, A represents the subset of $\mathbb{R}^2$ corresponding to the objective spatial frequencies and orientations making up possible Gabor patches, such as those used to generate the stimuli in Figure 1. Furthermore, since orientation is a measure of the angle of the grating, and an angle of 0 radians is identical to an angle of 2π radians, the manifold given by these coordinate functions does not have the topology of a plane, but rather a cylinder, which can be

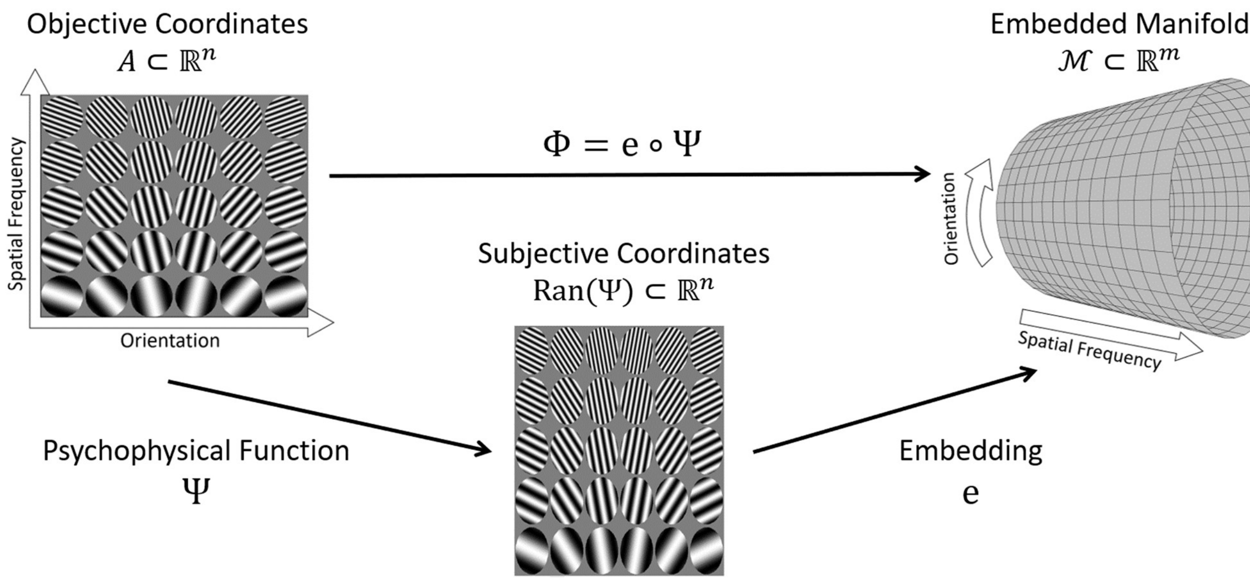

Figure 1. Schematic representation of the modeling approach.

embedded in $\mathbb{R}^3$. This treatment of psychological manifolds can be further generalized, and generalization may be necessary to study some stimuli or cognitive phenomena, but for simplicity we consider only this case.

With this formulation, we will find an expression for Φ in two parts. First, we will transform the objective coordinates of stimuli by incorporating information about which dimensions are currently being attended and the psychophysics of those dimensions. Then, to form the manifold $\mathcal{M}$, we will embed the transformed space into a typically higher-dimensional space so that stimuli that are perceived to change in continuous ways are assigned continuous curves on the manifold.

Incorporating Attention and Psychophysics. When perceiving or reasoning with stimuli, humans can selectively attend to the subset of the available information that is relevant to their goals (Carrasco, 2011; Driver, 2001). A key characteristic of selective attention is that, when dimensions are attended to, stimuli are more discriminable along those dimensions (Ho, Brown, Abuyo, Ku, & Serences, 2012; Ling, Liu, & Carrasco, 2009). In line with models such as the GCM, we assume that this process is well-modeled by a scaling mechanism, where each dimension is assigned a multiplicative attention weight between 0 and 1 that expands or compresses its scale. We will refer to these weights as $a_1, \ldots, a_n$ corresponding to the n dimensions in our stimuli. Additionally, we will assume that attention is a finite resource, so that $\sum_i a_i = 1$.

We will now specify the psychophysical transformation which will modify the objective coordinates of stimuli. Our strategy will be to give a coordinate transformation on the objective parameter space A under the constraint that distances in the new parameter space should coincide with known psychophysical laws when complete attention is paid to single dimensions. For simplicity, we restrict the discussion below to the orthogonal case, where changes in the value of one dimension do not necessitate changes in the values of other dimensions. This assumption may seem limiting; however, the method shown here can be generalized to any case where a researcher has an explicit hypothesis of interactions between dimensions.

With this assumption, consider what happens when all attentional resources are paid to a single dimension i. In this case, the discriminability of stimuli is only determined by their value on dimension i. Discriminability along a single attended dimension is a heavily studied area of psychophysics, and experimental methods for determining the psychophysical mapping for a single dimension are well developed (Gescheider, 1997). Assume that researchers know, either *a priori* or through experimentation, the form of the psychophysical functions $\psi_1, \ldots, \psi_n$ obeyed by each dimension in the current experimental context under total attention.

We will use these psychophysical mappings to induce a coordinate transformation on the objective parameter space. Let $\Psi: A \to \mathbb{R}^n$ be the transformation given by Equation 1.

$$\Psi(x_1, \ldots, x_n) = (a_1\psi_1(x_1), \ldots, a_n\psi_n(x_n)) \qquad (1)$$

Intuitively, Ψ may be thought of as the subjective coordinates used by an individual to represent stimuli under a particular attentional setting. Notice that this transformation will reproduce the above psychophysical functions ($\psi's$) as curves in $\mathbb{R}^n$ when all attention is paid to a single dimension, and where the distance along the curve indicates the subjective difference between two stimuli. Additionally, if the orthogonality assumption from above is violated, the form of Ψ may be substantially more complicated; but if dimensional interactions are hypothesized, a suitable coordinate transformation can still be derived.

Specifying the Manifold. To specify the manifold corresponding to an individual's subjective representation of the stimulus set, we must now embed the subjective coordinate space ($Ran(\Psi)$) into a typically higher-dimensional space where continuously changing stimuli are assigned continuous curves on the manifold. A stimulus that changes continuously is a sequence of stimuli where each stimulus is hard to discriminate from those that appear near it (Hénaff, Goris, & Simoncelli, 2019; Callahan-Flintoft, Holcombe, & Wyble, 2020). More rigorously, a continuous sequence of stimuli can be described for our purposes as a function $f: [0,\tau] \subset \mathbb{R} \to \mathbb{R}^n$ such that $\Phi \circ f$ is differentiable (and thus also continuous). For stimuli that change through time, the function f maps the time t to the objective coordinates of the stimulus at t, and τ represents the total duration of the sequence. Informally, this definition indicates that a continuous sequence is one where small changes in time result in small changes to the perceived values of stimuli, and the function f returns the objective coordinate representation of the changing stimulus at each time.

Having the subjective coordinates $\Psi(A)$ used by an individual, we now need to specify the manifold corresponding to an individual's subjective representations of stimuli. This must be done on a case-by-case basis such that continuous curves on the manifold correspond to continuous sequences of stimuli as defined above. For example, as we explained above, Gabor patches can be well described by a cylinder in $\mathbb{R}^3$, while other two-dimensional stimuli may be better modeled as a plane (e.g., spatial frequency and saturation) or a torus (e.g., orientation and hue).

In general, for orthogonal dimensions, a plausible manifold can be obtained from the quotient topology of the parameter space A, where the ends of dimensions are identified when that dimension is homeomorphic to a circle. Once a choice is made, an embedding $e: Ran(\Psi) \to \mathbb{R}^m$ can be chosen that equates or well-approximates distances along single dimensions in this new space to distances along the manifold in $\mathbb{R}^m$. (In the next section, we will see how to calculate the length of curves on the manifold once it is specified.) With this specification, we obtain the coordinate map $\Phi = e \circ \Psi$.

Arc Length of Curves

The most important assumption of our modeling approach is that the arc length of a continuous sequence of stimuli on the manifold is a measure of the magnitude that the stimulus is perceived to change during the sequence. In the previous section, we showed a method to obtain a representation of psychological space as a manifold embedded in $\mathbb{R}^m$ with a coordinate system $\Phi = e \circ \Psi$. With this complete, the arc length of curves can be determined in the standard way (Kreyszig, 1959). Recall that $f: [0,\tau] \to \mathbb{R}^n$ is a continuous sequence of stimuli, defined such that $\Phi \circ f$ is differentiable. Let a continuous sequence of stimuli f be represented with n component functions, so that $f(t) = \big(u_1(t), \ldots, u_n(t)\big)$. With this formulation, the line element of the manifold $\mathcal{M}$ can be determined via Equation. 2.

$$\mathrm{ds}^2 = \sum_{\alpha=1}^{n} \sum_{\beta=1}^{n} \left(\frac{\partial \Phi}{\partial u_\alpha} \cdot \frac{\partial \Phi}{\partial u_\beta}\right) \frac{du_\alpha}{dt} \frac{du_\beta}{dt} \qquad (2)$$

In this equation, the line element ds corresponds intuitively to the instantaneous distance traveled by f along $\mathcal{M}$ at t. With this specified, the arc length along $\mathcal{M}$ traversed by f from $t = 0$ to τ, is given in Equation 3.

$$s = \int_0^\tau \sqrt{ds^2}\, dt \qquad (3)$$

Typically, this integral cannot be solved exactly. However, it can be estimated with arbitrary precision for any continuous sequence of stimuli.

Arc Length for Two Dimensional Stimuli. Consider the example where a stimulus space is made up of Gabor-like stimuli varying in spatial frequency and saturation, such as those in Figure 2. Suppose that, in generating the stimuli, we used an objective parameter space $A \subset \mathbb{R}^2$, providing a nonnegative value for the spatial frequency and the saturation. Prior research has found that a power law provides an acceptable psychophysical representation of both frequency and saturation in relation to objectively measured units (Stevens, 1957). Thus, we can obtain the coordinate transformation:

$$\begin{aligned}\Psi(x_f, x_s) &= \big(a_f \psi_f(x_f), a_s \psi_s(x_s)\big) \\ &= \big(a_f x_f^{c_f}, a_s x_s^{c_s}\big)\end{aligned}$$

In this expression, the subscripts f and s correspond to frequency and saturation respectively, while the two c parameters correspond to power coefficients.

Since neither frequency nor saturation are homeomorphic to a circle, we can embed our transformed parameter space in $\mathbb{R}^2$ via the identity map. Thus, coordinates of our manifold can be given by $\Phi(x_f, x_s) = (a_f x_f^{c_f}, a_s x_s^{c_s})$.

Consider a continuous sequence of stimuli that follows an ellipse in objective coordinates:

$$f(t) = (r_f sin(2\pi t) + y_f, r_s \cos(2\pi t) + y_s) \qquad (4)$$

From Equations 2, 3, and 4, the arc length of this sequence is given in Equation 5.

$$s = 2\pi \int_0^\tau \sqrt{\begin{matrix}\left(a_f^2 c_f^2 r_f^2 x_f^{2c_f - 2} \cos^2 2\pi t\right) + \\ \left(a_s^2 c_s^2 r_s^2 x_s^{2c_s - 2} \sin^2 2\pi t\right)\end{matrix}}\, dt \qquad (5)$$

In this expression, the lengthy sum goes from the top line to the bottom line, all under the radical. The arc length s is hypothesized to correspond to an individual's representation of the total amount that the Gabor-like stimuli changes from $t = 0$ to τ.

From Arc Length to Distance. With a representation of the arc length of curves in psychological space, the distance between two points is canonically defined as the minimum arc length of a curve connecting them on the manifold, or the geodesic distance between those points. To test the model with continuous sequences of stimuli, we do not need to derive the geodesic distance. However, any attempt to define arc length will induce a geodesic distance between stimuli in this way. It can then be empirically determined whether the geodesic distance corresponds to the subjective difference between individual stimuli.

Testing the Model

In the previous section, we derived a model of psychological space that predicts the amount of change that an individual will perceive a stimulus to undergo throughout a transformation. To assess this model, we conducted an experiment to determine whether Equation 5 accurately predicts individuals' judgements of the average *rate of change* for two dimensional Gabor-like stimuli. We decided to assess the average rate of change for stimuli rather than the total distance because online pilot participants found instructions based on the speed that images were changing more intuitive than other formulations. Therefore, in what follows we assume additionally that the average rate of change perceived by a subject for a changing stimulus is proportional to its arc length. Furthermore, all stimuli have an identical duration of change (1.0 seconds).

Methods

Participants. This study utilized data collected online during the COVID-19 pandemic. Our sample was recruited from university students (*N*=24) taking an introductory psychology course, and participants received course credit for participation. All procedures used in this study were approved by the appropriate Institutional Review Board.

Stimuli. This study employed continuously changing two dimensional Gabor-like gratings varying in spatial frequency and saturation. All stimuli were presented 20° clockwise from a vertical orientation, while spatial frequency ranged from 0.3 to 11.2 cycles per stimulus (rescaled 0 to 1 for modeling analyses), and saturation values were determined on a 0 to 1 scale. Still images of these stimuli can be seen in Figure 2. Changing stimuli were created by saving static stimuli to a video file at a rate of 60 frames per second, where each frame was created using the GratingStim function in the Python 3 implementation of PsychoPy.

In order to change continuously, stimuli needed to follow a curve in its objective parameter space (the set A from above). For these curves, we chose to use ellipses parallel to the two coordinate axes, parameterized by center values radii as in Equation 4. Examples of these curves are shown in Figure 2.

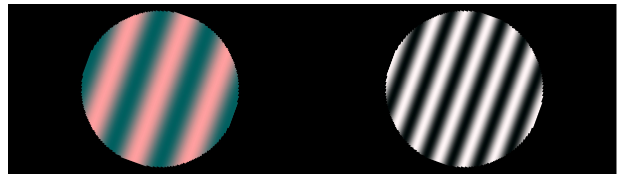

Figure 2. Two example stimuli varying in spatial frequency and saturation.

Two sets of stimuli were constructed: exemplars and generalization items. The ellipses followed by both sets of items were evenly distributed over the four centers seen in Figure 3. The radii of exemplar items were chosen from two extreme values for both dimensions. Taking all possible combinations of a center and two radii, there were sixteen total exemplar items. In contrast, generalization items were created with four intermediate radius values for both frequency and saturation, resulting in 64 stimuli.

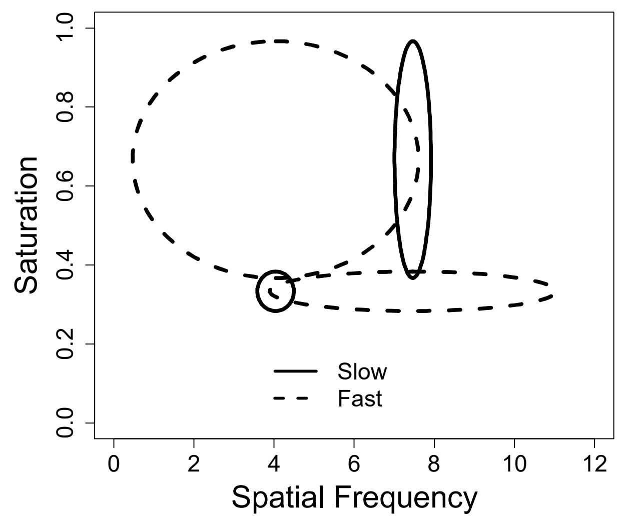

Figure 3. Four ellipses in parameter space corresponding to the paths taken by four changing stimuli. The four pictured items are each exemplar items and show one example from each center location and one example with each combination of radii. Note that "Slow" and "Fast" examples are distinguished by change in spatial frequency.

Procedure. This experiment was conducted online via the experimental platform Gorilla. Participants were first welcomed to the experiment and then introduced to static versions of the stimuli via text and images on their screen. They were informed that the stimuli could vary in two ways: "color" and "frequency" and were shown examples of static stimuli with varying saturation and frequency respectively. Then participants were introduced to changing stimuli and instructed that some were changing "Fast" while others were changing "Slow". Participants were informed that they should attend to the frequency dimension and judge whether a stimulus was changing fast or slow.

After the instructions, participants were given a multi-phase categorization task. During both phases, they were instructed on every trial to answer the question "Is this one changing fast or slow?" and presented with two response options labeled "Fast" and "Slow". A video stimulus was placed in the center of the screen, and participants could play the video as many times as desired by clicking on it.

During Phase 1, participants were shown each exemplar stimulus once and given feedback regarding their performance, determined only by the radius of the frequency dimension. Saturation was irrelevant for feedback, consistent with the instruction to pay attention to frequency. After eight and sixteen trials, participants were reminded via a text prompt to pay attention to changes in frequency.

In Phase 2, stimuli were randomly assigned into fifteen 8-trial blocks. Each block consisted of four exemplar stimuli chosen at random (with replacement) and 4 generalization stimuli chosen at random (without replacement). Feedback was provided as in Phase 1 for the exemplar stimuli. Additionally, between every block, participants were instructed to pay attention to changes in frequency.

Modeling. Responses from each participant in Phase 2 were predicted as a logistic regression of arc length given attention and power parameters. A response y was modeled as a Bernoulli variable using the relation

$$y \sim Bernoulli(logit^{-1}(ms + b))$$

Here, m and b are slope and intercept parameters fit for each subject individually, while s is the arc length calculated according to Equation 5. To estimate the integral in Equation 5 under different parameter values, we used the trapezoid rule, breaking the interval into 24 regions and taking the sum of the area of a trapezoid in each region. Finally, since $a_f + a_s = 1$, a single attention parameter a_f was used.

In order to estimate the values of latent parameters, a Bayesian modeling approach was implemented using RSTAN, the R implementation of the STAN modeling framework (STAN Development Team, 2017). For each participant, we gave latent parameters the following priors:

$$a_f \sim \text{Beta}(1, 1)$$
$$c_f, c_s \sim \text{Gamma}(1, 1)$$
$$m, b \sim \text{Normal}(0, 10)$$

Here, the Gamma distribution is parameterized by shape and rate, while the Normal distribution is parameterized with mean and standard deviation. Once the model was specified, STAN used Hamiltonian Monte Carlo sampling to estimate the joint posterior distribution over parameters.

Results

Behavioral. In aggregate, participants were able to learn the two categories ("Fast" and "Slow") by the end of the two training blocks, showing a mean accuracy of 0.81 over the final eight trials (Bayes Factor $\text{K} = 1.62 \times 10^7$ against a model with chance level responding; Rouder et al., 2009; Morey & Rouder, 2018). Accuracy continued to improve for exemplar items into Phase 2, where participants achieved a mean accuracy of 0.88 ($K = 2.61 \times 10^{17}$). This provides evidence that participants in aggregate learned the categories.

Additionally, we examined responses to generalization items based on their two radii and center parameters. If nonlinear psychophysical laws affect the perception of subjective change, we should expect both the radius of a curve as well as its center to predict subjective distance. However, if participants experience stimuli veridically, we would expect no effect of the ellipse's center.

To test this hypothesis, we converted responses to generalization items into a binary variable, where 1 stands for "Fast" responses. We then used a linear model with both radii and centers as predictors to see which aspects of a changing stimulus predicted participant responses. We found that all four predictors increased the Bayes Factor of the linear model, with strong evidence that the two radii and the center of the frequency dimension predict generalization responses (increase in $K > 10^7$ against a model with only an intercept), and only marginal evidence for the center of the saturation dimension (increase in $K = 1.81$). This pattern of results supports the need for a coordinate transformation such as Ψ above, as with veridical representations, there would be no reason for the center of the frequency dimension to predict the speed of change. Additionally, stronger evidence for psychophysical effects on the frequency dimension may be explained by the instructions and the category structure encouraging participants to attend to frequency.

Modeling. Our first modeling analysis aimed to determine whether arc length can predict subjective speed judgements. Fitting our model with RSTAN, we obtained samples from the joint posterior distribution of all parameters for each participant. A representation of these posterior distributions can be seen in Figure 4. From these posterior distributions, we found that the slope parameter (m) was positive with 95% or greater posterior confidence for 22/24 participants. This suggests that arc length, calculated as in Equation 5, is a reliable predictor of subjective speed of change judgements at the individual level.

Additionally, we examined whether the model could determine that participants were instructed to attend to the frequency dimension of stimuli (indicated by $a_f > 0.5$). We found that there was 95% or greater confidence that $a_f > 0.5$ for 18/24 participants, with no participants showing the opposite pattern (6 participants showed ambiguous evidence). This suggests that, for many participants, the model is able to capture the top-down attentional bias specified in our instructions. However, it is unclear if the model is unable to identify this parameter for the remaining six participants or if this reflects a meaningful individual difference, such as a difference in attentional control.

We also analyzed the two power parameters to examine the effect of psychophysics on participant responses. For the frequency dimension (c_f), there was 95% or greater confidence that nine participants had a power parameter less than one, corresponding to a negative concavity in their psychophysical function. For these participants, we have evidence that a nontrivial psychophysical transformation (Ψ) is important to predicting their judgements of changing stimuli. The remaining participants showed ambiguous evidence; however maximum a posteriori point estimates for 22/24 participants showed a power parameter less than 1. Together, these findings suggest that psychophysical parameters are important for some participants, though more trials may be needed to identify these parameters.

Finally, the power parameter for the saturation dimension (c_s) was more uncertain relative to its expected scale than other parameters. MAP estimates for this parameter found that 13/24 participants showed a value less than 1 (mean = 0.70), with 11 participants greater than one. This likely reflects the fact that participants attended to the frequency dimension. As a_f approaches a value of one, stimuli become less discriminable on the saturation dimension, and in the limit ($a_f = 1$) all values of c_s result in the same model predictions.

Discussion

In this paper, we presented a new method to assign psychological spaces a meaningful notion of distance as the arc length traversed by a stimulus through psychological space. We then derived an expression for arc length for 2-D stimuli (Equation 5) where each dimension is orthogonal and modeled psychophysically via a power law. We tested this prediction with an experiment where participants classified changing stimuli based on their rate of change. We found that arc length is a powerful predictor of the magnitude of change and found evidence that the attention parameter of the model accurately captures the dimension to which participants selectively attended. Finally, we found indirect evidence that a psychophysical model is needed to predict responses as well as explicit evidence that some participants deviate from a veridical representation of stimuli. Together, these results provide powerful preliminary evidence that arc length can be used as a geometric constraint for psychological spaces.

As discussed above, any model which can measure the arc length of continuously changing stimuli can also be used to induce a notion of geodesic distance, where the geodesic distance between two stimuli is the minimum length of all curves between the two points. Defined in this way, the geodesic distance is a kind of "transformational distance", measuring the magnitude of the shortest transformation that would convert one stimulus into another by continuously changing its attributes (Hahn, Chater, & Richardson, 2003). In the future we hope to explore commonalities and differences between geodesic distance and currently used distance functions in cognitive models, such as in the GCM.

Additionally, the results of our approach can be compared to other methods attempting to estimate the geometric properties of psychological spaces. For example, Hénaff, Goris, & Simoncelli (2019) provide a method to estimate the curvature of the curve corresponding to a continuous sequence of stimuli using a discrimination paradigm. By combining multiple geometric constraints such as arc length and curvature, future research can determine if these measures are in conflict and work toward a unified geometric framework for perception, attention, and representation.

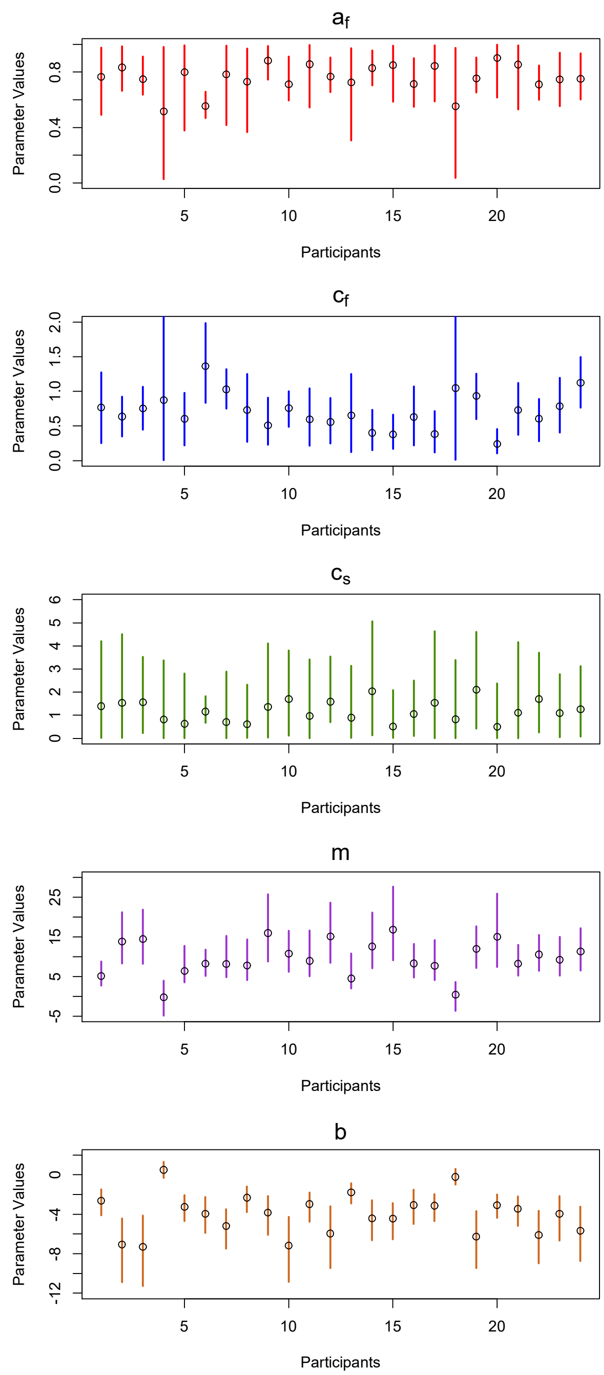

Figure 4. Mean posterior estimates with 95% credible intervals for each parameter. Each plot corresponds to an estimated parameter, with one posterior estimate for each participant. Note that two intervals for c_f are cut off because they extend far beyond the upper bound of the plot.

Acknowledgments

This work was supported in part by National Institutes of Health Grants R01HD078545 and P01HD080679 to Vladimir Sloutsky. The authors would also like to thank Zach Paitsel for helpful discussion and comments on an earlier draft.

References

Callahan-Flintoft, C., Holcombe, A. O., & Wyble, B. (2020). A delay in sampling information from temporally autocorrelated visual stimuli. *Nature communications, 11*, 1852.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research, 2011,* 1484-1525.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology, 92*, 53-78.

Gärdenfors, P. (2017). *The geometry of meaning: Semantics based on conceptual spaces.* The MIT Press.

Gescheider, G. A. (1997). *Psychophysics: The fundamentals.* Lawrence Erlbaum Associates, Inc.

Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. *Cognition, 87*, 1-32.

Hénaff, O. J., Goris, R. L. T., & Simoncelli, E. P. (2019). Perceptual straightening of natural videos. *Nature Neuroscience, 22*, 984-991.

Ho, T. C., Brown, S., Abuyo, N. A., Ku, E. j., & Serences, J. T. (2012). Perceptual consequences of feature-based attentional enhancement and suppression. *Journal of Vision, 12*, 1-17.

Kreyszig, E. (1959). *Differential geometry.* Dover Publications, Inc.

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research, 49*, 1194-1204.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111,* 309-332.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225-237

Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes factors for common designs. *R package version 0.9.12-4.2.* https://CRAN.R-project.org/package=BayesFactor

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology; General, 115*, 39-57.

Stan Development Team. (2017). Stan modeling language users guide and reference manual, version 2.17.0. Retrieved from http://mc-stan.org/.

Stevens, S. S. (1957). On the psychophysical law. *The Psychological Review, 64*, 153-181.

Turner, M. R. (1986). Texture discrimination by Gabor functions. *Biological Cybernetics, 55*, 71-82.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327-352.