

Perception-Driven Semi-Structured Boundary Vectorization

by

Shayan Hoshyari

B.Sc., Sharif University of Technology, 2015

M.A.Sc., University of British Columbia, 2017

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Computer Science)

The University of British Columbia
(Vancouver)

February 2019

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Perception-Driven Semi-Structured Boundary Vectorization

submitted by **Shayan Hoshyari** in partial fulfillment of the requirements for the degree of **Master of Science** in **Computer Science**.

Examining Committee:

Alla Sheffer, Computer Science
Supervisor

Leonid Segal, Computer Science
Examining Committee

Abstract

Artist-drawn images with distinctly colored, piecewise continuous boundaries, which are referred to as *semi-structured imagery* in this thesis, are very common in online raster databases and typically allow for a perceptually unambiguous mental vector interpretation. Yet, perhaps surprisingly, existing vectorization algorithms frequently fail to generate these viewer-expected interpretations on such imagery. In particular, the vectorized region boundaries they produce frequently diverge from those anticipated by viewers. This thesis proposes a new approach to region boundary vectorization that targets semi-structured inputs, and leverages observations about human perception of shapes to generate vector images consistent with viewer expectations. When viewing raster imagery, observers expect the vector output to be an accurate representation of the raster input. However, perception studies suggest that viewers implicitly account for the lossy nature of the rasterization process and mentally smooth and simplify the observed boundaries. The core algorithmic challenge is to balance these conflicting cues and obtain a piecewise continuous vectorization whose discontinuities, or corners, are aligned with human expectations.

This work centers around a simultaneous spline fitting and corner detection method that combines a learned metric, that approximates human perception of boundary discontinuities on raster inputs, with perception-driven algorithmic discontinuity analysis. The resulting method balances local cues provided by the learned metric with global cues obtained by balancing simplicity and continuity expectations. Given the finalized set of corners, the proposed framework connects those using simple, continuous curves that capture input regularities. The proposed method is demonstrated on a range of inputs, with its superiority validated over existing alternatives via an extensive comparative user study.

Lay Summary

Semi-structured images with visually distinct regions and perceptually clear region boundaries are natural candidates for vector representation. However, due to historic reasons, large numbers of such images are saved in raster form. Converting them to vector would enable more compact representations, artifact-free resizing, and other applications. Surprisingly, existing state-of-the-art vectorization methods produce multiple visible artifacts when processing such inputs. This thesis presents a perception-driven vectorization method designed for such semi-structured inputs that produces output vectorizations well-aligned with viewer expectation.

Preface

All the work presented in this thesis is a product of a working relationship between Shayan Hoshyari, Edoardo Alberto Dominici, Alla Sheffer, Nathan Carr, Duygu Ceylan, and Zhaowen Wang. The implementation of the methods and the data analysis has been done by Shayan Hoshyari. Edoardo Dominici implemented the method of Chapter 6, and helped with the conduction of the user study. Shayan Hoshyari was provided with invaluable guidance from Alla Sheffer, Nathan Carr, Duygu Ceylan, and Zhaowen Wang throughout the process.

The user study in Chapter 7 has been conducted with the approval of the University of British Columbia (UBC BREB Number H16-02320).

The ideas and algorithms described in this thesis has been published as the following research paper in ACM Transactions on Graphics (TOG) and was presented at the SIGGRAPH conference in August 2018.

Shayan Hoshyari, Edoardo Alberto Dominici, Alla Sheffer, Nathan Carr, Duygu Ceylan, Zhaowen Wang, and I-Chao Shen. Perception-driven semi-structured boundary vectorization. *ACM Transaction on Graphics*, 37(4), 2018. doi:10.1145/3197517.3201312. <http://www.cs.ubc.ca/labs/imager/tr/2018/PerceptionDrivenVectorization/>

Table of Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	vi
List of Tables	viii
List of Figures	ix
Acknowledgments	xi
1 Introduction	1
1.1 Motivation	1
1.2 Objective and Outline	2
1.3 Related Work	4
1.3.1 Natural Image Upscaling	4
1.3.2 Natural Image Vectorization	4
1.3.3 Pixel Art	5
1.3.4 Vectorization of Artist-Generated Cartoon Like Imagery	5
1.3.5 Curve Fitting	6
1.3.6 Corner Detection	7
2 Algorithm Overview	8
3 Initial Data-driven Corner Prediction	12
3.1 Learning Corner Likelihood	13
3.2 Training Data	15
3.3 Short Segments	16
4 Perception-Driven Corner Removal	18

4.1	Piecewise Smooth Vectorization	19
4.1.1	Fitting Energy	19
4.1.2	Constraints	20
4.1.3	Optimization	21
4.2	Corner Removal Iterations	23
4.2.1	Short segments	23
4.2.2	Computing Global Context Cues	25
5	Boundary Regularization	26
6	Multi-Color Inputs	28
7	Results and Validation	30
7.1	Perceptual Ground Truth Comparison	30
7.2	Algorithmic Alternatives	30
7.3	Qualitative Evaluation	33
7.4	Learning Statistics	36
7.5	Discussion on Black Box Learning	37
7.6	Resolution Dependence	37
7.7	Performance	37
8	Conclusions	38
	Bibliography	40
A	Curve Fitting	45
A.1	Projection of a Point on the Curve	45
A.2	Fitting a Single Curve	47
A.3	Fitting of Multiple Curves	48
B	A Fundamental Discrete-Continuous Optimization Problem	49
C	User Study Questions	52

List of Tables

Table 7.1	Comparison between neural network and random forest on corner detection accuracy	36
Table A.1	Curve degrees of freedom	45
Table A.2	Dimensions of the matrices involved in the computation of the derivative of the projection operator	46

List of Figures

Figure 1.1	Vectorization of an involved input	1
Figure 1.2	Comparison of various vectorization methods for a low resolution input	2
Figure 1.3	Trade off between simplicity and continuity	3
Figure 2.1	The impact of closure	9
Figure 2.2	Suppression of narrow features in rasterization.	9
Figure 2.3	Algorithm overview	10
Figure 3.1	The effect of misplaced corners	12
Figure 3.2	Precision-recall curves for random forest corner detectors	15
Figure 3.3	Examples of short segments	16
Figure 3.4	Training data examples	17
Figure 4.1	Iterative corner removal	19
Figure 4.2	Piecewise smooth vectorization.	20
Figure 4.3	Optimization stages	23
Figure 4.4	Effect of the half-space constraint and the short segment processing step	24
Figure 5.1	Effect of regularization	26
Figure 6.1	Removing ambiguities when multiple color regions meet at the same point	29
Figure 7.1	Comparison of single region vectorization methods on various resolutions	31
Figure 7.2	Vectorized multi-color examples	32
Figure 7.3	Vectorization of multi-color inputs	32
Figure 7.4	Comparison to artist generated vector images	33
Figure 7.5	Structured pixel-art vectorization comparison	34
Figure 7.6	Comparison to Kopf and Lischinski [2011] for non-structured pixel-art.	34
Figure 7.7	Comparison of our corner removal algorithm to using only a pure classifier	34
Figure 7.8	User study statistics	35
Figure 7.9	Failure cases	36

Figure 8.1	Effect of noise and anti-aliasing	39
Figure A.1	Schematic of the projection operator	46
Figure B.1	Schematic of the segmented least-squares problem	50
Figure C.1	User study question example for comparison against alternative algorithms	53
Figure C.2	User study question example for comparison against the work of an artist	54

Acknowledgments

First and foremost, I would like to thank my supervisor, Alla Sheffer, for her professional mentorship and continuous support. It has been a pleasure for me to work under her supervision. I would also like to thank Nathan Carr, Duygu Ceylan, and Zhaowen Wang for their insight and feedback.

A special thanks to Edoardo for his collaboration and tremendous help, especially during the SIGGRAPH submission deadline.

I would like to thank my labmates for their suggestions and discussions. In particular, Silver for his help with preparing the inputs, comparisons to alternatives algorithms, and rendering the outputs, and Nicholas Vining for proofreading the paper. Finally, thanks to Ilya Baran for sharing his code, Cornucopia, as an open-source library.

Thanks to Moritz Bächer, Stelian Coros, Hongyi Xu, and Espen Knoop for the great internship at Disney Research Zürich. It was a wonderful and enlightening experience.

Last but not least, I would like to thank my parents, Iraj and Mahboubbeh, and my sister, Leiana, for their love and support.

Chapter 1

Introduction

1.1 Motivation

Many artist-generated images frequently used in digital artwork (such as computer icons, comic book imagery, and simple graphic illustrations) consist of distinctly colored regions with piecewise continuous boundaries and visually pronounced corners (Figure 1.1). Such *semi-structured* images can be compactly represented in vector form enabling editing and other manipulation. For a range of historic reasons, many such images are still generated in raster, or bitmap formats; large legacy collections of semi-structured images and artwork exist in raster formats only, creating a demand for robust vectorization algorithms. For instance, the Adobe Stock Image database contains over nine million raster images labeled as icons or clip-art. Converting these images into a vector format would not only allow more effective editing, but would support re-interpretation of legacy designs that have been originally tuned to a particular resolution on modern high resolution displays.

The artist-generated images we target consist of a set of visually distinct regions, whose boundaries can be extracted using existing methods (Figure 1.1). The major vectorization challenge therefore is




Figure 1.1: Vectorizing artist-drawn semi-structured raster images (left) using existing methods (center) results in visible artifacts (highlighted in zoomed-in insets). The output generated by the perception-driven vectorization algorithm of this work (right) is consistent with viewer expectations. Please zoom in online to see details. Input image © aves – stock.adobe.com.




Figure 1.2: From left to right: input raster image, its vectorization results using Cornucopia (smooth polyline fitting) [Baran et al., 2010], Potrace [Selinger, 2003], Adobe Trace [Adobe, 2017], Vector Magic [Vector Magic, 2017], and our proposed method. Our result is more consistent with human perception than others. Input image adapted from Freepick – www.flaticon.com.

to convert these boundaries into piecewise free-form vector curves. While vectorization of generic image boundaries is often inherently ambiguous, human observers typically have a much more clear mental image of the expected vectorization outcome for semi-structured imagery such as the examples in Figures 1.1 or 1.2. However, existing methods for boundary vectorization are tuned toward natural imagery (i.e., vectorizing photographs), and frequently produce results inconsistent with viewer expectations on typical artist-drawn boundaries (Figures 1.1 and 1.2). These artifacts are most noticeable on lower resolution input, such as Figure 1.2. Fitting frameworks which target artist inputs [McCrae and Singh, 2008; Baran et al., 2010; McCrae and Singh, 2011] are similarly unsuitable for our needs, as they rely on input features such as tangents and curvature, which cannot be reliably computed on raster data (Figures 1.2 and 7.1).

1.2 Objective and Outline

Our goal is to explore the principles behind human perception of semi-structured raster imagery and to develop a computational algorithm that comes close to replicating the human mental process (Chapter 2). Since algorithmic rasterization of vector images results in raster boundaries that are close to their originating vector curves up to quantization error; we speculate that human observers expect the output vector curves to be similarly *accurate*, or close to their raster counterparts. We compute a solution consistent with human perception by combining accuracy with two key principles of Gestalt psychology: *simplicity* and *continuity* [Koffka, 1955; Wagemans et al., 2012]. Simplicity indicates that, given multiple possible interpretations for a 2D input, human observers prefer a simpler geometric interpretation. In our context, this points to a preference for lower-curvature vectorizations with few abrupt curvature changes [McCrae and Singh, 2008; Baran et al., 2010]. The *continuity* principle states that human observers have a tendency to group stimuli into continuous curves and patterns [Koffka, 1955; Wagemans et al., 2012]. Continuity plays a major role in the mental vectorization process. Observers consistently perceive jagged raster boundaries as piecewise smooth curves in their vector forms, and mentally segment

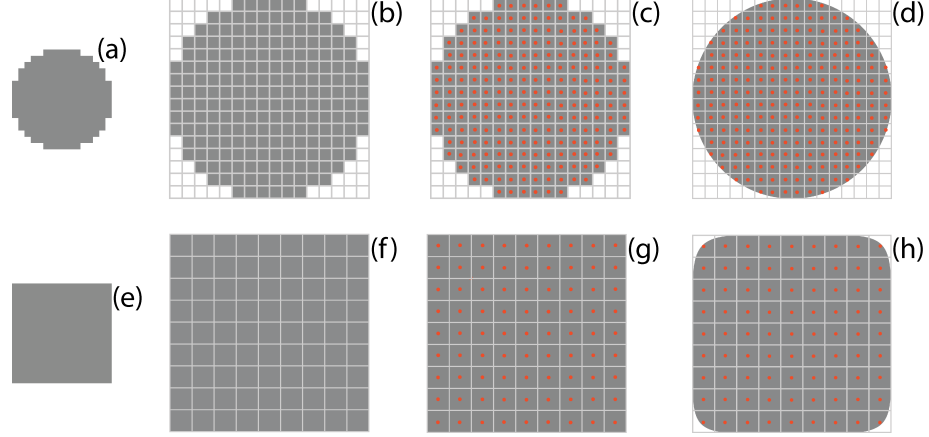


Figure 1.3: Given a raster input (a,e) with zoomed versions (b,f) an accurate and simple vectorization would simply preserve the axis-aligned input edges (c,g), while a purely continuous vectorization would generate an output continuous everywhere (d,h). Human observers and our algorithm balance the conflicting simplicity and continuity cues opting for vectorizations (d) and (g).

boundaries at *only* a small set of discontinuous (C^0) *corners* (Figure 1.3). We speculate that their choice of corners reflects a balance between continuity, accuracy, and simplicity.

We convert these observations into an actionable vectorization algorithm targeted at artist-generated raster inputs. Key to our method is the ability to detect input corners consistent with human perception. While perception literature provides some insights as to when humans perceive geometries as continuous [Hess and Field, 1999; Wagemans et al., 2012], these cannot be applied as is to our input raster data, which is locally discontinuous everywhere (Figure 1.3). We propose a method for corner detection that achieves results consistent with viewer expectations by combining learning from human data with algorithmic perception-driven discontinuity analysis. We note that the perception of corners is affected by both the local geometry surrounding a boundary vertex, and the overall boundary shape, specifically the locations of other corners. We learn to reason about local geometric surrounding from a collection of training raster images with manually annotated corners (Chapter 3). We then incorporate the learned classifier into a perception driven corner detection framework that accounts for global context (Chapter 4). This framework simultaneously locates the corners and fits simple G^1 continuous splines between them. Once the set of corners is finalized, we proceed to further optimize output simplicity without sacrificing continuity or accuracy (Chapter 5).

We validate our method in a number of ways: by comparing its outputs to manual vectorization; conducting a user study that compares our results against prior art; comparing our algorithmically detected corners to manually annotated ones; and generating a gallery of results on inputs with a diverse range of resolutions (Chapter 7). These validations confirm that our results are well aligned with human perception and are superior to those generated by alternative techniques. As shown by these comparisons, viewers prefer our results over the closest alternative vectorization method by a ratio of 5.58 (Figure 7.8).

Our proposed method is particularly effective for lower-resolution images where accuracy alone allows for a large set of solutions and earlier methods produce visually noticeable artifacts.

In summary, our contributions are as follows.

- We propose a novel vectorization approach targeted at artist-generated raster images. This method successfully balances learning and perception-driven algorithms to achieve desired vector interpretations without the need for large, potentially intractable, training sets.
- The effectiveness of this approach is demonstrated through a comprehensive set of user studies and state-of-the-art results.

1.3 Related Work

Vectorization of raster imagery is a long standing research problem. We discuss the different approaches below, classified by the type of inputs they target.

1.3.1 Natural Image Upscaling

Image upscaling or super-resolution methods increase the resolution of input images [Nasrollahi and Moeslund, 2014; Yang et al., 2014; Hou and Andrews, 1978; Li and Orchard, 2001; Fattal, 2007; Glasner et al., 2009; Dong et al., 2014; Wang et al., 2015; Kim et al., 2016; Dahl et al., 2017]. Classical approaches often work by applying filters, e.g., nearest-neighbor and bicubic, to reconstruct the underlying signal [Hou and Andrews, 1978; Li and Orchard, 2001]. More recent methods successfully hallucinate fine details using natural image statistics [Fattal, 2007] or self-similarity priors [Glasner et al., 2009]. Modern methods in this space successfully utilize deep learning frameworks that are trained on large datasets [Dong et al., 2014; Wang et al., 2015; Kim et al., 2016; Dahl et al., 2017]. Nevertheless, these methods are typically tuned to natural images and are poorly suited for fitting regions with distinct, piecewise continuous boundaries. Given such input they tend to produce jagged and wavy patterns, as shown in Figure 7.1.

1.3.2 Natural Image Vectorization

Natural image vectorization methods are widely employed in commercial products [Adobe, 2017; Vector Magic, 2017] and address two separate challenging problems: image segmentation and segment boundary fitting. Most methods concentrate on the segmentation stage [Orzan et al., 2008; Lecot and Levy, 2006; Sun et al., 2007; Xia et al., 2009; Wang et al., 2017; Favreau et al., 2017]. The boundaries of the computed segments are typically highly irregular; thus boundary fitting is mainly employed for smoothing rather than replicating or recovering [Kass et al., 1988; Caselles et al., 1997; Figueiredo et al., 2000]. Such lossy smoothing is undesirable for inputs with perceptually well-defined boundaries designed by

artists (Figures 1.2 and 7.1). Interactive vectorization approaches [Jun et al., 2017] can adjust different properties of region boundaries at the cost of user intervention. Some frameworks [Noris et al., 2013; Bessmeltsev and Solomon, 2018; Parakkat et al., 2018] are specifically designed for vectorization of artist strokes. Unlike the target images of our method, the inputs to such algorithms do not have boundary structure.

1.3.3 Pixel Art

Pixel art is characterized by extremely low resolution content (often 32×32 pixels or less) possessing only a few discrete regions of color. Processing pixel art requires dedicated methods [Kopf and Lischinski, 2011; Eagle, 1997; Mazzoleni, 2001; Stepin, 2003] as natural image processing approaches tend to undesirably blur this data. Upscaling methods for pixel art [Eagle, 1997; Mazzoleni, 2001; Stepin, 2003] are intended to run in real-time and are often based on simple strategies that result in quality degradation as resolution increases. In particular, the method of Stepin [2003], High Quality Magnification (HQX), is widely used in modern ports of legacy games. HQX offers magnification factors of $2\times$, $3\times$, and $4\times$ but does not produce an exact boundary vectorization. In this method, the 3×3 topological neighbours of each pixel are analyzed to be classified either as *distant* or *close* via predefined thresholds in the YUV color space. The classification of these 8 neighbours gives rise to 256 different combinations, each of which has a specific interpolation strategy saved in a lookup table.

Kopf and Lischinski [2011] proposed a dedicated method for vectorizing pixel art data. While they use the same thresholding strategy as Stepin [2003] for pixel region detection, their primary focus is on resolving the topological ambiguities that arise when regions with the same color touch across pixel corners. Arbitrary resolution of such ambiguities is undesirable, since it can cause thin features to become visually disconnected under large magnifications. Kopf and Lischinski associate the input image with a graph, where each pixel is represented by a vertex, and is connected to the pixels in its 3×3 surrounding region. They then recast the disambiguation process as removing two out of four edges connected to each vertex, based on perceptually motivated rules. Their rules allow for non-local solutions, yet are fast to compute. In a second stage, they apply a basic spline fitting approach to convert these pixel regions into vector form. Kreuzer et al. [2015] later designed a parallel implementation of this work as shaders in the OpenGL graphics pipeline. Our work focuses on the second stage of this process, and uses perceptual cues to infer piecewise smooth boundaries that better align with viewer expectations. We compare our method to the open-source implementation of Kreuzer et al. in Chapter 7.

1.3.4 Vectorization of Artist-Generated Cartoon Like Imagery

There have been few attempts to specifically target vectorization of artist-generated imagery. Sýkora et al. [2005] use Autotrace [Weber and Herzog, 2004] to vectorize boundaries on cartoon inputs. Zhang et al. [2009] similarly target cartoon animation; they fit smooth Bézier curves to region boundaries, and

only introduce discontinuities at corners where multiple regions meet. Fatemi et al. [2016] upsample binary images under the assumption that the originating boundaries are C^2 continuous everywhere. This assumption is clearly unsuited for raster boundaries which have visible discontinuities. Yang et al. [2016] fit Bézier splines to anti-aliased multi-region raster images. They assume that all the boundary transition vertices (C^0 or G^1) are given, and focus on optimizing the shape of individual Bézier segments connecting them. While we do not address anti-aliasing, our work complements Yang et al. in its focus on detecting the C^0 corners, the G^1 transition vertices, and the global regularities of vectorized boundaries.

A number of open-source and commercial software packages, e.g., [Selinger, 2003; Weber and Herzog, 2004; Vector Magic, 2017; ScanFont, 2017], specifically address region boundary vectorization. Excepting the work of Selinger [2003], however, they provide little details on the underlying algorithms. ScanFont [2017] is specifically designed for textual data. Vector Magic [2017] is based on a Bayesian method proposed by [Diebel, 2008]. Potrace [Selinger, 2003] is designed for vectorizing binary colored images. It extracts closed paths from the image, which are then fitted by polylines that strike a balance between fitting error and the number of line transitions. Finally, Potrace replaces the fitted polylines with smooth Bézier curves. While Potrace includes simple methods for the resolution of region connectivity ambiguities, they are only based on local information and are not suited for low resolution images. We compare the output of our method against the two publicly available programs, VectorMagic and Potrace (generalized version of Inkscape [2018] for color images), on a range of inputs (Figures 1.1 and 1.2). In our user study, participants preferred our results over VectorMagic and Potrace by a ratio of 5.58 and 6.89, respectively, as shown in Figure 7.8.

1.3.5 Curve Fitting

Classical curve fitting methods reconstruct originating curves from unordered noisy point samples by balancing fidelity to these samples versus curve fairness [Farin, 2002; Fleishman et al., 2005; Liu and Wang, 2008]. Our inputs are distinctly different – instead of dense samples, we deal with raster boundaries with axis-aligned edges; instead of denoising the input, we seek to recover its perceptual interpretation.

Our work is closer in spirit to methods for fitting artist intended curves to raw sketch strokes [McCrae and Singh, 2008; Baran et al., 2010; McCrae and Singh, 2011]. McCrae and Singh [2008] note that artists prefer to use curves with linearly changing curvature and avoid abrupt curvature discontinuities, consistent with the Gestalt simplicity cue. These methods consequently look for a compact set of simple curves, with gradually changing curvature, that jointly fit the dense raw input. They rely on curvature and fitting quality to identify possible corners, an approach which is poorly suited for raster inputs (Figure 1.2). In particular, local curvature estimation methods designed for noisy but generally smooth polylines [Baran et al., 2010; McCrae and Singh, 2011] are inapplicable to raster data which is locally discontinuous everywhere.

We employ a version of the fitting algorithm of Baran et al. [2010] with pre-specified corner locations as a local step within our boundary vectorization framework and use the fitting output to iteratively refine our choice of corners and guide our regularization decisions. As shown by the comparisons in Figure 1.2, our iterative framework significantly outperforms the original method of Baran et al. on typical semi-structured raster boundaries.

1.3.6 Corner Detection

Corner detection in 2D images is well studied in computer vision with both handcrafted [Harris and Stephens, 1988; Shi et al., 1994] and learning based methods [Rosten et al., 2010; Altwaijry et al., 2016; Yi et al., 2016]. Another line of research more related to our problem is to detect discontinuous or high curvature points on planar curves encoded in 1D chain format [Freeman and Davis, 1977; Beus and Tiu, 1987; Chetverikov and Szabo, 2003]. Medioni and Yasumoto [1986] and Langridge [1982] simultaneously detect corners and fit cubic splines to represent curves. Curve segmentations are iteratively refined via multi-resolution pyramids [Arrebola and Sandoval, 2005] and break point suppression [Carmona-Poyato et al., 2010]. Our work shares a similar goal; however, its focus is on raster images of much lower resolution than those typically considered in corner detection literature, and therefore perceptual cues are relied upon to resolve ambiguities.

Chapter 2

Algorithm Overview

Given a semi-structured input image, our method automatically vectorizes it by replacing each distinctly colored raster region with a vectorized equivalent delineated by a piecewise smooth boundary. Region detection on semi-structured imagery is straight forward, with local edge detection rules [Kopf and Lischinski, 2011; Stepin, 2003] producing the desired region segmentation. Consequently, our technical challenge is to vectorize the discrete region boundaries, consisting of vertical and horizontal pixel-side segments, into piecewise smooth curves.

The goal of classical boundary vectorization methods can be formulated as follows. Consider a rasterized region boundary $\tilde{S} = R(\tilde{S})$ where \tilde{S} denotes the original shape, in vector or parametric form with infinite precision, and R denotes the rasterization operator which takes curves and lines and converts them into a pixelized form. The goal of classical vectorization is to recover \tilde{S} from \tilde{S} . In general R is not invertible as multiple piecewise continuous curves are rasterized to the same raster polyline (Figure 1.3) making this goal unachievable.

Yet, given a rasterized boundary, human observers typically conjure an imaginary original shape S^P , which may or may not coincide with the actual originating shape \tilde{S} . We conjecture that, for semi-structured imagery, the shapes imagined by different observers are very similar. Our method is specifically designed for recovering these viewer perceived vectorizations S^P . We believe that for artist made inputs whose creators seek to clearly communicate their intent to viewers, recovering a result consistent with viewer perception is tantamount to recovering the original artist-intended vector image. We achieve this goal by leveraging the perceptual cues of accuracy, continuity, and simplicity identified in Chapter 1.

Accuracy. In general, we expect our reconstructed shape S^P , when rasterized, to produce the same pixelized content \tilde{S} , i.e., $R(S^P) \approx \tilde{S}$. In theory, the maximal distance between an originating curve and its rasterization can be unbounded, with the rasterization suppressing narrow but arbitrarily long features (Figure 2.2). However, following the WYSIWIG (What You See Is What You Get) principle, we note




Figure 2.1: The impact of closure. Left: perceptually consistent and locally smooth vectorization. Middle: vectorization inconsistent with the global context where the sides of the bottleneck should be perceived as part of the same imaginary segmentation contour. Right: a consistent vectorization on both sides produced by our framework correctly accounting for the closure cue. Input image adapted from Freepick – www.flaticon.com.




Figure 2.2: Suppression of narrow features in rasterization.

that viewers generally expect the vectorized output to be close to the actual boundary polyline and do not hallucinate imaginary features. This motivates us to enforce a hard bound on the distance between the reconstructed and raster boundaries, and to prioritize closer fits.

Continuity. We overcome the discrete nature of our input by looking for vectorizations that are G^1 continuous everywhere except at a small number of well-placed, discontinuous, C^0 only *corners*. In particular, our observations show that human observers expect the corners not to be clustered close to one another. We locate these corners by balancing continuity against other cues.

Simplicity. Following recent graphics literature [McCrae and Singh, 2008; Baran et al., 2010] we fit the simplest possible smooth curves to each corner-bounded input raster segment. We prioritize zero curvature (straight) or constant curvature (arc) vectorizations when satisfying the accuracy requirement, and employ clothoids or curves with linearly changing curvature when simpler solutions are inadequate. Following the same argument, we seek to minimize the number of fitted curve segments, and to reuse the same or similar segments along different portions of the same region boundary. The simplicity principle [Wagemans et al., 2012] further suggests more regular solutions which maximize symmetry and axis-alignment.

Closure. We consider another cue in addition to the core ones identified above. Research suggests that humans mentally segment objects into parts at points with negative minima of curvature [Wagemans et al., 2012] by mentally connecting these points to their closest opposite curvature minimum points (Figure 2.1). Consequently, they mentally pair potential concave corners, and expect them to share




Figure 2.3: Our framework consists of three major steps: (a) potential corner detection, (b) iterated corner removal, and (c) global regularization. Colors distinguish different curve types. Input image adapted from Freepick – www.flaticon.com.

similar continuity properties (Figure 2.1). Our algorithm enforces this expectation.

A major challenge in applying these principles to boundary vectorization is to efficiently and correctly detect the corners, or C^0 only discontinuities, that humans perceive along the input raster boundaries. Crafting a purely algorithmic corner classifier aligned with human perception is a daunting task: beyond the observations above, we do not have specific knowledge of how humans perceive corners. We are therefore motivated to employ machine learning on manually annotated corner data as a core part of the detection framework. However, learning corner classification entirely from human annotations is problematic. In our experiments, manual annotation required a non-negligible amount of time (5 to 10 minutes per input shape), making collection of large scale reliable training data, consisting of hundreds or thousands of shapes across multiple input resolutions, impractical. Participants also found it hard to distinguish between purely geometry-driven decisions, and content or recognition based ones, making extrapolation from a small set of inputs challenging. We develop an algorithm capable of obtaining perceptually consistent corner classification from limited training data by combining learned classifier prediction with insights about corner perception. An overview of our framework is shown in Figure 2.3.

We first compute local corner probabilities using a trained classifier (Chapter 3). We note that the classifier, while fairly accurate, can and does make occasional erroneous choices (Figure 3.1), both identifying false corners and missing true corners. Rather than addressing both types of errors, we start with a more lax set of potential, rather than definitively labeled, corners; we then gradually prune this set (Chapter 4). We use the obtained corners to compute an initial boundary vectorization that is discontinuous at these corners and that best satisfies our perceptual criteria of accuracy and simplicity. We then repeatedly compact this set by removing corners, if the vectorization produced following the removal better adheres to our measured perceptual criteria. Our removal decisions are dominated by accuracy and simplicity, since the set of potential corners produced by the classifier is typically small and sparse enough to satisfy the continuity principle as is.

Given the finalized set of corners, we further improve output simplicity (Chapter 5). We detect groups of fitted segments whose current fit can be further simplified by either using the same primitive across all segments (e.g., using arcs of the same circle or segments of the same straight line), or by using primitives with common parameters across them (e.g., parallel lines).

Hereinafter, we assume for simplicity that the input image contains a single closed region, whose boundary forms a closed polyline P . The ordered pixel corner points that lie on P are denoted as its vertices, $\{p_0, \dots, p_{m-1}, p_m = p_0\}$. We define the resolution of the region as the largest dimension of its bounding box rounded to the closest power of two. Chapter 6 details the minimal changes to this formulation necessary to account for images with multiple interacting regions (Figure 7.3).

Chapter 3

Initial Data-driven Corner Prediction

Given a boundary polyline P , we aim to identify a subset of boundary vertices $C = \{c_0, \dots, c_n = c_0\}$ that viewers perceive as C^0 discontinuities when mentally vectorizing this boundary. The local geometry around a polyline vertex plays a dominant role in determining whether this vertex is a corner. In many cases (Figure 1.3), this determination is crystal clear to human observers. For example, a vertex where long horizontal and vertical lines meet is clearly a corner, while a point in the middle of a long horizontal line is not. In other cases, however, the answer is less obvious (Figure 2.1). The first stage of our algorithm detects potential corners by using such local geometric information. Theoretically, one could perform training on complete raster images or complete, closed boundaries, and consequently exploit both global and local context. Nevertheless, successful and generalizable training on such data would require significantly more training input. We produce the desired final vectorization by combining the data-driven predictions with a subsequent perception-based corner processing step (Chapter 4).

To detect potential corners, we first collect a range of corner/non-corner labels on a representative set of binary raster images (Section 3.2). We then use the labeled dataset to train a predictor which computes the likelihood for each given polyline vertex of being a corner based on the shape of the boundary




Figure 3.1: Given the input in (a), a purely data-driven classifier produces corner labels that are misaligned with human corner annotations (b), leading to undesirable vectorization artifacts (c). Starting with a more lax set of possible corners (d), we produce output labels (e) that align with the manual annotations and result in better vectorization (f). Input image adapted from Freepick – www.flaticon.com.

polyline around it. We also learn a probability cutoff threshold that we use to determine which vertices have sufficient probability of being corners; we include those in our initial set of potential corners (Section 3.1).

3.1 Learning Corner Likelihood

We aim to predict corner likelihood based on local neighborhoods. Hence, we associate each vertex p_i with a feature vector \mathbf{f}_i that encodes the local neighborhood centered at p_i : $\mathbf{f}_i = [p_{i-s} - p_i, \dots, p_{i-1} - p_i, p_{i+1} - p_i, \dots, p_{i+s} - p_i]^T$. The local neighborhood size is denoted by s and is selected via cross validation as will be discussed later. We use counter-clockwise ordering of vertices with respect to the region interior to distinguish between convex and concave segments. In our experiments, and consistent with the closure principle, humans expect corners to show up in concave areas more frequently than in convex ones.

To predict the corner likelihood of each vertex p_i , we take a supervised learning approach. We learn from a set of manually labeled feature representations $\{\mathbf{f}_i, y_i\}$, where $y_i \in \{0, 1\}$ is a binary label indicating whether p_i is a corner or not (Section 3.2). In particular, we employ random forests [Breiman, 2001] to build a non-linear mapping $\hat{y} : \mathbb{R}^{4s} \rightarrow \mathbb{R}$, such that $\hat{y}(\mathbf{f}_i)$ approximates the probability of $y_i = 1$ for any training pair. A random forest is an ensemble of decision trees, where each tree is a decomposition of the feature space into labeled partitions. The mapping $\hat{y}(\mathbf{f}_i)$ is found by averaging the labels of the partitions containing \mathbf{f}_i in all the trees. We use the random forest implementation provided by the *scikit-learn* package [Pedregosa et al., 2011].

We expect the predicted corner likelihoods to behave symmetrically on symmetric inputs, where the feature vector \mathbf{f} is rotated or reflected (as long as the counterclockwise ordering is maintained). Thus, for each feature vector \mathbf{f}_i , we generate symmetric replicas via rotation by multiples of 90° , and reflection along X and Y axes around p_i and include them in the training data with the same labels. The details of generating these replicas are elaborated in Algorithm 3.1. We further note that data augmentation at training time, on its own, might not ensure invariance under symmetry at prediction time. Therefore, instead of using the output of the random forest \hat{y} as the corner likelihood, we use a symmetric mapping, \tilde{y} , constructed by averaging \hat{y} over all the possible replicas:

$$\tilde{y}(\mathbf{f}) = \frac{1}{|\mathcal{F}(\mathbf{f})|} \sum_{\mathbf{f}' \in \mathcal{F}(\mathbf{f})} \hat{y}(\mathbf{f}'), \quad (3.1)$$

where $\mathcal{F}(\mathbf{f})$ is the set of all replicas of \mathbf{f} and $|\mathcal{F}(\mathbf{f})| = 8$ is the number of replicas.

We use 10-fold cross-validation on the training set to determine the hyper-parameters of our feature representation and the random forest classifier: (1) the neighborhood size s , (2) the maximum number of features to be considered in decision tree splitting, (3) the minimum number of samples required for splitting one node, (4) the minimum number of samples required for a node to be considered as a leaf, (5)

Algorithm 3.1 Data augmentation

function GENERATEREPLICAS**Input:** \mathbf{f}

▷ Feature vector

Output: \mathcal{F} ▷ The set of all replicas of \mathbf{f} $\mathcal{F} = []$ $\mathcal{F}.\text{insert}(\mathbf{f})$ $\mathcal{F}.\text{insert}(\text{ToggleDirection}(\text{MirrorBisector}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{ToggleDirection}(\text{MirrorX}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{ToggleDirection}(\text{MirrorY}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{MirrorBisector}(\text{MirrorY}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{MirrorBisector}(\text{MirrorX}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{MirrorY}(\text{MirrorX}(\mathbf{f})))$ $\mathcal{F}.\text{insert}(\text{ToggleDirection}(\text{MirrorBisector}(\text{MirrorY}(\text{MirrorX}(\mathbf{f}))))$ **return** \mathcal{F} **function** MIRRORBISECTOR(\mathbf{f})

▷ Helper function

Swaps the X and Y entries in the feature vector \mathbf{f} . This is equivalent to mirroring all the points in \mathbf{f} with respect to the bisector of the first quarter of the coordinate system.

function MIRRORY(\mathbf{f})

▷ Helper function

Multiplies the X entries in the feature vector \mathbf{f} by -1 . This is equivalent to mirroring all the points in \mathbf{f} with respect to the Y axis.

function MIRRORX(\mathbf{f})

▷ Helper function

Multiplies the Y entries in the feature vector \mathbf{f} by -1 . This is equivalent to mirroring all the points in \mathbf{f} with respect to the X axis.

function TOGGLEDIRECTION(\mathbf{f})

▷ Helper function

Inverts the order of the points in \mathbf{f} .

decision tree node splitting rule [Safavian and Landgrebe, 1991], (6) whether to enable bootstrapping or not, and (7) the number of decision trees. Specifically, given a set of randomly sampled hyper-parameters within a plausible range, we select the optimum values based on the highest overall F_1 -score, calculated over all possible detection thresholds within $[0, 1]$. We split the training set of polygons into 10 distinct groups. For a set of randomly sampled hyper-parameters, we predict the corner probabilities of each vertex in a group using a classifier trained on the remaining 9 groups. We then select the optimum values of the hyper-parameters that result in the highest overall F_1 -score, calculated over all possible detection thresholds within $[0, 1]$.

Using the learned corner detector as is, we can obtain about 80~90% precision at 90% recall depending on input resolution (Figure 3.2). These numbers are comparable to many other learning frameworks that seek to mimic human perception, such as [Lun et al., 2015]. From an application perspective, however, the vectorizations produced with 10% missing corners and another 10% false corners are frequently less visually appealing (Figure 3.1). Our experiments show that alternative training strategies do not lead to accuracy improvement (Chapter 7). We believe that the inconsistency between the prediction and ground truth is due to a combination of additional cues employed by viewers when they mentally identify




Figure 3.2: Precision-recall curves for random forest corner detectors obtained by cross-validation. Results for resolutions 32 to 256 are shown. Solid dots represent the operating points at which we detect the initial corner sets.

corners.

As an initialization for the next step, we seek to detect all the vertices that viewers may potentially perceive as corners. We set a detection threshold to attain a recall of at least 95% in cross-validation (indicated by the solid dots in Figure 3.2), such that most of the true corners have scores above this threshold and are included in the initial corner set. In our experiments, we use a threshold value of 0.125 on unnormalized classifier output which works well for all the resolutions.

3.2 Training Data

We collected a total of 158 semi-structured raster images as our training/validation data set, consisting of 76 artist-generated images of recognizable objects (downloaded from flaticon.com and rasterized) and 82 synthesized images of basic geometric shapes including circles, ellipses, orthogonal polygons, and French curves. Figure 3.4-a shows a subset of the images used for training, and the full set of training models is provided in the project website [Hoshyari et al., 2018]. The object shapes designed by artists better reflect the actual content we seek to process, while the geometric shapes provide baseline corner annotations. In cross-validation, we only withhold the samples of actual objects while always keeping the augmented geometric shapes in training. The training images are across a range of representative resolutions, with 8/16/32/64/128/256 pixels on each side. This range of resolutions is chosen for two reasons: first, the space our method targets is dominated by low to mid-range resolution inputs; second, for higher resolution inputs, using the accuracy constraint alone is typically sufficient to produce visually pleasing results. In order to meaningfully reason about the differences in human perception across resolutions, our corpus consists of image series which depict nearly-identical shapes across resolutions (for instance, depicting a similar bell with varying levels of details). We train and validate models for




Figure 3.3: Examples of short segments: (a) corners corresponding to boundary edge midpoints; (b) a two-pixel-long segment

each input resolution independently. Note that only basic geometric shapes are used as training data for extra-low resolutions (8 and 16), since at such resolution humans largely expect shapes to be continuous everywhere and are much more selective in their corner annotation. We found that a small representative set of geometries is sufficient for training in extra-low resolutions. The detailed data composition for each resolution is specified in the project website [Hoshyari et al., 2018].

All the input images were annotated by a graphics graduate student, with no prior knowledge of the project internals, who marked the perceived corners on the inputs using a simple GUI. Figure 3.4-b shows a subset of the annotated images. The complete set of annotated images is provided in the project website [Hoshyari et al., 2018].

3.3 Short Segments

Our corner classifier operates on boundary vertices. The underlying implicit assumption is that viewer perceived discontinuities align with raster boundary discontinuities. While true in general, there are two scenarios in which a perceived corner may, in fact, lie on a flat raster segment. Specifically, on symmetric inputs (Figure 3.3) corners may naturally correspond to boundary edge midpoints where both boundary end-points are raster discontinuities (Figure 3.3-a). Actual or perceived vector corners may also be rasterized as two-pixel long segments (Figure 3.3-b). In these cases, the perceived corner corresponds to the middle vertex of the segment. To handle both cases, when collecting training data, we asked the annotator to mark both ends of such corner edges or segments as corners (rather than the middle). We found that our classifier produces more accurate results if we mark the end-points of “corner” segment as corners instead of its mid-point. While this choice requires some minor adaptation in the subsequent processing, we found that overall it results in better fitting outcomes. Our corner removal step accounts for these choices as discussed in Section 4.2.




Figure 3.4: Examples of training data: (a) semi-structured raster images; (b) manual annotations (red points are marked as corners).

Chapter 4

Perception-Driven Corner Removal

We expect the initial set of corners to include all prominent corners human observers are likely to perceive, but to also contain extra corners which are not aligned with human perception (Figure 4.1 left). We eliminate these redundant corners using an iterative corner removal algorithm. We first compute an initial vectorization of the boundary (Section 4.1) using all the initial corners as C^0 discontinuities. We then repeatedly assess the impact of removing a subset of corners on the perceptual consistency of the updated vector boundary and eliminate the subset of corners which improves our consistency metric the most (Section 4.2). The process is repeated until consistency can no longer be improved.

Our perceptual consistency assessment (Section 4.2) accounts for the perceptual cues we expect humans to employ when mentally vectorizing raster boundaries: simplicity, accuracy, continuity, and closure. The simplicity principle suggests that a corner is unnecessary if by eliminating it we can simplify the boundary vectorization. The simplified boundary may have fewer fitting primitives or have simpler primitives instead of more complex ones (Figure 4.1, steps 2 and 3). Human observers expect the vector boundary to deviate somewhat from its rasterized counterpart. We therefore prioritize simplicity over accuracy and remove corners when the resulting vectorization is simpler, as long as the deviation of the fitted spline from the raster input remains within a specified threshold (Section 4.1). We similarly remove corners if the resulting post-removal vectorization remains as simple as before but has improved accuracy, i.e., better approximates the input polyline, as shown in Figure 4.1, steps 4 and 5.

The continuity principle argues for avoiding close-by corners that form short boundary segments. We prioritize removing such segments in our iterative process. In selecting which of the corners to remove, we account for continuity in addition to the principles listed above (Section 4.2.1).

Lastly, closure argues for joint corner/non-corner labeling of pairs of vertices which bound viewer imagined contours (Figure 2.1). We enforce similar labeling by jointly considering paired corner vertices at each removal iteration; we remove them only if doing so improves simplicity or accuracy across all affected boundary segments. We use the raster geometry together with the fitted curves adjacent to each corner to determine the pairings (Section 4.2.2).




Figure 4.1: Iterative corner removal (left to right): initial corners and their induced vector boundary; incremental corner removal steps (with associated type and fit metrics); final corners and vector boundaries (unregularized). The red rectangle highlights the removed corner after each step. At step 1 short segments are processed, at steps 2 and 3 the type score decreases, and at steps 4 and 5 the fitting error decreases. Input image adapted from Freepick – www.flaticon.com.

4.1 Piecewise Smooth Vectorization

Given a set of corners $c_i, i \in [0 \dots n]$, that delineate raster boundary segments $L_i, i \in [0 \dots n]$, we seek a boundary vectorization that is continuous everywhere except at these corners, and that balances the number and complexity of the fitted primitives against fitting accuracy. In assessing accuracy, we both minimize the distance from the fitted curves to the raw input boundary and explicitly constrain the amount of allowable local deviation.

4.1.1 Fitting Energy

We measure fitting accuracy by assessing the distance between the fitted primitives and the corresponding boundary edge midpoints and segment corners. More formally, for each boundary segment L_i connecting consecutive corners, we seek to fit a set of $K_i + 1$ primitives $\{C_i^k(t)\}_{k \in [0 \dots K_i]}, t \in [0, 1]$, such that each primitive C_i^k approximates the midpoints of pixel edges in its corresponding segment L_i^k , and the combined spline curve consisting of these primitives is close to its corner points. We jointly optimize for all curve sections by minimizing the approximation error together with the fit simplicity, by minimizing:

$$\mathcal{E} = \mathcal{D} + \frac{1}{\alpha} \mathcal{R} \quad (4.1)$$

$$\mathcal{D} = \sum_{i,k} \sum_{e_j \in L_i^k} \min_{t \in [0,1]} \|C_i^k(t) - m_j\|_1 + \sum_i [\|C_i^0(0) - c_i\|_1 + \|C_i^{K_i}(1) - c_{i+1}\|_1] \quad (4.2)$$

$$\mathcal{R} = \sum_{i,k} r(\text{Type}_{C_i^k}) \quad (4.3)$$




Figure 4.2: Piecewise smooth vectorization.

The total energy \mathcal{E} consists of a data approximation term \mathcal{D} and a simplicity term \mathcal{R} . The first part of \mathcal{D} measures the shortest distances from the primitives to their corresponding midpoints, where e_j is the j -th segment edge and m_j is the corresponding midpoint (see Figure 4.2). The second part of \mathcal{D} measures the distances between the spline endpoints and the corresponding segment corners. We use the L_1 norm for measuring fitting accuracy as we seek to penalize large local deviations between the vector output and its raster input while being more tolerant to small deviations. In the simplicity term, $r(\cdot)$ is a discrete complexity cost based on primitive type, $Type_C$. We empirically set r to 1 for a straight line, 2 for an arc and 4 for a clothoid. These scores reflect the increase in curve complexity or fitting ability of each primitive – intuitively a clothoid can roughly fit the

same set of midpoints as two arcs with comparable accuracy, and an arc often can roughly approximate the midpoints spanned by two lines. The weight α is designed to balance the two terms. Intuitively, as image resolution (boundary length) increases, the fitting error will aggregate; thus, to keep the type versus accuracy balance constant, we use a weight inversely proportional to image resolution, setting $\alpha = 32/rs$ where rs is the input resolution.

4.1.2 Constraints

We impose three sets of constraints on the fitted splines: continuity, accuracy, and tangent. We require the curve sections to satisfy C^0 continuity at corners, i.e., between consecutive primitives that share common corner vertices $C_i^k(1) = C_{i+1}^0(0)$, and G^1 continuity between all other pairs of consecutive primitives: $C_i^k(1) = C_{i+1}^k(0)$, $\bar{C}_i^k(1) = \bar{C}_{i+1}^k(0)$ where $\bar{C}_i^k = \frac{dC_i^k}{dt} / \|\frac{dC_i^k}{dt}\|_2$.

Strictly enforcing the accuracy constraints requires evaluation at every pixel next to the processed boundary to test if its center is inside or outside the region enclosed by the vectorization – a non-trivial and non-local test. We approximate this constraint as follows. Accuracy implies that the vector boundary should generally lie in-between the pixel centers immediately inside and outside the raster boundary (see Figure 4.2). To encode this constraint, we define a line interval $I_m(t)$ centered at each midpoint m associated with L_i^k :

$$I_m(t) = m + (t - 0.5) \cdot u_m, \quad t \in [0, 1], \quad (4.4)$$


where u_m is a unit coordinate $(1, 0)$ or $(0, 1)$ depending on whether the pixel edge associated with m is oriented vertically or horizontally. The accuracy constraint requires each primitive C_i^k to intersect the intervals that correspond to the midpoints it spans:

$$\min_{s, t \in [0, 1]} \|C_i^k(s) - I_{m_j}(t)\|_2 = 0, \quad \forall e_j \in L_i^k. \quad (4.5)$$

In practice, we seek to allow small deviation from the strict intervals, as this leads to a simpler and

visually more appealing fitting output. Specifically we require the curve to either satisfy Equation (4.5) or to pass within ϵ -distance ($\epsilon = 0.1$) from one of its end points:

$$\min_{s \in [0, 1]} \|C_i^k(s) - I_{m_j}(1)\|_1 \leq \epsilon \text{ or } \min_{s \in [0, 1]} \|C_i^k(s) - I_{m_j}(0)\|_1 \leq \epsilon \quad (4.6)$$



Lastly, we note that human observers expect accuracy not only in terms of absolute proximity but also in terms of vector versus raster tangents. While tangents on raster data are clearly ill-posed, we note that at corners viewers expect the vectorized curve tangent to be in the same half-space (inset, green) as the polyline itself with respect to the polyline edge immediately emanating from the corner. Fits that violate this property (inset, red) appear

as counter-intuitive. We enforce this property by enforcing the segment tangents $\bar{C}_i^k(1)$ and $\bar{C}_i^{k+1}(0)$ to be within the relevant half-spaces. We compute the half-space by locating the first polyline edge away from the corner orthogonal to the edge emanating from it. Figure 4.4 shows the effect of enforcing this constraint.

4.1.3 Optimization

Our optimization goals require solving a discrete-continuous problem: we need to determine the number of primitives (K_i+1) for each segment L_i , associate each primitive C_i^k with its corresponding midpoints, select the primitive types, and optimize their parameters. We obtain an approximate solution to this challenging discrete-continuous optimization problem by adapting the Cornucopia fitting algorithm proposed by Baran et al. [2010]. This method is designed to compute fitted splines which similarly balance accuracy and simplicity, and supports the types of primitives we use. The main advantage of the framework is an efficient approximation of the discrete component of the problem that matches primitives to midpoints with a simple shortest cycle algorithm. Before further discussion of the solution method, we note that a similar but simpler discrete-continuous optimization problem is introduced in Appendix B for further clarifications.

We modify the framework to account for the accuracy constraints as follows. For every pair of midpoints m_i and m_j which do not have a corner between them, we first try to fit all the primitive types to approximate the sequence of points bounded by m_i and m_j (fitting a curve with known endpoints is described in Appendix A). We constrain the fitted primitives to satisfy our half-space constraints. We eliminate all primitives which do not satisfy the accuracy constraints. We then obtain the fitted spline by using the Cornucopia shortest cycle computation framework. We define a graph whose vertices are the different fitted primitives and whose edges denote the connectivity relations between these primitives. We assign a weight to each vertex equal to the sum of its L_1 fitting error and type cost. The edge weights are then initialized as the average weight of their endpoint vertices. An additional cost is added to each edge to approximate the extra fitting error caused by enforcing the C^0 or G^1 continuity between the

primitives connected by this edge [Baran et al., 2010]. For an edge representing a C^0 continuity between curves C_1 and C_2 , the additional cost is calculated as the distance from the endpoint of C_1 to the start point of C_2 , i.e.,

$$\text{additional cost} = \|C_1(1) - C_2(0)\|_1. \quad (4.7)$$

If the edge represents a G^1 continuity, the angles are also included in the additional cost to penalize discontinuities (Figure 4.3-b). In this case,

$$\text{additional cost} = \frac{\text{length } C_1 + \text{length } C_2}{2} (\|C_1(1) - C_2(0)\|_1 + \arccos(\bar{C}_1(1) \cdot \bar{C}_2(0))). \quad (4.8)$$

Once the shortest cycle is computed, we recompute the parameters of each selected primitive by enforcing continuity constraints (Appendix A describes the details of fitting multiple curves with transition constraints at prespecified locations). If any resulting primitive violates the accuracy constraints, we assign an infinitely large weight to all the edges associated with it and repeat the shortest cycle computation. In our experiments, this process typically converges in under 5 iterations. Algorithm 4.1 summarizes the optimization process, while Figure 4.3 visualizes its main steps.

Algorithm 4.1 Optimization pseudo-code

```

function FITCURVES
  Input: points                                ▶ Circular vector of ordered points to be fit with curves
  Input: isCorner                               ▶ Bit vector specifying if a point is a corner
  Output:  $\mathcal{C}$                                 ▶ The sequence of fitted curves
  n=size(points); edges=primitives=[]
  for i=1 to n, j=i to i+n, type in {line, arc, clothoid} do
    if sum(isCorner(i:j))==0 then
      primitives.pushBack( fitCurve(points(i:j), type) )    ▶ Fitting algorithm, Appendix A
  for all  $C_1$  and  $C_2$  in primitives do
    shouldAddEdge=(isCorner( $C_1$ .endPoint) and  $C_1$ .endPoint== $C_2$ .startPoint) or
                  (!isCorner( $C_1$ .endPoint) and  $C_1$ .endPoint== $C_2$ .startPoint+1)
    if shouldAddEdge then
      edge=( $C_1, C_2$ )
      findApproxCost(edge)                                ▶ Edge costs as described in Section 4.1.3
      edges.pushBack(edge)
  G=makeGraph(V=primitives, E=edges)
  repeat
     $\mathcal{C}$ =shortestCycleVertices(G)                        ▶ Shortest cycle algorithm from [Baran et al., 2010]
     $\mathcal{C}$ =enforceContinuities( $\mathcal{C}$ )                          ▶ Half-space,  $C^0$ , and  $G^1$ , Appendix A
    areConstraintsSatisfied=checkConstraintSatisfaction(G,  $\mathcal{C}$ )    ▶ Interval intersection
    updateEdgeCosts(G,  $\mathcal{C}$ )                                ▶ Assign an  $\infty$  cost to path edges of violating primitives
  until areConstraintsSatisfied
  return  $\mathcal{C}$ 

```




Figure 4.3: Optimization stages shown: (a) examples of fitted primitives; (b) examples of primitives connected by graph edges (high edge weights penalize discontinuities when G^1 continuity is enforced, see insets); (c) the shortest cycle in the graph; (d) final fitted primitives with the half-space constraint and the transition continuity enforced. Note that for illustrative purposes we only show a small selection of the total fitted primitives for parts (a) and (b). Input image adapted from Freepick – www.flaticon.com.

4.2 Corner Removal Iterations

At each iteration of corner removal we consider different subsets of current corners as removal candidates. We note that the impact of removing a sequence of consecutive corners is local – affecting only the segments they connect. Thus in general we can assess different consecutive sequences separately. The only exception to this observation is the processing of corners paired via closure (Section 4.2.2); when assessing the impact of removing a corner we need to simultaneously consider all corners paired with it.

To assess the impact of removing one corner, and its paired corners if these exist, we first compute a boundary vectorization with this corner vertex (and its pairs) excluded from the current corner set. We then compare the two resulting vectorizations. We add the corner to the removal candidate list if the vectorization complexity cost \mathcal{R} (Equation (4.3)) decreases following the removal, or if this cost remains the same but the fitting cost \mathcal{D} (Equation (4.2)) decreases (Figure 4.1). After iterating over all corners, we select the corner from the removal candidate list whose removal results in a vectorization with the lowest overall cost \mathcal{E} , and remove this corner vertex (and its paired corners) from the current corner list. We repeat this process for multiple iterations, until the removal candidate list is empty. Despite the greedy nature of this iterative process, we did not observe convergence to bad local minima in our experiments. Algorithm 4.2 summarizes the corner removal algorithm.

4.2.1 Short segments

The set of corners produced by our classifier is typically well-spaced, but can contain some close by corners – here we consider a pair of corners as close if the distance between them is shorter than 10% of the image resolution. In general viewers do not expect to see such short segments in their vectorization as they violate our expectation of visual continuity. Extra-short segments with one or two pixels are also a side effect of the discrete nature of corner detection, and the consequent corner labeling provided to training which replaces corners visually placed on edges or flat boundary segments with

Algorithm 4.2 Iterative corner removal

```
function CORNERREMOVAL(polygon, corners)
    curves = fitCurves(toMidPoints(polygon), toBitVector(corners))           ▶ Algorithm 4.1
    for all lines in curves with length  $\leq 3$  do
        corners = processSmallSegment(line, corners)           ▶ Short segment processing, Section 4.2.1
    cornerGroups = globalAnalysis(corners)           ▶ If required, group corners together, Section 4.2.2
    repeat                                           ▶  $\Delta$  denotes change after removing the corner group
         $g = \operatorname{argmin}_{\text{cornerGroups}}(\Delta\mathcal{R})$  subject to  $\Delta\mathcal{R} < 0$ 
        if  $g \neq \text{NULL}$  then remove  $g$ ; continue
         $g = \operatorname{argmin}_{\text{cornerGroups}}(\Delta\mathcal{E})$  subject to  $\Delta\mathcal{R} = 0$  and  $\Delta\mathcal{E} < 0$ 
        if  $g \neq \text{NULL}$  then remove  $g$ ; continue
        terminate
    until termination
    return remaining corners
```




Figure 4.4: Effect of the half-space constraint and the short segment processing step. Input image adapted from Freepick – www.flaticon.com.

those nearby non-flat boundary vertices (Section 3.3). We prioritize short segment removal and attempt to remove such segments before executing the main loop. We only consider segments that in the current vectorization are fit with individual line primitives, as we expect any higher complexity primitive to reflect actual necessary geometry. We measure the angles at the ends of the line primitive fitted ϕ_1 and ϕ_2 . We use these angles to assess if the primitive should be treated as symmetric or not, and classify it as symmetric if $|\phi_1 - \phi_2| < 2^\circ$. For non-symmetric primitives, we first apply the default removal test based on simplicity and accuracy, and remove the corner which improves \mathcal{R} the most. If neither corner is removed, we remove the corner with the more obtuse angle, if that angle is above 120° and \mathcal{R} increases by less than 1 (i.e., conceptually we replace a line by an arc). If one corner is convex and one concave, we only consider the convex one. For symmetric primitives, if the cost \mathcal{R} improves by removing both corners or if both angles are above 120° and the cost increase is less than 1 we remove both.


For length one segments we simply replace the two corners with one mid-edge corner. For length two segments we apply the same method but with relaxed thresholds, we classify the primitive as symmetric if $|\phi_1 - \phi_2| < 15^\circ$ and use an angle threshold of 100° instead of 120° . In addition, for primitives deemed symmetric we assess the cost of moving the corner to the middle, and perform this move if the cost improves. Figure 4.4 shows the effect of this process, and Algorithm 4.3 summarizes its main steps.

Algorithm 4.3 Short segment processing

```
function PROCESSSMALLSEGMENT(segment, corners)
  if seg.length  $\leq 1$  then replace the segment with a mid-edge corner; return
  else if seg.length  $\leq 2$  then thr =  $120^\circ$ ; tol =  $2^\circ$ 
  else if seg.length  $\leq 3$  then thr =  $100^\circ$ ; tol =  $15^\circ$ 
  else return
  if  $|\phi_1 - \phi_2| < \text{tol}$  then                                 $\triangleright$  Symmetric case ( $\phi$  values are the segment incident angles)
    if  $\Delta\mathcal{R} < 0$  when replacing the segment with a corner at the midpoint then
      Replace the segment with a corner at the midpoint; return
    if  $\Delta\mathcal{R} < 0$  after removing both sides then remove both sides; return
    if  $\Delta\mathcal{R} < 1$  after removing both sides and  $|\phi_1|, |\phi_2| > \text{thr}$  then remove both sides; return
  else                                                         $\triangleright$  Unsymmetric case
    i =  $\text{argmin}_{1,2}(\Delta\mathcal{R} \text{ after removing side } i)$  subject to  $\Delta\mathcal{R} < 0$ 
    if i  $\neq \text{NULL}$  then remove side i; return
    i =  $\text{argmin}_{1,2}(|\phi_i|)$  subject to  $\phi_i$  convex and  $\phi_i < \text{thr}$ 
    if i  $\neq \text{NULL}$  and  $(\Delta\mathcal{R} \text{ after removing side } i) < 1$  then remove side i; return
```

4.2.2 Computing Global Context Cues

There are multiple standard methods for the convex partitioning of closed curves, e.g., [Lien and Amato, 2004]. However, none are applicable to raster boundaries where the question of how to define convexity even on simple shapes (circle in Figure 1.3) is not evident. We use two criteria to identify pairs of concave corner vertices that viewers are likely to perceive as connected following the closure principle (Figure 2.1). First, we look for pairs of corner points which are closest vis-a-vis the model medial axis, i.e., ones with a common circle that touches both points and is entirely inside the processed region. We use a scaled version (by a factor of 0.90) of the common circle for the inside test to handle the inaccuracy inherent to raster data. We augment this test with a second closure test designed for corners which are visually paired due to local geometry but not necessarily closest vis-a-vis the medial axis.



This test is motivated by the continuation principle that suggests that viewers similarly partition shapes along invisible contours when these contours smoothly extend the outer contour (see inset). We apply this test to pairs of concave points that are internally visible from one another, i.e., ones that can potentially define an imaginary internal contour that partitions the shape. We pair the points if the angles between the line connecting the tangents of the fitted splines on the left or right of them are under 5° .

Chapter 5

Boundary Regularization

After computing the vectorization, we further regularize it by locating splines and individual primitives with similar but not identical characteristics and enforcing those to be identical. The set of characteristics we enforce include parallel and co-linear lines, and co-circular curves, similar to [Mehra et al., 2009; Li et al., 2011].

We apply the same greedy regularization scheme to each group of primitives (or entire splines) which can potentially be jointly regularized. Specifically, given a group of primitives that can potentially be fitted jointly, such as a set of near-parallel lines, we order them by degree of compatibility. We then start by fitting the closest pair, and then incrementally repeat the fitting each time adding the next closest candidate primitive. At each step, we test if the fitted result satisfies our accuracy constraints and stop the process if it fails. To facilitate stronger regularity, we relax the accuracy threshold ϵ to 0.2 (Equation (4.6)).

We enforce regularities in the following order. We first detect all nearly axis-aligned line segments, i.e., lines whose directions deviate from the vertical or horizontal axis no more than 10° , and force them to be strictly axis-aligned. We then consider all circle primitives used and assess if they can be used to fit other boundary sections currently described by straight lines – we note that by default our fitting method prefers lines to circles; thus a section that can be fitted either by an arc or by a line will be




Figure 5.1: Effect of regularization. Lines and circles connected by dashed lines share parameters. Input image adapted from Freepick – www.flaticon.com.

locally fitted by a line. However, from a simplicity perspective using the same circle across multiple sections is better than using different primitives. We re-fit each group of identified sections with the same circle parameters following the process above. We next consider all the segments defined by straight line primitives, and group those that can be described with the same line. Similarly, we proceed to handle near-parallel lines, and nearly orthogonal or almost identical corner tangents. Lastly, we improve continuity by considering corners with nearly identical tangents on either side and re-fitting them using a continuous vector spline. In our angle computation, we consider tangents to be nearly identical if they are within an angle difference of 20° , a threshold at which humans perceive two connecting curves with corresponding tangent angle as continuous [Hess and Field, 1999]. Figure 5.1 shows the before and after regularization results.

Chapter 6

Multi-Color Inputs

The work presented in this chapter has been developed by Edoardo Alberto Dominici under supervision of Alla Sheffer and is provided here for completeness.

Handling of multi-color inputs raises a few questions not addressed by the single boundary setup. First, we need to resolve topological ambiguities at corners [Kopf and Lischinski, 2011]. We use the framework by Kopf and Lischinski for this task. We also use their framework to extract the region boundary graph. Edges in this graph correspond to raster boundary segments shared by two regions and vertices correspond to raster boundary vertices shared by three or four regions. Our second task is to determine the location of corners along the boundaries. We use our corner detector trained at the closest power of two resolution to find the corner probabilities for each region boundary independently. We classify vertices along segments shared by two regions as corners if it is deemed as a corner with respect to both regions (we prefer to err on the side of more continuous solutions as viewers are more sensitive to redundant than to missing corners).

Lastly, we seek to determine the continuity at boundary vertices adjacent to multiple regions. We note that with accuracy constraints in place at vertices adjacent to three regions we can at most enforce one boundary to be continuous. At valence four vertices two crossing boundaries can also be continuous. We determine which continuity constraints to enforce as follows. If a multi-region vertex is considered a C^0 corner by all the boundaries, or all but one we use this classification as is. Similarly, if a vertex has a corner vertex within one pixel distance on any boundary, we classify it as a corner with respect to that boundary. We resolve conflicts where two or more boundaries have the vertex marked as non-corner using angles. As shown in Figure 6.1, we vectorize the conflicting regions independently, treating this vertex as corner on all. We then measure the angles between the tangents at this vertex and greedily mark the boundary with the most obtuse angle as the continuous one, and mark the vertex as a corner for the rest.

All resulting boundary segments are then fitted individually using a similar pipeline as the single color. The resulting vectorized boundary segments are not guaranteed to be connected. Our final step connects




Figure 6.1: Removing ambiguities when multiple color regions meet at the same point: (a) three regions meeting at the same vertex; (b) the shared point is classified as a corner on only one of the boundary regions; (c) each region is vectorized separately with the point of interest as a corner, and the angles between the tangents are measured; (d) the transition between the segments with the most obtuse angle is marked as G^1 , while the rest are marked as C^0 .

all segments by performing an additional global fitting step with positional constraints for the segment endpoints. If all segments meeting at a shared multi-region vertex are C^0 continuous at it, we simply constrain their end vertices to the average of their current positions. If one segment is G^1 continuous, we average end vertex positions of the other segments and project this average to the curve. We then refit the curves with this positional constraint.

Chapter 7

Results and Validation

We tested our method on a large set of close to 200 images, at different pixel resolutions ranging in size from 16×16 to 256×256 ; we also tested a range of multi-color inputs with region sizes as small as 2×2 . We show a range of representative results throughout the thesis (e.g., Figures 7.1 and 7.2). Numerous additional results are also provided in the project website [Hoshyari et al., 2018]. As our work focuses on boundary vectorization, our test sets are dominated by images that consist of two regions separated by a single closed boundary. While these inputs best illustrate the impact of our algorithmic choices, we also include various multi-region and multi-color examples (Figures 1.1 and 7.3). Our results look consistent with human expectations across the board. Besides visual inspection, we also perform a thorough evaluation by comparing our method to artist created results and algorithmic alternatives.

7.1 Perceptual Ground Truth Comparison

We asked an artist to manually vectorize 15 representative inputs. Figure 7.4 shows 2 of these artist produced vectorizations side-by-side with our results. Additional results are provided in the project website [Hoshyari et al., 2018]. While it took the artists on average 30 to 45 minutes to generate each example vectorization, our method automatically produces qualitatively similar results in a fraction of this time.

7.2 Algorithmic Alternatives

We also compared our results to those generated by a range of available methods. We compared our results to three publicly available vectorization methods: Adobe Trace [Adobe, 2017], Vector Magic [Vector Magic, 2017] and Potrace [Selinger, 2003]. Additionally, we tested against state-of-the-art upscaling [Stepin, 2003] and super-resolution [Wang et al., 2015] methods. Lastly, we compared against the polyline fitting method of [Baran et al., 2010]. Since this method fits a piecewise smooth curve

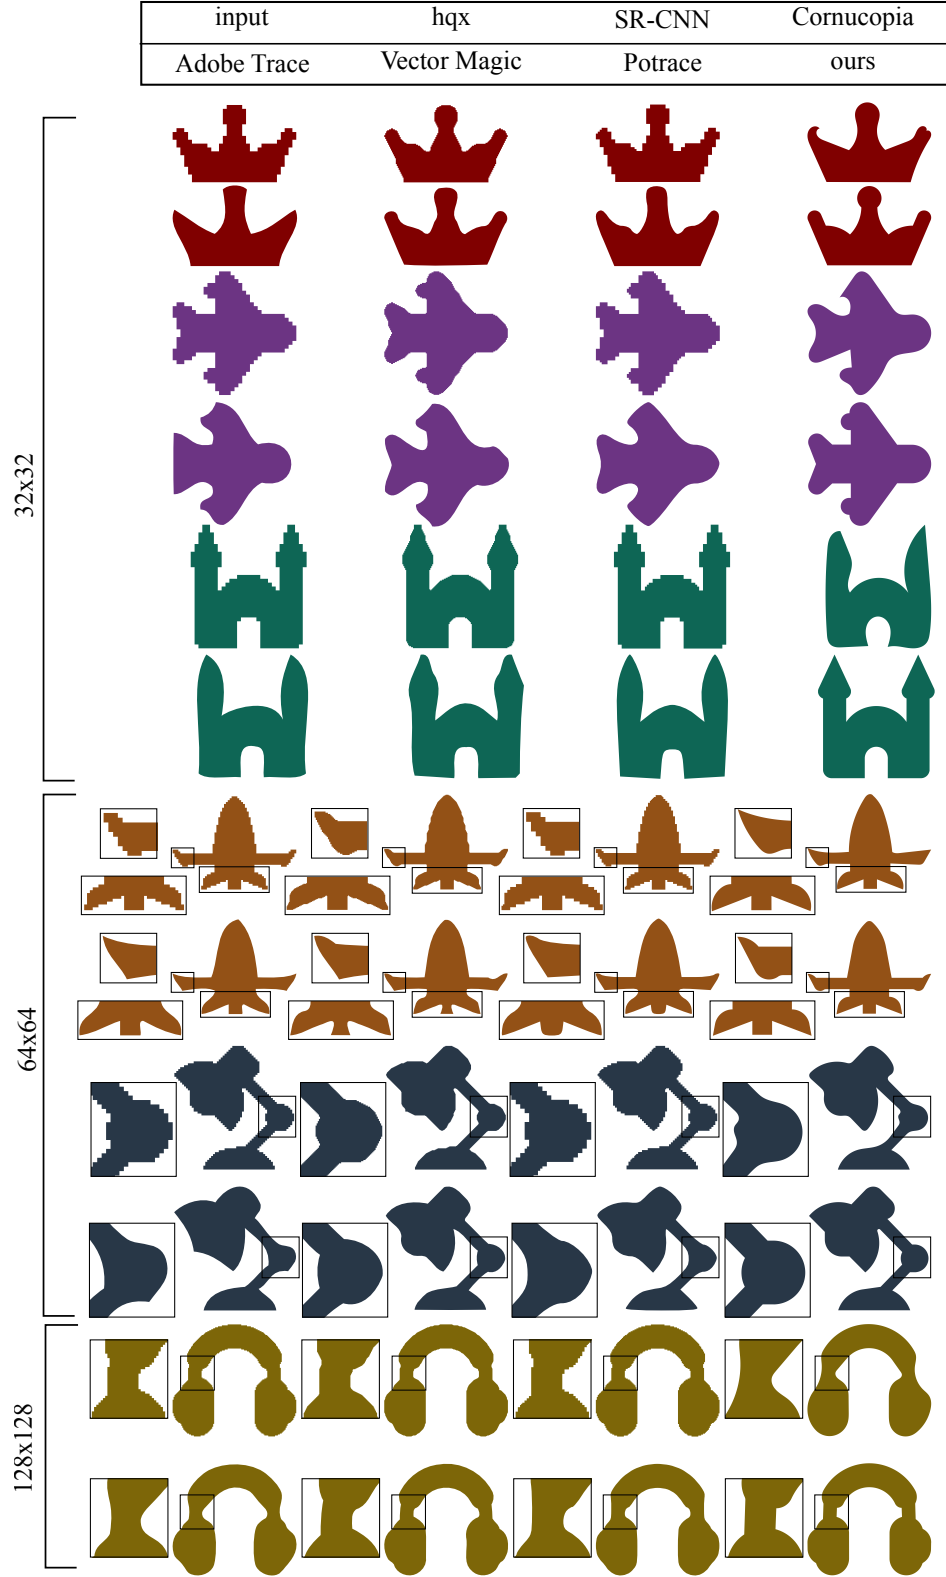


Figure 7.1: For inputs at various resolutions, we compare our method to an upscaling method (hqx [Stepin, 2003]), a super-resolution method (SR-CNN [Wang et al., 2015]), smooth polyline fitting [Baran et al., 2010], and three vectorization methods: Adobe Trace [Adobe, 2017], Vector Magic [Vector Magic, 2017], and Potrace [Selinger, 2003]. Input images adapted from Freepick – www.flaticon.com.



Figure 7.2: Vectorized multi-color examples. Bear, penguin, and giraffe images © pushnovaliudmyla – stock.adobe.com. Crocodile image © MarinaMays – stock.adobe.com.

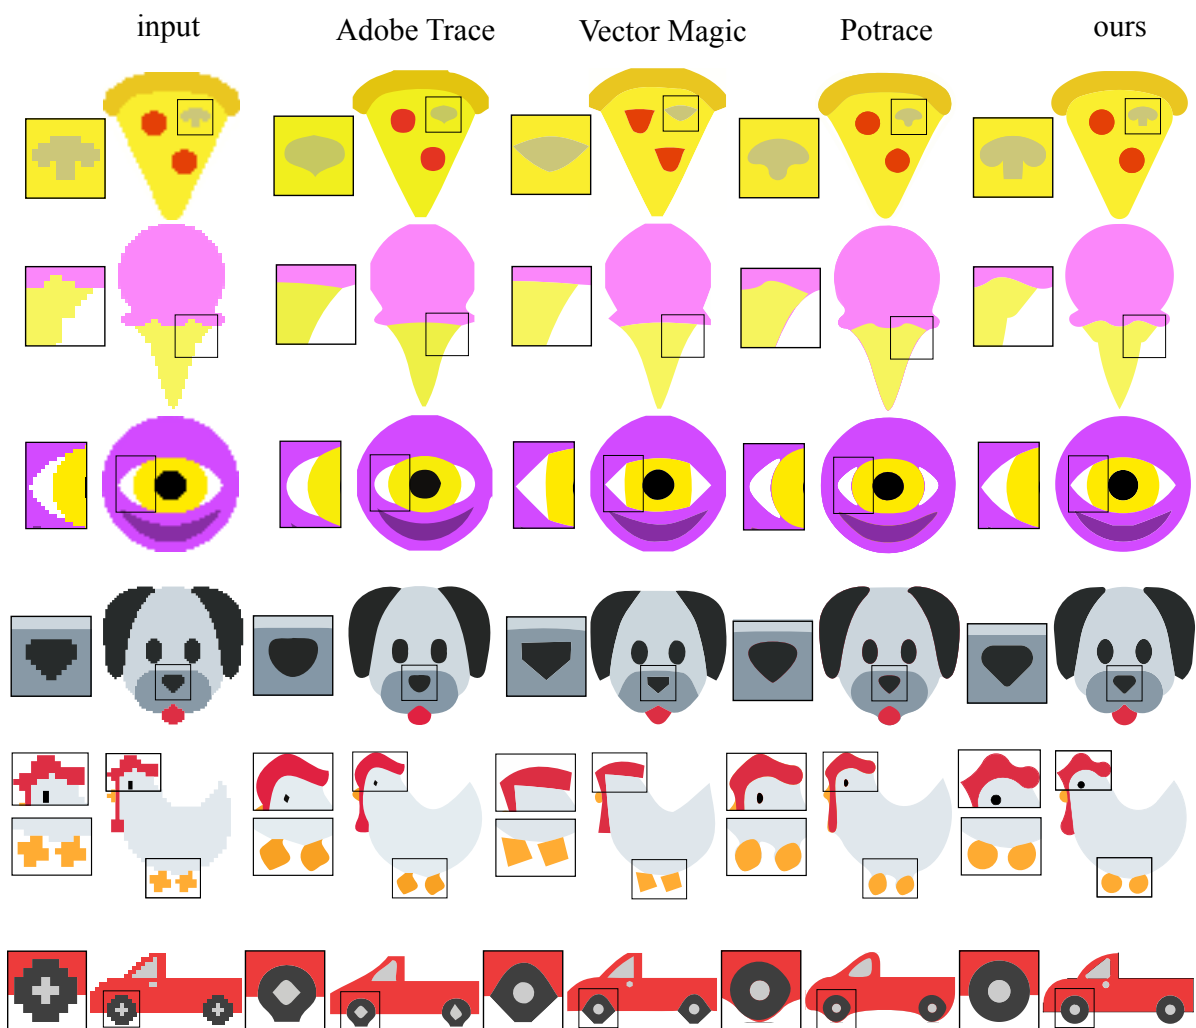


Figure 7.3: Vectorization of multi-color inputs. Dog and rooster images © Twitter, Inc.




Figure 7.4: Artist produced vectorizations shown (2 out of a total of 15) side-by-side with our results. In our user study, participants rated our outputs equal in quality to those manually produced by artists. Input images adapted from Freepick – flaticon.com.

to a polyline, we used boundary midpoints as input polyline vertices. In all cases, we adjusted method parameters to produce results that most closely resemble human expectations. For these comparisons, we focus on input resolutions of up to 128 and show representative examples at resolutions 32, 64 and 128 in Figure 7.1. We observe that, at higher resolutions, fitting accuracy dominates all other cues with all fitting and vectorization methods producing similar results. The small components of multi-region examples, as shown in Figure 1.1 and the additional results in the project website [Hoshyari et al., 2018], implicitly provide examples of lower resolution input. Our method produces outputs that are significantly more consistent with viewer expectations for these regions.

While we fundamentally focus on a different problem, we also provide comparisons to the method of Kopf and Lischinski [2011] in Figures 7.5 and 7.6. Kopf et al. focus on resolving topological ambiguities between multiple regions that touch at a pixel corner, and use standard spline fitting once these ambiguities are fixed. In contrast, we focus on vectorization of semi-structured region boundaries. While many of the inputs shown by Kopf and Lischinski have irregularly shaped regions only a few pixels large, we typically focus on resolutions where region shape or structure is more pronounced. On such inputs, our methods outperforms the method of Kopf and Lischinski [2011] (Figure 7.5). Small (1–2 pixel wide) irregularly shaped regions are best fitted using simple splines rather than our framework (Figure 7.6).

Lastly, we compare the results of our complete algorithm to those produced using the learned classifier alone, without the progressive corner removal (Figures 3.1 and 7.7).

7.3 Qualitative Evaluation

To further confirm that our results are well aligned with human expectation, we conducted a comparative user study. For each question, we showed participants the input, our result, and one of the algorithmic alternatives side by side. We then asked them to select the result that resembles the input the most.




Figure 7.5: Structured pixel-art vectorization comparison. User avatar image adapted from SimpleIcon – flaticon.com.



Figure 7.6: Comparison to Kopf and Lischinski [2011] for non-structured pixel-art.




Figure 7.7: Comparison of our corner removal algorithm (right) to using only a pure classifier (center) for detecting corners. Input image © VectorShopDesign – www.etsystudio.com.




Figure 7.8: The percentages of user preference to our method and alternative methods under different input rasterization resolutions.

Specifically, we showed the input image on top identified as “A” and the two vectorizations on the bottom identified as “B” and “C”. We then asked: “Which of the two images on the bottom “B” or “C” better represents image “A”? If both are equally good then select “Both”, and if neither represent “A” then select “Neither””. The answer options were “B”, “C”, “Both”, and “Neither”. Examples of the user study questions are provided in Appendix C, while the set of all questions are available in the project website [Hoshyari et al., 2018].

We included the artist-generated (ground truth) vectorizations as well as the results of all six baseline methods discussed above. All together, we composed a survey which had a total of 195 queries, 180 comparing against two alternative methods and 15 against the artist vectorizations. We conducted the survey via the Mechanical Turk interface, where each participant was shown 22 randomly selected queries with each query shown twice with “B” and “C” switched. For each participant, we discarded inconsistent answers where they chose different answers to the duplicated query, and discarded all answers from inconsistent participants who answered inconsistently over 60% of the queries. The results are summarized in Figure 7.8. As demonstrated, participants consistently preferred our result over the algorithmic alternatives, and rated our outputs as equally good as the manually produced results (Figure 7.4). These findings validate that our framework produces results on par with manually produced ones and superior to existing or potential alternatives.

In comparisons with alternative automatic schemes, we found only 3 queries (out of 180 comparisons across 30 inputs) where a plurality of respondents preferred an alternative output; one of those is shown in Figure 7.9-a, the others are included in the project website [Hoshyari et al., 2018]. On 2 (out of 15) inputs, viewers consistently prefer manually vectorized content to ours (Figures 7.9-b and -c). We hypothesize that humans often rely on semantic understanding of shapes to make judgments during vectorization, an interesting direction for future exploration.




Figure 7.9: Failure cases. Input images adapted from Freepick – flaticon.com.

Table 7.1: Comparison between neural network and random forest on corner detection accuracy. Precision/recall/ F_1 -score are reported for resolutions of 32, 64, and 128.

Resolution	Precision/Recall/ F_1	
	Neural network	Random forest
32	0.793/0.847/0.819	0.835/0.910/0.871
64	0.944/0.770/0.848	0.907/0.850/0.877
128	0.939/0.837/0.885	0.923/0.880/0.901

7.4 Learning Statistics

Finally, we validate one of the key components of our framework, the corner detection algorithm, using one-out cross validation on the training data of real object shapes. When evaluated against human annotation, our random forest classifier achieves F_1 -scores ranging from 0.87 to 0.90 on different resolutions (right column of Table 7.1). We also perform the same evaluation on an alternative learning algorithm – the multilayer perceptron neural network. We use the same setting to train a network model with two layers of hidden nodes for each resolution. Overall, the neural network achieves slightly worse results (middle column of Table 7.1) than the random forest employed in our framework.

7.5 Discussion on Black Box Learning

We train a corner classifier from a considerably small amount of training data, since it is unrealistic to generate a large collection of artist generated vectorizations of training raster images (in our experiments an artist took half an hour to vectorize a moderately complex input). We adjust the classifier parameters to avoid false negatives and focus on removing spurious corners in the subsequent stages of our algorithm. Theoretically, one could explore adding further training data to improve our results. A naïve source of training data for computation of R^P or its components could be pairs of vector shapes and their algorithmically generated rasterizations. However, such an inverse mapping approach would likely produce training data at odds with the vectorization solution perceived by humans.

7.6 Resolution Dependence

We use different classifiers for different region resolutions. Thus increasing image size while keeping region size the same has no impact on our results. We also experimented with using the classifier trained on the 32px data on higher resolution inputs. Surprisingly, the results were practically on par with those generated with the higher-res classifiers. In contrast, using higher-res classifier on coarser resolutions results in false negatives where perceived corners are not properly detected.

7.7 Performance

We measure the execution time of our method on a machine with AMD Ryzen 7 1800X 8 CPU @ 3.6Ghz with 16 logical cores. While predicting corner probabilities for all the vertices takes around 0.2–0.3 seconds, fitting different possible curve types to each segment between consecutive corner points and choosing the initial primitive type takes about 0.1–0.5 seconds. The corner removal execution time varies between 1–9 seconds for resolution 32, and 10–40 seconds for resolution 64. The final step of detecting and enforcing regularities takes 0.4–3 seconds for resolution 32 and 4–8 seconds for resolution 64. For input resolutions 128 and 256, fitting a piecewise smooth spline to the boundary takes 8 and 39 seconds, respectively. We note that our code is not optimized and there is a lot of room for improvement.

Chapter 8

Conclusions

We have presented a novel vectorization method specifically targeted at artist-generated semi-structured raster images, ones consisting of a small number of uniformly colored regions, was presented. Human observers consistently mentally vectorize such input and consequently expect algorithmic outputs to agree with this mental vectorization. However, these inputs pose a unique challenge for traditional vectorization methods, as contrary to other inputs. We presented the first method for vectorization of such inputs specifically designed to align with human perception. The highlights of our vectorization method include: conservative initial data-driven corner detection, perception-driven corner removal, minimization of a discrete-continuous fitting energy, and boundary regularization. As demonstrated by a range of studies our outputs are well aligned with human expectations.

This thesis opens several avenues for future work.

- The current implementation of the proposed algorithm is unoptimized, making it computationally expensive for large inputs. For example, the graph in the Cornucopia algorithm is constructed from scratch in every corner removal iteration when examining the removal of every corner. It is, however, possible to recycle the curves and the graph constructed during the first iteration. This can reduce the total time spent on graph construction by a factor of $O(n^2)$ with n representing the number of initially detected corners.
- Our approach is designed for vectorizing clean quantized images. Commercial software (Adobe Illustrator, Vector Magic) has dedicated image filters for quantizing antialiased and noisy data. While our framework can be applied to the outputs of these methods, artifacts they produce will impact our results, as shown in Figure 8.1. An interesting line of future work would be to modify our accuracy constraints and spline fitting energy to account for antialiased or other fuzzy boundaries, as well as to develop dedicated boundary extraction methods for such fuzzy semi-structured inputs.
- Humans are likely to rely on semantic understanding of shapes when vectorizing an image, as




Figure 8.1: Effect of noise and anti-aliasing. Input image adapted from Freepick – flaticon.com.

previously mentioned in Section 7.3. Exploration of such semantics can be an interesting direction of future work.

Bibliography

- Adobe. Adobe Illustrator 2017: Image Trace. <http://www.adobe.com/>, 2017.
- Hani Altwaijry, Andreas Veit, Serge J Belongie, and Cornell Tech. Learning to detect and match keypoints with deep architectures. In *BMVC*, 2016.
- Fabián Arrebola and Francisco Sandoval. Corner detection and curve segmentation by multiresolution chain-code linking. *Pattern Recognition*, 38(10):1596–1614, 2005.
- Ilya Baran, Jaakko Lehtinen, and Jovan Popović. Sketching clothoid splines using shortest paths. In *CGF*, volume 29,2, pages 655–664, 2010.
- Mikhail Bessmeltsev and Justin Solomon. Vectorization of line drawings via polyvector fields. *arXiv preprint arXiv:1801.01922*, 2018.
- H Lynn Beus and Steven SH Tiu. An improved corner detection algorithm based on chain-coded plane curves. *Pattern Recognition*, 20(3):291–296, 1987.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- A Carmona-Poyato, Francisco José Madrid-Cuevas, R Medina-Carnicer, and Rafael Muñoz-Salinas. Polygonal approximation of digital planar curves through break point suppression. *Pattern Recognition*, 43(1):14–25, 2010.
- Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *IJCV*, 22(1):61–79, 1997.
- Dmitry Chetverikov and Zsolt Szabo. A simple and efficient algorithm for detection of high curvature points in planar curves. In *CAIP*, volume 3, pages 746–753. Springer, 2003.
- Ryan Dahl, Mohammad Norouzi, and Jonathan Shlens. Pixel recursive super resolution. In *IEEE ICCV*, 2017.
- J. R. Diebel. *Bayesian image vectorization: The probabilistic inversion of vector image rasterization*. Ph.D. dissertation, Stanford Univ., 2008.
- Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, pages 184–199, 2014.
- Eagle. Eagle. http://everything2.com/index.pl?node_id=1859453, 1997.
- Gerald Farin. *Curves and Surfaces for CAGD: A Practical Guide*. Morgan Kaufmann Publishers Inc., 2002.

- M. Fatemi, A. Amini, L. Baboulaz, and M. Vetterli. Shapes from pixels. *IEEE TIP*, 25(3):1193–1206, 2016.
- Raanan Fattal. Image upsampling via imposed edge statistics. In *ACM SIGGRAPH*, 2007.
- Jean-Dominique Favreau, Florent Lafarge, and Adrien Bousseau Bousseau. Photo2ClipArt: Image Abstraction and Vectorization Using Layered Linear Gradients. *ACM TOG*, 36(6), 2017. URL <https://hal.inria.fr/hal-01581981>.
- M. A. T. Figueiredo, J. Leitão, and A. K. Jain. Unsupervised contour representation and estimation using b-splines and a minimum description length criterion. *IEEE TIP*, 9(6):1075–187, 2000.
- Shachar Fleishman, Daniel Cohen-Or, and Cláudio T. Silva. Robust moving least-squares fitting with sharp features. *ACM TOG*, 24(3):544–552, 2005.
- Herbert Freeman and Larry S. Davis. A corner-finding algorithm for chain-coded curves. *IEEE Transactions on computers*, 26(3):297–303, 1977.
- Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *IEEE ICCV*, 2009.
- Chris Harris and Mike Stephens. A combined corner and edge detector. In *Alvey vision conference*, volume 15, page 50, 1988.
- Robert Hess and David Field. Integration of contours: new insights. *Trends in cognitive sciences*, 3(12):480–486, 1999.
- Shayan Hoshyari, Edoardo Alberto Dominici, Alla Sheffer, Nathan Carr, Duygu Ceylan, Zhaowen Wang, and I-Chao Shen. Perception-driven semi-structured boundary vectorization. *ACM Transaction on Graphics*, 37(4), 2018. doi:10.1145/3197517.3201312. <http://www.cs.ubc.ca/labs/imager/tr/2018/PerceptionDrivenVectorization/>.
- Hsieh Hou and H Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(6):508–517, 1978.
- Xie Jun, Winnemöller Holger, Li Wilmot, and Schiller Stephen. Interactive vectorization. In *ACM SIGCHI*, 2017.
- Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *IJCV*, 1(4): 321–331, 1988.
- Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. pages 1646–1654, 2016.
- Jon Kleinberg and Eva Tardos. *Algorithm design*. Pearson Education India, 2006.
- K. Koffka. *Principles of Gestalt Psychology*. International library of psychology, philosophy, and scientific method. Routledge & K. Paul, 1955.
- Johannes Kopf and Dani Lischinski. Depixelizing pixel art. *ACM TOG*, 30(4):99:1–99:8, 2011.
- Felix Kreuzer, Johannes Kopf, and Michael Wimmer. Depixelizing pixel art in real-time. In *Proceedings of the 19th Symposium on Interactive 3D Graphics and Games*, pages 130–130. ACM, 2015.

- DJ Langridge. Curve encoding and the detection of discontinuities. *Computer Graphics and Image Processing*, 20(1):58–71, 1982.
- Gregory Lecot and Bruno Levy. Ardeco: Automatic region detection and conversion. In *EGSR*, pages 349–360, 2006.
- Xin Li and Michael T Orchard. New edge-directed interpolation. *IEEE TIP*, 10(10):1521–1527, 2001.
- Yangyan Li, Xiaokun Wu, Yiorgos Chrysanthou, Andrei Sharf, Daniel Cohen-Or, and Niloy J. Mitra. Globfit: Consistently fitting primitives by discovering global relations. *ACM TOG*, 30(4):52:1–52:12, 2011.
- Jyh-Ming Lien and Nancy M. Amato. Approximate convex decomposition of polygons. In *Proc. Symp. Computational Geometry*, pages 17–26, 2004.
- Yang Liu and Wenping Wang. A revisit to least squares orthogonal distance fitting of parametric curves and surfaces. In *Proc. Advances in Geometric Modeling and Processing*, pages 384–397, 2008.
- Zhaoliang Lun, Evangelos Kalogerakis, and Alla Sheffer. Elements of style: Learning perceptual shape style similarity. *ACM Trans. Graph.*, 34(4):84:1–84:14, 2015.
- Andrea Mazzoleni. Scale2x. <http://www.scale2x.it/>, 2001.
- James McCrae and Karan Singh. Sketching Piecewise Clothoid Curves. In *Proc. Sketch-Based Interfaces and Modeling*, 2008.
- James McCrae and Karan Singh. Neatening sketched strokes using piecewise french curves. In *Proc. EG Symposium on Sketch-Based Interfaces and Modeling*, pages 141–148, 2011.
- Gerard Medioni and Yoshio Yasumoto. Corner detection and curve representation using cubic b-splines. In *Robotics and Automation. Proceedings. 1986 IEEE International Conference on*, volume 3, pages 764–769. IEEE, 1986.
- Ravish Mehra, Qingnan Zhou, Jeremy Long, Alla Sheffer, Amy Gooch, and Niloy J. Mitra. Abstraction of man-made shapes. *ACM TOG*, 28(5):137:1–137:10, 2009.
- Kamal Nasrollahi and Thomas B. Moeslund. Super-resolution: A comprehensive survey. *Mach. Vision Appl.*, 25(6):1423–1468, August 2014.
- Gioacchino Noris, Alexander Hornung, Robert W Sumner, Maryann Simmons, and Markus Gross. Topology-driven vectorization of clean line drawings. *ACM Transactions on Graphics (TOG)*, 32(1):4, 2013.
- Alexandrina Orzan, Adrien Bousseau, Holger Winnemöller, Pascal Barla, Joëlle Thollot, and David Salesin. Diffusion curves: A vector representation for smooth-shaded images. *ACM TOG*, 27(3), 2008.
- Amal Dev Parakkat, Uday Bondi Pundarikaksha, and Ramanathan Muthuganapathy. A delaunay triangulation based approach for cleaning rough sketches. *Computers & Graphics*, 2018.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

- Ananth Ranganathan. The Levenberg-Marquardt algorithm. URL <http://users-phys.au.dk/jensjh/numeric/project/10.1.1.135.865.pdf>. Accessed: May 2018.
- Edward Rosten, Reid Porter, and Tom Drummond. Faster and better: A machine learning approach to corner detection. *IEEE transactions on pattern analysis and machine intelligence*, 32(1):105–119, 2010.
- S Rasoul Safavian and David Landgrebe. A survey of decision tree classifier methodology. *IEEE transactions on systems, man, and cybernetics*, 21(3):660–674, 1991.
- ScanFont. Font Lab, <http://old.fontlab.com/font-converter/scanfont/>, 2017.
- Peter Selinger. Potrace: a polygon-based tracing algorithm. In <http://potrace.sourceforge.net/potrace.pdf>, 2003.
- Jianbo Shi et al. Good features to track. In *Computer Vision and Pattern Recognition, 1994. IEEE Conference on*, pages 593–600. IEEE, 1994.
- Maxim Stepin. Hqx. <http://web.archive.org/web/20070717064839/www.hiend3d.com/hqx.html>, 2003.
- Jian Sun, Lin Liang, Fang Wen, and Heung-Yeung Shum. Image vectorization using optimized gradient meshes. In *ACM SIGGRAPH*, 2007.
- Daniel Šýkora, Jan Buriánek, and Jiří Žára. Sketching cartoons by example. In *Proc. Sketch-Based Interfaces and Modeling*, pages 27–34, 2005.
- The Inkscape team. Inkscape: A vector drawing tool. 2018. URL <https://gitlab.com/inkscape/>. Accessed: May 2018.
- Vector Magic. Cedar Lake Ventures <http://vectormagic.com/>, 2017.
- J. Wagemans, J. H. Elder, M. Kubovy, S. E. Palmer, M. A. Peterson, M. Singh, and R von der Heydt. A century of gestalt psychology in visual perception i. perceptual grouping and figure-ground organization. *Psychological Bulletin*, 138(6):1172–1217, 2012.
- C. Wang, J. Zhu, Y. Guo, and W. Wang. Video vectorization via tetrahedral remeshing. *IEEE TIP*, 26(4):1833–1844, April 2017.
- Zhaowen Wang, Ding Liu, Jianchao Yang, Wei Han, and Thomas Huang. Deep networks for image super-resolution with sparse prior. In *IEEE ICCV*, pages 370–378, 2015.
- M. Weber and B. Herzog. Autotrace. <http://autotrace.sourceforge.net>, 2004.
- Tian Xia, Binbin Liao, and Yizhou Yu. Patch-based image vectorization with automatic curvilinear feature alignment. *ACM TOG*, 28(5), 2009.
- Chih-Yuan Yang, Chao Ma, and Ming-Hsuan Yang. Single-image super-resolution: A benchmark. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *ECCV*, pages 372–386. Springer International Publishing, 2014.
- M. Yang, H. Chao, C. Zhang, J. Guo, L. Yuan, and J. Sun. Effective clipart image vectorization through direct optimization of bezigons. *IEEE TVCG*, 22(2):1063–1075, 2016.
- Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. Lift: Learned invariant feature transform. In *European Conference on Computer Vision*, pages 467–483. Springer, 2016.

Song-Hai Zhang, Tao Chen, Yi-Fei Zhang, Shi-Min Hu, and Ralph R. Martin. Vectorizing cartoon animations. *IEEE TVCG*, 15(4):618–629, 2009.

Appendix A

Curve Fitting

As previously mentioned, the optimization method of Section 4.1.3 relies on fitting curves to a series of points while enforcing various constraints on them at prespecified transition points. For completeness, this appendix presents the method of Baran et al. [2010] along with our modifications for achieving this task. Hereinafter, it is assumed that a curve C has type $Type_C$ (line, arc, or clothoid), is fully described by a set of degrees of freedom χ_C , and is parameterized by $0 \leq s \leq 1$. The full set of degrees of freedom χ used in this thesis for each curve type are shown in Table A.1.

A.1 Projection of a Point on the Curve

Curve fitting would first require the introduction of the projection operator, i.e., a function that returns the closest point on the curve to another input point:

$$Proj(\chi_C; Type_C, m) = C(s^*) \text{ where } s^* = ArgProj(\chi_C; Type_C, m) = \underset{s}{\operatorname{argmin}}(\|C(s) - m\|_2), \quad (A.1)$$

where $Proj$ and $ArgProj$ are projection operators (the former finds the closest point, while the latter finds the corresponding curve parameter), and m is the input point. For lines and arcs, the projection can be carried out analytically, while for clothoids, numerical methods such as Newton or bisection can be used.

Another value of importance is the Jacobian of the projection operator with respect to curve degrees of

Table A.1: Curve degrees of freedom

Curve type	Degrees of freedom χ
Line segment	start point X and Y coordinates, start angle, length
Arc	start point X and Y coordinates, start angle, length, start curvature
Clothoid	start point X and Y coordinates, start angle, length, start curvature, end curvature




Figure A.1: Schematic of the projection operator and perturbations to the curve parameters.

Table A.2: Dimensions of the matrices involved in the computation of the derivative of the projection operator

Matrix	Dimension	Matrix	Dimension	Matrix	Dimension
$\partial_\chi \text{ArgProj}(\chi_C; \text{Type}_C, m)$	$1 \times \chi $	C	2×1	m	2×1
$\partial_\chi \text{Proj}(\chi_C; \text{Type}_C, m)$	$2 \times \chi $	$\partial_s C$	2×1	$\partial_{ss}^2 C$	2×1
$\partial_\chi C$	$2 \times \chi $	$\partial_{s\chi}^2 C$	$2 \times \chi $		

freedom $\partial_\chi \text{Proj}(\chi; \text{Type}, m)$. One way to evaluate this Jacobian matrix is the use of finite difference formulas. As depicted in Figure A.1, each entry of χ can be perturbed by a small amount of ϵ and then the corresponding column in $\partial_\chi \text{Proj}(\chi; \text{Type}, m)$ can be approximated as $(\text{Proj}(\chi + \Delta\chi; \text{Type}, m) - \text{Proj}(\chi; \text{Type}, m)) / \epsilon$.

Alternatively, an analytical formula is available. Note that the position of every point on the curve is a function of the curve parameter s and the curve degrees of freedom χ . Using the partial derivatives of this function ($\partial_\chi C$, $\partial_s C$, $\partial_{ss} C$, and $\partial_{\chi s} C$) and the chain rule, the Jacobian $\partial_\chi \text{Proj}(\chi; \text{Type}, m)$ can be evaluated as

$$\partial_\chi \text{ArgProj}(\chi_C; \text{Type}_C, m) = \frac{(\partial_s C)^T (\partial_\chi C) - (C - m)^T (\partial_{\chi s}^2 C)}{(\partial_s C)^T (\partial_s C) + (C - m)^T (\partial_{ss}^2 C)}, \quad (\text{A.2})$$

$$\partial_\chi \text{Proj}(\chi_C; \text{Type}_C, m) = \partial_\chi C + (\partial_s C)^T (\partial_\chi \text{ArgProj}(\chi_C; \text{Type}_C, m)).$$

The dimensions of the matrices involved in Equation (A.2) are given in Table A.2, where $|\chi|$ is the number of degrees of freedom.

A.2 Fitting a Single Curve

Given a set of points $L = \{m_1, m_2, \dots, m_n\}$, the goal is to find a curve C with type $Type_C$, that minimizes the least square error

$$\mathcal{D} = \sum_{j=1}^n \|Proj(\chi_C; Type_C, m_j) - m_j\|_2^2. \quad (\text{A.3})$$

In other words, the curve parameters of Table A.1 have to be solved for to minimize \mathcal{D} .

Initial Guess. First, an initial guess is made, depending on the curve type:

- A line segments must pass through the centroid of the points, $\bar{m} = \sum_{j=1}^n m_j / n$, and have a tangent equal to the largest eigenvector of the covariance matrix, $I \in \mathcal{R}^{2 \times 2}$, where $I_{11} = \sum_{s=1} (m_s - \bar{m})_x (m_s - \bar{m})_x$, $I_{12} = I_{21} = \sum_{s=1} (m_s - \bar{m})_y (m_s - \bar{m})_x$, $I_{22} = \sum_{s=1} (m_s - \bar{m})_y (m_s - \bar{m})_y$. The line segment is finally clamped between the projections of the points m_1 and m_n .
- An initial guess is made for arcs by fitting an ellipse to the lifted points $z_j = ((m_j - m_1)_x, (m_j - m_1)_y, \|m_j - m_1\|_2^2)$, and subsequently projecting the ellipse back into the XY -plane to create a circle. Finally, the circle is clamped between the projections of m_1 and m_n .
- For a clothoid, the curvature varies linearly with the arc-length. Therefore, the tangential angle would be a quadratic function of the arc-length. Thus, to find an initial guess for a clothoid, a quadratic approximation of the tangent angle, $\theta = as^2 + bs + c$, is fitted to the tuples (s_i, θ_i) , where $s_i = \left(\sum_{j=2}^i \|m_j - m_{j-1}\|_2 \right) / \left(\sum_{j=2}^n \|m_j - m_{j-1}\|_2 \right)$, and θ_i is the angle between the vectors $m_{i+1} - m_i$ and $m_i - m_{i-1}$. Here, s_i and θ_i are approximations of the arc-length and tangent angle at the projections of the points on the clothoid, respectively. The clothoid parameters are then extracted from this quadratic angle fit, and additionally requiring the clothoid to pass through centroid \bar{m} .

Gauss-Newton Iterations. The initial guess is not accurate enough, and has to be refined in the case of circles and clothoids. Defining a vector function, $F \in \mathcal{R}^{2n}$, as

$$F_i = \begin{cases} \text{x component of } Proj(C; m_j) - m_j & i = 2j \\ \text{y component of } Proj(C; m_j) - m_j & i = 2j + 1 \end{cases}, \quad (\text{A.4})$$

we can write the objective function as $\mathcal{D} = \|F\|_2^2$. Gauss-Newton iterations can then be performed to minimize \mathcal{D} in terms of the curve degrees of freedom χ . Therefore, the curve degrees of freedom are updated as

$$\begin{aligned} [(\partial_\chi F)^T (\partial_\chi F) + \lambda I] \delta \chi &= -(\partial_\chi F)^T F, \\ \chi &\leftarrow \chi + \delta \chi. \end{aligned} \quad (\text{A.5})$$

Here, I is the identity matrix, the Jacobian matrix $\partial_\chi F$ is found using the procedure of Section A.1, and λ is a damping term added to prevent divergence. The addition of the damping term is also known as

the Marquardt-Levenberg method [Ranganathan].

A.3 Fitting of Multiple Curves

Fitting multiple curves and imposing certain constraints on them can be tackled in a similar way. In this problem, the goal is to fit curves C_k to separate series of points $L_k = \{m_1^k, \dots, m_{n_k}^k\}$ ($1 \leq k \leq K$), while enforcing transition continuities between them (in the absence of any constraint, the problem will reduce to multiple instances of fitting a single curve). By defining a vector function F as

$$F_i = \begin{cases} \text{x component of } Proj(C_k; m_j^k) - m_j^k & i = 2(j + \sum_{k'=1}^k n_{k'}) \\ \text{y component of } Proj(C_k; m_j^k) - m_j^k & i = 2(j + \sum_{k'=1}^k n_{k'}) + 1 \end{cases}, \quad (\text{A.6})$$

the objective function can again be written as $\mathcal{D} = \|F\|_2^2$. Enforcing a C^0 transition between two curves k and k' can then be achieved by requiring $C_k(1) - C_{k'}(0) = 0$. The transition can further be promoted to G^1 by adding an additional constraint $\bar{C}_k(1) - \bar{C}_{k'}(0) = 0$. Gathering all the constraints in a single vector function G , we can write the minimization problem as

$$\chi = \underset{\chi}{\operatorname{argmin}}(\mathcal{D}) \text{ subject to } G(\chi) = 0, \quad (\text{A.7})$$

where χ is now a vector containing the degrees of freedom of all the curves involved. Rather than solving Equation (A.7) exactly, a penalty method can be used, i.e., we can solve

$$\chi = \underset{\chi}{\operatorname{argmin}}(\mathcal{D} + \beta \|G(\chi)\|_2^2), \quad (\text{A.8})$$

where β is a large number (in this thesis $\beta = 100$ has been used). In this case, Gauss-Newton iterations would take the form of

$$[(\partial_\chi F)^T (\partial_\chi F) + \beta (\partial_\chi G)^T (\partial_\chi G) + \lambda I] \delta \chi = -(\partial_\chi F)^T F. \quad (\text{A.9})$$

Note that the Jacobian matrices $\partial_\chi F$ and $\partial_\chi G$ can be assembled easily by realizing the fact that the projection of a point on the curve C_k is independent of the degrees of freedom of the other curves.

Appendix B

A Fundamental Discrete-Continuous Optimization Problem

In Chapter 4, our discrete-continuous objective function for boundary vectorization and its optimization based on the work of Baran et al. [2010] were introduced. The goal of this appendix is to provide further intuition by introducing the simpler but less complicated problem of *segmented least-squares*. The segmented least-squares problem can be found in many algorithm design books, e.g., [Kleinberg and Tardos, 2006].

The problem is as follows. A series of n data points, $L = \{m_1, m_2, \dots, m_n\}$, is given, arranged such that $(m_1)_x < (m_2)_x < \dots < (m_n)_x$. The goal is to find a set of transition points $\mathcal{T} = \{m_{t_1} = m_1, m_{t_2}, \dots, m_{t_{k-1}}, m_{t_k} = m_n\}$, and a set of lines $\mathcal{C} = \{C_1, C_2, \dots, C_{k-1}\}$, such that the objective \mathcal{E} , given as

$$\begin{aligned}\mathcal{E}(\mathcal{C}, \mathcal{T}; L) &= \mathcal{D} + \frac{1}{\alpha} \mathcal{R}, \\ \mathcal{D}(\mathcal{C}, \mathcal{T}; L) &= \sum_{i=1}^{k-1} \sum_{j=t_i}^{t_{i+1}} |(m_j)_y - C_i((m_j)_x)|^2, \\ \mathcal{R}(\mathcal{C}, \mathcal{T}; L) &= k,\end{aligned}\tag{B.1}$$

is minimized. Here, it is assumed that the equation of each line is $C_i(x) = a_i x + b_i$, and α is a hyper-parameter. Similar to Equations (4.1)–(4.3), this objective function is composed of two terms. The first term \mathcal{D} ensures that the lines stay close to the points, and the second term \mathcal{R} controls the number of lines (similar to simplicity). When $\alpha = 0$, the objective term \mathcal{R} will dominate. Therefore, the solution will include a single line only, as shown in Figure B.1-b, whose equation can be found using the procedure given in Section A.2. When $\alpha = \infty$, the objective term \mathcal{D} will dominate, and the solution will consist of $n - 1$ lines passing through all the points, as shown in Figure B.1-c. A suitable selection of α can result in a trade off between accuracy and simplicity (Figure B.1-d).




Figure B.1: Schematic of the segmented least-squares problem for various values of α : (a) input points; (b) $\alpha = 0$; (c) $\alpha = \infty$; (d) $\alpha = 1$.

The key to the solution is the observation of the following recursive relation

$$\min_{\mathcal{C}, \mathcal{T}} \mathcal{E}(\mathcal{C}, \mathcal{T}; L) = \min_{1 \leq i < n} \left(\min_{\mathcal{C}, \mathcal{T}} \mathcal{E}(\mathcal{C}, \mathcal{T}; L(1:i)) + \min_{a,b} \sum_{j=i+1}^n |(m_j)_y - a(m_j)_x - b|^2 \right), \quad (\text{B.2})$$

where the notation $L(1:i)$ denotes the set containing the first i points. This would allow for a dynamic programming algorithm to find the optimal solution, as shown in Algorithm B.1. For simplicity, only the calculation of the minimum value of the objective function \mathcal{E} is shown (the corresponding transition points \mathcal{T} and line parameters \mathcal{C} can be found in a similar fashion). Note that the combinatorial nature of the problem has resulted in a high computational cost of fitting $O(n^2)$ lines during the solution process.

Algorithm B.1 Solution to the segmented least square problem

```

function MIN ▷ Minimum attainable objective function
  Input:  $L$  (the array of input points with length  $n$ ) and  $\alpha$ 
  Output: Array  $\mathcal{E}^*$  where  $\mathcal{E}^*(i) = \min_{\mathcal{C}, \mathcal{T}} \mathcal{E}(\mathcal{C}, \mathcal{T}; L(1:i))$  for  $1 \leq i \leq n$ .
   $\mathcal{E}^*$  = empty array of size  $n$ 
   $\mathcal{E}(1) = 0$ 
  for  $i = 2:n$  do
     $\mathcal{E}^{**}$  = empty array of size  $i - 1$ 
    for  $j = 1:(i - 1)$  do
       $\mathcal{E}^{**}(j) = \mathcal{E}^*(j) + \text{MinSingle}(L(j:i), \alpha)$ 
     $\mathcal{E}^*(i) = \min(\mathcal{E}^{**})$ 
  return  $\mathcal{E}^*$ 

function MINSINGLE ▷ Helper function for fitting a single line (a special case of Appendix A)
  Input:  $L$  (the array of input points with length  $n$ ) and  $\alpha$ 
  Output:  $\mathcal{E}^* = \min_{a,b} \sum_{m_j \in L} |(m_j)_y - a(m_j)_x - b|^2$ 
   $a = (n \sum (m_j)_x (m_j)_y - \sum (m_j)_x \sum (m_j)_y) / (n \sum (m_j)_x^2 - (\sum (m_j)_x)^2)$  ▷  $\Sigma$  denotes summing over all  $L$ 
   $b = (\sum (m_j)_y - a \sum (m_j)_x) / n$ 
   $\mathcal{E}^* = \sum |(m_j)_y - a(m_j)_x - b|^2 + \alpha$ 
  return  $\mathcal{E}^*$ 

```

The solution to the vectorization problem of Section 4.1.3 has a similar nature. It is, however, innately

more complex, as it has to satisfy the extra constraints and handle the addition of curve types to the search space. Traces of the segmented least squares problem can also be found in other curve fitting frameworks. For example, McCrae and Singh [2008] solve the segmented least square problem in the curvature space. Likewise, the method of Selinger [2003] is quite similar to Algorithm B.1.

Appendix C

User Study Questions

Figure C.1 depicts an example of the questions used in the user study conducted to compare our vectorization method to the other existing methods. Figure C.2 depicts an example of the questions used in the user study conducted to compare our vectorization method to the work of an artist. The full list of user study questions (15 and 180 questions for comparison against alternative algorithms and the work of an artist, respectively) are available on the project website [Hoshyari et al., 2018].

Question 32-castle: our: 24 vectormagic: 2 both: 0 neither: 2

Which of the two images on the bottom (B) or (C) better represents image (A)? If both are equally good then select 'Both', and if neither represents (A) then select 'Neither'.




Figure C.1: User study question example for comparison against alternative algorithms. (A), (B), and (C) represent the input, Vector Magic result, and our result, respectively.

Question 64-headphone2: our: 20 artist: 2 both: 1 neither: 5

Which of the two images on the bottom (B) or (C) better represents image (A)? If both are equally good then select 'Both', and if neither represents (A) then select 'Neither'.




Figure C.2: User study question example for comparison against the work of an artist. (A), (B), and (C) represent the input, artist's work, and our result, respectively.