Factors in Perceptual Shape Completion

Tenzin Chosang

A THESIS SUBMITTED TO THE FACULTY OF GRADUATE STUDIES IN PARTIAL

FULFILLMENT OF THE REQUIREMENTS FOR

THE DEGREE OF MASTER OF ARTS

GRADUATE PROGRAM IN PSYCHOLOGY YORK UNIVERSITY

TORONTO, ONTARIO

June 2024

ABSTRACT

While we live in a 3D world, each 3D object in our visual field projects to the retina as a 2D region bounded by a 1D closed contour. Humans are known to rely profoundly on these bounding contours to segment the scene and to detect and recognize objects. This task is complicated by interpositions of objects in the visual field that lead to occlusions: partial blocking of one object by another. The impact of occlusions on human perception is mitigated by our ability to perceptually complete partially occluded bounding contours, i.e., to fill-in missing contour fragments.

Prior research has evaluated both local and global methods for shape completion, in terms of their objective accuracy against ground truth. However, these methods have not been evaluated against human perceptual completion. In this thesis, observers viewed a series of partial bounding contours where an interval of each contour has been erased, simulating occlusion. A virtual line passing through the missing interval is identified and observers are asked to adjust the location of a dot along this line until it appears to lie where the contour would be, were it visible. The gap length of the missing interval varied from 10 to 50% of the total shape. After analysis of the objective error and bias, it was found that humans employ more than local cues in the process of shape completion, but it is unclear if this directly suggests the use of global cues.

ACKNOWLEDGEMENT

First and foremost, I must acknowledge my supervisor and teacher, Dr. James H. Elder. If it is your logical and lightening sharp mind that makes you stand out, it must certainly be your kindness, and patience that makes you truly remembered. Patience is one of the greatest treasures, and it is a treasure not easily acquired by many. And the amount you have shown me during my masters is otherworldly. Even with my ridiculously slow progress these past two years of the pandemic, you have always been there for me when I really got stuck on a problem. And whenever I get nervous and overthink about some aspect, you always put my mind at ease. You have a very calming presence and a calm manner of speech, which makes you so cool and professional. I appreciate all the advice and the space and freedom you give me. This entire thesis would not be possible without you.

I would also like to acknowledge Professors Richard Murray and Laurence Harris for writing me reference letters for graduate scholarships -without which my research would not be possible. Thank you so much for all your kindness and guiding me during my undergraduate research years. I appreciate you both. I also want to thank Professor Joe Desouza, for giving me a lot of research opportunities during undergrad, thank you! Next, I would like to acknowledge the psychology graduate office members Lori, Freda and Barbara, and the graduate program director Dr. Suzanne Macdonald. Your kindness will be remembered and thanks for always responding quickly to my emails- even on weekends! I would also like to thank Dr. Kevin Lande for his questions and emailing me with papers to look at- thanks. I would also like to thank my fellow CVR trainees, lab members-especially Shaiyan for his helpful clarifications- and friends, you make the environment more welcoming and fun. I also must thank the individuals who participated in my experiment.

Next, I must acknowledge some other teachers, both near and far, both in the past and present. It is your purity and integrity that are my source of comfort, strength, and hope.

Lastly, I would like to acknowledge my parents. Without you two, my world would be grayscale.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Shape Completion

Shape completion is a process that allows for the perception of whole objects from visual fragments (Ben-Yosef & Ben-Shahar, 2012). This is accomplished through a 'filling-in' of the shape when parts of the shape are occluded (Singh, 2004). This phenomenon of completion involves two variations, amodal and modal completion (Figure 1). Modal completion (shown as the dashed red line) separates the two elephants, where contrast falls to near zero. Amodal completion (dashed cyan line) occurs when objects are occluded; in this particular example, the amodal contour completes the boundary of the elephant on the right where it is occluded by the tree (Wagemans et al., 2012).

Figure 1. Contour completion. Red dashed line: Modal completion. Cyan dashed line: Amodal completion. Figure taken from Wagemans et al. (2012).

Humans have been known to use geometric properties of orientation and curvature, among others, to perform contour interpolation and extrapolation (Wagemans et al., 2012). Humans perceive the continuation of contours as smooth rather than jagged, even when there is no cue to suggest that percept (Singh & Fulvio, 2005; Geisler & Perry, 2009), thus suggesting that the visual system could potentially be using priors based on regularities in the natural environment to inform completion.

1.2 Local and Global Cues

For shape completion, we define a local cue as a feature of the visible contour defined at a point. This cue could include position, orientation, and curvature. Typically, the values of these local features at the endpoints of the visible contour bounding the missing contour interval are posited to co-determine the completion (Figure 2).

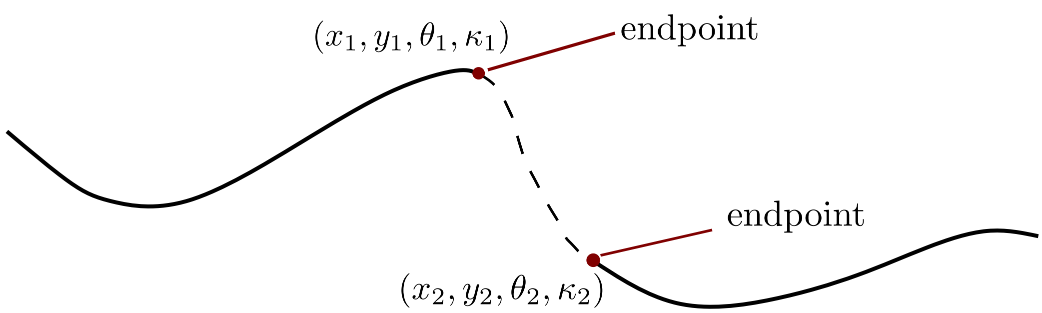

Figure 2. Local cues to contour completion: 2D position (x, y), orientation θ and curvature defined at the endpoints of the visible contour.

In addition to local cues, non-local or even global shape cues could potentially influence shape completion. Global cues involve properties such as closure, convexity, and symmetry amongst others (as cited in Elder, 2018). Here, we will consider two global shape models that could be used for completion: Shapelets (Dubinskiy & Zhu, 2003) and Formlets (Grenander et al. 2007, Oleskiw et al. 2010, Elder et al. 2013).

1.3 Local Methods: Linear Interpolation and Elastica

The linear interpolation model uses only endpoint position information directly linking the two endpoints of the visible part of the shape using a line segment (Figure 3).

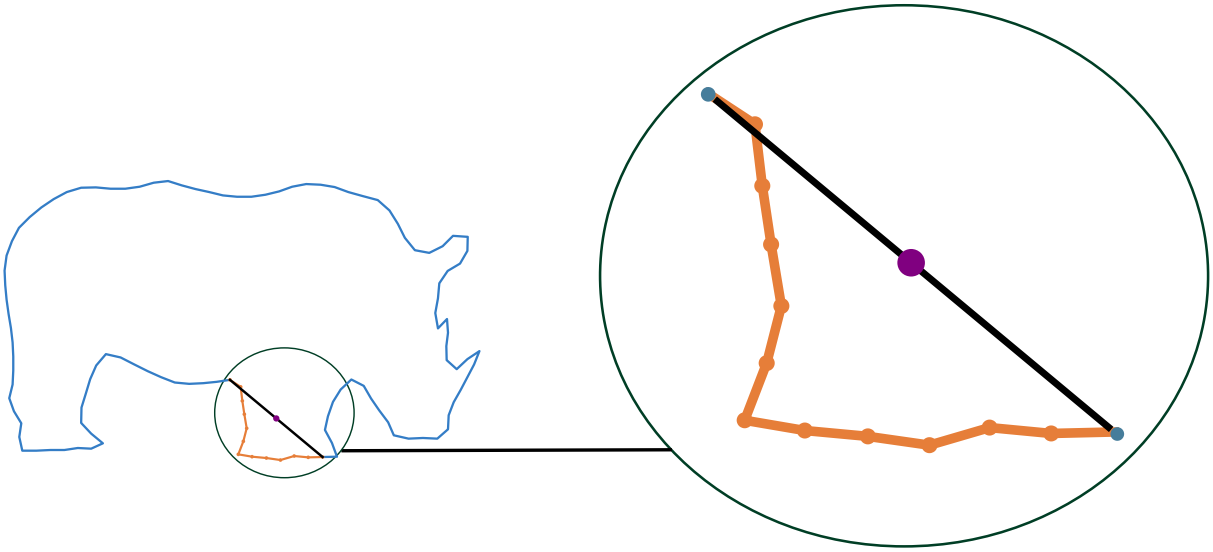

Figure 3. Linear interpolation model for natural shape at 10% gap level. The linear interpolation (black) connects the endpoints of the visible contour (in blue). Purple: midpoint of the linear interpolant.

This linear approach does not generate smooth curves. An alternative local model called elastica does yield smooth curves. The concept of elastica dates back to the works of the mathematician Euler (Mumford,1994, pp.491). Elastica are curves $\gamma(s)$ that minimize a weighted sum S of squared curvature κ^2 and arc length while ensuring that the tangent orientations at the endpoints align with the visible curve:

$$S = \int_{s_1}^{s_2} (\kappa^2(s) + \lambda)\, ds \tag{1}$$

Subject to $\gamma'(s_1) = \overrightarrow{T_1}$ and $\gamma'(s_2) = \overrightarrow{T_2}$, where s is the unit arc length parameter and $\overrightarrow{T_1}$ and $\overrightarrow{T_2}$ are the tangent vectors of the visible curve at the two endpoints. λ is a free parameter that controls the relative importance of curvature and arc length in the energy functional

(Yakubovich, 2014). If the curve $\gamma(s)$ is approximated discretely as a polygon, the objective function can be minimized by gradient descent. I optimized the lambda parameter by searching through a range of values to find the value that minimizes the average Euclidean distance between the ground truth and the model estimate of an interpolated point.

1.4 Global Model: Shapelets

Shapelets are contour-based shape representations that represent shapes by a linear combination of basis functions (Dubinskiy & Zhu, 2013; Figure 4). The basis functions are localized and capture information at different scales. The shapelet basis consists of Gabor-like planar curves that map arclength (t) to the image and can be represented in the complex plane $\psi:[0, 1] \rightarrow C$. Each shapelet is controlled by two parameters: the arc length position parameter (μ), which determines the position of the shapelet along the curve, and the scale parameter (σ), which affects the spatial extent of the shapelet. A shapelet $\gamma(t;\sigma,\mu)$, has the specific form:

$$\gamma(t;\sigma,\mu) = \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right)\left(\cos\big(2\pi\sigma(t-\mu)\big) + i\sin\big(2\pi\sigma(t-\mu)\big)\right) \qquad (2)$$

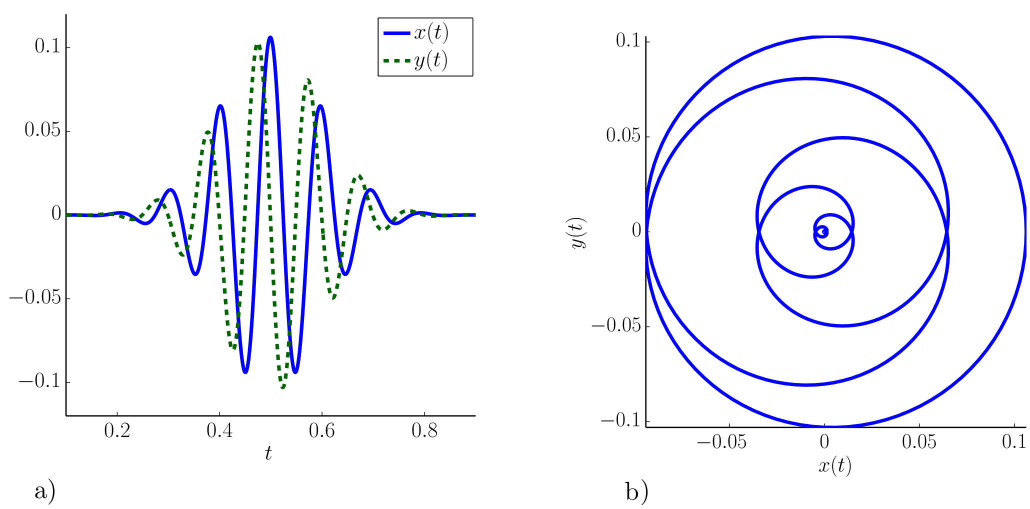

Figure 4. Example Shapelet. a) Coordinate functions and b) Trace. Figure taken from Elder et al., (2013).

The basis functions are subject to an affine transformation effected by a 2 x 2 matrix of basis coefficients A_k:

$$A_k = \begin{pmatrix} a_k & b_k \\ c_k & d_k \end{pmatrix} \tag{3}$$

This affine transformation is applied to each shapelet in image space prior to linear combination, and captures different actions such as anisotropic scaling, rotation and shearing (Figure 5). An over-complete dictionary of basis functions is formed by grid-sampling over position μ, and extent σ (Dubinskiy & Zhu, 2013).

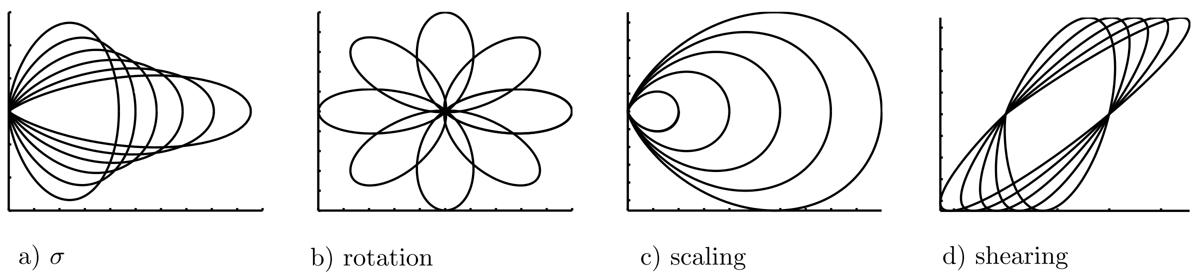

Figure 5. Instances of bases at different scales subject to transformations. Each shape base is a lobe-shaped curve. Figure taken from Dubinskiy & Zhu (2003).

To learn the representation of a shape, Dubinskiy and Zhu (2013) apply a matching pursuit algorithm, a greedy algorithm that at each iteration selects the shapelet that minimizes reconstruction error, measured by the L^2 norm of the residual:

$$E(\Gamma^{obs}, \Gamma^K) = |\Gamma^{obs} - \Gamma^K|_2^2 \tag{5}$$

Where Γ^{obs} is the observed or target shape and Γ^K is the approximation after adding K shapelets. Unfortunately, because this shapelet model does not explicitly capture regional properties of shapes, it tends to yield non-simple i.e., self-intersecting curves when used as a generative model (Elder et al., 2013; Figure 6).

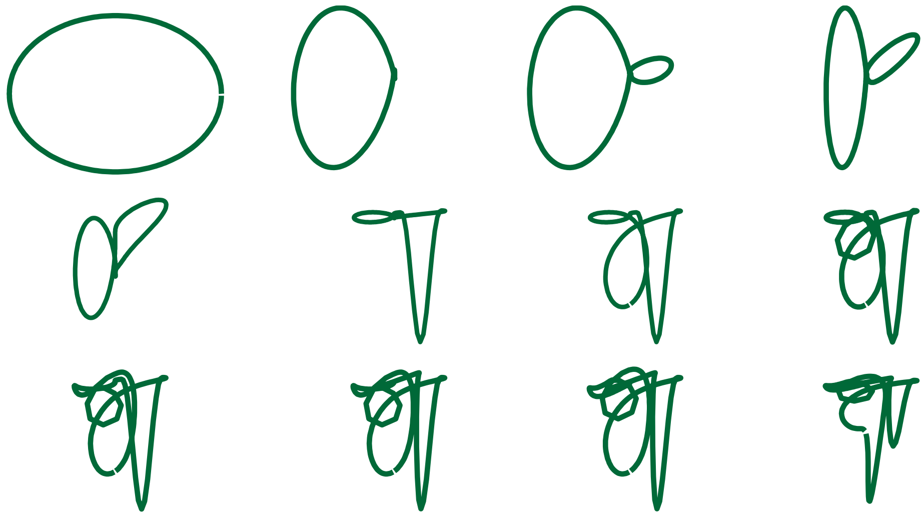

Figure 6. Sampling from the shapelet model produces non-simple, i.e., self-intersecting curves. Figure taken from Elder et al., (2013).

To apply the shapelet model to contour completion, the shapelet representation is fit only to the visible portion of the curve. Since each shapelet is a complete closed curve, the resulting representation will also be a complete closed curve, interpolating the missing interval of the curve it is fit to.

1.5 Global Model: Formlets

Formlets build on the Growth by Random Iterated Diffeomorphisms (GRID) model, which models growth as a sequence of local and radial deformations (diffeomorphisms) (Grenander et al., 2007), and has been used to track tumour growth in medical images. Elder et al. (2013) called these diffeomorphisms Formlets. Formlets are localized in both scale and space. They are defined using a Gabor-like deformation function f centered at the point ζ in the complex plane and deforming space within a standard deviation σ about ζ:

$$f(z;\zeta,\sigma,\alpha) = \zeta + \frac{z-\zeta}{|z-\zeta|}\rho(|z-\zeta|;\sigma,\alpha),$$

$$\rho(r;\sigma,\alpha) = r + \alpha \sin\left(\frac{2\pi r}{\sigma}\right)\exp\left(\frac{-r^2}{\sigma^2}\right) \tag{4}$$

A Formlet is characterized by its location ζ , scale σ, and gain α, which represents the magnitude of the deformation. To preserve the topology, the authors place constraints on this gain parameter α. A dictionary of formlets is constructed by grid sampling over position ζ and scale σ (Figure 7). Figure 8 shows the effect of scale and gain parameters in terms of the deformation applied to a grid.

To represent a shape, one begins with a simple embryonic shape generated by sampling a unit circle and then applying an affine transformation to generate an elliptical shape Γ^0 that minimizes the L^2 error between the target shape Γ^{obs} and the ellipse. To then learn the formlet representation $\{f_1 ... f_k\}$ of the shape Γ^{obs}, matching pursuit is employed to iteratively find the formlet minimizing the reconstruction error when applied to the current model Γ^{k-1}. The objective function for formlets is the same as for shapelets. The gain value α is optimized analytically, subject to the diffeomorphism constraints.

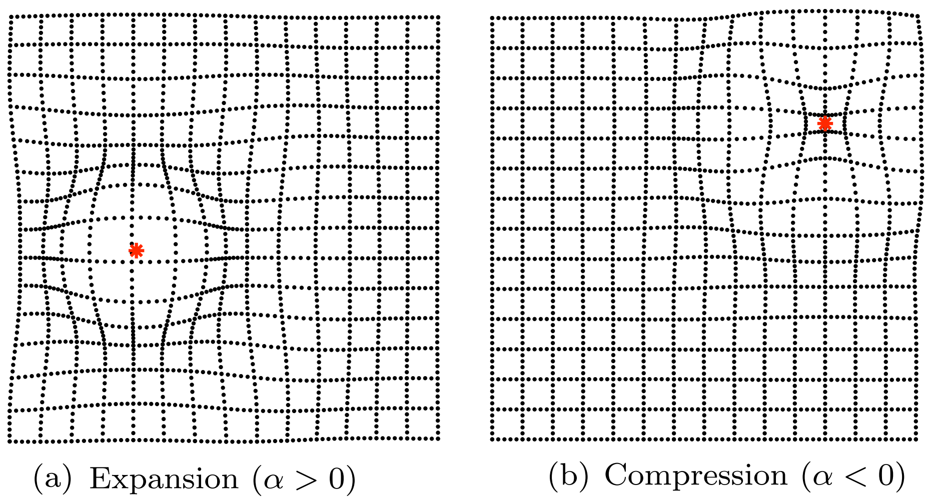

Figure 7. Example formlet deformation. Positive and negative gains α produce expansion and contraction, respectively (Elder et al., 2013). The location ζ of the formlet is indicated by the asterisk.

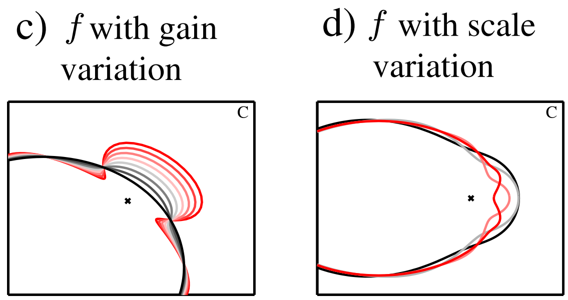

Figure 8. Figures (a) and (b) display several formlets applied to an ellipse. Figure taken from Elder et al., (2013).

As for shapelets, we can apply formlets to shape completion by fitting them only to the visible portion of the contour. Since it is based on a complete closed embryonic shape, the resulting representation will interpolate the missing interval of the contour.

1.6 Accuracy of Formlets and Shapelets for Contour Completion

While neither the local nor the global models reviewed above have been assessed as models of human shape completion for naturalistic shapes, Elder et al. (2013) did evaluate the accuracy of Shapelet and Formlet models relative to ground truth. In particular, they assessed the accuracy of completion for 10% and 30% intervals of animal shapes and found that Formlets were more accurate (Figure 9). Formlets also preserved the topology of the shape, as can be seen in Figure 10, while Shapelets failed to preserve the topology of the shape. Both Formlets and Shapelets serve as shape models for which completions are determined by non-local factors. The use of these models is to determine if they would make more accurate predictions for natural shapes, which have global regularities that the models might capture. As well as to determine whether these models would therefore predict human judgements.

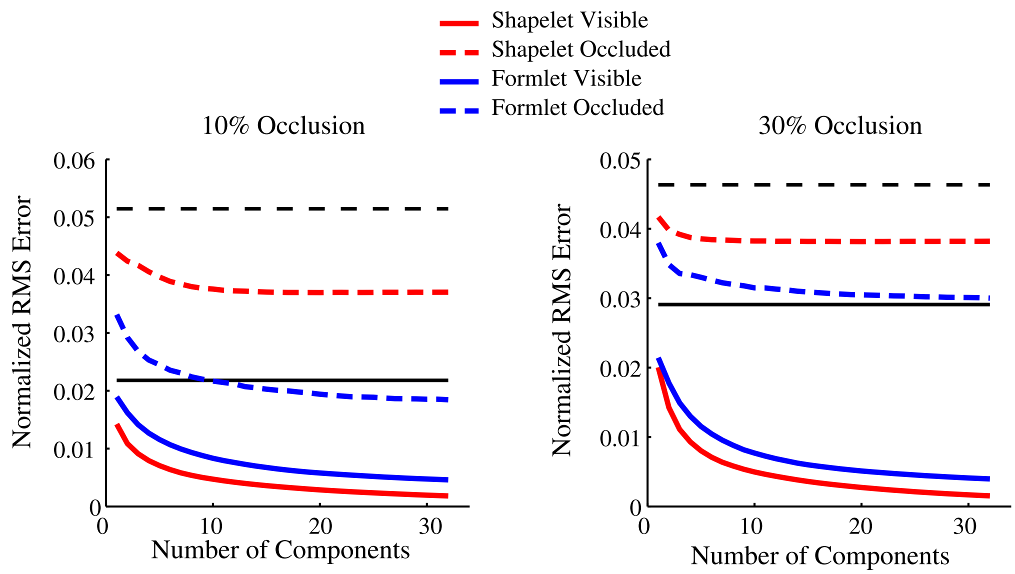

Figure 9. Error defined using $\boldsymbol{L^2}$ Hausdorff distance of occlusion pursuit evaluation of Shapelets (red) and Formlets (blue). Figure taken from Elder et al., (2013). Solid and dashed horizontal black lines indicate the errors of the initial affine fit ellipse for visible and occluded intervals, respectively.

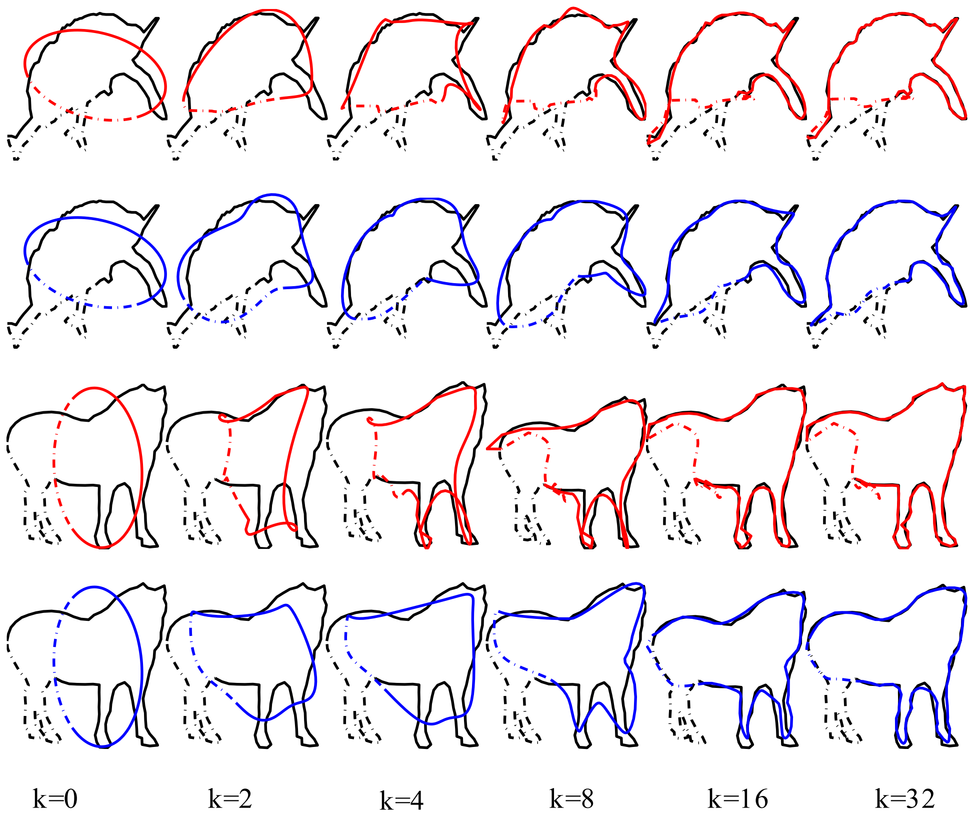

Figure 10. Examples of pursuit on 30% occluded shapes. Formlets (blue) and Shapelets (red) in increasing iterations from left to right (k = 0,2,4,8,16,32). Solid lines indicate the visible contour and dashed lines indicate the occluded portion of the contour. Figure taken from Elder et al. (2013).

1.7 Aims of Present Study

The impact of occlusion on human perception is mitigated by the human ability to perceptually complete partially occluded bounding contours, i.e., to fill in missing shape information. This study aims to investigate the degree to which humans use local and non-local cues to perform this perceptual completion task. The study analyzes the completion ability of the four models reviewed above and compares them to human judgements in terms of accuracy and bias. In addition to natural shapes, we will assess human and model shape completion on synthetic shapes that match the curvature statistics of the natural shapes but are otherwise maximum entropy and are thus lacking the global regularities of the natural shapes (Elder et al., 2018). This will allow us to test whether these global regularities play a role in human or model shape completion.

CHAPTER 2

Methods

2.1 Human Psychophysics

2.1.1 Participants

All 12 participants were recruited from York University. Participants gave their informed consent to participate in the study. All participants were naïve to the purpose of the experiment. All experiments were approved by the York University Human Participants Review Committee (HPRC) prior to data collection. The participants were all graduate student volunteers; no monetary incentive was given to participate in the experiment.

2.1.2 Apparatus

The experiment was conducted in a laboratory setting, using a PC desktop computer with the Windows 10 operating system. The stimuli were displayed on a monitor with a resolution of 2560 x 1440 pixels, at a viewing distance of 60 cm, maintained with a chin rest. Experimental stimuli were drawn and displayed using python and PsychoPy3 (Peirce, 2007). All shape stimuli were scaled so that the maximum Euclidean distance of a point on the contour from its centroid was 4 degrees. A keyboard and mouse were used to capture the behavioral responses of participants. All statistical analyses were conducted using Python (version 3.7).

2.1.3 Stimuli

We employed 300 natural and 300 synthetic shapes, sampled at 120 points, roughly uniformly spaced. The statistics of the turning angles (discrete curvatures) of the synthetic shapes matched the turning angle statistics of the natural shapes. Otherwise, these synthetic shapes were maximum entropy and lacked the global regularities of the natural shapes (Elder et al., 2018). To avoid biasing results, we did not interpolate between the points to form a continuous contour, but rather displayed the points as dots on the screen. Each dot subtended 1.2 arcmin. Figure 11 shows some sample natural and synthetic shape stimuli. To study shape completion, we removed a random, contiguous sequence of 11 (10 %), 29 (25%) or 59 (50%) dots (Figure 12).

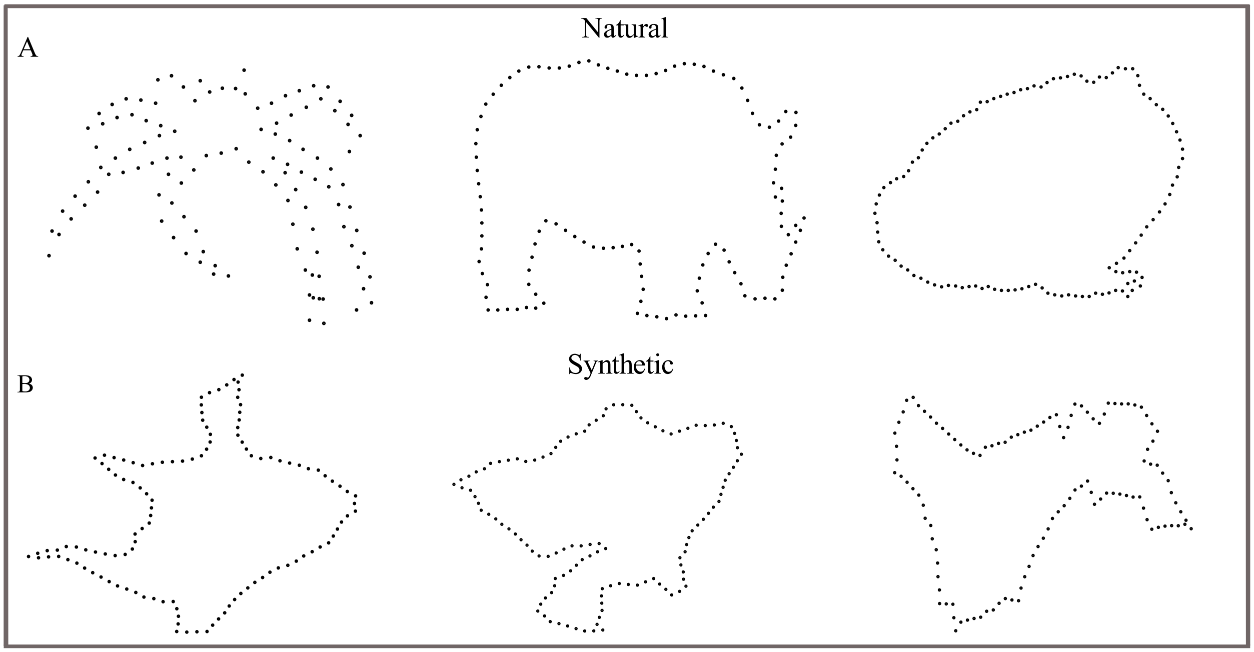

Figure 11. Sample natural (A) and synthetic (B) shape stimuli.

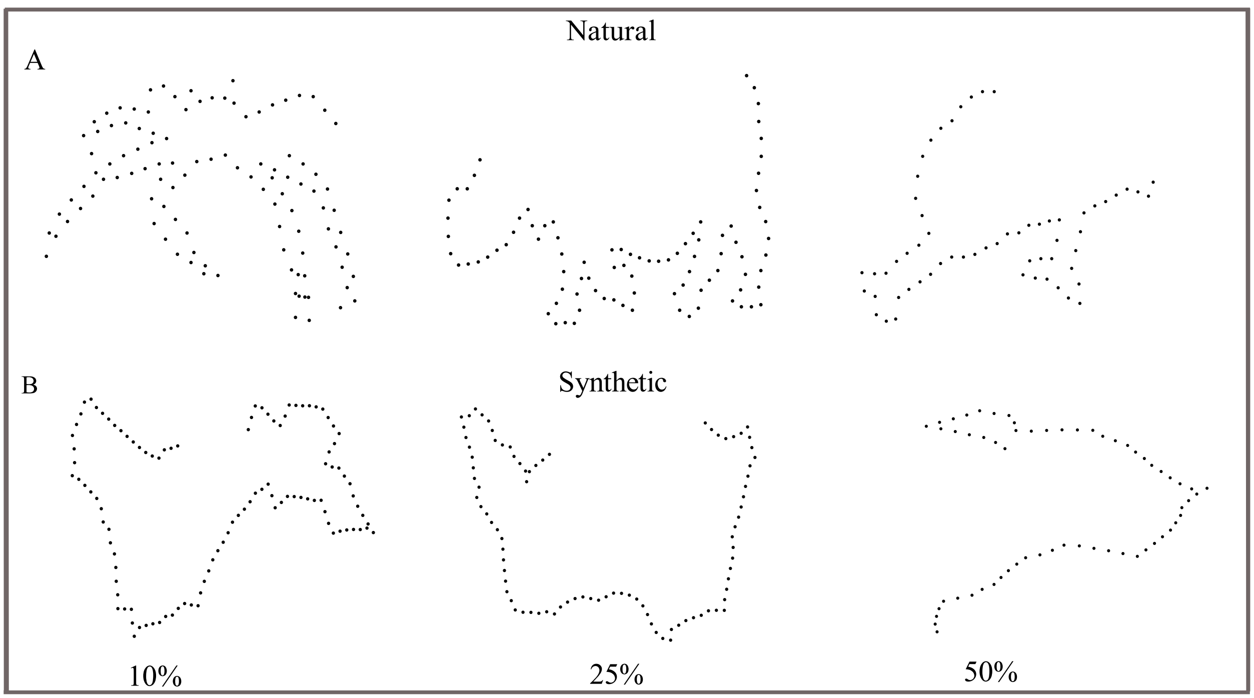

Figure 12. Sample natural (A) and synthetic (B) shapes with occlusion intervals employed in this study.

2.1.4 Defining the Completion Task

Linear Interpolant

For any given shape, first an interval of either 10 %, 25% or 50% of the shape was removed (see Figure 12). Next, the midpoint of the linear interpolant connecting the endpoints of the visible part of the shape was determined (Figure 13).

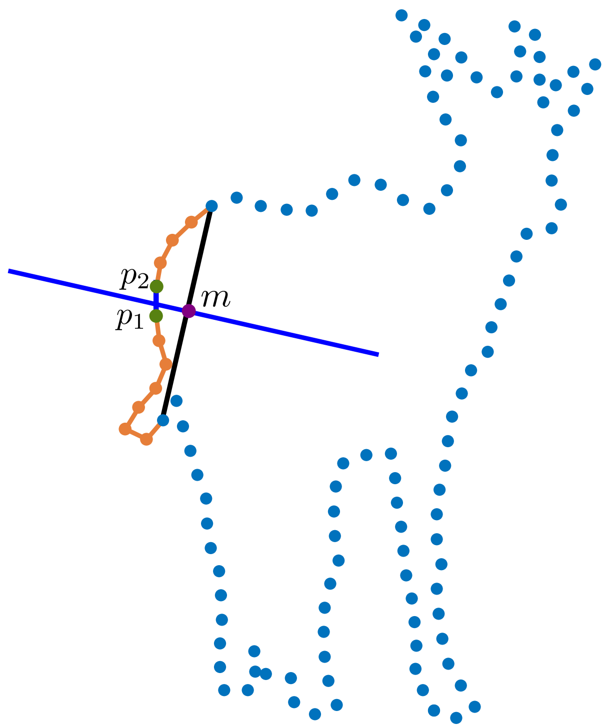

Figure 11. Blue dots: visible contour. Orange dots: missing interval. Linear interpolant (black line), with the midpoint (m) in purple. Probe line (blue line) is orthogonal to the linear interpolant and passes through the midpoint (m). $\overline{\boldsymbol{p_1 p_2}}$ is the line segment of the missing interval that intersects with the probe line (blue). The completion contour (in orange) connects the missing dots.

Probe Line

The participant will adjust the position of a dot on a virtual probe line until it appears to align with the missing contour. The probe line is orthogonal to the linear interpolant and passes through its midpoint. In homogenous coordinates, it can be represented by the 3-vector $l_o = (a_o, b_o, c_o)^T$ where $a_o = n_y$, $b_o = -n_x$ and $c_o = n_x y_m - x_m n_y$, and $(\mathrm{x_m}, \mathrm{y_m})$ and $(\mathrm{n_x}, \mathrm{n_y})$ are the midpoint of and normal to the linear interpolant, respectively.

Ground Truth

To obtain the intersection of the probe line with the completion line (Figure 14), I consider in turn each line segment that has been removed. Consider the i^{th} segment with endpoints $p_1 = (x_1^i, y_1^i)$ and $\mathrm{p}_2 = (x_2^i, y_2^i)$. For example, in Figure 13, $\overline{p_1 p_2}$ represents the 8th segment out of the 12 segments that make up the completion line. This line $\overline{p_1 p_2}$ can be represented by the 3-vector $l_i = (a_i, b_i, c_i)$ such that every point (x, y) on the line satisfies $a_i x + b_i y + c_i = 0$, where $a_i = y_2^i - y_1^i$, $b_i = x_1^i - x_2^i$ and $c_i = x_2^i y_1^i - x_1^i y_2^i$. The point of intersection $\tilde{p} = (x^*, y^*, \lambda^*)$ of this line with the probe line l_o is given by the cross product $\tilde{p} = l_i \times l_o$.

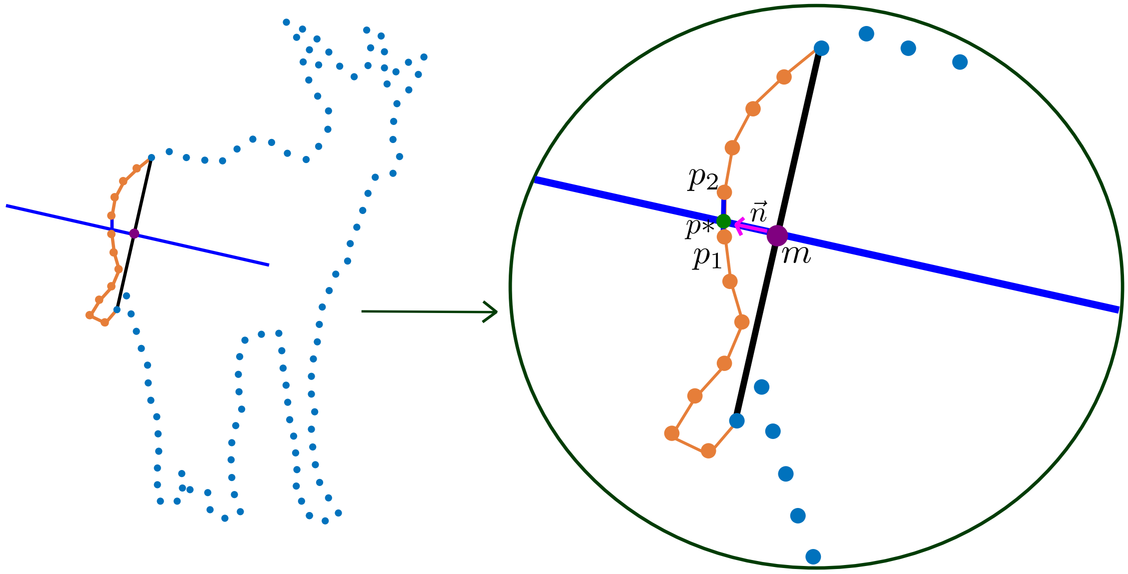

Figure 12. Ground truth calculation. The intersection of the completion contour (orange line) with the blue probe line is denoted by $\boldsymbol{p}^*$. Black: Linear interpolant. Pink arrow: unit normal $\vec{\boldsymbol{n}}$ to linear interpolant.

The Euclidean representation $p^* = (x^*, y^*)$ of this point of intersection given by $p^* = (x^*/\lambda^*,\ y^*/\lambda^*)$. For p^* to be between the endpoints of the i^{th} segment, two conditions must be satisfied.

$$(1)\ |p^* - p_1| < |p_2 - p_1|$$

$$(2)\ (p^* - p_1)^T(p_2 - p_1) > 0$$

At least one of the deleted segments must satisfy these conditions. If more than one segment satisfies the conditions, the task is ambiguous, so we discard this construction.

2.1.5 Range of Probe Line

The range of the probe line was calculated by taking the maximum deviation (Euclidean distance) of the ground truth intersection p^* of the completion line from the linear interpolant over all stimuli. This was done separately for each gap level and shape class. I then set the range of the probe line to be 30% beyond this value. I initialize the probe dot randomly over this range. Figure 15 displays the range of the probe line.

2.1.6 Procedure

Participants sat with their chin resting on a chin rest to ensure a viewing distance of 60 cm. They were informed about the procedure for the experiment and went through five practice trials before beginning the experiment. No feedback was given during the practice trials or the actual experiment, and the participants had unlimited time to complete each trial. After the practice trials, participants began the real experiment, which consisted of 100 trials per condition. There was a total of six conditions (natural/synthetic shapes x 3 levels of occlusion), resulting in 600 trials per human participant, taking about one hour to complete. The order of conditions was randomized. Figure 15 displays an example of the stimuli and task.

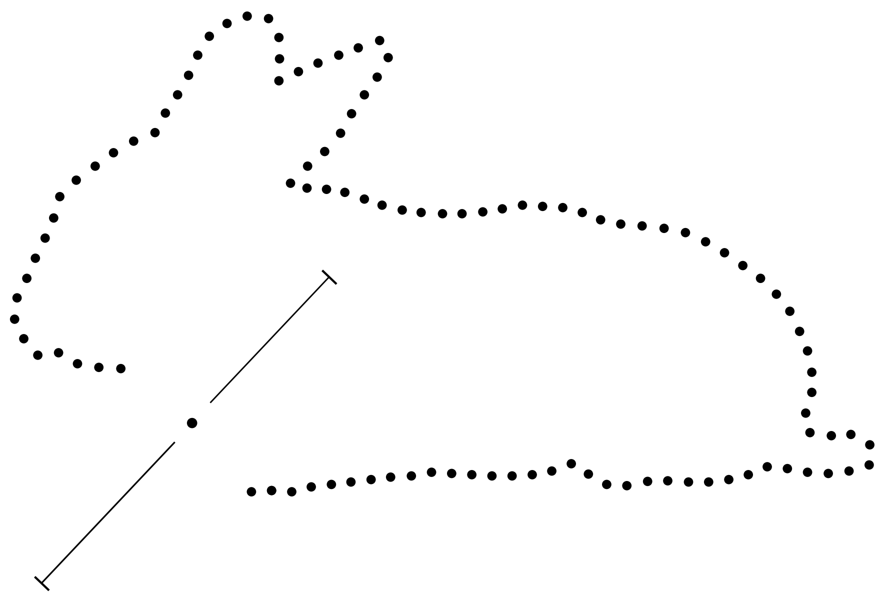

Figure 13. Experimental Paradigm. Black dots: visible shape. The participants move the black probe dot using arrow keys to where they think it would lie on the missing contour interval. The black bar represents the range of the probe dot for this shape.

2.1.7 Error Measurements

Performance was measured by comparing the probe location selected by the participant with the ground truth location (the intersection of the probe line with the completion contour). Error is considered to be negative if the probe setting is displaced inward relative to the ground truth and positive if displaced outwards (see Figure 16).

The results were analyzed using a linear mixed model where the predictors (gap level, shape class and system) were a mix of categorical and continuous variables. P-values were adjusted using the Bonferroni adjustment for multiple comparisons, which controls the overall Type I error rate.

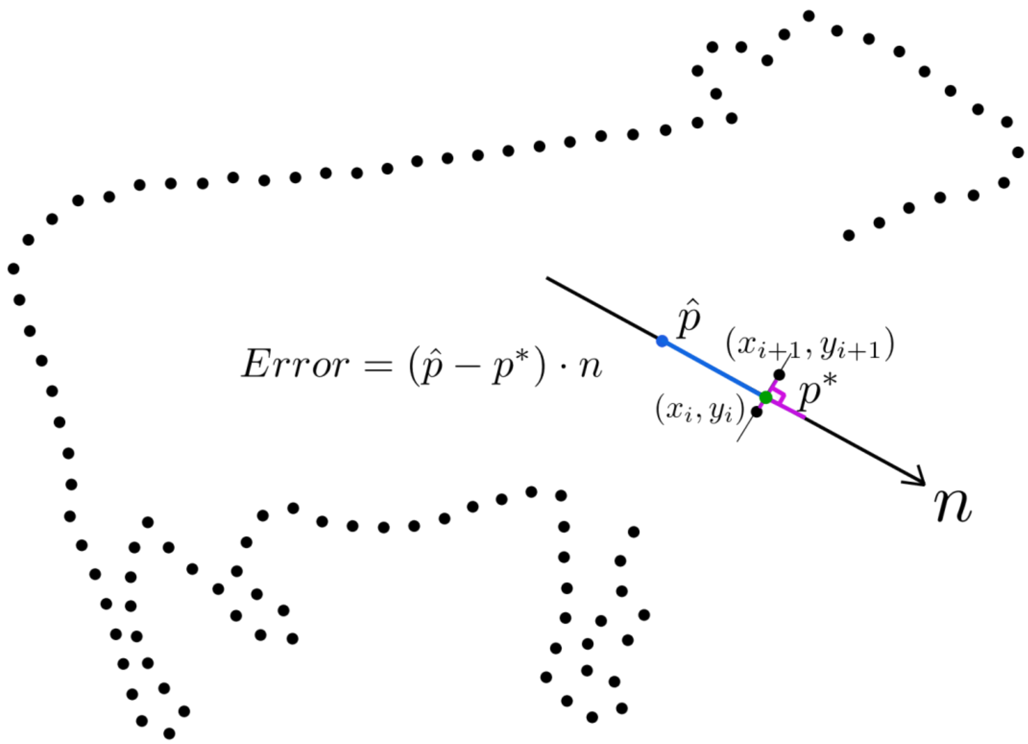

Figure 14. Error Calculation. We define the 11 displacements of the completion estimate $\hat{p}$ from the ground truth p^* as $(\hat{p}$ - $p^*) \cdot\ n$, where n is the outward normal to the linear interpolant.

2.2 Model Predictions

The human completion judgements can be compared to model predictions, defined as the intersection of their completions with the probe line (Figure 14).

2.2.1 Linear Interpolation

Figure 17 shows examples of completions predicted by linear interpolation. Note that errors can be large.

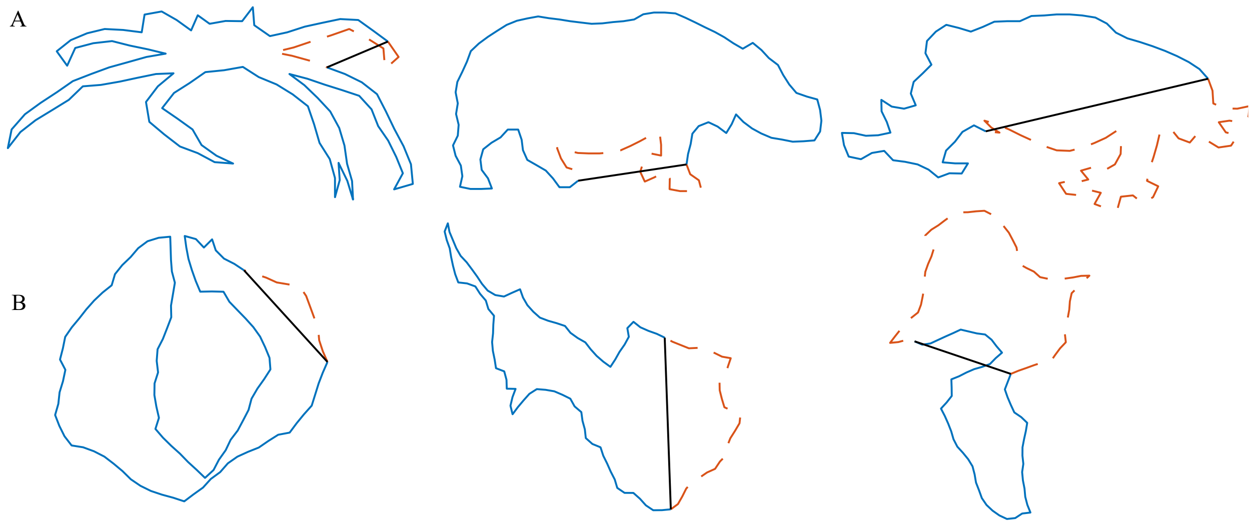

Figure 15. Linear completions for (A) Natural and (B) Synthetic Shapes. Blue: visible shape. Dashed: occluded contour. Black: Linear completion.

2.2.2 Elastica

The Elastica model has a free parameter λ determining the balance between curvature and length. To optimize this parameter, we employed as training stimuli 90 natural and 90 synthetic shapes not used in the experiment, and randomly selected 10 occlusion intervals for each shape and level of occlusion. We then swept over a range of λ values to determine the λ value that minimized the mean absolute error of the prediction relative to the ground truth for both the natural shapes (Figure 18), as well as the synthetic shapes (Figure 19). Optimal λ values are shown in Table 1, where λ values increased from 10% to 25% occlusion levels. Sample shape reconstructions at optimal λ values are shown in Figure 20.

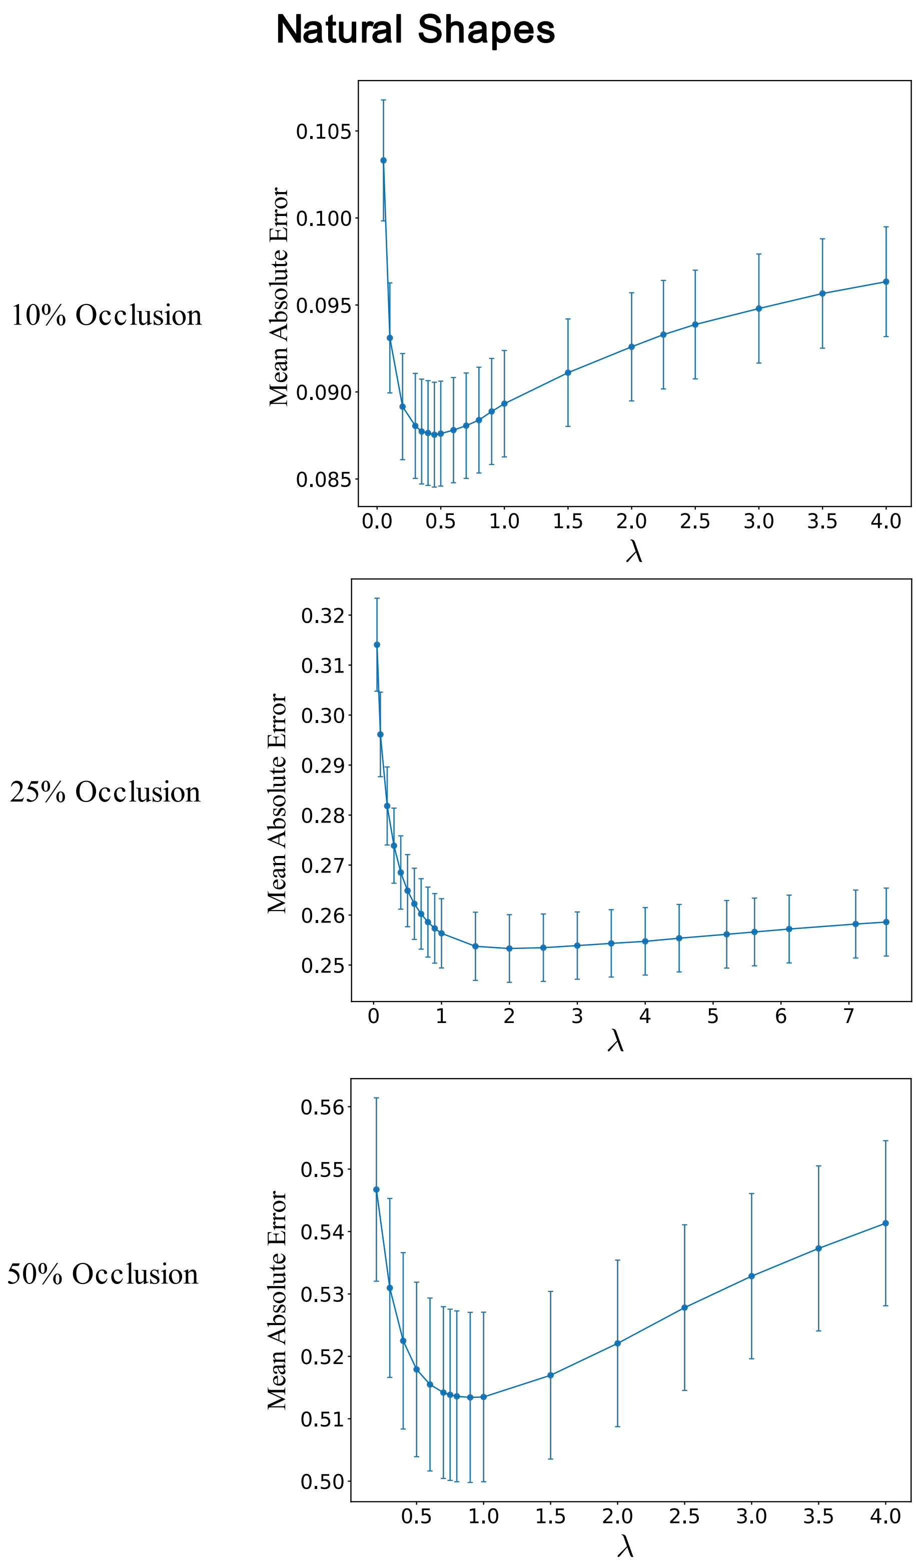

Figure 16. Elastica MAE on natural training shapes as a function of the free parameter $\boldsymbol{\lambda}$. Error bars represent SEM.

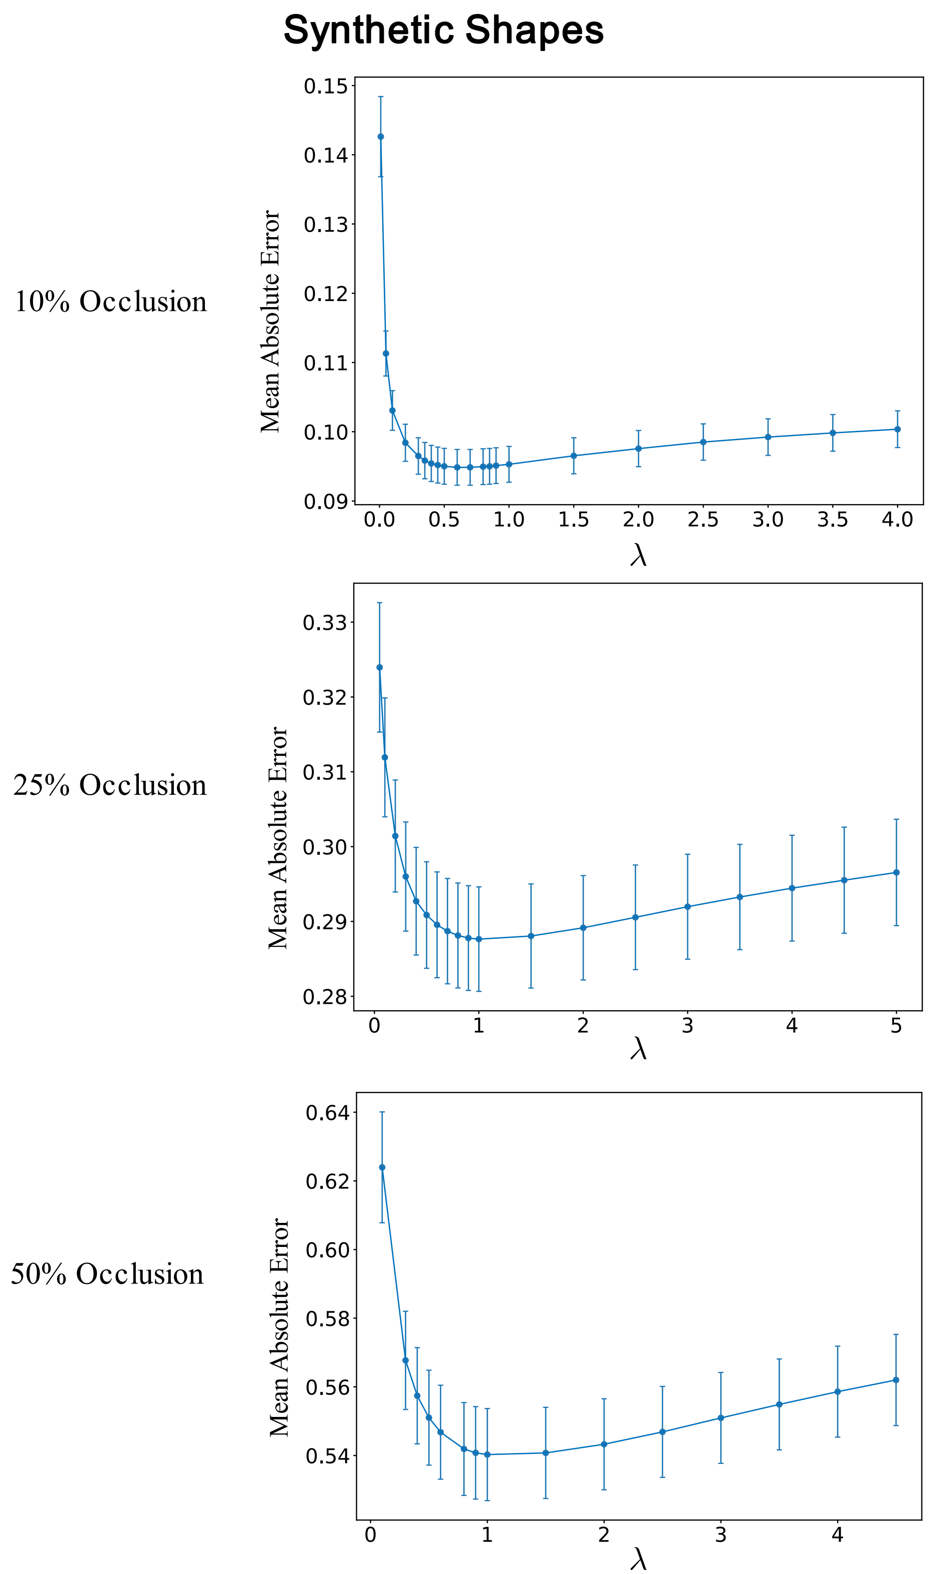

Figure 17. Elastica MAE on synthetic training shapes as a function of the free parameter $\boldsymbol{\lambda}$. Error bars represent SEM.

Table 1. Optimal λ on training shapes

Level of Occlusion	λ Natural Shapes	Synthetic Shapes
10%	0.45	0.6
25%	2.0	1.0
50%	0.9	1.0

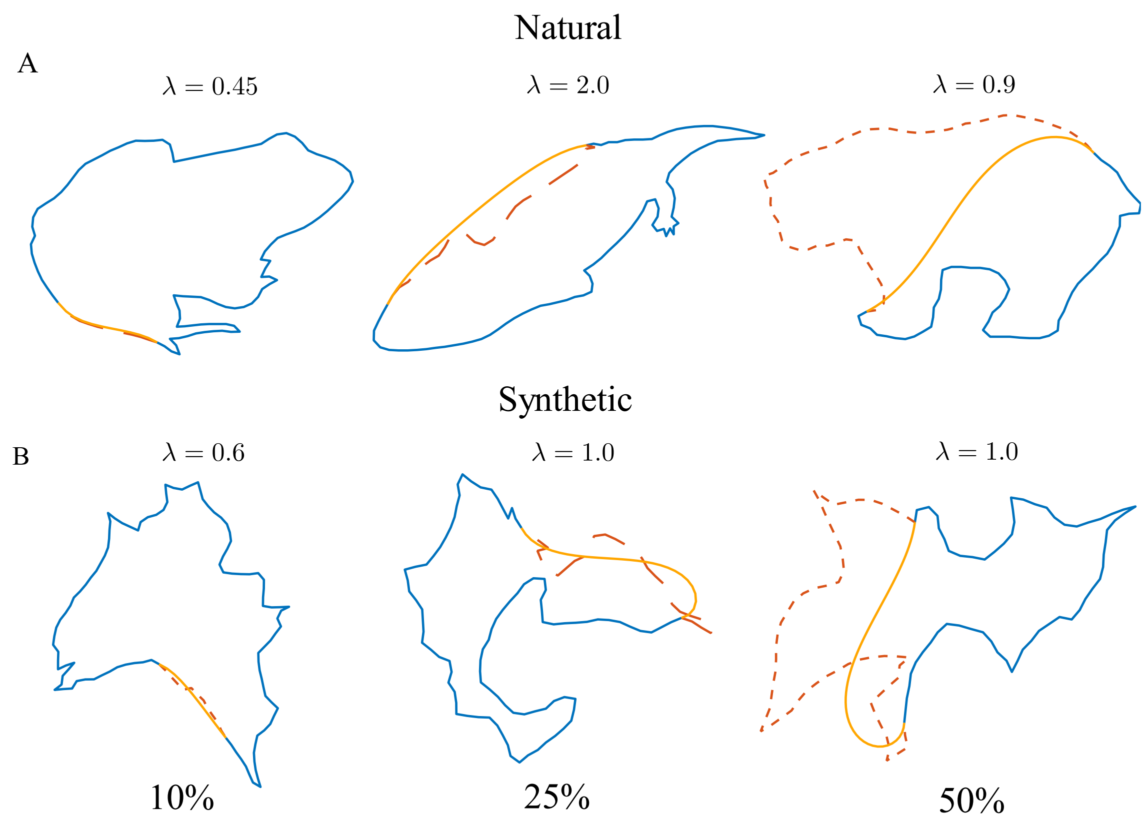

Figure 18. Sample Elastica completions with optimal λ. Blue: visible contour. Dashed orange: occluded contour. Yellow: Elastica completion.

2.2.3 Shapelet and Formlet Models

I used 64 shapelets and 64 formlets to form completion predictions for Shapelet and Formlet models (Figure 21 and Figure 22). Figure 46 shows the different iterations of formlets, starting with the initial elliptical shape Γ^0 that minimizes the L^2 error between the target shape Γ^{obs} and the ellipse.

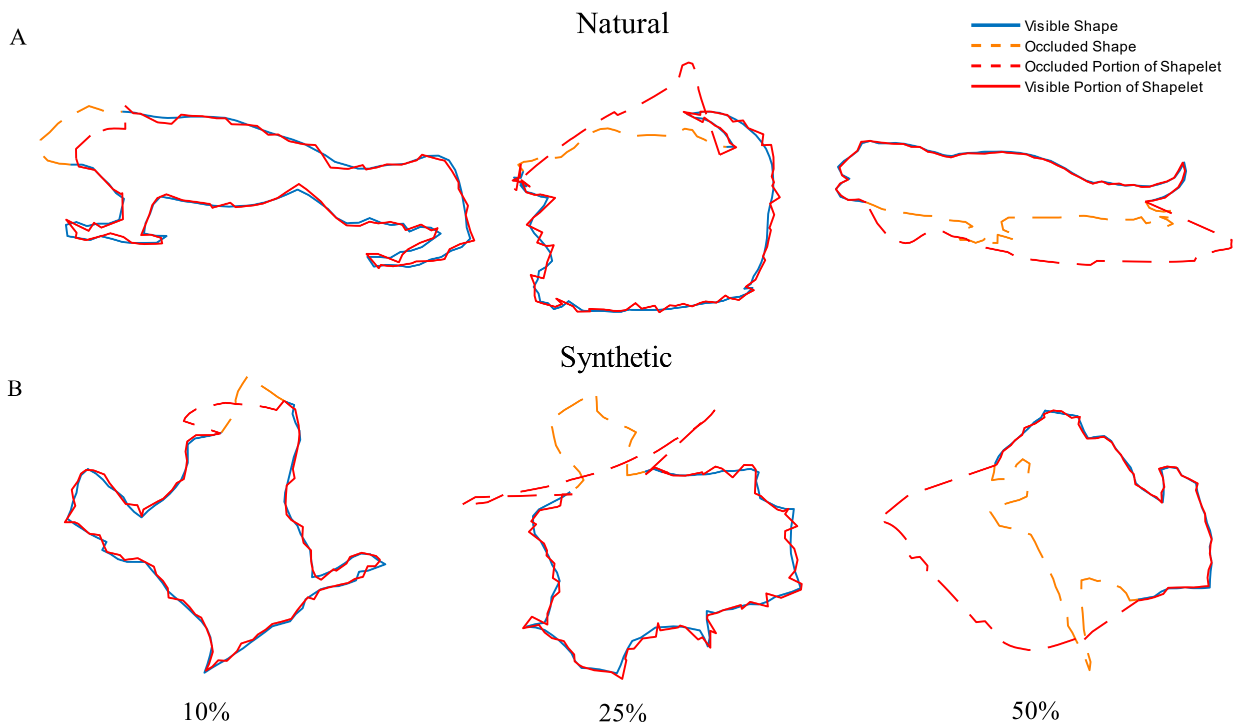

Figure 19. Sample Shapelet completions. Blue: visible contour. Orang**e**: occluded contour. Red: Shapelet completion.

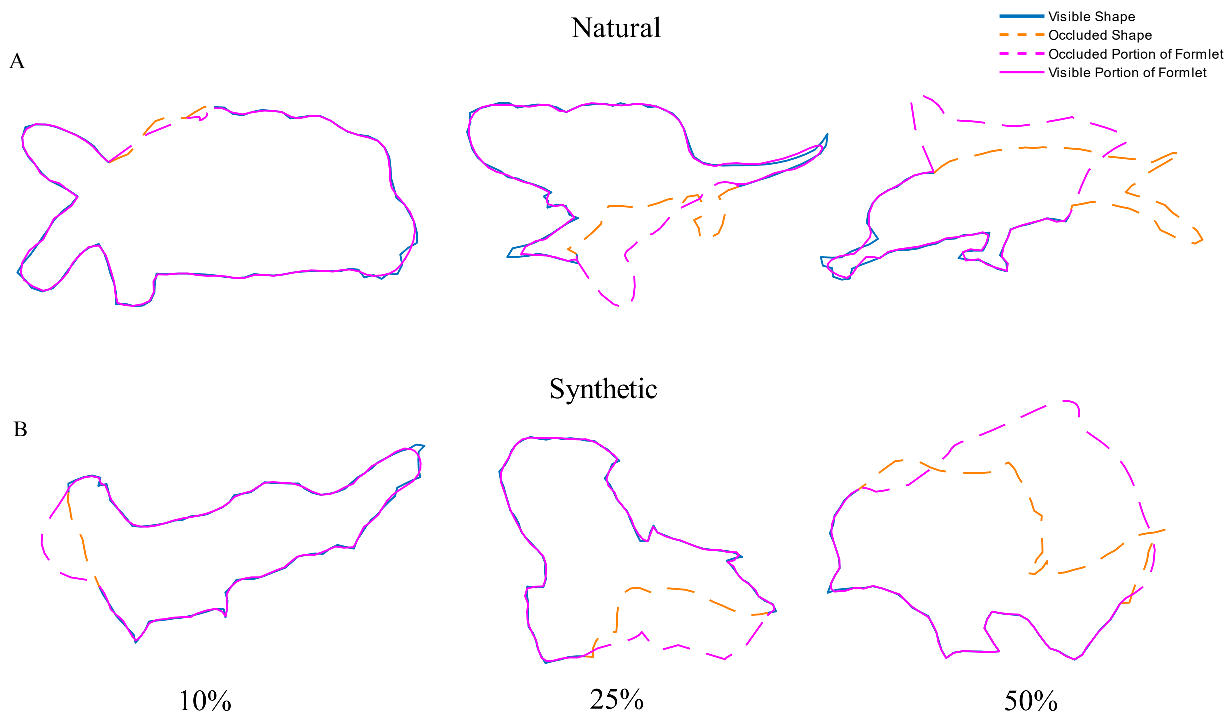

Figure 20. Sample Formlet completions. Blue: visible contour. Orange: occluded contour. Magenta: Formlet completion.

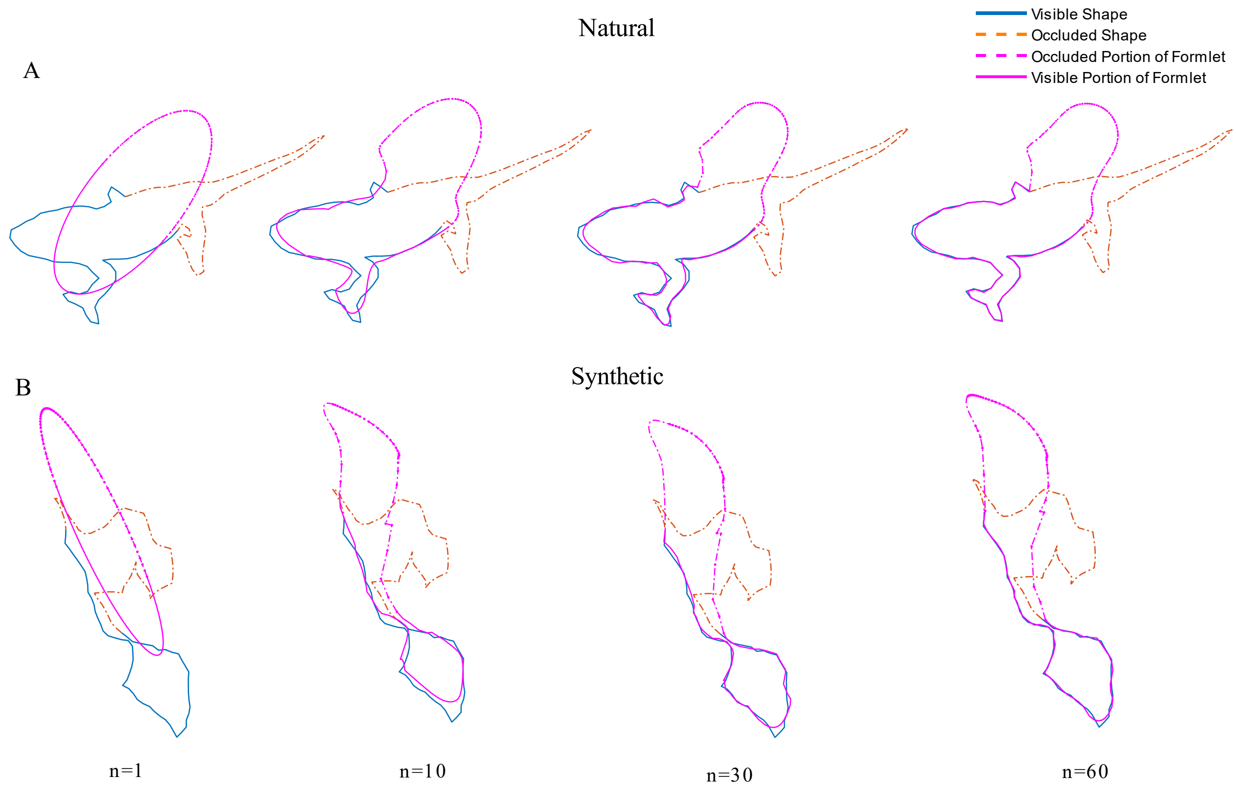

Figure 21. Formlet completions. Examples of pursuit on 50% occluded shapes. A) Natural Shape B) Synthetic shape. Formlets (magenta) in increasing iterations from left to right (n=0,10,30,60).

CHAPTER 3

Results

3.1 Mean Signed Error (Bias)

The mean signed error (bias) of humans and models was evaluated and compared for both natural and synthetic shapes (Figure 24). This analysis identifies whether on average, human and model judgements tend to be biased inward or outward relative to the actual shape, and by how much. Figure 25 compares bias for natural and synthetic shapes within each system (human or model).

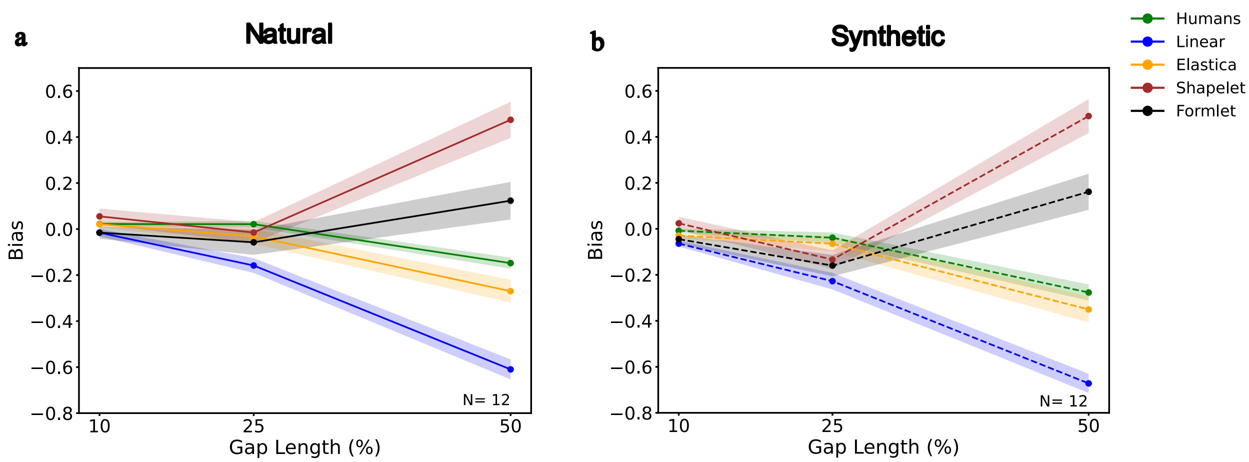

Figure 22. Comparison of bias for humans and models: (a) Natural shape (b) Synthetic shape. For the human data, shaded bars represent standard error calculated over all participants (N=12). For the model results, shading represents standard error over stimuli (N = 100). Bias is expressed in 'shape units', where 1 is the average Euclidean distance of a point on the shape from its centroid.

A linear mixed effects model was used to analyze how system type (Human, Linear, Elastica, Shapelet or Formlet), shape class (natural or synthetic shapes) and gap level (10%, 25% or 50%) influence signed error. All the analysis was done in RStudio software. I fit the following model with mean signed error as the dependent variable:

$$y_{signed\_error} = \beta_0 + \beta_g X_g + \beta_s X_s + \beta_c X_c + \beta_{gs}(X_g \cdot X_s) + \beta_{sc}(X_s \cdot X_c) + \beta_{gc}(X_g \cdot X_c) + \beta_{gsc}(X_g \cdot X_s \cdot X_c) + \epsilon \tag{5}$$

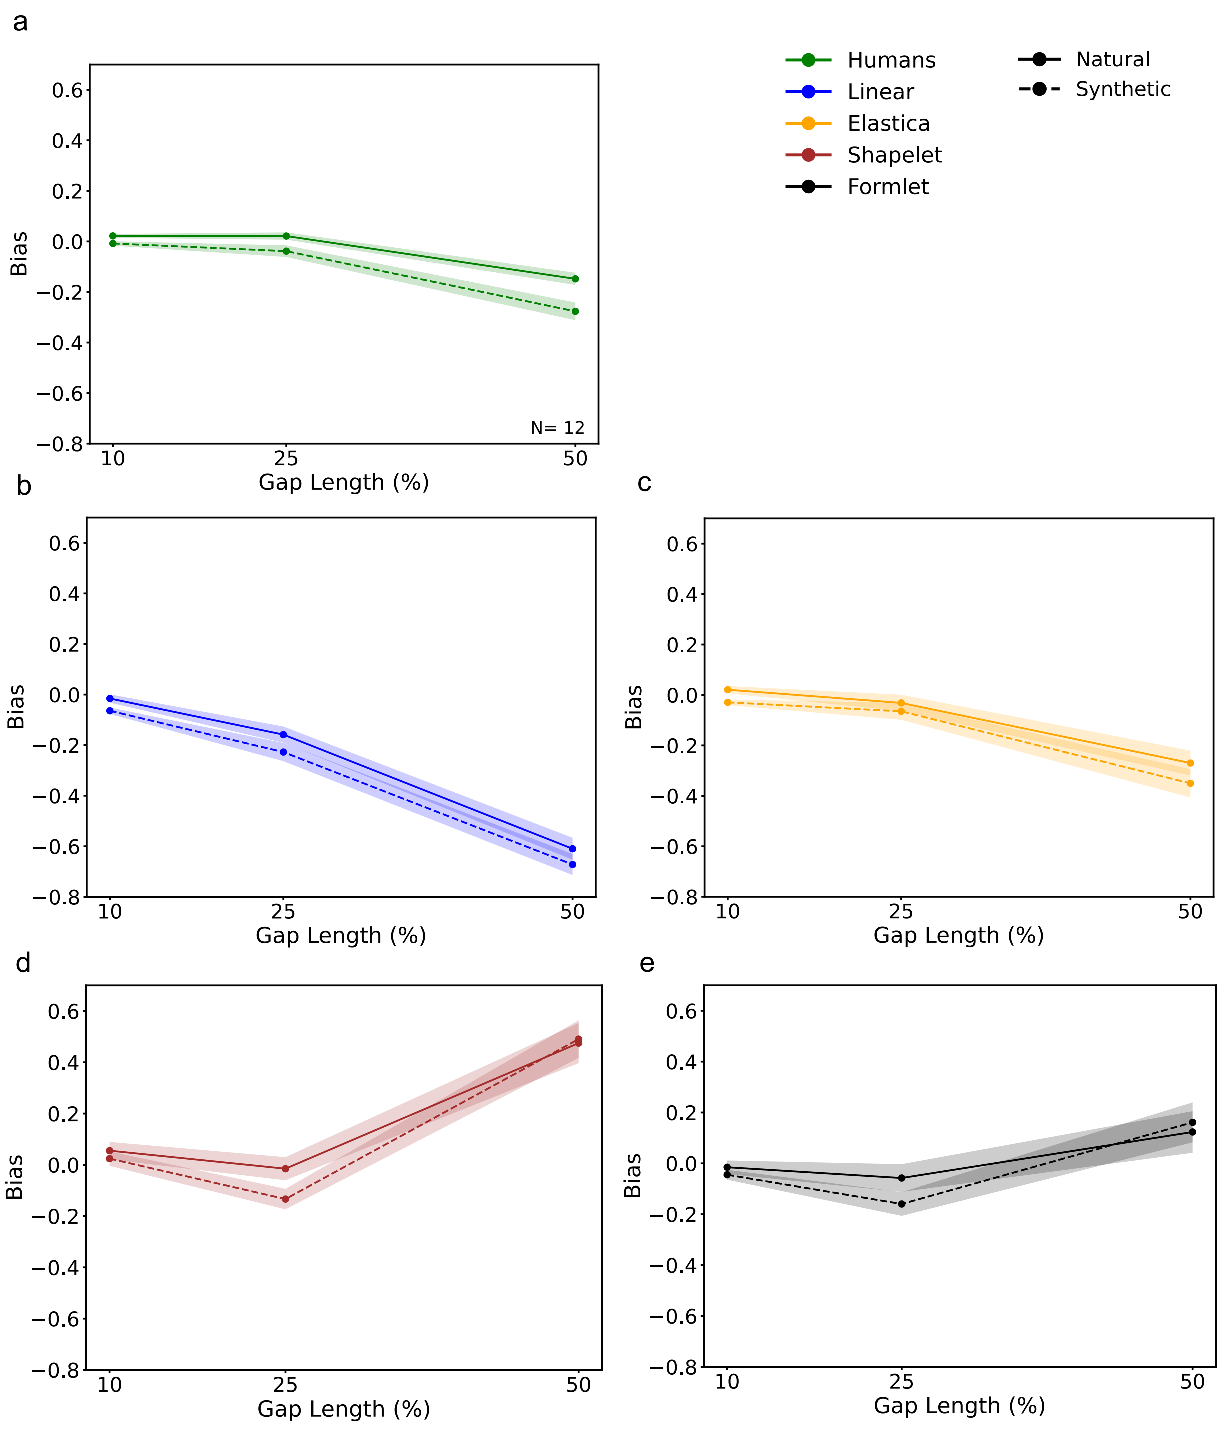

Figure 23. Mean bias for (a) Human (b) Linear model (c) Elastica model (d) Shapelet model and (e) Formlet model. Bias is expressed in 'shape units', where 1 is the mean Euclidean distance of a point on the shape from its centroid. Shading represents standard error over subjects for human data and over shapes for all models.

The predictor variables include X_g, representing the gap level treated as a categorical variable, X_s representing the system type (Human, Linear, Elastica, Shapelet or Formlet) and X_c representing shape class (natural, synthetic). We treated the gap as categorical as bias appeared to be nonlinearly related to gap. 2- and 3-way interaction terms were also included. Pairwise comparisons were done, and p-values were adjusted using the Bonferroni adjustment for multiple comparisons.

The results of the interaction tests are shown in Table 2, where β_0 represents the intercept, β_g represents the main effect of gap level, β_s is the main effect of system type and β_c is the main effect of shape class. β_{gs} is the interaction between gap level and system, β_{gc} is the interaction between gap level and shape class, while β_{sc} is the interaction between system and shape class. The three-way interaction is represented by β_{gsc}.

Table 2. Main Effects and Interactions of Linear Regression Model Coefficients for Bias.

Main Effects & Interactions	*p-level*
β_0 : 0.021, t(35970)=1.710	0.087
β_g : F(2, 35970) =178.431	<0.001
β_s : F(4, 35970)=961.933	<0.001
β_c : F(1, 35970) = 127.498	<0.001
β_{gs} : F(8, 35970)=700.015	<0.001
β_{gc} : F(2, 35970)= 6.692	<0.01
β_{sc} : F(4, 35970) =2.268	0.115
β_{gsc} : F(8, 35970)=8.118	0.227

Sign of Bias

The bias for each gap level was tested for statistical significance against zero across both shape categories and systems, as detailed in Table 3.

Table 3. Bias Significance for each level of gap at $p<0.05$. 0: No significant bias, + : Significantly positive bias, – : Significantly negative bias.

		10%	*25%*	*50%*
Natural Shape	Human	0	0	−
	Linear	0	−	−
	Elastica	0	−	−
	Shapelet	+	0	+
	Formlet	0	−	+
Synthetic Shape	Human	0	−	−
	Linear	−	−	−
	Elastica	−	−	−
	Shapelet	0	−	+
	Formlet	−	−	+

Bias is generally non-significant or negative for small gaps (10%, 25%). For large gaps (50%) bias is negative for humans and local models but positive for global models.

Bias as a Function of Gap

Pairwise comparisons were done to assess how bias changes as gap increases difference (Table 4). As gap increases from 10% to 25%, bias either decreases or remains (statistically) constant. But as gap increases further from 25% to 50%, bias decreases for humans and local models but increases for global models.

Table 4. Bias change with gap level. ↑: increases significantly; ↓ : decreases significantly; → : no significant difference. Significance at $p<0.05$.

		10% - 25%	*25% - 50%*
Natural Shape	Human	→	↓
	Linear	↓	↓
	Elastica	↓	↓
	Shapelet	↓	↑
	Formlet	→	↑
Synthetic Shape	Human	→	↓
	Linear	↓	↓
	Elastica	→	↓
	Shapelet	↓	↑
	Formlet	↓	↑

Bias as a Function of Shape Class

The mean bias averaged over gap, is shown in Figure 26. Bias is, overall negative for humans and local models, positive for shapelets and non-significant for formlets. Pairwise comparisons were conducted to examine the differences in bias across shape classes for each system, averaged over gap levels. The results of these comparisons are summarized in Table 5, indicating that for all systems bias was more positively (or less negative) for natural than synthetic shapes.

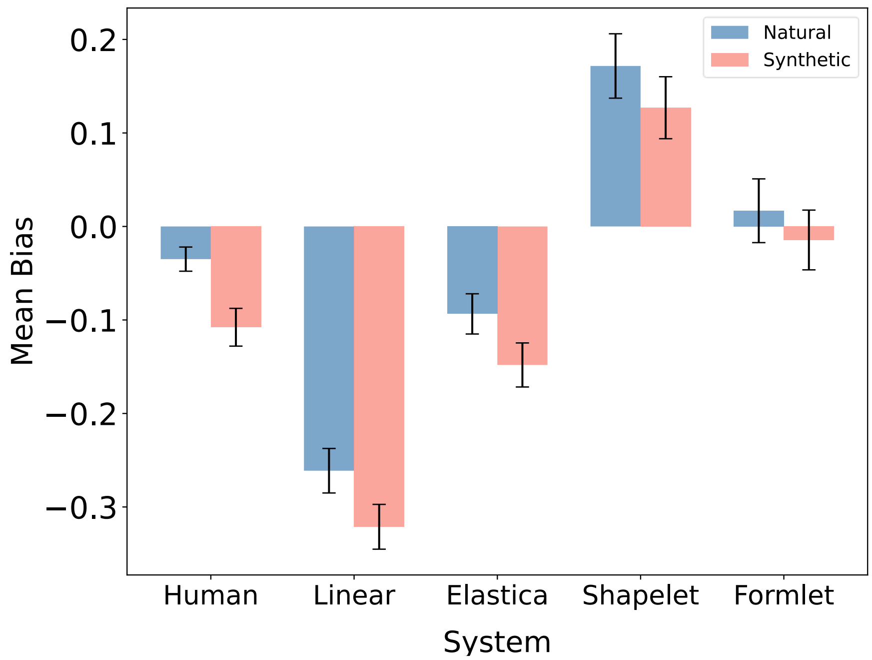

Figure 24. Mean bias averaged over gap. Error bars represent standard error across shapes for models, and across participants for human.

Table 5. Shape class differences across systems for bias. Shape Class separated by one or more inequality symbols are significantly different at the $p< 0.05$ level.

System	*Shape Class Comparison* *(+ve bias) > (-ve bias)*	*p-value*
Human	Natural > Synthetic	<0.0001
Elastica	Natural > Synthetic	<0.0001
Linear	Natural > Synthetic	<0.0001
Shapelet	Natural > Synthetic	<0.0001
Formlet	Natural > Synthetic	0.003

Difference in Bias between Humans and Models

Figure 26 considered the effect of system. However, it would be convenient to examine the effect of shape class. To address this, Figure 27 displays the mean bias, averaged over gap for natural and synthetic shapes. Among the models, shapelets generate the most pronounced positive bias, whereas the linear model produces the greatest negative bias. The local linear and elastica models, as well as humans, all demonstrate a tendency towards negative bias. This trend is consistent across both natural and synthetic shapes.

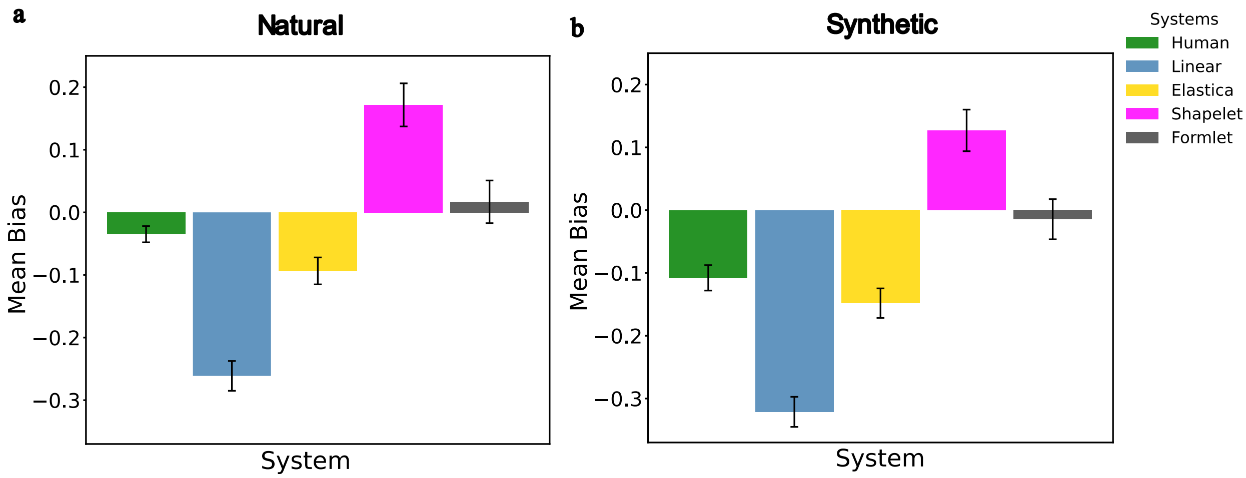

Figure 25. Mean bias for (a) natural and b) synthetic shapes. Error bars represent standard error over shapes for models, and over participants for the human system.

Analysis of pairwise differences between systems were carried out, and the results are summarized in Table 6. The shapelet model demonstrated a significantly higher positive bias than the formlet model, followed by humans, and then the local models. The linear model produced the most negative bias, and these trends were consistent across both natural and synthetic shapes.

Table 6. Differences in bias between humans and models. Systems separated by one or more inequality symbols are significantly different at the $p < 0.05$ level.

	System *(-)* ⟷ *(+)*
Natural Shapes	Linear > Elastica > Human > Formlet > Shapelet (-) (+)
Synthetic Shapes	Linear > Elastica > Human > Formlet > Shapelet (-) (+)

3.2 Standard Deviation

For each participant, the standard deviation of their signed errors across all shapes was calculated. Then the mean and standard error of these standard deviations across all participants (N=12) was computed. This provides a complement to the analysis of bias in Section 3.1, identifying how variable the signed error is across trials. The means and standard error of these standard deviations is plotted in Figure 28.

Figure 28 shows that bias increases with gap across humans ($p<0.05$), as well the standard deviation is similar for humans and local models, though a bit higher for the models at 50% occlusion ($p<0.05$). The standard deviation is higher for global models as well. Similarly, the standard deviation of signed errors across all shapes was computed for each model in Figure 29.

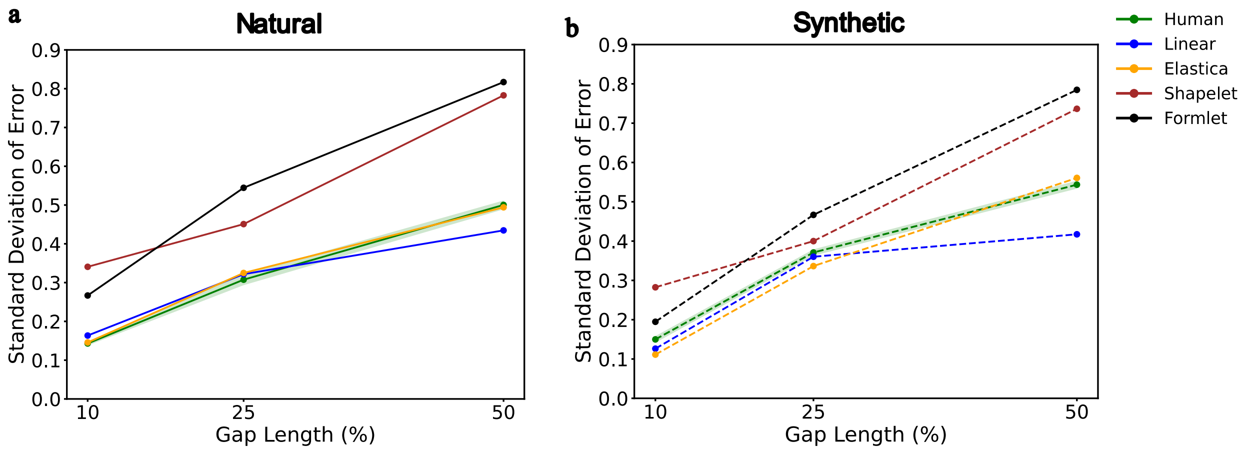

Figure 26. Comparison of standard deviations of signed error for natural and synthetic shapes: (a) Natural shape condition (b) Synthetic shape condition. Mean standard deviation was across participants for human data.

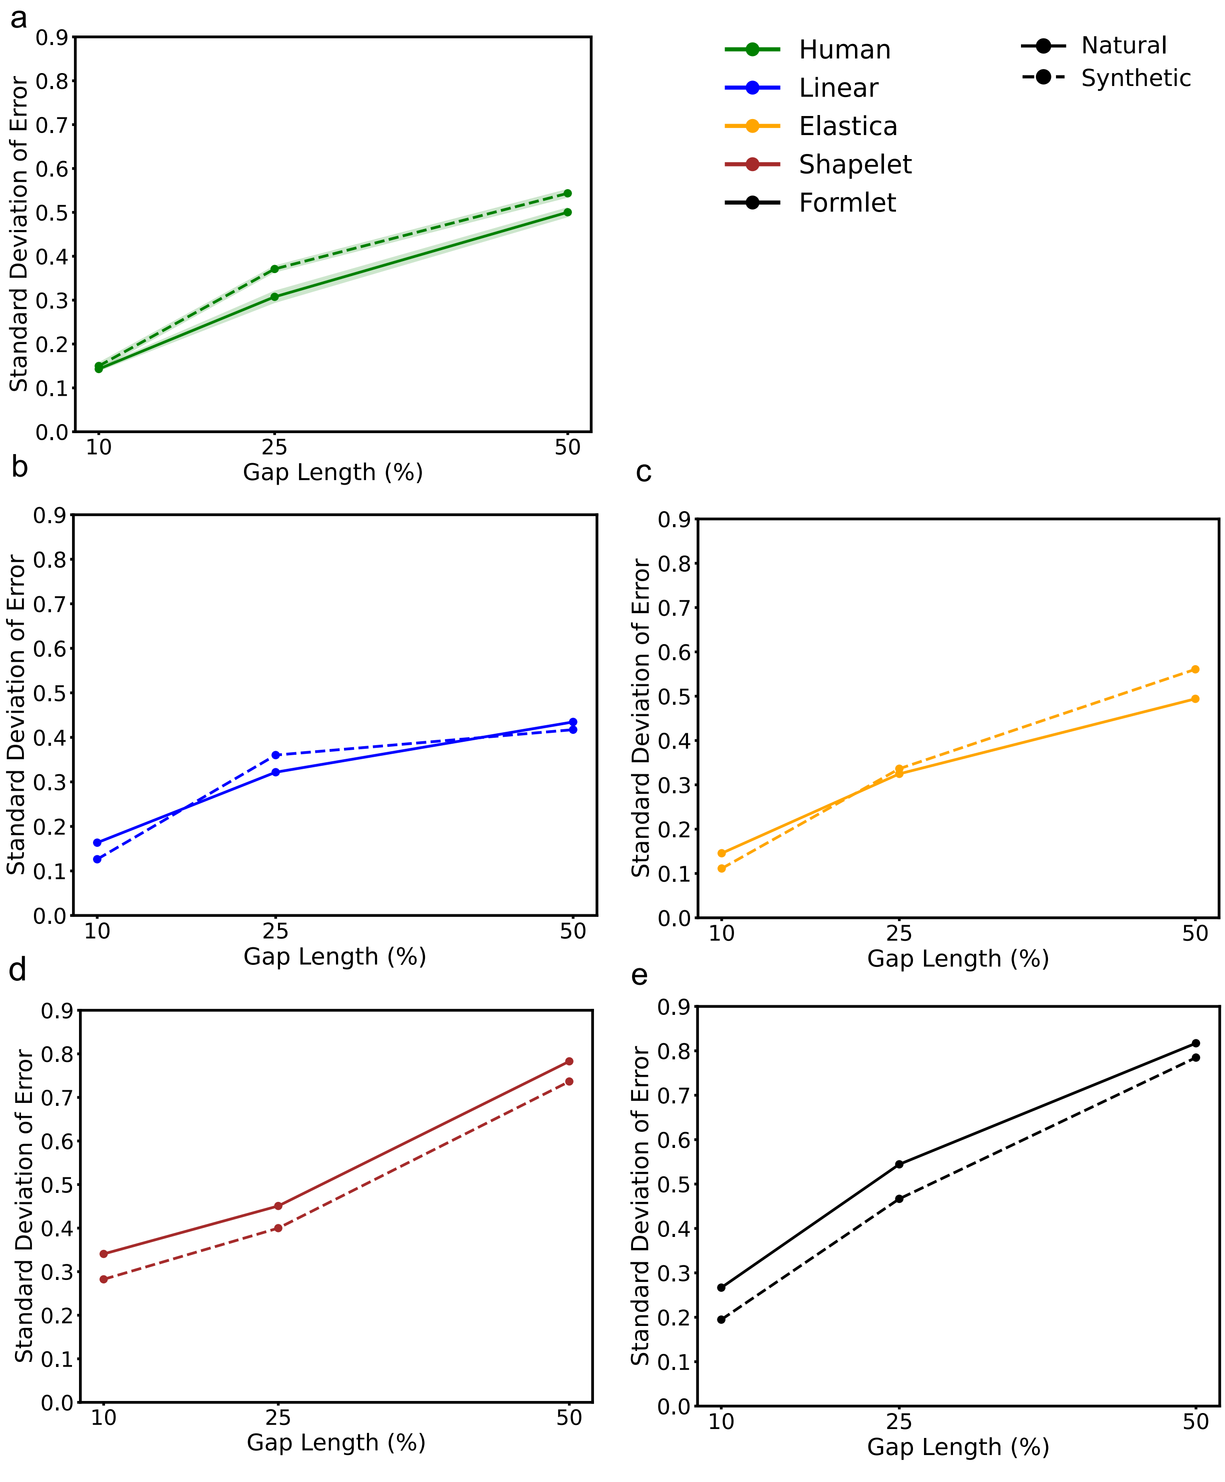

Figure 27. Standard deviations of signed error for (a) Human (b) Linear model (c) Elastica model (d) Shapelet model and (e) Formlet model.

3.3 Mean Absolute Error

The mean absolute error (MAE) of humans and models relative to the ground truth was evaluated and compared (Figure 30). A linear mixed model was employed to analyze how system (human, linear, elastica, shapelet or formlet), shape class (natural shape or synthetic shape) and gap level (10%, 25% and 50%) influence mean absolute error (accuracy). The type of system s and shape class c, were treated as categorical variables, but gap level g was treated as a scalar variable, since MAE appeared roughly linear over gap. I fit the following model with mean absolute error as the dependent variable:

$$y_{m.a.e} = \beta_0 + \beta_s X_s + \beta_c X_c + \beta_g X_g + \beta_{gs}(X_g \cdot X_s) + \beta_{sc}(X_s \cdot X_c) + \beta_{gc}(X_g \cdot X_c) + \beta_{gsc}(X_g \cdot X_s \cdot X_c) + \epsilon \quad (6)$$

The predictor variables include X_g, representing the gap level treated as a continuous variable, X_s representing the system type and X_c representing shape class. P-values were adjusted using the Bonferroni adjustment for multiple comparisons.

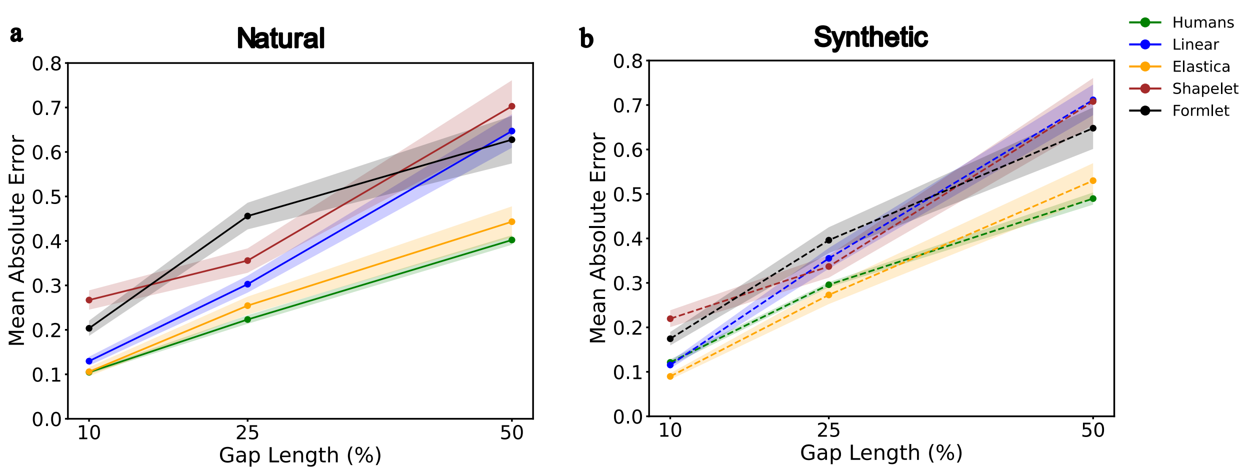

Figure 28. Comparing MAE across humans and models. (a) Natural shapes (b) Synthetic shapes. Mean absolute error is expressed in ‘shape units’, where 1 is the mean Euclidean distance of a point on the shape from its centroid. Shading represents standard error over subjects for human data (N=12) and over shapes for all models.

The results of the interaction tests are shown in Table 7, where β_0 represents the intercept, β_g represents the main effect of gap level, β_s is the main effect of system type and β_c is the main effect of shape class. β_{gs} is the interaction between gap level and system, β_{gc} is the interaction between gap level and shape class, while β_{sc} is the interaction between system and shape class. The three-way interaction is represented by β_{gsc}.

Table 7. Main Effects and Interactions of Linear Regression Model Coefficients for MAE.

Main Effects & Interactions	*p-level*
β_0 = 0.033, t(35980) = 3.380	<0.001
β_g : F(1,35980) = 13149.5	<0.001
β_s : F(4, 35980)=441.272	<0.001
β_c : F(1,35980) = 22.275	<0.001
β_{gs} : F(4, 35980)=103.435	<0.001
β_{gc} : F(1, 35980)= 91.248	<0.001
β_{sc} : F(4, 35980) =28.272	<0.001
β_{gsc} : F(4,35980)=1.135	0.227

Statistical Tests

Gap level was treated as a continuous variable to determine its effect on MAE within systems and shape classes. This involved testing the statistical significance of slope and intercepts against zero (Table 8).

Table 8. Significance of intercepts where p<0.05. + : Error is significantly greater than 0, − : Error is significantly less than 0, 0 : Error is not significantly different from 0.

		Intercept	*p-value*
Natural Shape	Human	+	0.001
	Linear	0	0.357
	Elastica	+	0.002
	Shapelet	+	<0.001
	Formlet	+	<0.001
Synthetic Shape	Human	+	<0.001
	Linear	−	0.008
	Elastica	0	0.218
	Shapelet	+	<0.001
	Formlet	+	<0.001

The results indicate that for all shape classes and systems, the slope is significantly greater than zero ($p<0.05$). Synthetic shapes are associated with steeper slopes for all systems compared to natural shapes (Figure 31, Table 9).

Table 9. Slope contrasts across shape classes. Shape Class separated by one or more inequality symbols are significantly different at the p< 0.05 level.

System	*Slope Comparison* *(steep slope) > (less steep)*	*p-value*
Human	Synthetic > Natural	0.0001
Elastica	Synthetic > Natural	<0.0001
Linear	Synthetic > Natural	<0.0001
Shapelet	Synthetic > Natural	0.0013
Formlet	Synthetic > Natural	0.0001

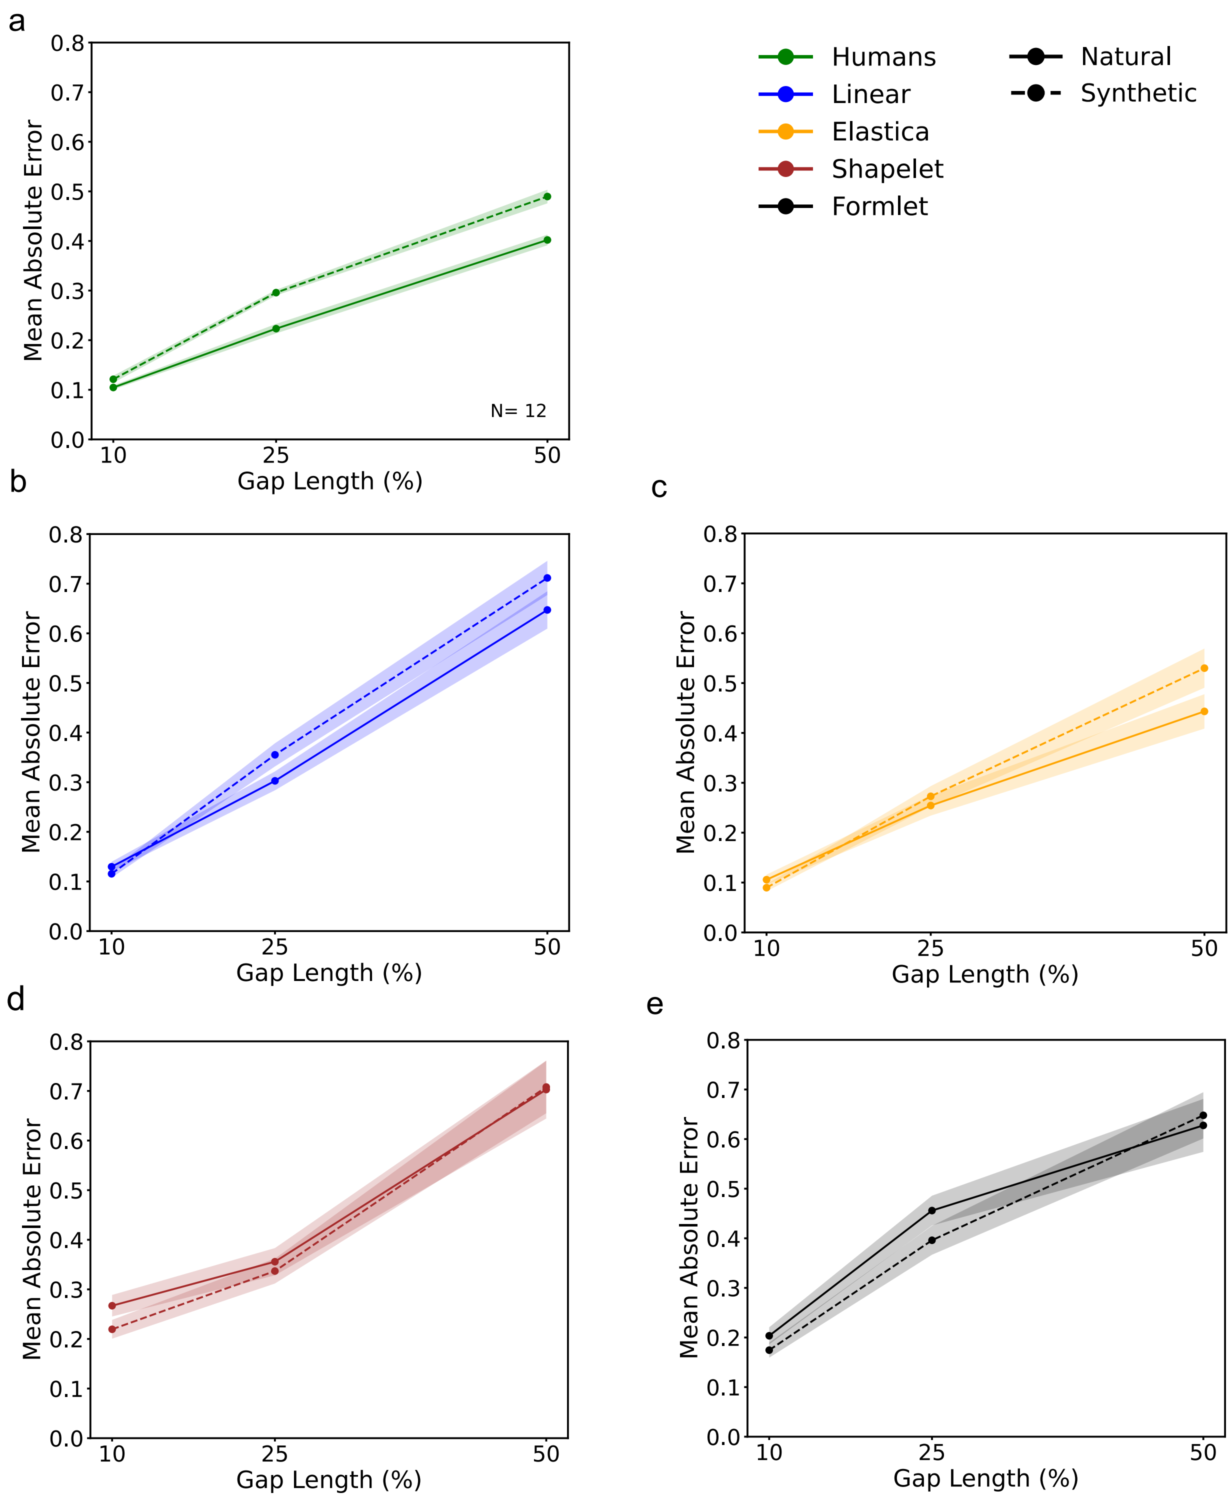

Figure 29. Mean Absolute Error for (a) Human (b) Linear model (c) Elastica model (d) Shapelet model and (e) Formlet model. Error is expressed in 'shape units', where 1 is the mean Euclidean distance of a point on the shape from its centroid. Shading represents standard error over subjects for human data (N=12) and over shapes for all models.

Intercepts for natural and synthetic shapes were not significantly different for human and linear systems. For all other systems, natural shapes produced significantly greater intercepts than synthetic shapes (Table 10).

Table 10. Intercepts of MAE across natural and synthetic shapes. Shape Class separated by one or more inequality symbols are significantly different at the $p < 0.05$ level. 0 for non-significance.

System	*Intercept Comparison* *(larger error) > (small error)*	*p-value*
Human	0	0.4786
Elastica	Natural > Synthetic	0.0182
Linear	0	0.7475
Shapelet	Natural > Synthetic	<0.0001
Formlet	Natural > Synthetic	<0.0001

Effect of System: Slopes and Intercepts

Table 11 summarises the differences in slope across systems. The linear system showed the steepest increase in MAE with increasing gap relative to all other systems, and humans and elastica showed the smallest. Humans were not significantly different in steepness of slope with elastica model for natural shapes but were significantly different for synthetic shapes.

Table 11. Slopes of MAE across systems. Systems separated by one or more inequality symbols are significantly different at the $p < 0.05$ level.

System Slopes

(high +) ⟷ *(low -)*

Natural Shapes	Linear > Formlet, Shapelet > Elastica, Human
	(steep slope) (less steep)
Synthetic Shapes	Linear > Shapelet > Elastica > Human
	Linear > Formlet > Human

Table 12 presents a summary of results from the pairwise comparisons of intercepts between systems within each shape class.

Table 12. System intercepts of MAE for natural and synthetic shapes. Systems separated by one or more inequality symbols are significantly different at the $p < 0.05$ level.

System Intercepts

(high +) ⟷ *(low -)*

Natural Shapes	Shapelet, Formlet > Elastica, Human > Linear
	(high error) (low error)
Synthetic Shapes	Human, Shapelet, Formlet > Elastica, Linear

3.4 Mean Absolute Deviation (MAD) between models and humans.

The mean absolute deviation of model estimates from human estimates was evaluated and compared. This was done to evaluate how well each model could serve as a model of human judgements, as opposed to their accuracy relative to the ground truth. The type of system s and shape class c, were treated as categorical variables, while gap level g was treated as a scalar variable, since MAD was observed to be roughly linear over gap. I fit the following model with mean absolute deviation as the dependent variable:

$$y_{deviation} = \beta_0 + \beta_s X_s + \beta_c X_c + \beta_g X_g + \beta_{gs}(X_g \cdot X_s) + \beta_{sc}(X_s \cdot X_c) + \beta_{gc}(X_g \cdot X_c) + \beta_{gsc}(X_g \cdot X_s \cdot X_c) + \epsilon \quad (7)$$

P-values were adjusted using the Bonferroni adjustment for multiple comparisons. The results of the interaction tests are shown in Table 13.

Table 13. Main Effects and Interactions of Linear Regression Model Coefficients for MAD.

Main Effects & Interactions	*p-value*
β_0 = 0.048, t(35980) = 5.409	<0.01
β_g : F(1, 35980) =11914.593	<0.001
β_s : F(4, 35980)=1417.828	<0.001
β_c : F(1,35980) = 5.255	<0.01
β_{gs} : F(4, 35980)=443.691	<0.001
β_{gc} : F(1, 35980)= 150.508	<0.001
β_{sc} : F(4, 35980) =10.56	<0.001
β_{gsc} : F(4,35980)=39.113	< 0.001

Mean Absolute Deviation as a Function of Gap.

The statistical significance of the slope and intercepts of the Mean Absolute Deviation (MAD) as a function of gap was analyzed. Except for the linear system, synthetic shapes consistently demonstrated a steeper slope than natural shapes across all systems (Table 14, Figure 32).

Table 14. Slope of MAD contrasts across shape classes. Shape Class separated by one or more inequality symbols are significantly different at the $p < 0.05$ level.

System	*Slope Comparison*	*p-value*
Human	Synthetic > Natural	0.0001
Elastica	Synthetic > Natural	<0.0001
Linear	Natural > Synthetic	<0.0001
Shapelet	Synthetic > Natural	0.0013
Formlet	Synthetic > Natural	0.0001

All intercepts are positive or zero with the exception of the intercept for the shapelet model on synthetic shapes (Table 15), where the large positive curvature in the data has driven the intercept to a negative value.

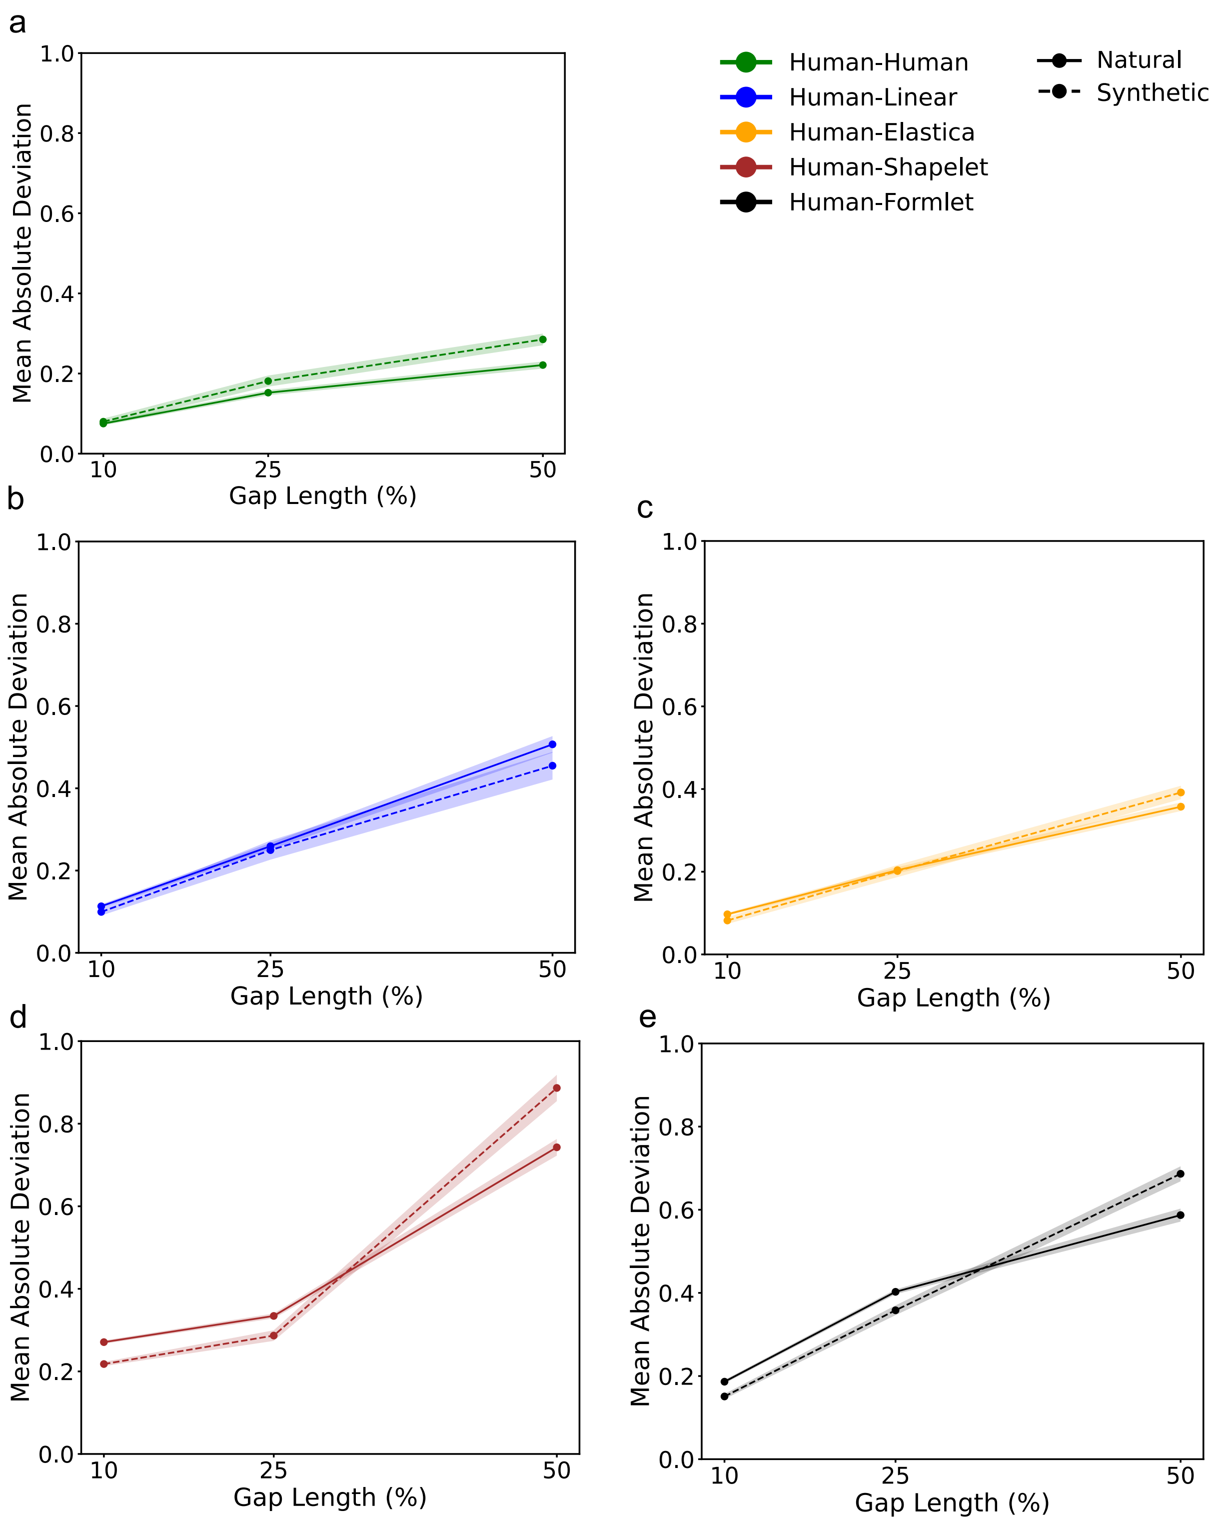

Figure 30. Mean absolute deviations of human estimates from: (a) Mean of the other human estimates (b) Linear model estimates (c) Elastica model estimates (d) Shapelet model estimates (e) Formlet model estimates.

Table 15. Significance of intercepts where $p < 0.05$. + : Error is significantly greater than 0, − : Error is significantly less than 0, 0 : Error is not significantly different from 0.

		Intercept	*p-value*
Natural Shape	Human	+	0.001
	Linear	0	0.357
	Elastica	+	0.002
	Shapelet	+	<0.0001
	Formlet	+	<0.0001
Synthetic Shape	Human	+	<0.0001
	Linear	+	0.008
	Elastica	0	0.218
	Shapelet	−	<0.0001
	Formlet	+	<0.0001

Mean Absolute Deviation as a Function of Shape Class

The slope of MAD as a function of gap was smaller for synthetic shapes than for natural shapes for humans and all models, except for the linear model for which slope was higher for synthetic shapes (Table 14).

In terms of the intercepts, human and linear deviations from human estimates were found to be not significantly different across shape classes. For the other three system deviations, natural shapes produced larger intercepts than synthetic shapes (see Table 16).

Table 16. Intercepts of MAD across natural and synthetic shapes. Shape Class separated by one or more inequality symbols are significantly different at the p< 0.05 level. Non-significance = 0.

System	*Intercept Comparison*	*p-values*
Human	0	0.389
Elastica	Natural > Synthetic	0.002
Linear	0	0.220
Shapelet	Natural > Synthetic	<0.0001
Formlet	Natural > Synthetic	<0.0001

MAD as a Function of System

For both natural and synthetic shapes, slopes were highest for shapelets and lowest for humans (Figure 33, Table 17). Of the models, elastica had the smallest slopes for natural shapes and linear and elastica models had the smallest slopes for synthetic shapes (Table 17).

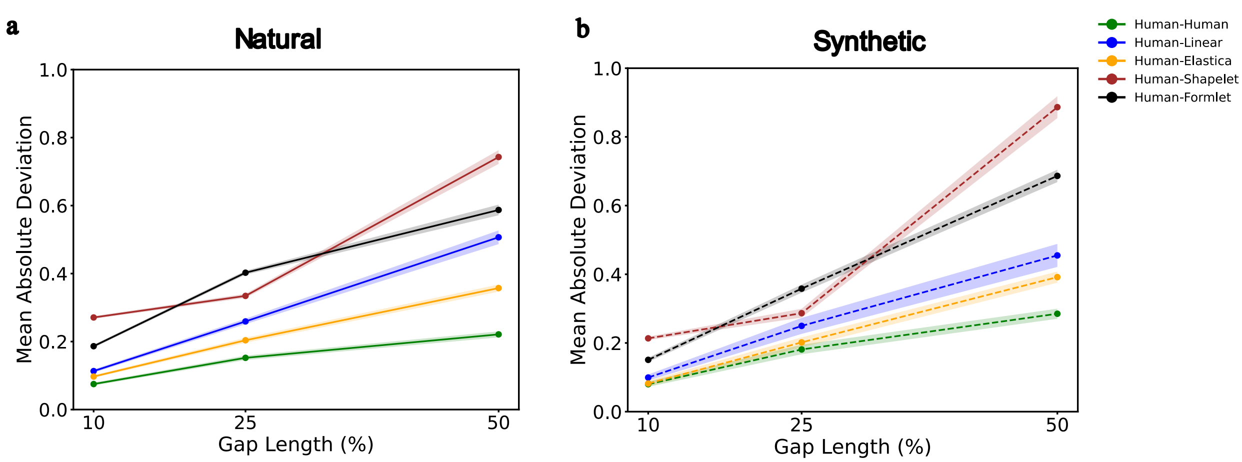

Figure 31. Mean absolute deviation of individual human participants from the mean judgements of other participants (Human-Human) together with the mean absolute deviation of each model's judgement from mean human judgements.

Table 17. Slopes of MAD across systems. Systems separated by one or more inequality symbols are significantly different at the p< 0.05 level.

	System Slopes *(high +)* ⟷ *(low -)*
Natural Shapes	Shapelet > Linear, Formlet > Elastica > Human
Synthetic Shapes	Shapelet > Formlet > Linear, Elastica > Human * all other slope contrasts were non-significant.

For both natural and synthetic shapes, intercepts were largest for the shapelet model and smallest for humans. Of the models, intercepts were smallest for the elastica model (Table 18).

Table 18. Intercepts of MAD across natural and synthetic shapes. Systems separated by one or more inequality symbols are significantly different at the p< 0.05 level.

	System *(high +)* ⟷ *(low -)*
Natural Shapes	Shapelet > Formlet, Linear > Elastica > Human
Synthetic Shapes	Shapelet > Formlet > Linear > Elastica > Human

3.5 Outlier Model Predictions

To gain further insight into the quantitative results presented above, I checked whether any particular shapes & completion tasks generated outlier or extreme data for either humans or models. Figure 34 displays the signed error across all trials for the first condition (Natural Shapes and 10% Gap level). There is one shape (Shape 13- Figure 35) that generates an extreme negative error for the formlet model.

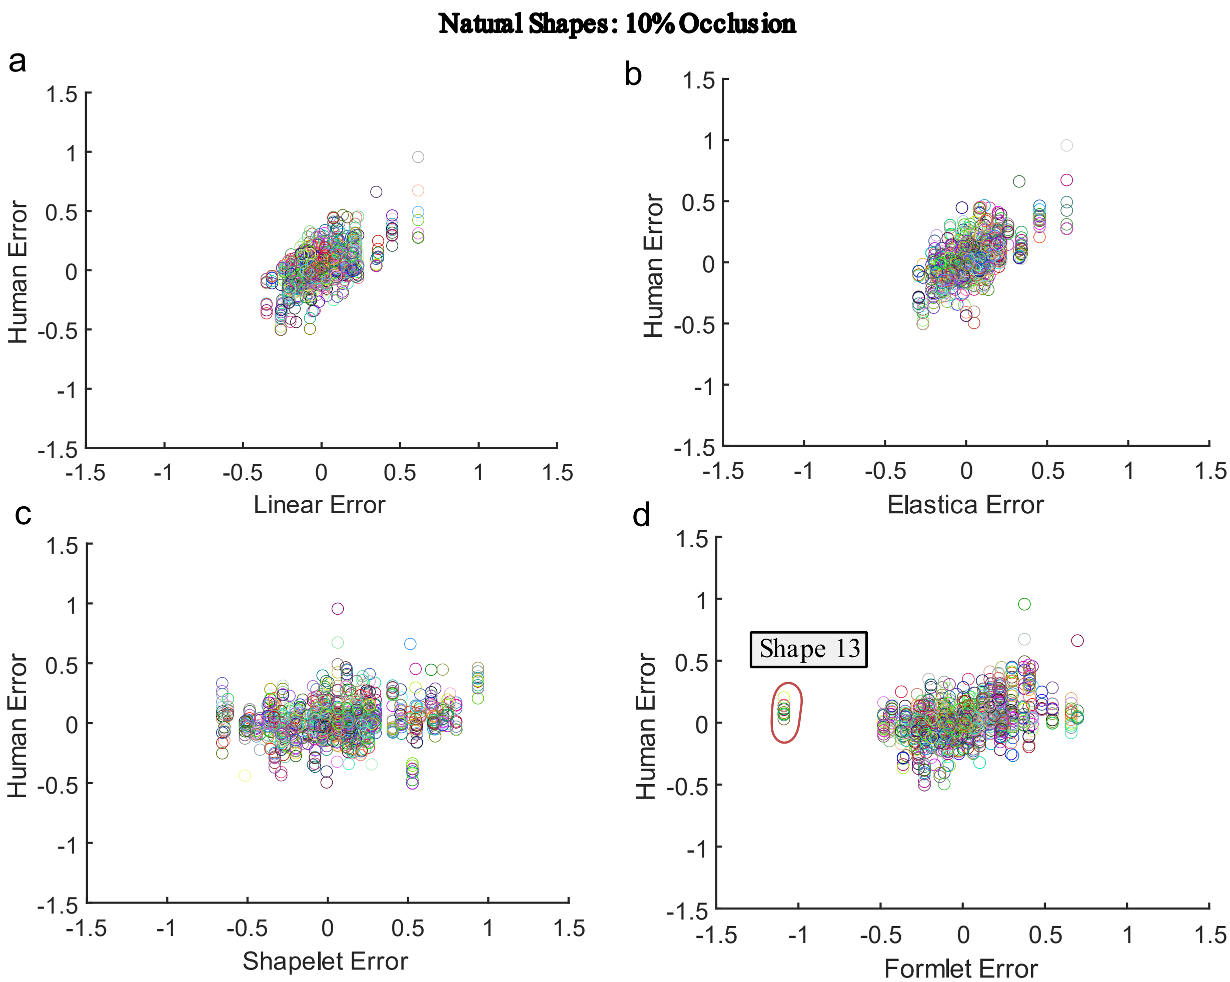

Figure 32. Scatterplot of signed deviation for all participants and models.

Note that error is also negative and large in magnitude for the shapelet model. These errors are caused by large amplitude oscillations in the completion contours. These are likely produced by large amplitude formlets deployed to fit the legs of the giraffe.

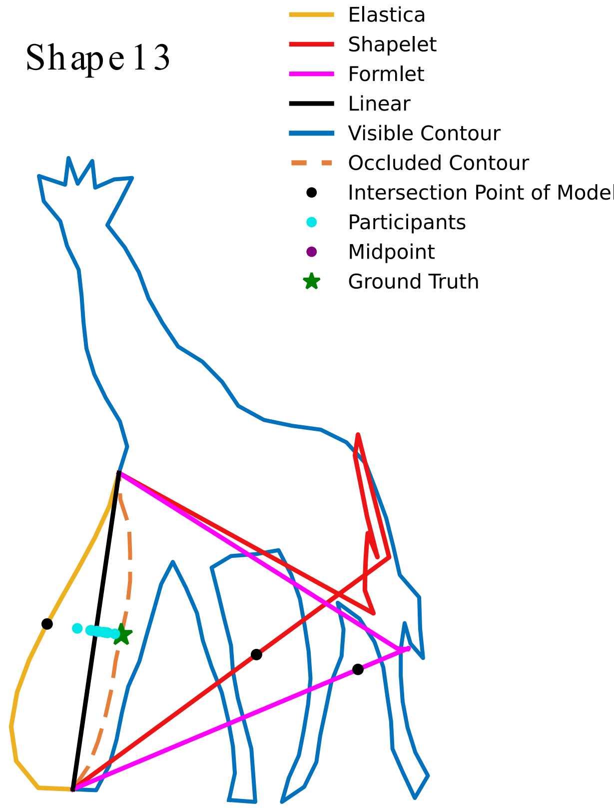

Figure 33. Outlier formlet prediction for a natural shape at 10% occlusion.

Figure 36 displays the signed error for natural shapes, at 25% gap level. We observed one shape (Shape 24) for which the Elastica model generated an unusually high positive error. This shape is shown in Figure 37. We note that the endpoint tangent directions on which the elastica prediction is based are misleading in this case. None of the models are accurate and human estimates are broadly dispersed.

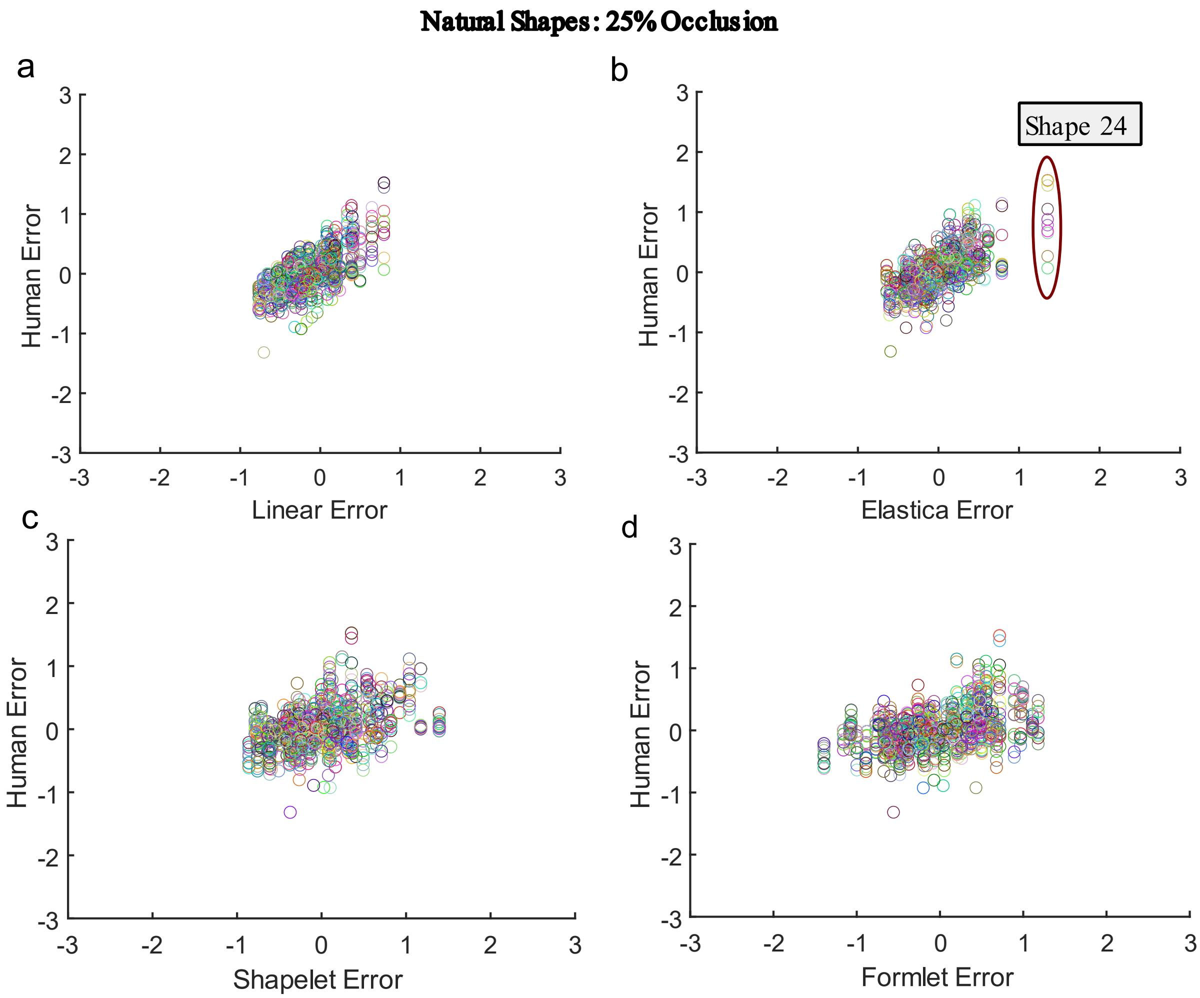

Figure 34. Scatterplot of signed deviations from ground truth for all participants and models. Red Outline: outlier in elastica completion and formlet completion.

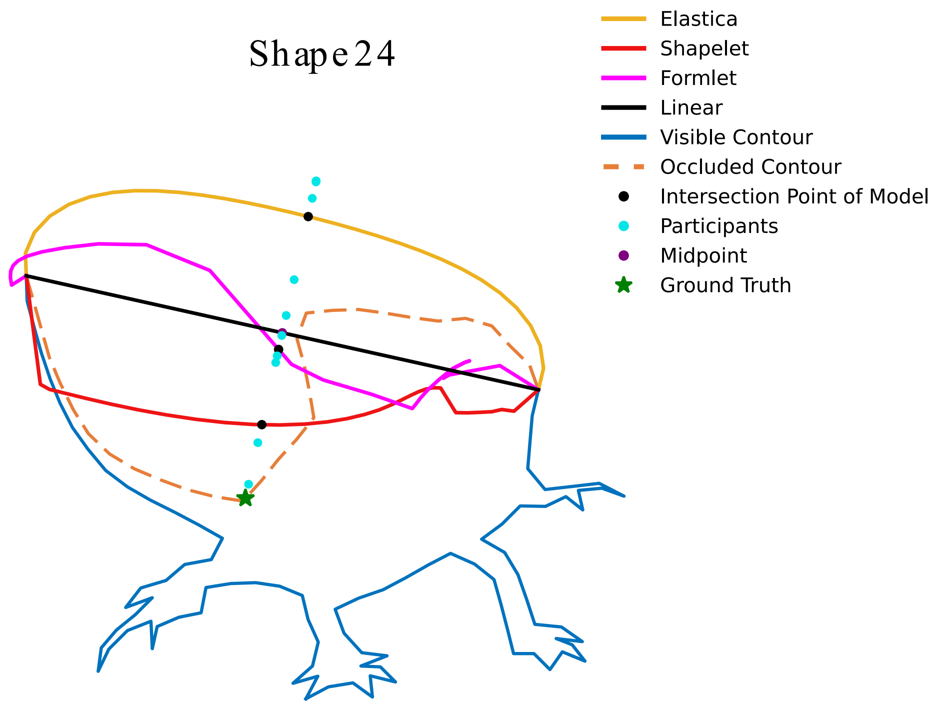

Figure 35. Outlier elastica prediction for a natural shape at 25% occlusion.

The scatterplots for natural shapes at 50% gap level is shown in Figure 38. The shapelet and formlet models generate negative errors for Shape 48. In Figure 39, the majority of human participants experience a figure-ground reversal. Interestingly, this experience is shared by the global models, whereas by paying attention to only local cues at the endpoints, the local models do much better.

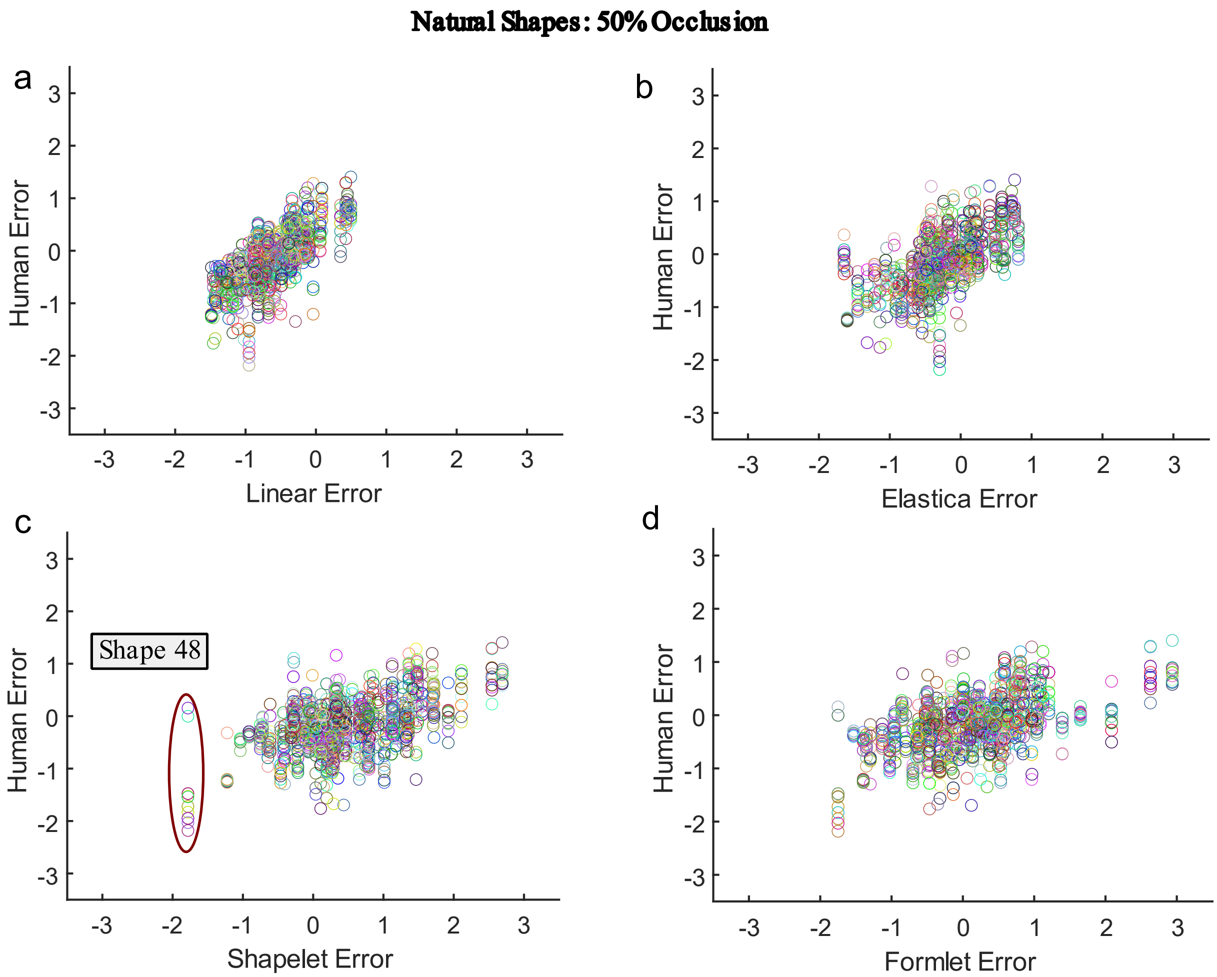

Figure 36. Scatterplots of signed deviations for natural shapes at 50% occlusion.

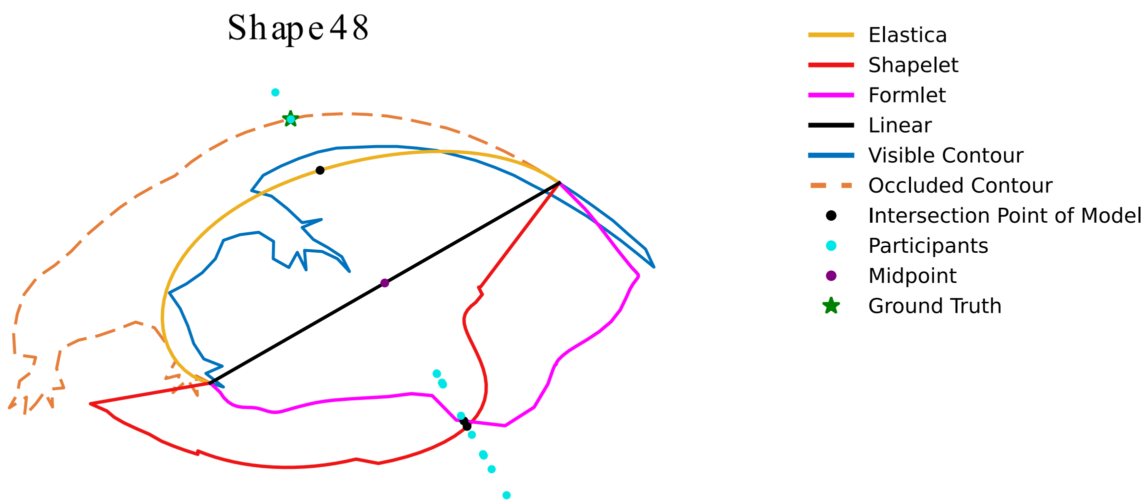

Figure 37. Outlier shapelet and formlet predictions for a natural shape at 50% occlusion.

Figure 40 represents the signed error across all trials for synthetic shapes at the 10% gap level. No clear outliers are apparent.

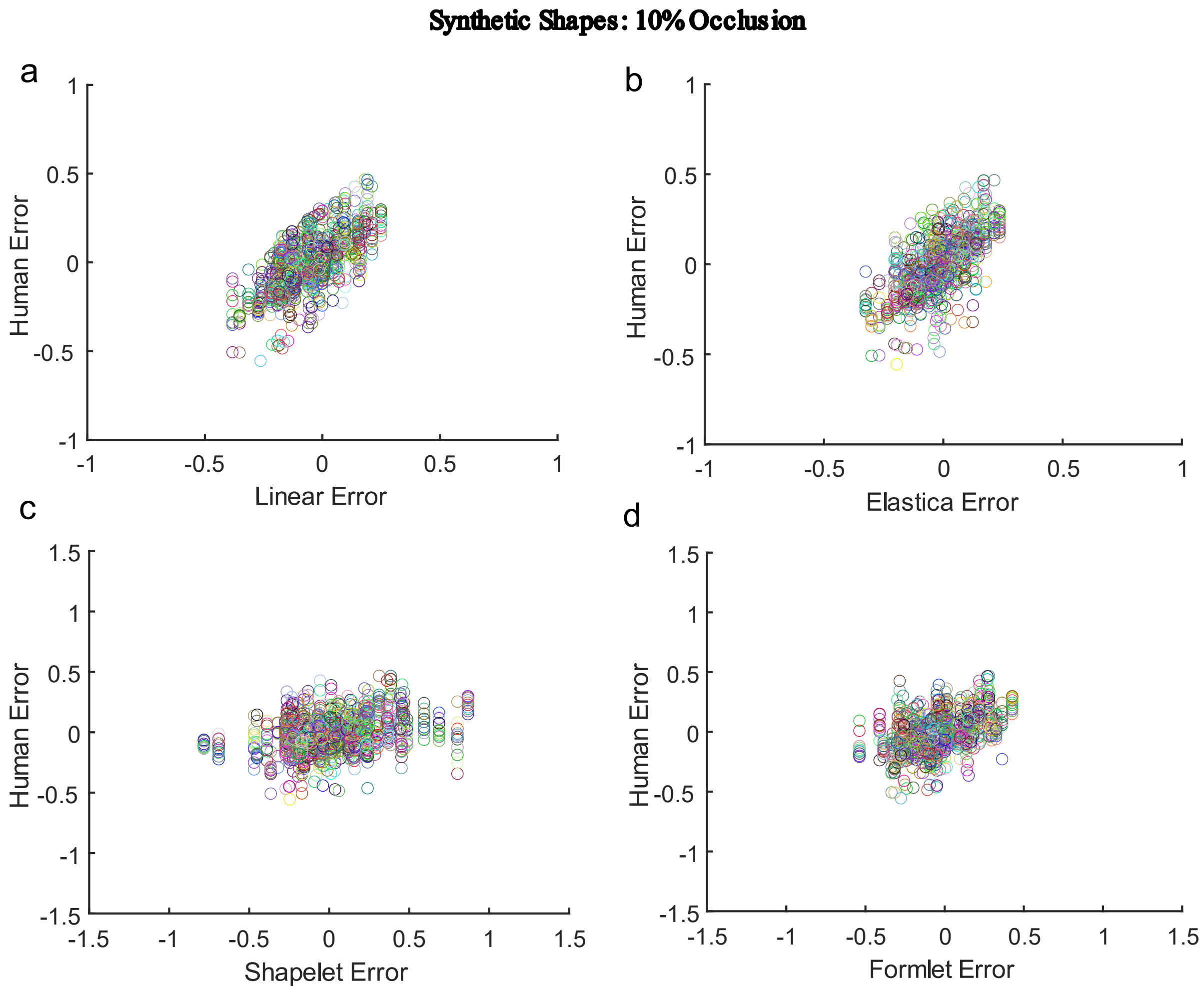

Figure 38. Scatterplot of signed deviations for synthetic shapes at 10% occlusion.

Figure 41 represents the signed error across all trials for synthetic shapes at the 25% gap level. One notable extreme case is shape 83, where the shapelet completion has very large negative error, while humans are much less biased (Figure 42).

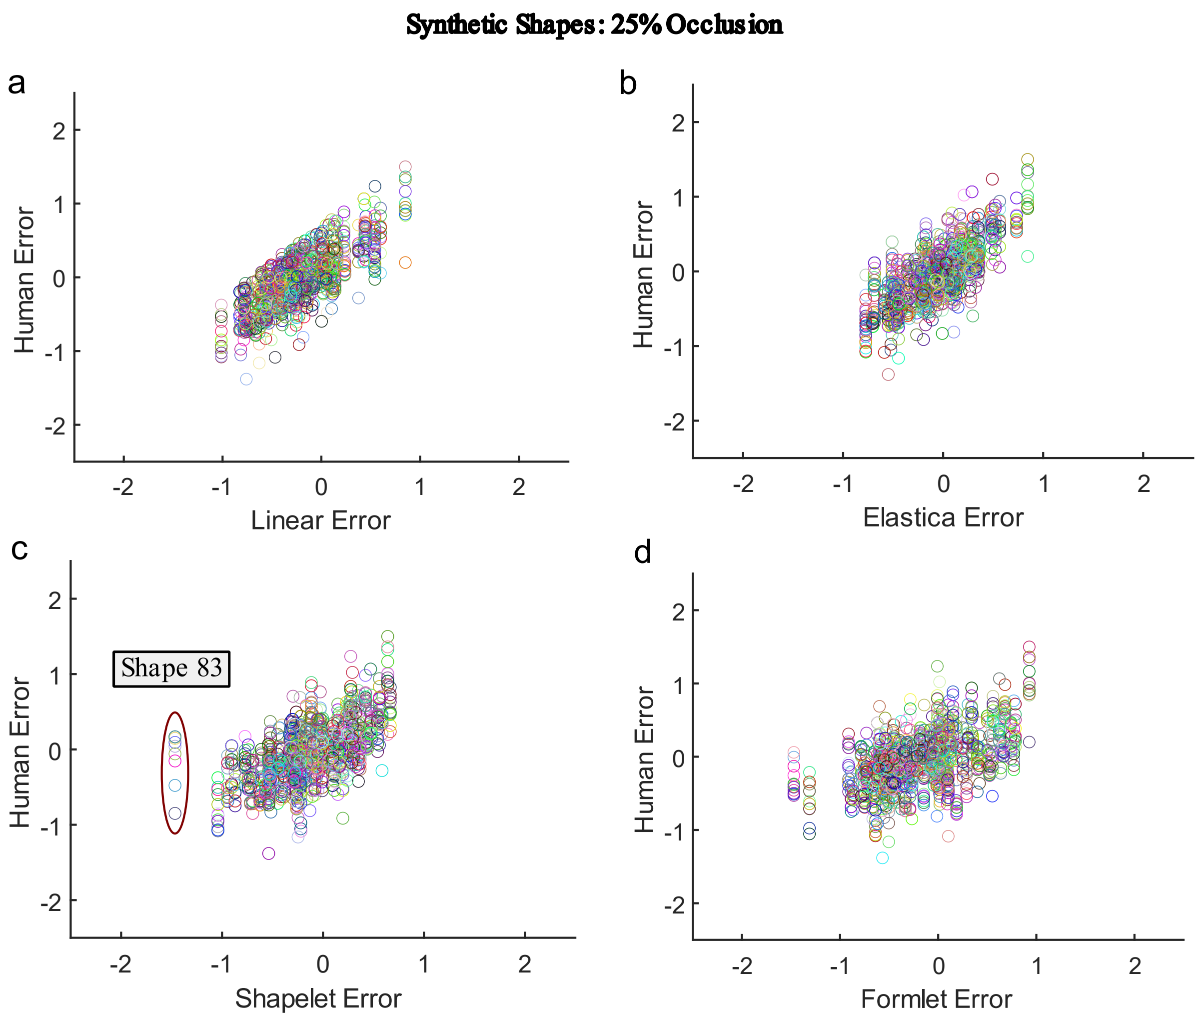

Figure 39. Scatterplot of signed deviations for synthetic shapes at 25% occlusion.

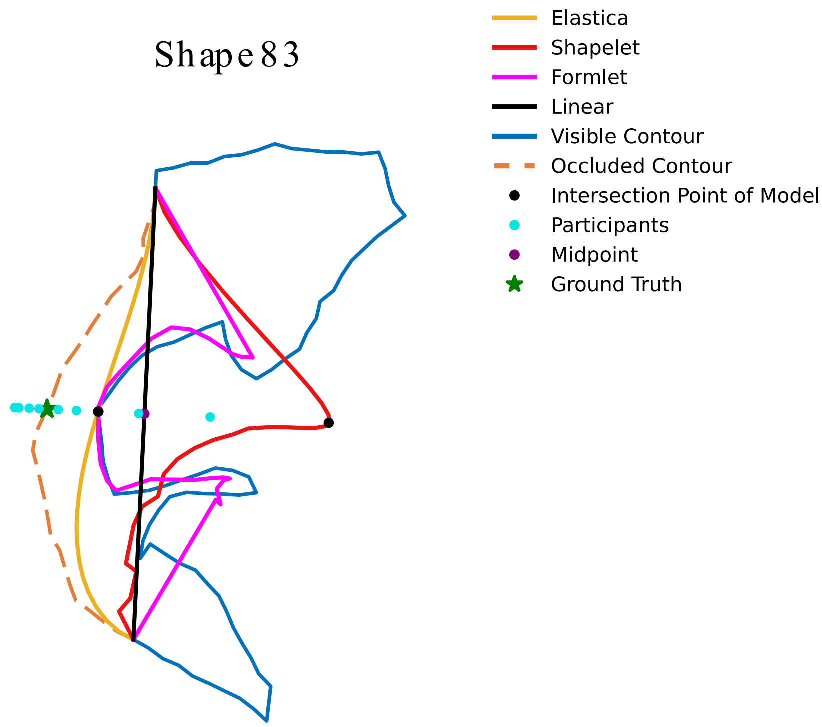

Figure 40. Outlier shapelet prediction for a synthetic shape at 25% occlusion.

Figure 43 shows the signed error for synthetic shapes at 50% gap level. The shapelet model generated an unusually large negative error whereas humans were more spread out in their estimation (Figure 44), some generating large negative errors while others making predictions that were fairly accurate. In this case, the endpoint tangents are quite predictive, and so the elastica model performs quite well.

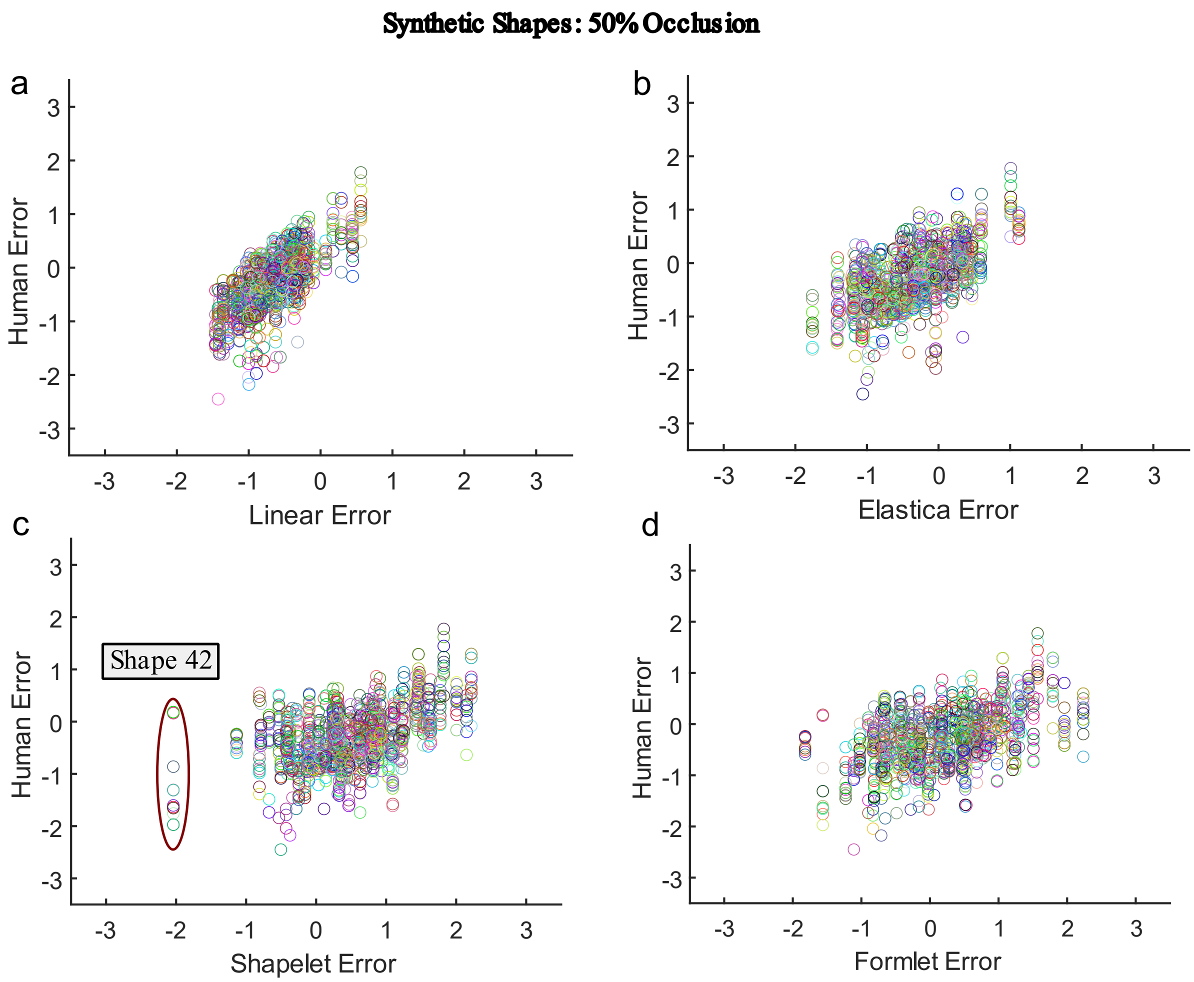

Figure 41. Scatterplot of signed deviations for synthetic shapes at 50% occlusion. Red outline: Shapelet and formlet completions, corresponding to shape 42.

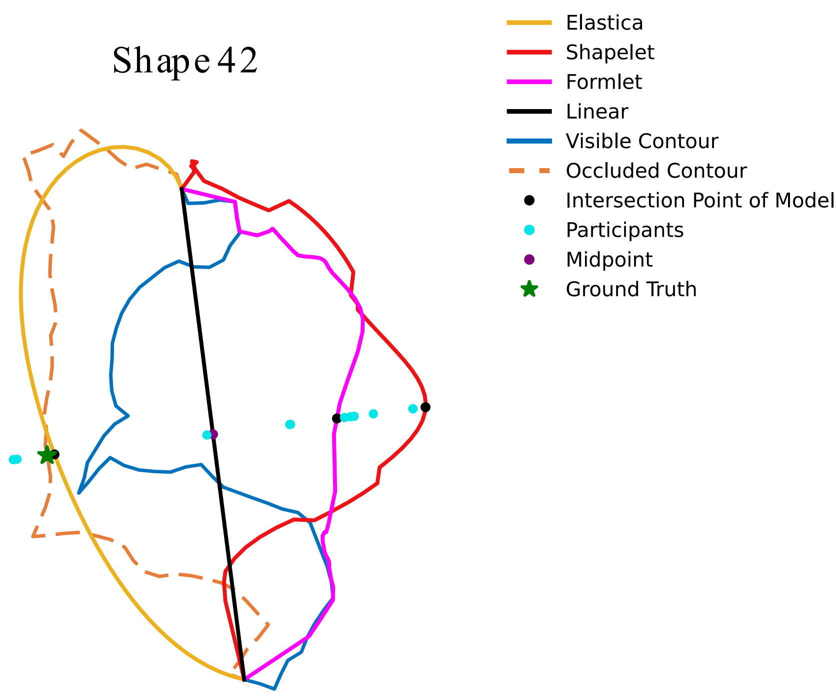

Figure 42. Outlier shapelet prediction for a synthetic shape at 50% occlusion.

Shapelet Self-Intersection

As noted previously, one disadvantage of the shapelet model is that it can generate self-intersecting shapes. To assess the potential impact on completion we computed the proportion of times self intersections were generated for 50% completion task. We found that the shapelet model generated self-occlusions for 14% of natural shapes and 19% of synthetic shapes self-intersect. Figure 45 displays the positive bias of the shapelet model.

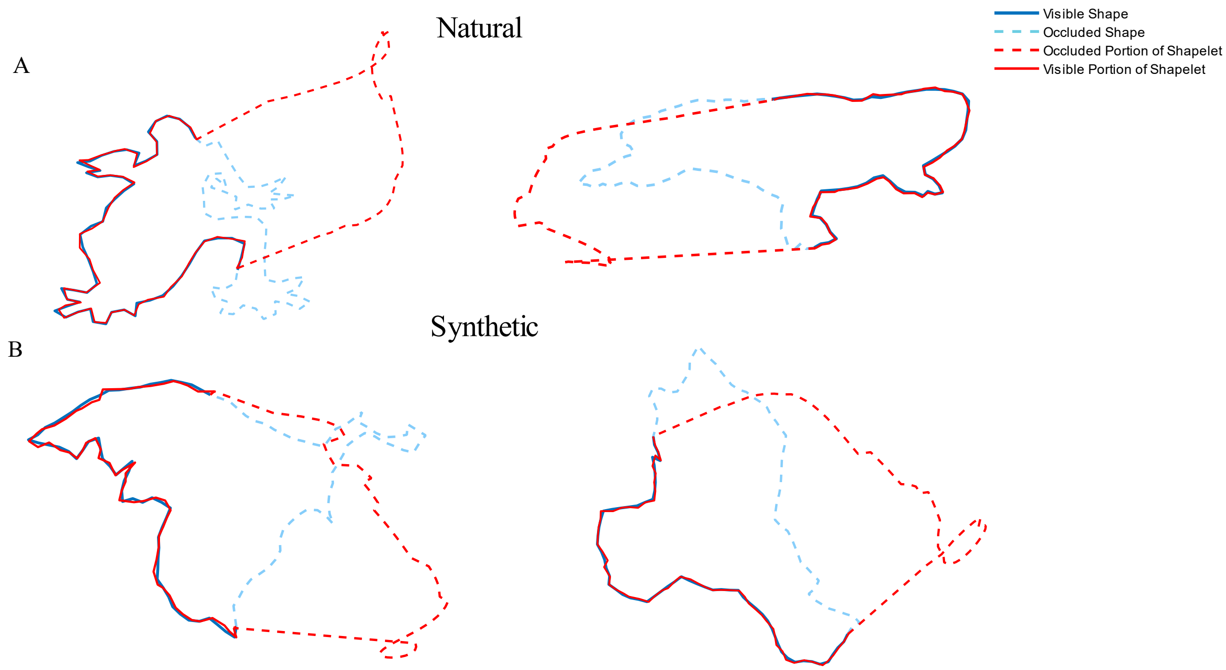

Figure 43. Example shapelet self-intersections.

CHAPTER 4

Conclusion

4.1 Summary

One of the most notable and challenging aspects of object perception is dealing with occlusions. Typically, we see parts of objects, as light reflects into our eyes from these parts, but we still perceive the objects as whole. To investigate one aspect of this perception of partly occluded objects, this thesis looked at how humans and shape models perform on a shape completion task.

Out of all the models, the one that most closely resembles human performance on the various metrics (bias, MAE and MAD), is the elastica model. This result could suggest that humans prefer smooth curves, and this coincides with prior work on shape completion whereby a bias towards smooth curves in humans as opposed to jagged or sharp angles was indicated (Geisler & Perry, 2009; Singh & Fulvio, 2005). But it still fails to capture aspects of the human data, suggesting that humans may rely on non-local or higher order cues. Notably, with 50% occlusion, where half the shape is missing, human performance surpasses that of the other models. Synthetic shapes tend to produce greater error and more pronounced bias for humans and local models, perhaps due to the higher-order regularities in natural shapes that lead to increased smoothness. Also, humans seem to be better for synthetic shapes as well. It could be that there are still some global regularities in the synthetic shapes due to the non-intersection constraint and humans are sensitive to these regularities.

It is interesting to note that formlets were more accurate than shapelets for 10% occlusion and 50% occlusion, but less accurate for 25% occlusion. This is not consistent with the Elder et al. (2013) results, where they found that Formlets were more accurate at 10% and 30% occlusion. One difference is that they measured the Hausdorff error between the entire estimated completion curve and the ground truth, whereas in my analysis, I measure the error only along the probe line.

While humans are less consistent for synthetic shapes, global models show greater consistency with humans for synthetic shapes. This could be because natural shapes have more features like legs that produce extreme formlets and shapelets. Formlets and shapelet MAE intercepts are much higher than for human and local models. This could reflect the fact that these models do not approximate the visible portion of the shape exactly.

There are several limitations of this study, notably the formlet model produces intersections, which casts doubt on the validity of the errors it generates. This could be further explored in future research. Additionally, the optimal $\boldsymbol{\lambda}$ value does not progress smoothly from shorter to longer gap levels, and the reason for this remains unclear.

Another limitation is that the participant pool consisted of CVR trainee students, and no monetary incentive was given to participate in the experiment. This could have impacted their motivation and, consequently, their performance on this completion task.

4.2 Future Directions

It is interesting that the optimal λ coefficient that balances the smoothness and length of completions does not vary monotonically with gap length. This should be further investigated in future research.

It would also be interesting to study the effect of familiarity by comparing human completion performance for upside down shapes VS right side up shapes. This kind of inversion is known to lower recognition performance. Thus, any differences in completion for upright and inverted shape stimuli would suggest an effect of semantics on completion processes.

It might also be interesting to employ a 2-Alternative Forced Choice (2AFC) paradigm rather than the method of adjustment, to determine whether the mechanics of adapting the probe interferes in any way with the judgements.

REFERENCES

Ben-Yosef, G., & Ben-Shahar, O. (2012). A Tangent Bundle Theory for Visual Curve Completion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *34*(7), 1263–1280.

Dubinskiy & Zhu. (2003). A multi-scale generative model for animate shapes and parts. *Proceedings Ninth IEEE International Conference on Computer Vision*, 249–256 vol.1.

Elder, J. H., & Zucker, S. W. (1996). Computing contour closure. In B. Buxton & R. Cipolla (Eds.), *Computer Vision—ECCV '96* (pp. 399–412). Springer.

Elder, J. H., Krupnik, A., & Johnston, L. A. (2003). Contour grouping with prior models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *25*(6), 661–674.

Elder, J. H., Oleskiw, T. D., Yakubovich, A., & Peyré, G. (2013). On growth and Formlets: Sparse multi-scale coding of planar shape. *Image and Vision Computing*, *31*(1), 1–13.

Elder, J. H. (2018). Shape from Contour: Computation and Representation. *Annual Review of Vision Science*, *4*(1), 423–450.

Elder, J. H., Oleskiw, T. D., & Fruend, I. (2018). The role of global cues in the perceptual grouping of natural shapes. *Journal of Vision*, *18*(12), 14.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field." *Vision Research*, *33*(2), 173–193.

Geisler, W. S., & Perry, J. S. (2009). Contour statistics in natural images: Grouping across occlusions. *Visual Neuroscience*, *26*(1), 109–121.

Grenander, U., Srivastava, A., & Saini, S. (2007). A pattern-theoretic characterization of biological growth. *IEEE Transactions on Medical Imaging*, *26*(5), 648–659.

The MathWorks Inc. (2022). Optimization Toolbox version: 9.4 (R2022b), Natick, Massachusetts: The MathWorks Inc. https://www.mathworks.com

Mumford, D. (1994). Elastica and Computer Vision. In *Algebraic Geometry and its Applications* (pp. 491–506). Springer, New York, NY.

Oleskiw, T., Elder, JH., Peyr´e, G. (2010). On growth and formlets: sparse multi-scale coding of planar shape. In *2010 IEEE Conference on Computer Vision and Pattern Recognition*. Los Alamitos, CA: IEEE, 459–66.

Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods.*

Python Software Foundation. Python Language Reference, version 3.7. Available at http://www.python.org.

Singh, M. (2004). Modal and Amodal Completion Generate Different Shapes. *Psychological Science, 15*(7), 454–459.

Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences*, *102*(3), 939–944.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A Century of Gestalt Psychology in Visual Perception I. Perceptual Grouping and Figure-Ground Organization. *Psychological Bulletin*, *138*(6), 1172–1217.

Yakubovich, A. (2014). *Extensions of the Formlet Model of Planar Shape*. [Master's thesis, York University].

Yakubovich, A., & Elder, J. H. (2014). Building Better Formlet Codes for Planar Shape. *2014 Canadian Conference on Computer and Robot Vision*, 84–91.