Tangent Bundle Elastica and Computer Vision

Ohad Ben-Shahar, *Senior Member, IEEE* and Guy Ben-Yosef, *Member, IEEE*

Abstract—Visual curve completion, an early visual process that completes the occluded parts between observed boundary fragments (a.k.a. inducers), is a major problem in perceptual organization and a critical step toward higher level visual tasks in both biological and machine vision. Most computational contributions to solving this problem suggest desired perceptual properties that the completed contour should satisfy *in the image plane*, and then seek the mathematical curves that provide them. Alternatively, few studies (including by the authors) have suggested to frame the problem not in the image plane but rather in the unit tangent bundle $\mathbf{R}^2 \times \mathcal{S}^1$, the space that abstracts the primary visual cortex, where curve completion allegedly occurs. Combining both schools, here we propose and develop a biologically plausible theory of *elastica in the tangent bundle* that provides not only perceptually superior completion results but also a rigorous computational prediction that inducer curvatures greatly affects the shape of the completed curve, as indeed indicated by human perception.

Index Terms—Visual completion, curve completion, tangent bundle, elastica

1 Introduction

Visual completion is a basic visual mechanism which facilitates the experience of whole objects from visual fragments. Critical in this process is the completion of occluding boundaries, contours, and other curve-like structures, a process typically known as visual curve completion. Completion behind real occlusions, also called *amodal completion*, is one form of this process, while the perception of illusory objects and subjective contours (as in the Kanizsa triangle [30]) is considered another form, often called *modal completion* [30].

Both modal and amodal completion are widely considered part of a visual process which is mostly bottom up and local (in terms of the visual information it uses). Indeed, now-classical demonstrations by Kanizsa [30, pp.84-88,97] show that perceptual results are often dominated by local continuity (or good continuation) rather than by seemingly strong top-down factors such as contextual information or previous visual experience (see Fig. 1). While it would be naive to think that higher level vision is not involved in the completion process at all [33], ample experimental and theoretical studies from the last three decades have consistently argued that it is a fundamental stage of early vision which facilitates higher level tasks like recognition (e.g., see [24], [30], [34], [35], [37], [40], [61], [63], [67], to name but a few). As such, visual curve completion has been studied from all aspects of vision sciences, including the computational, perceptual, and neurophysiological perspectives.

Suppose we are given an image region where parts of boundaries are missing due to (amodal or modal) occlusion. Suppose further that also given are two *boundary positions* between which the boundary curve should be completed (e.g., black dots in Fig. 1E). Clearly, there are infinitely many ways to complete a curve between these two positions (e.g., yellow curves in Fig. 1E), from which the visual system chooses a unique one (say, the red curve in Fig. 1E). To narrow down the set of possible sensible completions our visual system is believed to exploit additional information that is "measured" from the *visible* part of the contour around the points of occlusion. This information together with the boundary positions on both sides of the occlusion induce the completed curve, and thus often referred to as the *inducers*. The detection and localization of the inducers in the image is a task in itself, but it is left outside the scope of this paper where we assume this information to be given.

It is widely accepted that beyond their position, our visual system also uses the *boundary orientations* as part of the inducers [30], [34], [61]. (e.g., the white oriented segments in Fig. 1E). In fact, many computational solutions to the curve completion problem have focused exclusively on producing a unique curve from two given pairs of boundary positions and orientations (e.g., [5], [28], [34], [35], [44], [61], and many others). However, since some perceptual evidence suggests that the shape of completed curves are affected by yet additional inducer features (see below), we first define our computational problem without committing to any particular type of inducer information:

Problem 1. *Given two inducers, i.e., two boundary positions $q_0 = [x_0, y_0]$ and $q_1 = [x_1, y_1]$ and additional information associated with the visible part of the contour around these two points, find the shape of the perceptual curve that is completed between q_0 and q_1.*

Obviously, this definition is general and partial, as it does not formalize the "additional information" or what is the (rather illusive) perceptual outcome. Pursuing a computational theory that seeks the *perceptual* outcome is of course not new in the context of curve completion (e.g., [61, page 1], [67, page 838], or [35, page 161]), though unfortunately little was done in practice to validate any computational theories

• *The authors are with the Department of Computer Science, Ben-Gurion University of the Negev, PO Box 653, Beer-Sheva 84105, Israel. E-mail: {ben-shahar, guybeny}@cs.bgu.ac.il.*

Manuscript received 16 Oct. 2013; revised 5 June 2014; accepted 12 June 2014. Date of publication 25 July 2014; date of current version 5 Dec. 2014. Recommended for acceptance by C. Tang. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TPAMI.2014.2343214

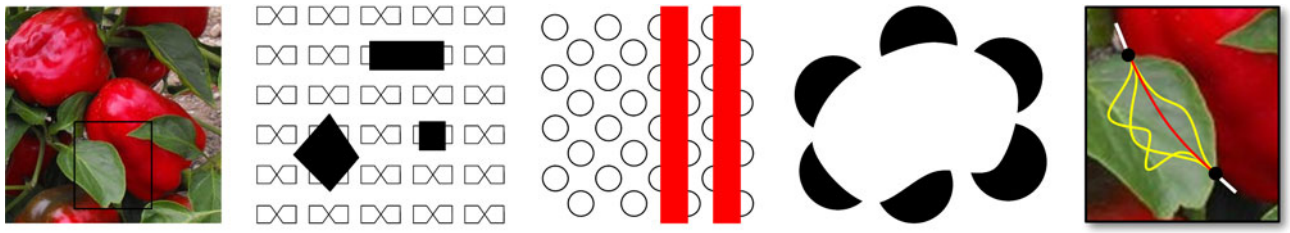

Fig. 1. (A) A natural visual scene which triggers amodal completion. (B) Reproduced from Kanizsa [30]), the gray occluders trigger the perception of rectangles, although the global regular pattern specifically suggest different objects. Here, local continuity cues override contextual information. (C) When this array of circles is occluded in part by the red rectangles, the perception below the occluders changes qualitatively to a wavy pattern, again in contrast to the strong contextual cues in the scene. (D) An example of modal completion where the visual system constructs an illusory objects whose boundaries are clearly observed even where no pictorial cues exist. (E) Illustrated using the black ROI from panel A, the abstract shape completion problem (Problem 1) requires the construction of the single *perceptual shape* (in red) between two occlusion points (in black), while considering also additional information from the visible part of the contour around these points (e.g., orientation, as marked in white). This problem is severely underconstrained, as normally there would be infinitely many curves (a few shown in yellow) that would satisfy these conditions.

against perceptual and psychophysical findings. While we further elaborate on this issue below, it is worth mentioning already at this point that outside the scope of the prior art that is relevant to our research are the numerous computer vision algorithms that fill in missing pieces in images using variational, Bayesian, or non parametric synthesis methods (e.g., [10], [12], [15]). Since they seek different goals, these algorithms are typically apathetic to perceptual results (except for seeking a visually pleasing result that has as few visual artifacts as possible), are often trained to a particular image or class of images, and are typically incapable of handling those synthetic stimuli where perception is the strongest (such as those in Figs. 1B, 1C, and 1D). Hence, except for this last short comment we will consider this type of studies outside the scope of our prior art.

2 Related Work

Ever since Ullman's original paper on computational curve completion [61], work in the field has been either axiomatic or mechanistic [4]. Axiomatic approaches define desired perceptual or geometrical properties that completed boundaries are expected to satisfy in the image (i.e., retinal) plane, and then explore these constraints in a formal fashion to derive the curve (or family of curves) that satisfy them (e.g., [8], [28], [35], [44], [53], [61], [64]). Over the years many axioms have been postulated, including isotropy of the completed curve, smoothness, total minimum curvature, extensibility, scale invariance, roundedness, total minimum change of curvature, etc. Unfortunately, while most axioms have intuitive appeal, few were motivated (or for that matter, confronted and confirmed retrospectively) by rigorous perceptual findings. Unfortunately still, some axioms contradict each other, leaving the computational community with the debate of which are the correct axioms to use.

Unlike axiomatic approaches, mechanistic models of curve completion do not hypothesize about the perceptual or geometrical properties of the completed curve but rather theorize about the operation (or principles of operation) of the mechanism that generates it (e.g., [4], [5], [11], [16], [18], [26], [47], [51], [52], [67]). Imposing assumptions on the mechanism rather than on the final result may be more difficult to analyze theoretically, but it also bears several advantages, and in particular, such approaches are readily testable for biological plausibility and they are typically simpler (in the spirit of Occam Razor) as a scientific explanation for the corresponding computation in the visual cortex.

In this paper we follow our general framework for completion in the tangent bundle [5] and suggest to combine ideas from both the axiomatic and mechanistic schools toward a new computational theory that is unique not only in its perceptually superior completion results, but also in its rigorous computational consequence that inducer curvatures are both necessary and significant in shaping the completed curve, as indeed indicated by several perceptual studies.

2.1 Elastica in Computational Vision

One celebrated geometrical axiom explored extensively in axiomatic methods is the one of minimum total curvature. Employed loosely as early as Ullman [61], studied explicitly and rigorously by Horn [28] and later by Mumford [44] (a study to which our title pays tribute), and extended in several ways by other researchers in the last two decades (e.g., [8], [53], [64]), the class of curves that satisfy this property is known as *elastica*.[1] Borrowing from mathematical mechanics, elastica models the visual curve by the shape of a thickless physical rod that minimizes its potential (or bending) energy [38] after its endpoints have been fixed at the inducers. To solve it, one typically applies calculus of variation to the corresponding functional $\mathcal{E}_{2D} = \int k(s)^2 ds$.

Requiring the completed curve to conform to elastica in the image plane is very appealing since these structures are visually and analytically smooth, they are naturally related to the Gestalt principle of good continuation [65], and they can be associated with the maximum likelihood reconstruction of a plausible stochastic mechanism [44]. At the same time, elastica curves cannot be expressed in closed form [28], [44], they violate other popular axioms (e.g., roundedness and scale invariance [53], [64]), and no clear perceptual evidence has been provided to support them as a completion model.

In our theory we indeed suggest to employ an elastica principle, but not in the 2D image plane and not as a

1. Elastica is a centuries-old mechanical problem which generalizes *Hooke's law* (the latter known mostly in the context of springs [36], [38]) and studies the shape of elastic physical structures when they undergo large-scale deflections.

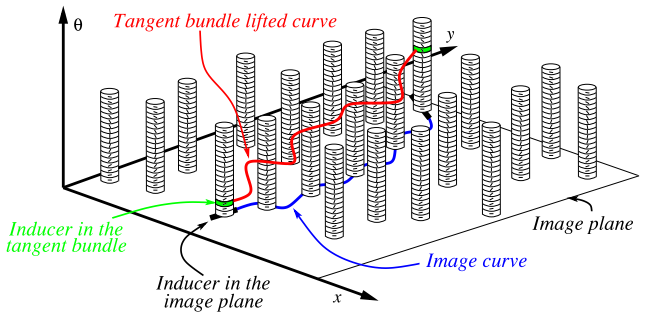

Fig. 2. The primary visual cortex can be abstracted as the space $T(I) = \mathbf{R}^2 \times \mathcal{S}^1$, the unit tangent bundle associated with the image plane I. In this space the curve completion problem requires the construction of a single continuous curve (in red) which abstracts a population of active neurons between two active inducer cells (in green). The perceived completed curve is the projection of the constructed tangent bundle curve down to the image plane (in blue).

perceptual axiom. Rather, we propose to employ it as a mechanistic principle in the visual cortical tissue that is responsible for the computation of completed contours. In formal terms, this means applying elastica not in the 2D image (retinal) plane $\mathbf{R}^2$ but in the 3D unit tangent bundle associated with the image plane, i.e., $\mathbf{R}^2 \times \mathcal{S}^1$.

2.2 The Tangent Bundle Approach to Curve Completion

The seminal work by Hubel and Wiesel [29] on the organization of the primary visual cortex (V1) around orientation hypercolumns, and the good many evidences showing that V1 cells participate in the curve completion process (e.g., [37], [40]), has led several computational vision scientists, including ourselves [5], to consider the contour completion process not in the image plane but rather *directly in the mathematical space that abstracts V1*. A proper abstraction that is well accepted is the unit tangent bundle $T(I) \triangleq \mathbf{R}^2 \times \mathcal{S}^1$ of the image plane $I = \mathbf{R}^2$ (e.g., [1], [2], [11], [27], [47], [50]), where each point represents explicitly both position and orientation, as suggested by the response of simple cell receptive fields [29]. An image curve $\alpha(t)$ (e.g., blue curve in Fig. 2) is represented in V1 as an activation pattern of all those cells that correspond to the oriented tangents along the curve's arclength [29]. Hence, in the continuous abstraction of V1 as the unit tangent bundle $T(I)$, this activation pattern forms a "lifted" curve $\beta(t)$ (e.g., red curve in Fig. 2), where both position and tangent orientation are represented explicitly along the path.

Thinking about the completion problem this way, one can clearly observe that the goal of the computational process now becomes the construction of curves between boundary *points* in $T(I)$ [5]. This immediately entails that the "additional information" from Problem 1 cannot avoid but include, at the very least, the orientation of the curve at the point of occlusion (since the simplest geometrical structure in $T(I)$, a point, represents both position and orientation in the image plane). Furthermore, working with tangent bundle representation of image curves incorporates an additional complication since these curves cannot be arbitrary. Rather, as we argued in the past, one should consider only those curves $\beta(t) = [x(t), y(t), \theta(t)] \in T(I)$ whose third coordinate corresponds to the tangential angle of their projection to I, i.e., curves $\beta(t)$ that satisfy $tan\,\theta(t) = \frac{\dot{y}(t)}{\dot{x}(t)}$ and called *admissible* [5]. Once an admissible three-dimensional curve $\beta(t)$ is constructed between the inducer points in the tangent bundle, the corresponding image curve is then obtained by properly projecting $\beta(t)$ "down" to the image plane.

The advantage of modeling the curve completion problem in $T(I)$ is not limited to considering it in a space more reminiscent of the visual area that solves this perceptual problem. It also provides the opportunity to avoid an axiomatic approach and predefined perceptual constraints on the "desired" perceived shape, and to employ physical or biological constraints on the *generation process* that reconstructs it. For example, Williams and Jacobs [67] suggested a stochastic completion process, Citti and Sarti proposed an area minimization process [11], and we explored a principle of minimum neural energy consumption expressed via the length of the constructed admissible curve [4], [5]. All these principles were applied directly in the unit tangent bundle.

In this paper we follow this general approach, but suggest a new theory that not only provides perceptually superior results, but also entails qualitatively different type of (verifiable) predictions that no previous computational work could, namely that inducer *curvatures* are an essential factor in determining the shape of the completed curve, as indeed indicated by several perceptual studies. As we shall see, this emerges from the unique application of 3D elastica (constrained via admissibility) directly in the unit tangent bundle.

3 Tangent Bundle Elastica (TBE) for Visual Curve Completion

We theorize that the completed curve between two given image inducers is the 2D image projection of the 3D admissible elastica curve between the lifted inducer points in the unit tangent bundle. Note that when it is used like so, tangent bundle elastica serves in our theory as a *mechanistic* principle (that operates in early visual cortical areas) rather than a perceptual axiom. While computational theories need not necessarily justify their elements, one might wonder if and how this principle is biologically plausible vis-a-vis the functional organization of the early visual cortex. While this issue is outside the scope of our largely theoretical and computational paper, here we nevertheless discuss it shortly before moving on to the main theoretical and formal aspects.

3.1 The Biological Plausibility of Tangent Bundle Elastica

The biological motivation to consider TBE as a mechanistic principle in early vision is derived from the *connectivity* between orientation selective cells in the primary visual cortex (e.g., [6], [22], [23], [41], [42], [48]). As mentioned already in Section 2.2, image contours (either real or completed) are represented in V1 as an activation pattern of all those cells that correspond to the oriented tangents along the curve's arclength [29]. The long range horizontal connections between these cells facilitate contextual modulations that affect their response, an assertion supported perceptually, computationally, and neuroanatomically (e.g., [17], [31],

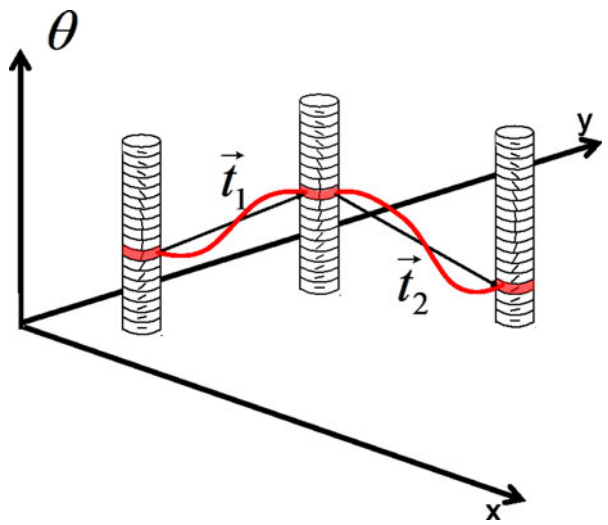

Fig. 3. The links from and to a cell in V1 can be thought of as links between points in the unit tangent bundle, hence abstracted as tangent (or difference) vectors [45, page 6] in $T(I)$.

[46]). Hence, these connections must be considered an intrinsic part of representation of contours in early vision. Here we ask if and how a certain condition on these connections can lead us to consider elastica as the mechanistic principle in the tangent bundle.

We hypothesize that for the curve completion problem the activation pattern of active cells that is formed to represent the completed cells is not optimized by the number of active contours as suggested in the past [4], [5], but rather by how *similar* are active consecutive links between cells along the activation chain (i.e., the degree to which the link to the "next" active cell in the pattern is similar to the link from the "previous" cell). As Fig. 3 alludes to, each such link can be abstracted as a *tangent vector* [45, page 6] in $T(I)$. In this way, similarity between two consecutive links along the path of cells that represent the completed curve is abstracted as the rate of change between two consecutive tangent bundle vectors (see Fig. 4). More formally, let us denote as $\vec{t}_1$ and $\vec{t}_2$ two such adjacent vectors, and let them be the tangent vectors at consecutive points $\beta(s_1)$ and $\beta(s_2)$ of some curve $\beta(s)$ in $T(I)$, respectively. The rate of change between these vectors is then

$$\frac{\vec{t}_2 - \vec{t}_1}{s_2 - s_1} \triangleq \frac{\Delta\vec{t}}{\Delta s},$$

which we require, for similarity of $\vec{t}_1$ and $\vec{t}_2$, to be as small as possible. Finally, taking $\beta(s_1)$ and $\beta(s_2)$ to be infinitesimally close to each other, i.e., at the limit as $\Delta s \to 0$, we get that our requirement becomes one on the curvature κ_β of β at s_1

$$\lim_{\Delta s \to 0} \frac{\Delta\vec{t}}{\Delta s} = \kappa_\beta \to \min.$$

Similarity between links is thus abstracted as curvature in $T(I)$, and striving for similarity along the path as much as possible is represented in the continuous case as a tangent bundle curve with total *minimum curvature*, namely *elastica in the tangent bundle*.

Is similarity between links a plausible organizational principle for early visual cortical areas? Various models have argued in favor of reciprocal and symmetrical connectivity fields (e.g., [3], [17], [46] in a way that permits colinear or cocircular contextual facilitation that runs in agreement with natural image statistics (e.g., [20]). Tangent bundle elastica takes these insights to a next level of specificity and implies that the cells that become active to represent the completed curve are those whose "incoming" and "outgoing" active connections attempt to maximize similarity along the pattern, all while maintaining admissibility of the pattern as a whole. While strict similarity is impossible for all cells simultaneously (unless the image curve is a straight line), TBE attempts to globally maximize this measure as much as possible (Fig. 4).

3.2 The Computational Problem

Hypothesizing that perceptually completed contours are admissible elastica curves in the unit tangent bundle implies that the energy of the completed curve is determined not by its tangent bundle length (i.e., by the number of active cells that represent the curve in V1) as suggested in our earlier studies [4], [5], but rather by its total *tangent bundle curvature*.[2] In other words, Problem 1 can now be described more concretely as follows:

Problem 2. *Given inducers $p_0 = [x_0, y_0, \theta_0]$ and $p_1 = [x_1, y_1, \theta_1]$, find the curve $\beta(s) = [x(s), y(s), \theta(s)] \in T(I)$ that minimizes the bending energy*

$$\mathcal{E}(\beta) = \int_{t_0}^{t_1} \kappa_\beta(s)^2 ds, \tag{1}$$

subject to admissibility $\tan\theta(t) = \frac{\dot{y}(t)}{\dot{x}(t)}$ and additional information associated with the visible part of the contour around the inducers.

At this point we still leave open the specification of the "additional information" but as we shall see shortly, the mere requirement for elastica will dictate certain boundary conditions that imply not only inducers' position and orientation in the image, but also their curvature. To get there, however, one first needs to realize that the curvature κ_β in the functional is *not* the curvature of the perceptual curve in the image (retinal) plane, but rather the curvature of the lifted curve in the unit tangent bundle. Not only that this curvature is not "visual", it is also a curvature of a 3D space curve (as opposed to a plane curve). Hence, to facilitate the analysis of tangent bundle elastica we first study the meaning of this abstract property. In particular, to assess the stationary (minimum) point of the functional in Eq. (1), it is desirable to describe the curvature κ_β in terms of the curvature of the image curve α.

3.3 The Curvature of a Tangent Bundle Curve

Let $\alpha(s) = [x(s), y(s)]$ be an arclength parameterized curve in I, and let

2. We note that elastica in 3D space is the combination of a bending energy term (proportional to the curvature squared) and a twisting energy term (proportional to the torsion squared) [36, pp. 62, 70]. In this paper we consider only the former, though an extension is possible and left for future work.

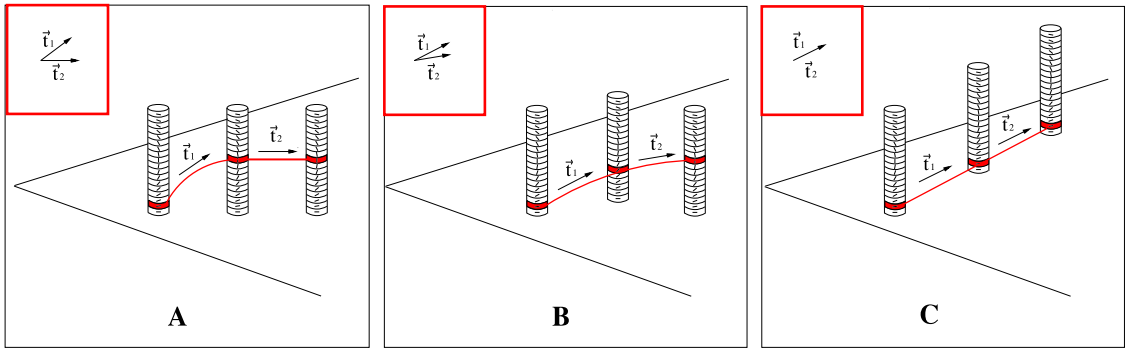

Fig. 4. Connection similarity as a possible substrate for elastica in the tangent bundle. With the abstraction of each connection as a tangent bundle tangent vector [45, page 6], the TBE mechanism hypothesizes that each two consecutive links $\vec{t}_1$ and $\vec{t}_2$ along the activation pattern that represents the completed curve strive to be as similar as possible, i.e., such that their corresponding tangent vectors in the tangent bundle (see insets) are as least different as possible. In the limit, this rate of change becomes the curvature (or bending energy) of the activation pattern as a whole.

$$\beta(s) = [x(s), y(s), \theta(s)] \qquad \theta(s) = \tan^{-1}\frac{\dot{y}(s)}{\dot{x}(s)}$$

be its corresponding lifted curve in $\mathbf{R}^2 \times \mathcal{S}^1$, where $\tan^{-1}$ is the 4 quadrant arctangent of its argument. To make the curve's dimensions commensurable, we follow our previous work [5] and incorporate a proportionality constant $\hbar$ into the description of β

$$\beta(s) = [x(s), y(s), \hbar\theta(s)]. \tag{2}$$

Since both α and β are parameterized by s, we can rewrite β as

$$\beta(s) = x(s)\hat{x} + y(s)\hat{y} + \hbar\theta(s)\hat{z} = \alpha(s) + \hbar\theta(s)\hat{z}, \tag{3}$$

where $\hat{z}$ is the orientation basis vector in $T(I)$ that is orthogonal to the spatial coordinates $\hat{x}$ and $\hat{y}$. Since this regular curve is not necessarily parametrized by arclength, its curvature function (using Newton's notation for differentiation) is given by

$$\kappa_\beta(t) = \frac{|\dot{\beta} \times \ddot{\beta}|}{|\dot{\beta}|^3} \ . \tag{4}$$

Having these notations, the following proposition describes the relationship between the curvature of β and the curvature of α:

Proposition 1. *Let κ_α be the curvature of some arclength-parametrized curve $\alpha \in I$, and let κ_β be the curvature of its corresponding lifted curve $\beta \in T(I)$. Then*

$$\kappa_\beta^2 = \frac{\kappa_\alpha^2 + \hbar^2 {\kappa_\alpha}^4 + \hbar^2 \dot{\kappa_\alpha}^2}{\left(1 + \hbar^2\kappa_\alpha^2\right)^3} \ , \tag{5}$$

where $\dot{x}$ represents differentiation with respect to the arclength of α.

Proof. Deriving Eq. (3) and considering the numerator of Eq. (4) we obtain

$$\begin{aligned}\dot{\beta} \times \ddot{\beta} &= (\dot{\alpha} + \hbar\dot{\theta}\hat{z}) \times (\ddot{\alpha} + \hbar\ddot{\theta}\hat{z}) \\ &= (\dot{\alpha} \times \ddot{\alpha}) + \hbar\ddot{\theta}(\dot{\alpha} \times \hat{z}) + \hbar\dot{\theta}(\hat{z} \times \ddot{\alpha}) + \hbar^2\dot{\theta}\ddot{\theta}(\hat{z} \times \hat{z}).\end{aligned}$$

Note that $\dot{\alpha}, \ddot{\alpha}$ are planar on XY, and thus perpendicular to $\hat{z}$. Furthermore, since α in arclength parametrization $\dot{\alpha}, \ddot{\alpha}$ are also perpendicular to each other, thus $<\dot{\alpha}, \ddot{\alpha}, \hat{z}>$ forms an orthogonal coordinate system (i.e., a frame). Hence the magnitude of the product vector in this formulation becomes rather simple:

$$\begin{aligned}|\dot{\beta} \times \ddot{\beta}|^2 &= (\dot{\beta} \times \ddot{\beta}) \cdot (\dot{\beta} \times \ddot{\beta}) \\ &= |\dot{\alpha} \times \ddot{\alpha}|^2 + \hbar^2\ddot{\theta}^2|\dot{\alpha}|^2 + \hbar^2\dot{\theta}^2|\ddot{\alpha}|^2 \ .\end{aligned} \tag{6}$$

Now, since α is in arclength parametrization, we obtain

$$\begin{aligned}\dot{\theta} &= \kappa_\alpha \\ \ddot{\theta} &= \dot{\kappa}_\alpha \\ |\dot{\alpha}|^2 &= 1 \\ |\ddot{\alpha}|^2 &= \kappa_\alpha^2 \\ |\dot{\alpha} \times \ddot{\alpha}|^2 &= \kappa_\alpha^2\end{aligned}$$

(the last row is an immediate result of Eq. (4)), and we get that

$$|\dot{\beta} \times \ddot{\beta}|^2 = \kappa_\alpha^2 + \hbar^2\dot{\kappa}_\alpha^2 + \hbar^2\kappa_\alpha^4 \tag{7}$$

and

$$|\dot{\beta}|^2 = |\dot{\alpha}|^2 + \hbar^2\dot{\theta}^2 = 1 + \hbar^2\kappa_\alpha^2 \ , \tag{8}$$

Substituting Eqs. (7) and (8) into Eq. (4) finalizes our proof. □

3.4 Elastica Completion in the Tangent Bundle

With the curvature of tangent bundle curves spelled out, we are now able to define formally the curve completion problem according to our tangent bundle elastica principle. We first observe that elastica in space, as elastica in the plane, demands *Cauchy* boundary conditions (as opposed to the *Dirichlet* boundary conditions that are required for models involving just length [5]). In other words, to solve for specific elastica curve one needs boundary *derivatives* in addition to boundary positions [28], [36], [38]. Since the derivatives of tangent bundle curves involve the quantity $\dot{\theta} = \kappa_\alpha$, by construction we must also supply the image plane *curvature* of the inducers in order to solve for the unique TBE curve between them. Hence, the "additional information" in Problem 2 must now include not only inducer *orientation* (which is already incorporated in p_0 and p_1) but also inducer *curvature*. Interestingly, this very first

(and somewhat technically driven) outcome of the TBE theory matches important findings from the perceptual literature [55], [57], [58]. We discuss this aspect of our theory in more details later on.

Based on the above, and expressing all quantities (energy functional and inducer properties) in the image plane (in particular, by substituting Eq. (5) into Eq. (1)), our curve completion via TBE (i.e., Problem 2) now becomes as follows:

Problem 3. *Find the curve* $\beta(s) = [x(s), y(s), \theta(s)] : [0, l] \rightarrow T(I)$ *that minimizes the tangent bundle bending energy*

$$\mathcal{E}(\beta) = \int_0^l \frac{\kappa_\alpha^2 + \hbar^2 {\kappa_\alpha}^4 + \hbar^2 \dot{\kappa_\alpha}^2}{\left(1 + \hbar^2 \kappa_\alpha^2\right)^3} ds \tag{9}$$

subject to the admissibility constraint $\tan\theta(t) = \frac{\dot{y}(t)}{\dot{x}(t)}$ *and the boundary conditions (i.e., inducers)* $\beta(0) = [x_0, y_0, \theta_0]$, $\kappa_\alpha(0) = k_0$, *and* $\beta(l) = [x_1, y_1, \theta_1]$, $\kappa_\alpha(l) = k_1$.

Note that in this formulation s is the arclength parameter of the corresponding image curve α and that l is its length. Importantly, l is unknown and the search for an optimal curve should be done among curves of all possible lengths. With this (now fully formalized) problem description, the following is our main theoretical result of the TBE theory:

Theorem 1. *Of all admissible curves in* $T(I)$ *that project to non-inflectional image curves, those that minimize the functional in Eq. (9) belong to a two parameter* (c, ϕ) *family defined by the following differential equation:*

$$\begin{aligned}\ddot{\kappa_\alpha} = &\frac{\kappa_\alpha^2 - 2\hbar^2\kappa_\alpha^4 - 3\hbar^4\kappa_\alpha^6 + \hbar^2\dot{\kappa_\alpha}^2 + 7\hbar^4\kappa_\alpha^2\dot{\kappa_\alpha}^2}{2\hbar^2\kappa_\alpha\left(1 + \hbar^2\kappa_\alpha^2\right)} \\ &- \frac{c\sin(\theta + \phi)\left(1 + \hbar^2\kappa_\alpha^2\right)^3}{2\hbar^2\kappa_\alpha}.\end{aligned} \tag{10}$$

Proof. Since l, the total length of α, is unknown, a further change of representation to our elastica functional is needed. Since α has no inflection points, the orientation $\theta(s)$ of α with respect to its arclength is monotonic. Without loss of generality let us assume that $\theta(s)$ is monotonically increasing. With this, θ may serve as a parameter of the integral using the following two identities:

$$\begin{aligned}\kappa_\alpha &= \frac{d\theta}{ds} \\ \dot{\kappa_\alpha} &= \frac{d\kappa_\alpha}{ds} = \frac{d\kappa_\alpha}{d\theta} \cdot \frac{d\theta}{ds} = \acute{\kappa_\alpha} \cdot \kappa_\alpha,\end{aligned}$$

where we denote by $\acute{X}$ differentiation with respect to θ. Using these identities, Eq. (9) becomes

$$\mathcal{E}(\beta) = \int_{\theta_0}^{\theta_1} \left[\frac{\kappa_\alpha + \hbar^2{\kappa_\alpha}^3 + \hbar^2\acute{\kappa_\alpha}^2\kappa_\alpha}{\left(1 + \hbar^2\kappa_\alpha^2\right)^3}\right] d\theta. \tag{11}$$

By using this form to describe the curve, we are at risk of ignoring the boundary conditions on positions $[x_0, y_0]$ and $[x_1, y_1]$ that must be introduced back into the problem.[3] This can be done by adding certain constraints that force the projection of the induced curve to pass through these two end points. For that, note that one can derive the following identities for the "span" of the end points:

$$\begin{aligned}\Delta x \triangleq x_1 - x_0 &= \int_0^l \dot{x} = \int_{\theta_0}^{\theta_1} \frac{\cos\theta}{\kappa_\alpha} d\theta \\ \Delta y \triangleq y_1 - y_0 &= \int_0^l \dot{y} = \int_{\theta_0}^{\theta_1} \frac{\sin\theta}{\kappa_\alpha} d\theta,\end{aligned}$$

from which we can rewrite the following integral constraints on the desired κ_α and θ functions:

$$\begin{aligned}\int_{\theta_0}^{\theta_1} \left[\frac{\cos\theta}{\kappa_\alpha} - \frac{\Delta x}{\Delta\theta}\right] d\theta = 0 \\ \int_{\theta_0}^{\theta_1} \left[\frac{\sin\theta}{\kappa_\alpha} - \frac{\Delta y}{\Delta\theta}\right] d\theta = 0,\end{aligned}$$

where $\Delta\theta = \theta_1 - \theta_0$. These additional constraints can now be incorporated in our new functional in Eq. (11) using two arbitrary Lagrange multipliers λ_x and λ_y. The result is the following minimization problem in terms of θ:

$$\begin{aligned}\mathcal{E}(\beta) = \int_{\theta_0}^{\theta_1} \Bigg[&\frac{\kappa_\alpha + \hbar^2{\kappa_\alpha}^3 + \hbar^2\acute{\kappa_\alpha}^2\kappa_\alpha}{\left(1 + \hbar^2\kappa_\alpha^2\right)^3} \\ &+ \lambda_x\left(\frac{\cos\theta}{\kappa_\alpha} - \frac{\Delta x}{\Delta\theta}\right) + \lambda_y\left(\frac{\sin\theta}{\kappa_\alpha} - \frac{\Delta y}{\Delta\theta}\right)\Bigg] d\theta.\end{aligned} \tag{12}$$

Note that although θ is measured in radians, $\tan\theta = \frac{\dot{y}}{\dot{x}}$, $\cos\theta$, and $\sin\theta$ are dimensionless and therefore do not unbalance the units in the functional.

Eq. (12) is a functional of the type $\int F[\theta, \kappa_\alpha(\theta), \acute{\kappa_\alpha}(\theta)]$. Thus, the application of the Euler-Lagrange equation (EL) becomes

$$\frac{d}{d\theta}\left(\frac{\partial F}{\partial \acute{\kappa_\alpha}}\right) = \frac{\partial F}{\partial \kappa_\alpha} \tag{13}$$

while,

$$\begin{aligned}\frac{d}{d\theta}\left(\frac{\partial F}{\partial \acute{\kappa_\alpha}}\right) &= \frac{2\hbar^2(1 + \hbar^2\kappa_\alpha^2)\kappa_\alpha\acute{\kappa_\alpha} + 2\hbar^2\acute{\kappa_\alpha}^2 - 10\hbar^4\kappa_\alpha^2\acute{\kappa_\alpha}^2}{\left(1 + \hbar^2\kappa_\alpha^2\right)^4} \\ \frac{\partial F}{\partial \kappa_\alpha} &= \frac{1 + \hbar^2\acute{\kappa_\alpha}^2 - 2\hbar^2\kappa_\alpha^2 - 5\hbar^4\kappa_\alpha^2\acute{\kappa_\alpha}^2 - 3\hbar^4\kappa_\alpha^4}{\left(1 + \hbar^2\kappa_\alpha^2\right)^4} - \frac{c\sin(\theta + \phi)}{\kappa_\alpha^2},\end{aligned} \tag{14}$$

where the symbols c and ϕ rename the free parameters such that $\lambda_x = c\sin\phi$ and $\lambda_y = c\cos\phi$. Putting these terms in the EL equation and simplifying we get

$$\begin{aligned}\acute{\kappa_\alpha} = &\frac{1 - 2\hbar^2\kappa_\alpha^2 - \hbar^2\acute{\kappa_\alpha}^2 + 5\hbar^4\kappa_\alpha^2\acute{\kappa_\alpha}^2 - 3\hbar^4\kappa_\alpha^4}{2\hbar^2\kappa_\alpha\left(1 + \hbar^2\kappa_\alpha^2\right)} \\ &- \frac{c\sin(\theta + \phi)\left(1 + \hbar^2\kappa_\alpha^2\right)^3}{2\hbar^2\kappa_\alpha^3}.\end{aligned} \tag{15}$$

Finally, it is left to substitute back to $\acute{\kappa_\alpha} = \frac{\dot{\kappa_\alpha}}{\dot{\kappa_\alpha}}$, and to

$$\acute{\kappa_\alpha} = \frac{d\acute{\kappa_\alpha}}{d\theta} = \frac{d\acute{\kappa_\alpha}}{ds} \Big/ \frac{d\theta}{ds} = \frac{1}{\kappa_\alpha}\frac{d}{ds}\left(\frac{\kappa_\alpha}{\dot{\kappa_\alpha}}\right) = \frac{\ddot{\kappa_\alpha}\kappa_\alpha - \dot{\kappa_\alpha}^2}{\kappa_\alpha^3} \tag{16}$$

in order to get Eq. (10). □

3. Note that the initial conditions on the orientation of the two inducers become very explicit in this form and are embedded directly in the limits of the integral in Eq. (11).

We conjecture that the same property expressed in Eq. (10) is valid for general inflectional curves as well. Indeed, if one assumes that the admissible TBE curve β between the two endpoints in the tangent bundle projects to an image curve α with a finite number of $n-1$ inflection points $p_1, p_2, \ldots, p_{n-1}$ (this is a reasonable assumption since we do not expect the optimal curve to wiggle too much, certainly not an infinite number of times), then α can be thought of as a chain of n non-inflectional curve segments, each of which is described by Eq. (10) and the proper boundary conditions (at p_i and p_{i+1}), with a possible change of the sign of κ_α between them. Hence, by Theorem 1, Eq. (10) is valid to all curve segments and thus to the entire length of the inflectional curve as a whole. This conjecture is further strengthened by our empirical results below.

Eq. 10 is a highly non linear second order differential equation that is unlikely to assume a closed from analytical solution. Still, this equation, which we denote as the *tangent bundle elastica differential equation* (TBE-DE), constrains the behavior of image curvature of admissible tangent bundle elastica curves, hence, may be used to solve for these curves numerically. Next we explore this possibility.

3.5 Numerical Solution

To solve Eq. 10 numerically we apply a nonlinear optimization process for solving ordinary differential equations (ODE) with boundary value conditions. In particular, we numerically seek the values of the parameters c and ϕ such that the resultant curve satisfies both the TBE-DE and the position, orientation, and curvature of the inducers with which it should meet. To do so, define the auxiliary function $g(s)$ as the derivative of image curvature with respect to arclength, i.e., $g(s) \triangleq \dot{\kappa_\alpha} = \frac{d\kappa_\alpha}{ds}$. Initializing the arclength parameter to $s_0 = 0$ and starting at point $p_0 = [x_0, y_0, \theta_0]$ and the given initial curvature $\kappa(0) = k_0$, we apply an Euler-based iteration as follows:

$$
\begin{aligned}
&s_{n+1} \triangleq s_n + \delta \\
&g_{n+1}(s_{n+1}) = g(s_n + \delta) \approx g(s_n) + \delta \cdot \dot{g}(s_n) \\
&\quad = g(s_n) + \delta \cdot \ddot{\kappa}(s_n) \\
&\quad = g_n + \delta \cdot \left(\frac{\kappa_n^2 - 2\hbar^2\kappa_n^4 - 3\hbar^4\kappa_n^6 + \hbar^2 g_n^2 + 7\hbar^4\kappa_n^2 g_n^2}{2\hbar^2\kappa_n\left(1+\hbar^2\kappa_n^2\right)} \right. \\
&\qquad \left. - \frac{c\sin(\theta+\phi)\left(1+\hbar^2\kappa_n^2\right)^3}{2\hbar^2\kappa_n} \right). \\
&\kappa_{n+1} \triangleq \kappa(s_{n+1}) = \kappa(s_n + \delta) \approx \kappa(s_n) + \delta \cdot \dot{\kappa}(s_n) \\
&\quad = \kappa_n + \delta \cdot g_n. \\
&\theta_{n+1} \triangleq \theta(s_{n+1}) = \theta(s_n + \delta) \approx \theta(s_n) + \delta \cdot \kappa(s_n) \\
&\quad = \theta_n + \delta \cdot \kappa_n. \\
&y_{n+1}(s_{n+1}) \approx y(s_n) + \delta \cdot \dot{y}(s_n) \\
&\quad = y_n + \delta \cdot \sin\theta_n. \\
&x_{n+1}(s_{n+1}) \approx x(s_n) + \delta \cdot \dot{x}(s_n) \\
&\quad = x_n + \delta \cdot \cos\theta_n,
\end{aligned}
\tag{17}
$$

where δ is a pre-selected step size and the approximation obtained is of order $O(\delta)$. The curve $\beta_i(x) = [x(s_i), y(s_i), \theta(s_i)]$ computed by this step is then evaluated at $s_n = l$, where l is defined as the total length of the projection of β in I, (i.e., after $n = l/\delta$ iterations) to obtain the point $[x_{end}, y_{end}, \theta_{end}] = [x(l), y(l), \theta(l)]$. Importantly, since Eq. (17) solves for $\dot{\kappa_\alpha}$ (the *derivative* of κ_α) based on $\ddot{\kappa_\alpha}$ from Eq. (10), the process is indifferent to the emergence of inflection points (which occur when κ_α changes sign during the iteration) and permits their existence in the resultant completed curve without any additional or ad-hoc computational steps.

Clearly, the construction of a curve from the family specified by Eq. (10) requires a starting point in $T(I)$ and an initial curvature, as well as knowledge of the unknown parameters c, ϕ, g_0, and l. However, since the starting point is given as $p_0 = [x_0, y_0, \theta_0]$ endowed with k_0, we are required to match the unknown parameters with the rest of the four boundary conditions $p_1 = [x_1, y_1, \theta_1]$ and k_1. This can be done via gradient descent as follows:

1) Guess initial values for the parameters c, ϕ, g_0, and l.
2) Apply Eq. (17) to construct a tangent bundle curve of length l starting from $p_0 = [x_0, y_0, \theta_0]$ and boundary curvature k_0.
3) Evaluate the correctness of the parameters by assessing the error between the four-tuple $[x(l), y(l), \theta(l), \kappa(l)]$ (i.e., the *constructed* end point and its curvature) and $[x_1, y_1, \theta_1, k_1]$ (i.e., the *desired* end point and its curvature):

$$E(c, \phi, g_0, l) \triangleq \|[x_1, y_1, \theta_1, \kappa_1] - [x(l), y(l), \theta(l), \kappa(l)]\|^2$$

4) Improve the guess for c, ϕ, g_0, and l via gradient descent on E, and iterate back to step 2.

Convergence of this optimization process guarantees a correct curve. Like most nonlinear processes, however, this depends on a "good" initial guess. In our implementation the initial guess for c, ϕ, g_0, and l was decided via a quick and coarse brute force sampling of the search domain for minimum, where each of the four axes is sampled between 5 to 10 values. While finding a better way to make an initial guess is left for future research, we note that our simple search approach has led to excellent results in virtually all cases tested.

4 Perceptual Implications and Predictions

Unlike axiomatic approaches, mechanistic models do not make any prior assumption on the perceptual shape that the visual system completes between inducers. It is therefore expected that one would be able to generate perceptual predictions from the mechanism. While a comprehensive psychophysical evaluation of the TBE theory is forthcoming, in what follows we derive such predictions in the context of the axioms commonly used in axiomatic models. Indeed, several of these are straight forward to make. For example, we immediately observe that since the solution for elastica in the tangent bundle is not linked to any specific frame, it is clearly *isotropic*, i.e. invariant to rigid transformations in the image plane. Since the completed curves are solutions of a differential equation (Eq. (10)), they clearly possess *smoothness*, though this property is already implicit in the need to have a well defined curvature along the tangent bundle curve β, and therefore a well defined curvature *derivative* along the corresponding image curve α (cf. Eq. (5)). Furthermore, since the model minimizes total curvature in $T(I)$, it must be *extensible* [61] in that space and hence in the image plane also, i.e.,

any two arbitrary inducers along a completed curve α would generate the same shape as the shape of the portion of α connecting them. Naturally, here we extend Ullman's original notion of extensibility to also include boundary curvatures.

While smoothness, isotropy, and extensibility are also found in many previous models, the signature behavior of TBE is revealed in the context of three other perceptual properties.

4.1 Roundedness?

A very popular property that has been sought-after in many axiomatic models of contour completion is roundedness, i.e., the assertion that cocircular inducers are perceptually completed by a circular arc (e.g., [8], [35], [61]). To test if tangent bundle elastica predicts roundedness, we examine if and when Eq. (10) is satisfied for $\kappa_\alpha = const$.

Substituting $\kappa_\alpha = K$, $\dot{\kappa_\alpha} = 0$, and $\ddot{\kappa_\alpha} = 0$ in Eq. (10) yields

$$0 = \frac{K^2 - 2\hbar^2K^4 - 3\hbar^4K^6 + 0 + 0}{2\hbar^2K(1+\hbar^2K^2)} - \frac{c\sin(\theta+\phi)(1+\hbar^2K^2)^3}{2\hbar^2K} . \quad (18)$$

This equation can be satisfied only for $c = 0$ (or else θ must be constant, which contradicts the desired behavior along circular arcs), and hence the roundedness condition reduces to a vanishing numerator of the left component, i.e.,

$$0 = K^2 - 2\hbar^2K^4 - 3\hbar^4K^6 = K^2(1 - 2\hbar^2K^2 - 3\hbar^4K^4) . \quad (19)$$

By solving this equation we get that circular paths are possible only for $K = 0$ (i.e., for straight lines) or for

$$K = \pm\frac{1}{\sqrt{3}\hbar} . \quad (20)$$

Note that this very specific curvature value is associated with both the inducers and the (circular) completed curve as a whole.

We conclude that in general, even if *all* boundary conditions (i.e., all of inducer positions, orientations, and curvatures) correspond to the same circle, TBE would *not* generate a circular arc between the inducers except for a singular case.[4] In other words, in general our new model does *not* satisfy the axiom of roundedness, a prediction that fits nicely with several psychophysical findings (e.g., [19], [25]). Examples of non-circular and (the singular) circular paths are shown in Fig. 5.

4.2 Scale Invariance?

Perhaps even more popular than roundedness is the axiom of scale invariance, i.e., the assumption that perceptually completed contours are invariant to viewing distance (e.g., [8], [35], [64]). Unfortunately, this property is no longer considered desired as perceptual studies concluded that it does not hold for human vision [16], [19], [21], [54], [56]. To investigate how our model behaves with change of scale we first recall that the classical meaning of changing scale of a curve completion problem involves changing the inducers' distance but *not* their orientation. Since our new formulation also involves image curvature, changing scale now requires changing the boundary curvatures as well. Indeed, if we now think of the inducer as a thickless circular arc (as opposed to an oriented segment in classical formulation), then increasing viewing distance decreases the observed radius of curvature and hence increases the observed curvature of that inducer (and vice versa). Hence, to investigate scale dependency we must examine the model's output for two pairs of inducers that are scaled properly in both their distance and curvatures (while keeping their orientation fixed).

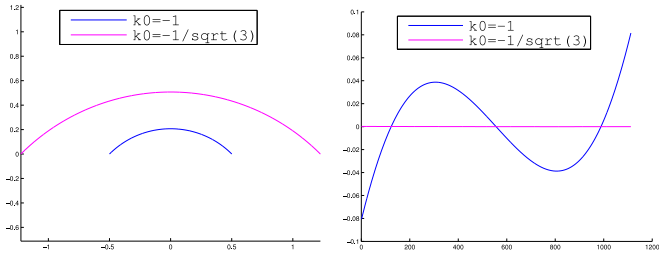

Fig. 5. On the left are two completed curves for two pairs of co-circular inducers at two scales. Blue curve corresponds to $p_{0,1} = [\mp 0.5, 0, \pm 45°], k_{0,1} = -1/\sqrt{2}$. Magenta corresponds to $p_{0,1} = [\mp\sqrt{6}/2,\ 0, \pm 45°]$, $k_{0,1} = -1/\sqrt{3}$. If TBE was scale invariant, the two completed curves should be scaled versions of each other. However, their rate of change of curvature $g(s)$ plotted on the right reveals circular behavior in one case (which happens to be the corresponding singular case from Eq. (20)) but not in the other. Hence, the model provides neither scale invariance nor roundedness, as indeed reported in psychophysical experiments.

While such direct theoretical examination of scale invariance is possible, the singular behavior of *circular* completions from Section 4.1 immediately proves that tangent bundle elastica is *not* scale invariant. Indeed, Eq. (20) suggests that a circular completion is possible for a singular value of inducer curvatures. Since scaling changes inducer curvature, the completion of a scaled version of the singular circular inducers no longer satisfies Eq. (20) and therefore is not circular (see Fig. 5). Importantly, this scale-dependent behavior is consistent with human perception not only by the mere fact that it is scale dependent. Indeed, as Fig. 6 shows, in agreement with the accumulating psychophysical findings (e.g., [16], [21]), TBE predicts flatter completions for larger scales. Although this qualitative consistency strengthen the validity of our theory, it should be taken with care since unlike our theory, all perceptual work on scale invariance (or lack thereof) in curve completion is yet to take into account inducer curvature. Clearly, a final confirmation will require quantitative evaluation against human data, once the latter become available.

4.3 Perceptual Sensitivity to Inducer Curvature

As mentioned in our background section, virtually all computational curve completion research has focused on inducers defined by the position and orientation of the

4. The fact that this value K depends on $\hbar$ provides a unique opportunity for perceptual calibration of the latter. Indeed, Since $\hbar$ represents the relationship between the spatial and angular axes in the unit tangent bundle that abstracts the primary visual cortex, it is in fact a constant of the system that should be determined and calibrated once and for all for each observer (or perhaps for all observers, though this is less likely). But how? The fact that the theory predicts a unique case where completions are circular (Eq. (20)) implies that findings this case (via perceptual measurements) will reveal the value of $\hbar$.

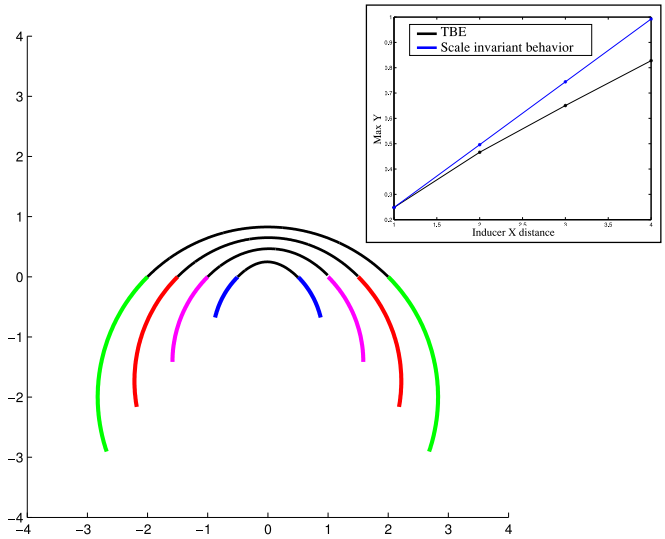

Fig. 6. TBE is scale dependent in a way that is qualitatively consistent with perceptual finding. Shown here are a pair of inducers whose distance and curvatures are scaled for different viewing distances. The results indicate flatter completions with increasing scale, as indicated by perceptual findings. The inset compares the TBE behavior (black) compared to the one predicted by scale-invariant (blue) one, in terms of the peak height as a function of inducers distance. The flatter TBE completions for increased scale are indicated by the progressively lower peaks.

observed part of the contour at the point of occlusion (Notable exceptions are the axiomatic theory due to Takeichi et al. [59] and the mechanistic model by Fantoni and Gerbino [16]). Although doing so follows much of the perceptual literature (e.g., [19], [25], [30], [34]), it unfortunately fails to explain our perceptual experience since often inducer curvature appears to affect the shape of completed contours (Figs. 7A and 7B). Psychophysical reports of this effect are not prevalent but have been accumulating slowly, both directly [58] and indirectly [55]. Although in our theory it emerges as a technical property (the need for *Cauchy* boundary conditions), such an effect is explicitly predicted by our tangent bundle elastica model. Computationally, the smoothness of the solution dictates that boundary curvatures would determine how quickly the orientation of the completed curve would deviate from inducer orientation near the occlusion point, and therefore predicts that smaller inducer curvature (in absolute value) will induce "sharper" completions. Fig. 7C exemplifies such predictions graphically.

5 Experimental Results

We applied our TBE theory and numerical procedure to various synthetic and natural scenes. Fig. 8 illustrates several tangent bundle elastica completions for various pairs of inducers. We emphasize again that the information now carried by inducers includes not only position and orientation at the point of occlusion, but boundary curvature also, hence we depict them as short circular arcs with the corresponding radius of curvature. A more explicit demonstration of the effect of inducer curvature is shown in Fig. 9 using classical synthetic images employed in previous psychophysical studies (e.g., by Takeichi [58]). A selected range of shapes induced by changing only the curvature of the inducers was already shown in Fig. 7C.

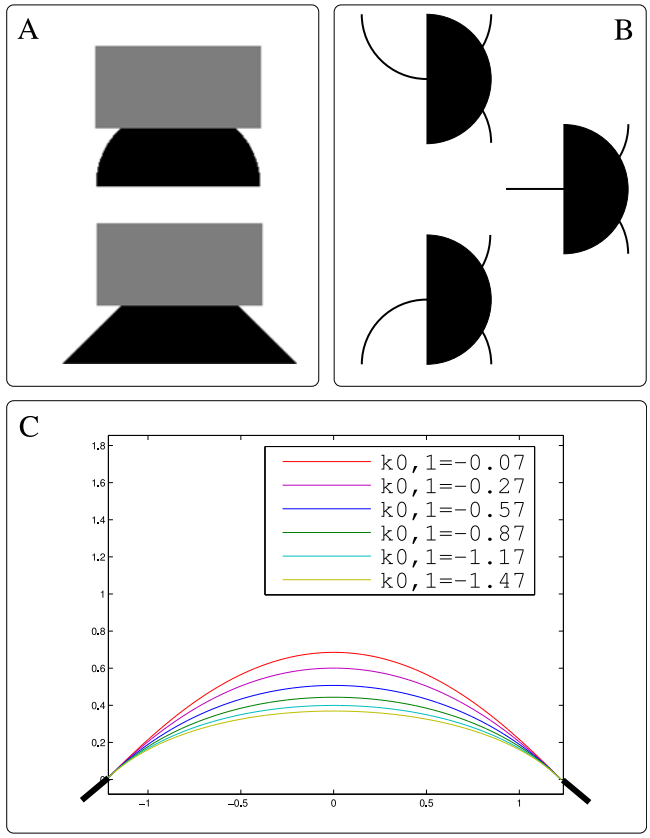

Fig. 7. The perceptual and computational effect of boundary curvatures. **(A)** Despite having identical inducers in terms of position and orientation, human observers tend to perceive different completions in these stimuli. **(B)** The curve penetrating the half disk occluder from the left appears to continue in different ways in these three stimuli, although its orientation and position is fixed in all of them. A plausible explanation for this perceptual outcome may be the different curvature at the point of occlusion. **(C)** Effect of boundary curvature on elastica in $T(I)$. We show here the shape of the completed curve for fixed boundary positions and orientation ($\pm 45^\circ$), but varying values of boundary curvatures. $\hbar = 1$ in all cases.

To apply our theory and numerical computations to given natural and synthetic images, one first needs to measure inducer orientations and curvatures from image data. Such measurements, especially of curvature, may be challenging, but various approaches have been devised in the past (e.g., [60], [66]) to facilitate the countless computer vision tasks that incorporate curvature, from perceptual organization (e.g., [39], [46]) through virtually all active contour techniques (e.g., [9], [32]), to object representation and recognition (e.g., [14], [43]). In our case, we measured inducers' orientation and curvature by fitting a circular arc to three manually marked points on the observed part of the contour at or near the occlusion point. The three dots (one of which is the occlusion point itself), the circle they form, and its tangent line at the point of occlusion, are all plotted for selected examples in Fig. 10. Since we argued that our TBE theory is biologically plausible it is also worth mentioning that the biological measurement of curvature by cells in the visual cortex has also been discussed in the neurophysiological literature, and models have been proposed both in terms of the tuning properties of their receptive fields [13], [62] and the distribution of the long range horizontal connections that connect them [3].

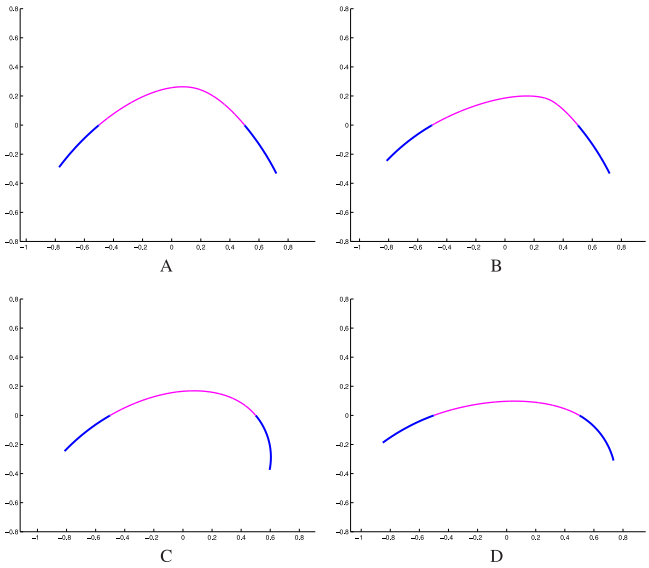

Fig. 8. Examples of tangent bundle elastica for different inducers. Each inducer $<x_i, y_i, \theta_i, \kappa_i>$ is plotted as a circular arc to express also the curvature information. The positions of the two inducers were fixed for all examples at $x_0 = -0.5, x_1 = 0.5, y_{0,1} = 0$. Orientation and curvature are: (A) $\theta_0 = 40°, \kappa_0 = 0.57$, $\theta_1 = -50°, \kappa_1 = 0.57$ (B) $\theta_0 = 30°, \kappa_0 = 0.70$, $\theta_1 = -50°, \kappa_1 = 0.57$ (C) $\theta_0 = 30°, \kappa_0 = 0.70$, $\theta_1 = -50°$, $\kappa_1 = 23$ (D) $\theta_0 = 20°, \kappa_0 = 0.70, \theta_1 = -30°, \kappa_1 = 2$. $\hbar = 1$ in all examples.

Following the above, Figs. 11, 12, and 13 demonstrate results on natural and synthetic images on which we overlaid artificial occluders to trigger a completion scenario. A canonical scale was used in all images such that distance between adjacent pixels was set to 10^{-2}. When executed on a low-end PC computer with 8GB RAM and non optimized Matlab code, running time of the numerical solution from Section 3.5 was approximately $T = 1$ sec for a given initial guess and step size $\delta = 0.01$. The non linear optimization was obtained with Matlab's `fminsearch` function with a tolerance value $\tau = 0.005$ on $E(c, \phi, g_0, l)$. While not shown here, excellent results were obtained with significantly more liberal values for δ and τ, which led to T as small as

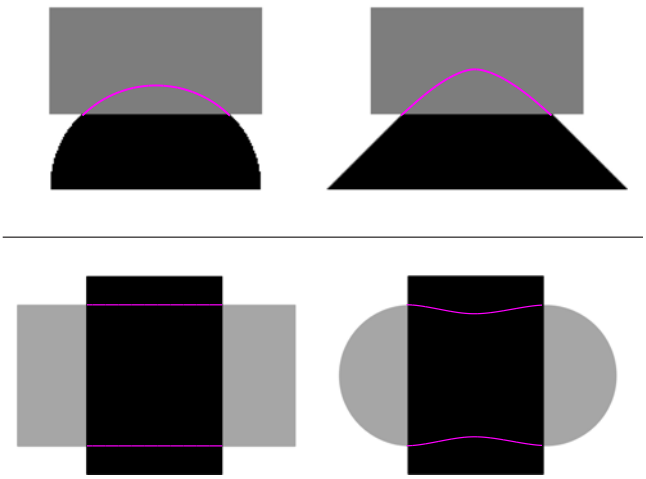

Fig. 9. Tangent bundle elastica and the effect of boundary curvature. While the inducers in each of these two pairs of stimuli are identical in terms of position and orientation, note the effect of inducer curvature on the predicted completed contour.

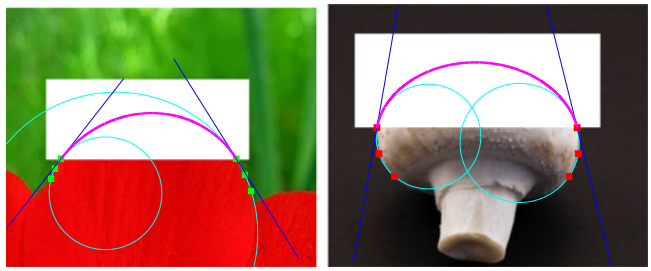

Fig. 10. Measuring inducer information. Shown for two examples are the image, the manually marked points on the observed part of the contour (in red or green), the osculating circle (cyan) fitted to these points and the inducer's tangent line (blue). Also shown is the tangent bundle elastica completion (in magenta).

50 ms. However, since in our implementation we used a naive brute force sampling procedure for finding a good initial guess (see end of Section 3.5), total running time depended on the number of samples tested serially. In our case, using the platform and the conservative selection of δ and τ mentioned above, this lasted up to 45 minutes. Naturally, both an informed scheme for selecting an initial guess or the parallel execution on a GPU or many independent nodes/cores could set the total running time as low as T.

Figs. 11, 12, and 13 also show comparison to various completion models from the literature. Note that in some cases all models produce comparable results, but in others the results can be quite disparate. Note in particular how the TBE completions of the two cars are qualitatively different and how they match human perception significantly better than previous models. Fig 14 brings another important example where continuity with inducer curvatures (a property of TBE) induces an inflectional completion which is more perceptually consistent and matches better the physical curve. Comparison to the minimum-length-in-the-tangent-bundle [5] and the image elastica models [28], [44]

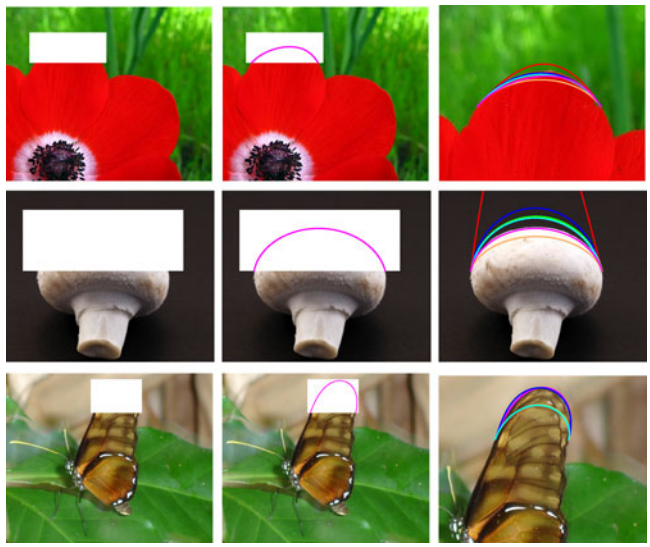

Fig. 11. Examples of curve completion via tangent bundle elastica applied in different settings. left: Image with a synthetic occluder. Middle: Tangent bundle elastica completions for $\hbar = 1$ and pixel size (i.e., viewing scale) of 0.01. Right: For comparison we show the unoccluded contour and the completions due to other models from the literature: the biarc model (green) [49], [61], cubic interpolation (red) [7], image plane elastica (blue) [28], scale invariant elastica (yellow) [64], Euler spiral (cyan) [35], and minimum length in the tangent bundle (orange) [5].

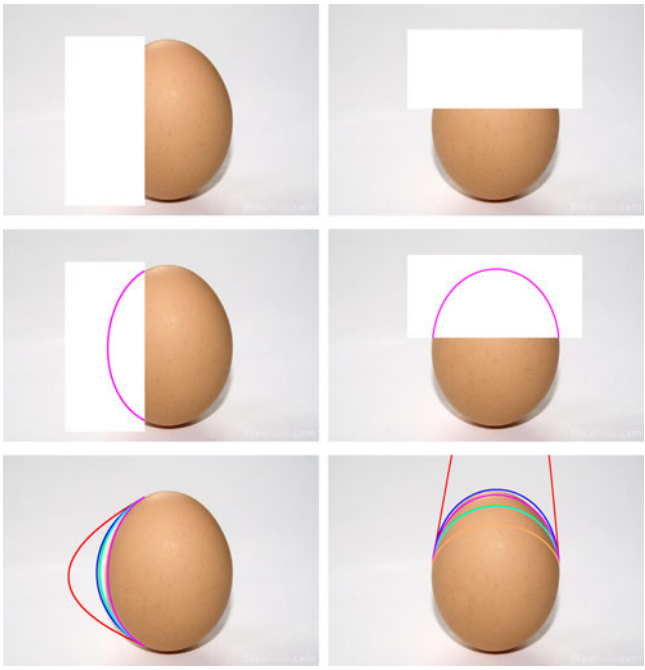

Fig. 12. Additional curve completion examples. Top row shows the stimulus, middle row the tangent bundle elastica, and the bottom row also shows the completions due to other models from the literature: the biarc model (green) [49], [61], cubic interpolation (red) [7], image plane elastica (blue) [28], scale invariant elastica (yellow) [64], Euler spiral (cyan) [35], and minimum length in the tangent bundle (orange) [5].

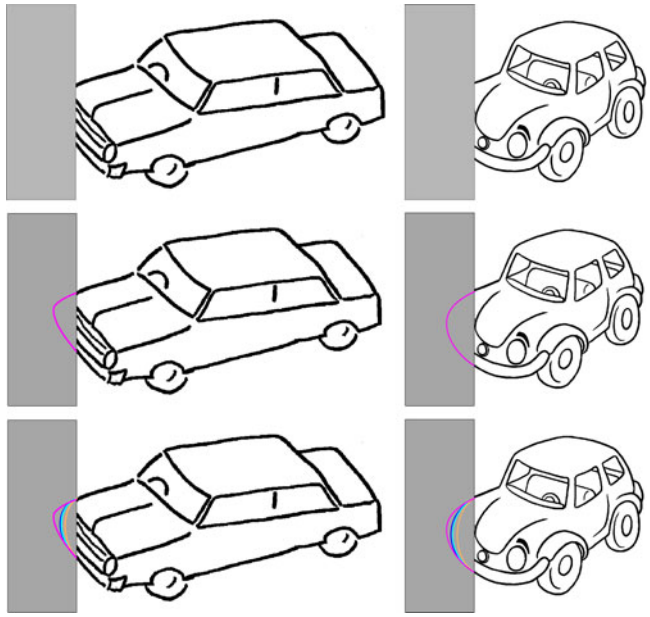

Fig. 13. Additional curve completion examples. Top row shows the stimulus, middle row the tangent bundle elastica, and the bottom row also shows the completions due to other models from the literature: the biarc model (green) [49], [61], cubic interpolation (red) [7], image plane elastica (blue) [28], scale invariant elastica (yellow) [64], Euler spiral (cyan) [35], and minimum length in the tangent bundle (orange) [5]. Note how the completions of the two cars are qualitatively different even though the two pairs of inducers are very similar in terms of position and orientation. Indeed, note how the TBE completions match human perception significantly better than previous models.

indicates qualitative differences. We emphasize that comparison for the real contour of the objects (where applicable, e.g., in the mushroom, butterfly, egg, and mountain scenes) is brought as a reference only and one should keep in mind that it bears little significance in terms of the *perceptual* correctness of the computational completions.[5]

When devising computational procedures of any sort it is inevitable to consider the consequences of noise, errors, or uncertainty in the initial measurements. While such considerations were never before addressed in the curve completion literature, they seem particularly desired in the context of our TBE theory due to the dependency on inducer curvature—a second order differential property that is more susceptible to measurement noise than just the position or the orientation. We have therefore examined empirically how such measurement errors may affect the shape of the completed TBE curve by introducing additive noise to the orientation, curvature, or both properties of the inducers simultaneously. Fig. 15 depicts the results of this test graphically, where the error bars are ± 1 standard deviation (STD) in the vertical position of the shape at several sample point along the curve. In the worst case, where both inducer properties are corrupted together, the mean and max STD along the completed shape are ~ 0.009 and ~ 0.012, respectively, which represent 0.5 and 0.7 percent relative error in the vertical position of the curve at its peak. The corresponding STD in the tangent bundle bending energy of these noisy completed shapes is approximately 6 percent. When only orientation or curvatures are corrupted, this relative error drops to 3.8 and 2.2 percent, respectively. Taken together, such results indicate no outstanding sensitivity of the computation to errors or uncertainty in inducer properties.

Finally, it is worth emphasizing again that without additional perceptual insights the value of $\hbar$ is not known yet and is set in all our examples to a fixed value $\hbar = 1$. Since this parameter affects the completed shape, it may be important to understand in what way it does so and what is the expected sensitivity of the results to its exact value (for the mere reason that the future perceptual procedure discussed in footnote 4, like any other psychophysical measurement, will not be able to provide an exact or definite evaluation of $\hbar$). Toward that end we examined how changing $\hbar$ affects the completed shape. Qualitatively, as shown in Fig. 16 for two randomly selected inducers, as $\hbar$ turns smaller the resultant completed shape becomes flatter, as indeed would be expected from the increasing dominance of the spatial dimension over the angular one (cf., Eq. (2)). At the same time, it is evident that the effect of changing $\hbar$ is both gradual and modest, thus indicating small sensitivity to errors or uncertainty in the experimentally measured value. Note that in the limit, when $\hbar \to 0$, Eq. (5) provides

$$\lim_{\hbar \to 0} \kappa_\beta^2 = \lim_{\hbar \to 0} \frac{\kappa_\alpha^2 + \hbar^2 {\kappa_\alpha}^4 + \hbar^2 \dot{\kappa_\alpha}^2}{\left(1 + \hbar^2 \kappa_\alpha^2\right)^3} = \kappa_\alpha^2,$$

implying that the TBE functional in Eq. (9) converges to the standard image elastica functional. That said, the two

5. It is likely that our visual system has evolved in correspondence to the statistical properties of our visual world. In this sense physical curves *are* important reference, but more in the statistical sense (i.e., on average) rather than for individual cases. The future experimental evaluation vis-a-vis natural image statistics (see the Summary) is designed to address exactly this possibility.

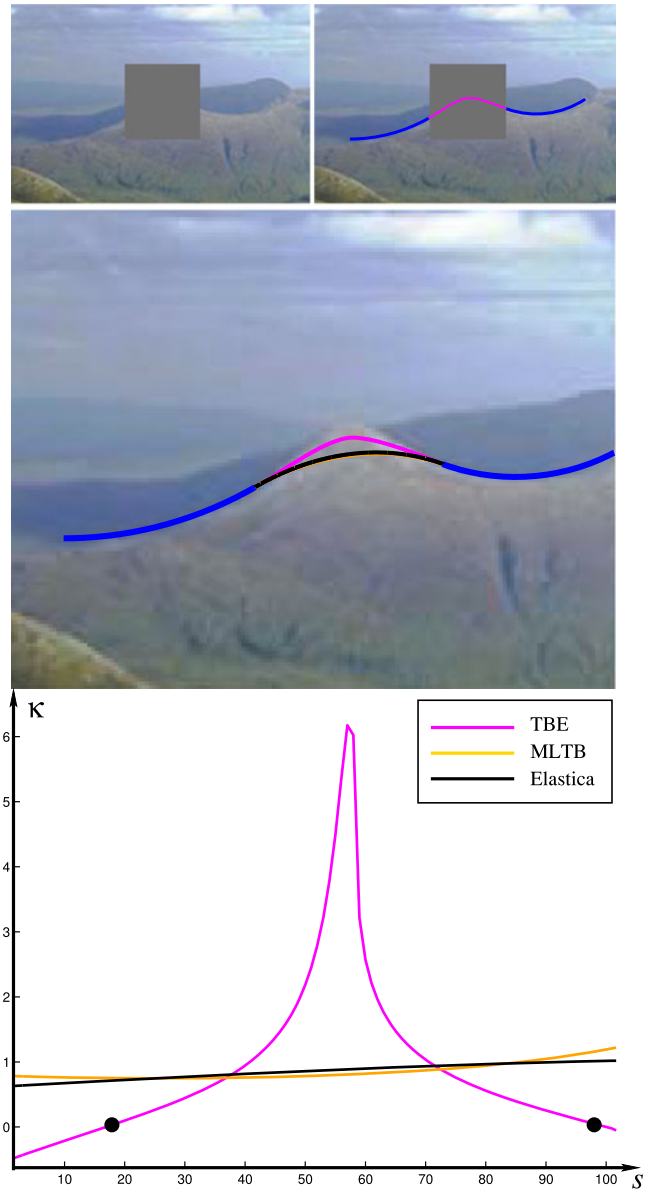

Fig. 14. TBE curve completion with inducer curvatures that induce inflectional curve (magenta). Note how the completion is more perceptually consistent, much sharper than competing models for the same inducers, and also corresponds better to the physical curve. In particular, note how image elastica (black) and minimum-length-in-the-tangent-bundle (gold, almost completely covered by the black curve) are much flatter and induce discontinuities of curvature at the inducers. The graph at the bottom shows the radically different behavior of the curvature of the three completed curves and emphasizes the emergence of inflection points (black points, where $\kappa = 0$) in the TBE completion.

models may still provide different completions even in this degenerate case due to the curvature constraints that TBE imposes on the completed curve near the inducers regardless of the value of $\hbar$.

6 Summary and Conclusions

We propose a new theory for visual curve completion, a fundamental process for biological and machine vision systems and a critical step towards the perception and recognition of whole objects from visual fragments. Combining classical and contemporary ideas from both the axiomatic and mechanistic approaches to contour completion, our tangent bundle elastica theory employs a *single* parameter-free mechanistic principle into a framework that generates novel predictions in a rigorous fashion. One of the most important outcomes of our theory is that inducer curvature is both necessary and significant in shaping the perceptually completed curve, as indeed reported in several perceptual studies. While a thorough comparison and a full quantitative psychophysical evaluation vis-à-vis human perception and natural image statistics is forthcoming, the possibilities that our theory allows with respect to the effect of boundary curvatures offers and the flexibility of the proposed framework for incorporation of other unexplored factors (such as shape prior and biases, perceptual anisotropies, etc.) may open new directions for visual completion research.

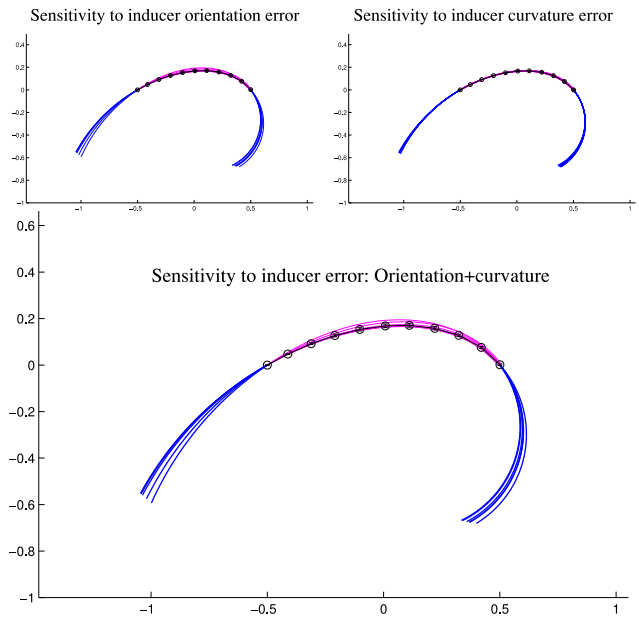

Fig. 15. Sensitivity to inducer noise. While the insets show the resultant range of shapes for orientation and curvature noise separately, the main figure depicts the shape variance for the more realistic case when both orientation and curvature suffer uncertainty. In all examples the base orientation of the left and right inducers are 30° and -50°, while their curvatures are -0.7 and -2.4, respectively. Orientation was then corrupted by additive Gaussian noise of $\sigma = 3^\circ$ with noise samples lying effectively in the range $\pm 5^\circ$ (to resemble the 10° orientation resolution of simple cell receptive fields [29]). In the absence of similar biological inspiration, curvature was corrupted by additive Gaussian noise of $\sigma = 0.05$. Blue circular arcs depict the inducers and their corresponding individual completion in cyan. Average shape is shown in black and errors bars are ± 1 STD. When both properties are corrupted simultaneously (as might happen in realistic scenarios), shape variance is still smaller than its arclength counterpart in the observed part of the inducer indicating no outstanding sensitivity to measurement or computation errors.

Acknowledgments

This work was supported in part by the National Institute for Psychobiology in Israel (grant no. 9-2012/2013) founded by the Charles E. Smith Family and the Israel Science Foundation (ISF grant no. 259/12). We also thank the Frankel Fund, the ABC Robotics initiative at BGU, and the Zlotowski Center for Neuroscience at Ben-Gurion University for their generous support. GB present address is the Faculty of Mathematics and Computer Science, The Weizmann Institute of Science, POB 26 Rehovot 76100, Israel.

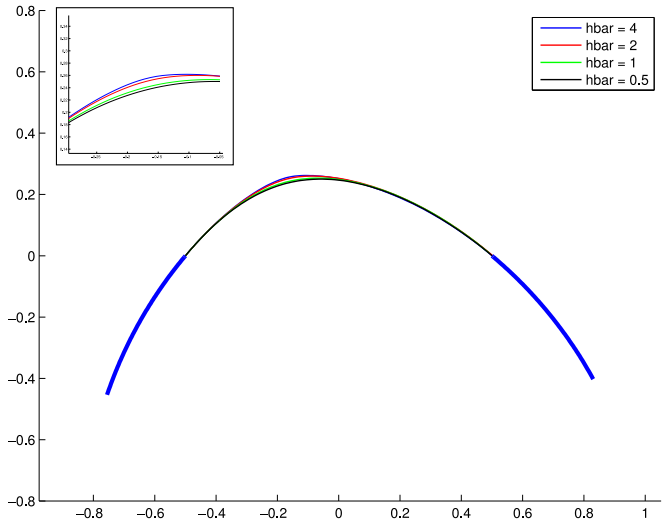

Fig. 16. Effect of $\hbar$ on the TBE completed shape. Note how smaller $\hbar$ values generate flatter (and therefore shorter) completed shapes, as indeed would be expected from the increasing dominance of the spatial dimension over the angular one. The inset zooms in on the differences in the most sensitive region.

REFERENCES

[1] J. August and S. Zucker, "Sketches with curvature: The curve indicator random field and Markov processes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 4, pp. 387–400, Apr. 2003.

[2] O. Ben-Shahar and S. Zucker, "The perceptual organization of texture flows: A contextual inference approach," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 4, pp. 401–417, Apr. 2003.

[3] O. Ben-Shahar and S. Zucker, "Geometrical computations explain projection patterns of long range horizontal connections in visual cortex," *Neural Comput.*, vol. 16, no. 3, pp. 445–476, 2004.

[4] G. Ben-Yosef and O. Ben-Shahar, "Tangent bundle curve completion with locally connected parallel networks," *Neural Comput.*, vol. 24, no. 12, pp. 3277–3316, 2012.

[5] G. Ben-Yosef and O. Ben-Shahar, "A tangent bundle theory for visual curve completion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 7, pp. 1263–1280, Jul. 2012.

[6] W. Bosking, Y. Zhang, B. Schofield, and D. Fitzpatrick, "Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex," *J. Neurosci.*, vol. 17, no. 6, pp. 2112–2127, 1997.

[7] M. Brady, W. Grimson, and D. Langridge, "Shape encoding and subjective contours," in *Proc. Assoc. Adv. Artif. Intell.*, 1980, pp. 15–17.

[8] A. Bruckstein and A. Netravali, "On minimal energy trajectories," *Comput. Vis., Graph. Image Process.*, vol. 49, no. 3, pp. 283–296, 1990.

[9] V. Caselles, F. Catte, T. Coll, and F. Dibos, "A geometric model for active contours in image processing," *Numerische Mathematik*, vol. 66, no. 1, pp. 1–31, 1993.

[10] T. Chan and J. Shen, "Variational image inpainting," *Commun. Pure Appl. Math.*, vol. 58, no. 5, pp. 579–619, 2005.

[11] G. Citti and A. Sarti, "A cortical based model of perceptual completion in the roto-translation space," *J. Math. Imaging Vis.*, vol. 24, no. 3, pp. 307–326, 2006.

[12] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1200–1212, Sep. 2004.

[13] A. Dobbins, S. Zucker, and M. Cynader, "Endstopped neurons in the visual cortex as a substrate for calculating curvature," *Nature*, vol. 329, no. 6138, pp. 438–441, 1987.

[14] G. Dudek and J. Tsotsos, "Shape representation and recognition from multiscale curvature," *Comput. Vis. Image Understanding*, vol. 68, no. 2, pp. 170–189, 1997.

[15] A. Efros and T. Leung, "Texture synthesis by non-parametric sampling," in *Proc. IEEE Int. Conf. Comput. Vis.*, 1999, pp. 1033–1038.

[16] C. Fantoni and W. Gerbino, "Contour interpolation by vector-field combination," *J. Vis.*, vol. 3, pp. 281–303, 2003.

[17] D. Field, A. Hayes, and R. Hess, "Contour integration by the human visual system: Evidence for a local 'association field'," *Vis. Res.*, vol. 33, no. 2, pp. 173–193, 1993.

[18] K. Fukushima, "Neural network model for computing occluded contours," *Neural Netw.*, vol. 23, pp. 528–540, 2010.

[19] J. Fulvio, M. Singh, and L. Maloney, "Precision and consistency of contour interpolation," *Vis. Res.*, vol. 48, pp. 831–849, 2008.

[20] W. Geisler, J. Perry, B. Super, and D. Gallogly, "Edge co-occurrence in natural images predicts contour grouping performance," *Vis. Res.*, vol. 41, no. 6, pp. 711–724, 2001.

[21] W. Gerbino and C. Fantoni, "Visual interpolation in not scale invariant," *Vis. Res.*, vol. 46, pp. 3142–3159, 2006.

[22] C. Gilbert, "Horizontal integration and cortical dynamics," *Neuron*, vol. 9, pp. 1–13, 1992.

[23] C. Gilbert and T. Wiesel, "Clustered intrinsic connections in cat visual cortex," *J. Neurosci.*, vol. 3, no. 5, pp. 1116–1133, 1983.

[24] D. Grosof, R. Shapley, and M. Hawken, "Macaque V1 neurons can signal 'illusory' contours," *Nature*, vol. 365, pp. 550–552, 1993.

[25] S. Guttman and P. Kellman, "Contour interpolation revealed by a dot localization paradigm," *Vis. Res.*, vol. 44, pp. 1799–1815, 2004.

[26] F. Heitger, R. von der Heydt, and O. Kubler, "Simulation of neural contour mechanisms: Representing anomalous contours," *Image Vis. Comput.*, vol. 16, pp. 407–421, 1998.

[27] W. Hoffman, "The visual cortex is a contact bundle," *Appl. Math. Comput.*, vol. 32, pp. 137–167, 1989.

[28] B. Horn, "The curve of least energy," *ACM Trans. Math. Softw.*, vol. 9, no. 4, pp. 441–460, 1983.

[29] D. Hubel and T. Wiesel, "Functional architecture of macaque monkey visual cortex," *Proc. Royal Soc. London, Ser. B*, vol. 198, pp. 1–59, 1977.

[30] G. Kanizsa, *Organization in Vision: Essays on Gestalt Perception*. West Port, NY, USA: Praeger Publishers, 1979.

[31] M. Kapadia, G. Westheimer, and C. Gilbert, "Spatial distribution of contextual interactions in primary visual cortex and in visual perception," *J. Neurophysiol.*, vol. 84, pp. 2048–2062, 2000.

[32] M. Kass, A. Witkin, and D. Trezopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1988.

[33] P. Kellman, "Interpolation processes in the visual perception of objects," *Neural Netw.*, vol. 16, pp. 915–923, 2003.

[34] P. Kellman and T. Shipley, "A theory of visual interpolation in object perception," *Cogn. Psychol.*, vol. 23, pp. 141–221, 1991.

[35] B. Kimia, I. Frankel, and A. Popescu, "Euler spiral for shape completion," *Int. J. Comput. Vis.*, vol. 54, no. 1–3, pp. 159–182, 2003.

[36] L. Landau and E. Lifshitz, *Theory of Elasticity*. New York, NY, USA: Pergamon, 1986.

[37] T. Lee and M. Nguyen, "Dynamics of subjective contours formation in the early visual cortex," *Proc. Nat. Acad. Sci. USA*, vol. 98, no. 4, pp. 1907–1911, 2001.

[38] A. Love, *A Treatise on the Mathematical Theory of Elasticity*. New York, NY, USA: Dover Publications, 1944.

[39] D. Lowe, "Organization of smooth curves at multiple scales," *Int. J. Comput. Vis.*, vol. 3, pp. 119–130, 1989.

[40] M. Maertens, S. Pollman, M. Hanke, T. Mildner, and H. Moller, "Retinotopic activation in response to subjective contours in primary visual cortex," *Frontiers Human Neurosci.*, vol. 2, no. 2, pp. 1–7, 2008.

[41] R. Malach, Y. Amir, M. Harel, and A. Grinvald, "Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex," *Proc. Nat. Acad. Sci. USA*, vol. 90, pp. 10469–10473, 1993.

[42] G. Mitchison and F. Crick, "Long axons within the striate cortex: Their distribution, orientation, and patterns of connection," *Proc. Nat. Acad. Sci. USA*, vol. 79, pp. 3661–3665, 1982.

[43] F. Mokhtarian and A. Mackworth, "A theory of multiscale, curvature-based shape representation for planar curves," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 8, pp. 789–805, Aug. 1992.

[44] D. Mumford, "Elastica and computer vision," in *Algebric Geometry and Its Applications*, B. Chandrajit, Ed. New York, NY, USA: Springer-Verlag, 1994.

[45] B. O'Neill, *Elementary Differential Geometry*. San Francisco, CA, USA: Academic Press, 1966.

[46] P. Parent and S. Zucker, "Trace inference, curvature consistency, and curve detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 8, pp. 823–839, Aug. 1989.

[47] J. Petitot, "Neurogeometry of V1 and Kanizsa contours," *Axiomathes*, vol. 13, nos. 3/4, pp. 347–363, 2003.

[48] K. Rockland and J. Lund, "Widespread periodic intrinsic connections in the tree shrew visual cortex," *Science*, vol. 215, no. 19, pp. 1532–1534, 1982.
[49] W. Rutkowski, "Shape completion," *Comput. Vis., Graph. Image Process.*, vol. 9, pp. 89–101, 1979.
[50] A. Sarti, G. Citti, and J. Petitot, "Functional geometry of the horizontal connectivity in the primary visual cortex," *J. Physiol. (Paris)*, vol. 103, pp. 37–45, 2009.
[51] A. Sarti, R. Malladi, and J. Sethian, "Subjective surfaces: A method for completing missing noundaries," *Proc. Nat. Acad. Sci. USA*, vol. 97, no. 12, pp. 6258–6263, 2000.
[52] E. Saund, "Perceptual organization of occluding contours generated by opaque surfaces," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 1999, pp. 624–630.
[53] E. Sharon, A. Brandt, and R. Basri, "Completion energies and scale," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 10, pp. 1117–1131, Oct. 2000.
[54] M. Singh, "Modal and amodal completion generate different shapes," *Psychol. Sci.*, vol. 15, no. 7, pp. 454–459, 2004.
[55] M. Singh and J. Fulvio, "Visual extrapolation of contour geometry," *Proc. Nat. Acad. Sci. USA*, vol. 102, pp. 939–944, 2005.
[56] M. Singh and J. Fulvio, "Bayesian contour extrapolation: Geometric determinants of good continuation," *Vis. Res.*, vol. 47, pp. 783–798, 2007.
[57] M. Singh and D. Hoffman, "Completing visual contours: The relationship between relatability and minimizing inflections," *Perception Psychophisics*, vol. 61, pp. 943–951, 1999.
[58] H. Takeichi, "The effect of curvature on visual interpolation," *Perception*, vol. 24, pp. 1011–1020, 1995.
[59] H. Takeichi, H. Nakazawa, I. Murakami, and S. Shimojo, "The theory of the curvature-constraint line for amodal completion," *Perception*, vol. 24, pp. 373–389, 1995.
[60] D. Tsai and M. Chen, "Curve fitting approach for tangent angle and curvature measurements," *Pattern Recognit.*, vol. 27, no. 5, pp. 699–711, 1994.
[61] S. Ullman, "Filling-in the gaps: The shape of subjective contours and a model for their creation," *Biol. Cybern.*, vol. 25, no. 1, pp. 1–6, 1976.
[62] M. Versavel, G. Orban, and L. Lagae, "Responses of visual cortical neurons to curved stimuli and chevrons," *Vis. Res.*, vol. 30, no. 2, pp. 235–248, 1990.
[63] R. von der Heydt, E. Peterhans, and G. Baumgartner, "Illusory contours and cortical neuron responses," *Science*, vol. 224, pp. 1260–1262, 1984.
[64] I. Weiss, "3d shape representation by contours," *Comput. Vis., Graph. Image Process.*, vol. 41, pp. 80–100, 1988.
[65] M. Wertheimer, "Laws of organization in perceptual forms," in *A Source Book of Gestalt Psychology*, W. Ellis, Ed. Evanston, IL, USA: Routledge & Kegan Paul, 1955, pp. 71–88.
[66] D. Williams and M. Shah, "A fast algorithm for active contours and curvature estimation," *Comput. Vis., Graph. Image Process.: Image Understanding*, vol. 55, no. 1, pp. 14–26, 1992.
[67] L. Williams and D. Jacobs, "Stochastic completion fields: A neural model of illusory contour shape and salience," *Neural Comput.*, vol. 9, no. 4, pp. 837–858, 1997.

Ohad Ben-Shahar received the BSc and MSc degrees in computer science from Technion, Israel Institute of Technology, and the PhD degree in computer science from Yale University, Connecticut, in 2003. An associate professor of computer science at Ben Gurion University of the Negev (BGU), he is also the director of the BGU interdisciplinary Computational Vision Lab. His main area of research is in computational vision and image analysis, where he is focusing primarily on issues related to the differential geometrical foundations of perceptual organization and early vision. His work is multidisciplinary and his computational research is often endowed by investigations into human perception, visual psychophysics, and computational neuroscience of biological vision. He is a senior member of the IEEE and the IEEE Computer Society.

Guy Ben-Yosef received the BSc degrees in physics and in computer science and the MSc and PhD degrees in computer science from Ben Gurion University of the Negev, Israel, in 2005, 2007, and 2011, respectively. He is currently a post-doctoral fellow at the computer science and applied mathematics department at the Weizmann Institute of Science, Israel. His main research interests include computational and biological vision, perceptual organization, and object recognition. He is a member of the IEEE and the IEEE Computer Society.