Curve Reconstruction Via Line Intersection Tomography

by

Sekhar Chandra Tatikonda

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

January 1995

Author ....
Department of Electrical Engineering and Computer Science
January 30, 1995

Certified by ....
Sanjoy Mitter
Professor of Electrical Engineering
Thesis Supervisor

Accepted by ....
Frederic R. Morgenthaler
Chairman, Departmental Committee on Graduate Students

Curve Reconstruction Via Line Intersection Tomography

by

Sekhar Chandra Tatikonda

Submitted to the Department of Electrical Engineering and Computer Science
on January 30, 1995, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

Richardson [Richardson, 1993] has shown that a large class of one-dimensional sets in the plane can be reconstructed from the set's line intersection function. Herein we give a summary of this work, provide a detailed description of the line intersection function, provide a reconstruction algorithm, and give some basic heuristics for implementing it.

Thesis Supervisor: Sanjoy Mitter
Title: Professor of Electrical Engineering

Acknowledgments

I would like to thank Professor Sanjoy Mitter for his continued advice, patience, and belief in me throughout this project. Also I would like to thank Tom Richardson for setting the groundwork for LIT and for his helpful comments and corrections.

Thanks go to my parents for their unwavering support and to my brother Mohan for his preternatural ability to make any situation a better one.

Finally I would like to thank the National Science Foundation for providing me with an NSF Graduate Fellowship.

Contents

Chapter 1

Introduction

The techniques of tomography have been used extensively in computer vision and medical imaging. Here we discuss a specific type of tomography called line intersection tomography (LIT). LIT allows us to reconstruct one-dimensional sets in the plane from the set's line intersection function. This line intersection function counts the number of times, including multiplicity, a given line intersects our image set. Richardson [Richardson, 1993] has shown that this reconstruction can be carried out with arbitrary precision for a class of sets called the $\mathcal{K}$-sets. These are sets which can be described as the union of a finite number of closed curves in the plane. Note an important subset of the $\mathcal{K}$-sets is the class of curves in the plane.

Herein we give a summary of Richardson's work, describe in detail the line intersection function, prove a reconstruction algorithm, and provide some general heuristics for implementing this algorithm.

1.1 Problem Background

Traditionally LIT has been used for the detection of boundaries of objects in an image. Thirion [Thirion, 1992] has shown that LIT is very good for detecting boundaries of highly contrasted objects like bones and blood vessels. In this way both internal and external boundaries can be detected. LIT, though, works very poorly on occluded objects. Unlike other transforms, Radon for example, that take the integral of the density over a slice of the image, LIT just counts intersections; so dense occluded objects cannot be reconstructed.

LIT, from a historical point of view, derives its mathematical basis in integral geometry. Specifically it concerns itself with the problem of determining when a one-dimensional set in the plane is completely determined by its projections. Richardson has shown that the $\mathcal{K}$-sets are completely determined by their projections.

Finally LIT is related to the so called Hough transform of computer vision. The Hough transform is used to determine line segments in an image. We will show in the next section, though, that the Hough transform is limited to reconstructing objects that are easily parameterized. The work described herein amounts to the inversion of the Hough transform.

Until now there has been no efficient algorithm proposed for general curve reconstruction using LIT.

1.2 Hough Transform

The Hough transform was developed to detect lines in computer images. It consists of parameterizing all lines that lie in the plane by two parameters: namely the angle it makes with some fixed coordinate system and the distance between the line and the origin of said coordinate system. The parameter space is $S^1 \times R^1$. We partition this space into disjoint regions, usually boxes. We pick some representative point for each region in the parameter space. For every point in the image and for every region in the parameter space we check to see if the point lies on one of the lines in the parameter region. If it does we mark a tally for that region. We say a representative line is in the image if the value for its region is greater than some threshold.

We can generalize this transform to arbitrary feature identification [Leavers, 1992]. The features, though, must be parameterized be a small number of parameters. For example the feature circle can be parameterized by two coordinates for the origin and one for the radius. Thus the parameter space is three dimensional. One sees that for complicated shapes the number of computations increases exponentially with the number of parameters.

Though the Hough transform works well for lines and other simple features it does not do well on detecting arbitrary curves in images. The reason is that arbitrary curves require many parameters to encode them. Kulkarni [Kulkarni, et. al. page 939] shows that it takes an exponential, in the accuracy parameter, number of representative points to encode a curve of finite length and turn. (Turn is a measure of the absolute total curvature of a curve.) Hough transform type algorithms have not been used for general curve reconstruction.

1.3 Stochastic Approaches

Though Richardson provides a theory for reconstruction in the exact case, any algorithm we devise will have to operate on partial information. Kulkarni provides one such reconstruction in a stochastic setting. We provide a deterministic algorithm.

Kulkarni shows that all curves of finite length and finite turn can be reconstructed to arbitrary accuracy by using a uniform sampling of the projection lines in the PAC sense. Unfortunately there are two major problems with this reconstruction: the number of sampling points is exponential in the accuracy parameter and we do not reconstruct the original curve but instead its line intersection function.

What follows is a brief summary of the methods Kulkarni used that will be of help to us in the following chapters.

Kulkarni uses the idea of metric entropy. Specifically let $[X, d]$ be a metric space. We say a set X^ϵ is an ϵ cover of X if every member $x \in X$ is less than ϵ away from some member of X^ϵ. The metric entropy is then the $log_2|X^\epsilon|$. The metric entropy idea is similar to the idea of picking representative points discussed in previous section.

Let the line intersection function be $n_c(l)$. This equals the number of times the line l intersects the curve c.

We define a distance over the space of curves to be $d_P(c_1, c_2) = E|n_{c_1}(l) - n_{c_2}(l)|$, where the expectation is taken over the uniform distribution over the space of lines $l \in S_1 \times R_1$. Though we are using d_P to measure how close two curves are, it is really a measure of how close the curve's respective line intersection functions are. Richardson [Richardson, theorem 3.1] proves that if $d_P(c_1, c_2) = 0$ then $c_1 = c_2$ $\mathcal{H}^1$-almost everywhere. (Where $\mathcal{H}^1$ is the one-dimensional Hausdorff measure.)

Kulkarni shows that the space of curves of finite length and finite turn has a metric entropy that is polynomic in ϵ.

What Kulkarni's PAC algorithm does is sample the $n_c(l)$ function at many different points l. Then it finds the closest fit of the data to one of the $n_c(\cdot)$ functions in our ϵ-cover. We use a preexisting table to match our best fitting $n_c(\cdot)$ to its underlying curve. This table, though, will have an exponential number, in the accuracy parameter, of entries. PAC arguments have been mainly used for proving existence of algorithms. Usually effienciy does not come into play. This is the case with Kulkarni's PAC algorithm.

1.4 Summary of Thesis

In chapter two we give a review of Richardson's work. We give a reconstruction algorithm in chapter three. Then we give a detailed description of the line intersection function in chapter four. In chapter five we discuss what "good approximation" means and give some basic heuristics for implementing an LIT algorithm. Finally in chapter six we conclude and discuss some open problems.

Chapter 2

Background

Richardson has constructed a rigorous mathematical theory for the reconstruction of certain one-dimensional sets in the plane. What follows is a summary of Richardson's results and noted definitions and theorems that will be useful in the following chapters. The basic idea is to model an image set A by a number of curves $C = \{c_i\}$. Each curve $c \in C$ induces an extended indicatrix (tangent curve) in the space of lines $\mathcal{G}$. These extended indicatrices turn out to be closely related to the jumpset of the line intersection function n_A. Thus from n_A one can reconstruct A.

2.1 Background

A *one dimensional set* in the plane is any set $E \subset R^2$ such that $0 < \mathcal{H}^1(E) < \infty$.

2.1.1 Curves

The following are the standard definitions for parameterized curves and curves. A *parameterized curve* is a continuous function $c : [a, b] \to R^n$ which is non-constant on any open subinterval. We say another curve $c' : [c, d] \to R^n$ is *equivalent* to c if there exists a homeomorphism $\Psi : [c, d] \to [a, b]$ such that $\Psi(c) = a, \Psi(d) = b$, and $c' = c \circ \Psi$. An equivalence class over the parameterized curves is a *curve.* A parameterized curve is *closed* if $c(a) = c(b)$. We denote $c|_{[t_1,t_2]}$ to equal the curve c

restricted to $[t_1, t_2]$. A non-closed curve is *simple* if it does not cross itself. The *trace* of a curve c is: $tr(c) := \{x : c^{-1}(x) \neq \emptyset\}$. An *inflection point* is a point where a convex and a concave segment meet. A *cusp* is a point where two concave or two convex segments meet with common tangent. A *corner* is a point where the incoming and outgoing tangents to the point differ.

We now define total length and total absolute curvature. First we define length and absolute curvature for piecewise linear curves. Let $v_0, ..., v_n$ be a set of $n+1$ points in R^n. Let a_i be the line segment connecting v_{i-1} to v_i. Call the set of line segments $P = \{a_i\}$. The length of the piecewise linear curve P is defined as $\mathcal{L}(P) := \sum_{i=1}^{n} |v_i - v_{i-1}|$. Let the exterior angle between line segments a_i and a_{i+1} be φ_i for $i = 1, ..., n-1$. Then the absolute curvature is $\kappa(P) = \sum_{i=1}^{n-1} \varphi_i$. For a general curve c we define the length of c to be $\mathcal{L}(c) := \sup \{\mathcal{L}(P) : P \text{ is inscribed in } c\}$. The absolute curvature is $\kappa(c) := \sup \{\kappa(P) : P \text{ is inscribed in } c\}$. Note that when c is continuously differentiable then the above definition of length corresponds to $\int |\dot{c}| dt$. When c is twice differentiable then the above definition of curvature corresponds to $\int_0^{\mathcal{L}(c)} |\kappa(t)| dt$, where $c(t)$ is arclength parameterized and $\kappa(t)$ is the curvature at the point $c(t)$.

We say a curve is *rectifiable* if $\mathcal{L}(c) < \infty$. All rectifiable curves can be uniquely parameterized by arclength. We let $T_c(t)$ denote the unit tangent vector wherever it exists.

2.1.2 Lines and Angles

Let S^1 be the unit circle and let θ and $*\theta$ be two elements that form a basis for S^1. That is $\theta \wedge *\theta = e_1 \wedge e_2$ where e_1, e_2 are the standard orthonormal basis elements in R^2. We let $\angle(\theta_1, \theta_2)$ equal the smaller of the two arcs of S^1 with θ_1, θ_2 for endpoints. Let $|\angle(\theta_1, \theta_2)|$ equal the angle between the vectors θ_1 and θ_2. Finally let $\mathcal{R}(\theta_1, \theta_2) := sign(\langle \theta_1, *\theta_2 \rangle) |\angle(\theta_1, \theta_2)|$. This is just the amount θ_2 must rotate to coincide with θ_1.

Since we are interested in the number of intersections a given line has with our curve we need an analytic way to describe lines. As stated before a line can be parameterized by a tangent angle and a distance from the origin. Specifically $\mathcal{G} =$

$R^1 \times S^1$. Let $G = (\rho, \theta) \in \mathcal{G}$. Then let $\underline{G} = \{x \in R^2 : \langle x, *\theta \rangle = \rho\}$. Note that while G is an oriented line its image in R^2 is not. We can define a measure on $\mathcal{G}$ as follows: $d_{\mathcal{G}}(G_1, G_2) = ((\rho_1 - \rho_2)^2 + |\angle(\theta_1, \theta_2)|^2)^{\frac{1}{2}}$.

2.1.3 Generalized Sets

Since it is possible for a curve to cross itself we need a notion of multiplicity. We define a *generalized set* to be any set with multiplicity. If A is a set in R^n then we denote $A(x)$ to be the multiplicity of the member x in A. The multiplicity can be infinite. Richardson has shown how to generalize the notions of Hausdorff measure and trace to deal with generalized sets. Specifically $\mathcal{H}^0(A) = \int_{R^2} A(x) d\mathcal{H}^0$ and $\mathcal{H}^1(A) = \int_{R^2} A(x) d\mathcal{H}^1$ and $g\_trc(x) := |\{c^{-1}(x) \bmod \mathcal{L}(c)\}|$. We let $n_A(G) = \int_{\underline{G}} A(x) d\mathcal{H}^0$. For a given curve c we denote $n_c := n_{g\_trc}$.

2.1.4 $\mathcal{K}$-Sets

We call a generalized set A a $\mathcal{K}$-set if there exists a finite set of closed curves with finite total absolute curvature $C = \{c_1, ..., c_n\}$ such that $\mathcal{H}^1(|2A - g\_trC|) = 0$ where $trC = \cup_{c \in C} trc$ and $g\_trC = \sum_{c \in C} g\_trc$. We say in this case that C *represents* A. If $2A = g\_trC$ then we say C represents A *exactly.* Two sets C_1, C_2 are *equivalent* if $\mathcal{H}^1(|g\_trC_1 - g\_trC_2|) = 0$. The factor of two is required because we are trying to approximate a set A by closed curves. For example we cannot approximate a simple nonclosed curve by a closed curve.

We define $\kappa(C) = \sum_{c \in C} \kappa(c)$. Now we can define the curvature of an arbitrary $\mathcal{K}$-set A. Define $\kappa^*(A) := \frac{1}{2} \inf \{\kappa(C) : C \text{ represents } A\}$.

Note that κ^* does not correspond to our general notion of curvature κ for curves. For example take a simple non-closed curve c with curvature $\kappa(c)$. Then $\kappa^*(c) = \kappa(c) + \pi$. This is because C can represent the image set A but it cannot represent the "total turn" of A. In general for a given set A our intuitive notion of curvature will equal κ^* minus $\frac{1}{2}\pi$ for every endpoint in A.

2.1.5 Tangency and Regularity

Finally we discuss tangency and regularity. There are three different but related notions of tangency: tangency to a plane set, parametric tangency to a curve, and generalized tangency.

To define tangency to a plane set we need the notion of density. For any set $E \subset R^2$ let the density of E at x be $\Theta(E,x) = lim_{\rho \to 0^+} \frac{1}{2\rho} \mathcal{H}^1(B_\rho(x) \cap E)$. We say x is a *regular point* of E when $\Theta(E,x) = 1$.

We define a sector at x in the direction of y with width ϵ to be $S(x,y,\epsilon) := \{z \in R^2 : \langle z-x, y\rangle \geq (1-\epsilon)|y||z-x|\}$. We say y is a *tangent* to E at x if $\Theta(E \cap S(x,y,\epsilon)) > 0$ and for every $\epsilon > 0$ we have $\Theta(E \backslash S(x,y,\epsilon) \backslash S(x,-y,\epsilon)) = 0$. We say E has a *tangent line* at x if both y and $-y$ are tangent to E at x.

Note because $\mathcal{G}$ acts locally like R^2 we can use the above definitions for sets on $\mathcal{G}$.

We say a curve c has *unit left tangent* $T_c(t-)$ at $x = c(t)$ if $T_c(t-) = lim_{t_i \to t^-} \frac{c(t_i)-c(t)}{|c(t_i)-c(t)|}$ exists and has *unit right tangent* $T_c(t+)$ at $x = c(t)$ if $T_c(t+) = lim_{t_j \to t^+} \frac{c(t_j)-c(t)}{|c(t_j)-c(t)|}$ exists. We say a curve c has a *tangent line* at $x = c(t)$ if $T_c(t-) = T_c(t+)$.

Note that if y is a tangent line to *trc* at $x = c(t)$ then $\frac{y}{|y|} = \pm T_c(t)$.

We define a *generalized tangent* to a point $x = c(t)$ to be any vector which is a positive linear combination of $T_c(t-)$ and $T_c(t+)$.

2.1.6 Richardson's Main Theorem

We define the following pseudo-metric on the space of generalized sets: $d_{\mathcal{I}}(A_1, A_2) = \int_{\mathcal{G}} |n_{A_1}(G) - n_{A_2}(G)| dG$.

Theorem 1 *If A_1 and A_2 are $\mathcal{K}$-sets, then $d_{\mathcal{I}}(A_1, A_2) = 0$ if and only if $\mathcal{H}^1(|A_1 - A_2|) = 0$.*

Proof: Richardson, theorem 3.1.

If we can approximate n_A in the $d_{\mathcal{I}}$ sense then we can approximate the underlying set A in the $\mathcal{H}^1$ - measure sense. Note $d_{\mathcal{I}}$ is just a deterministic version of Kulkarni's d_P.

2.2 Curves in R^2

Since we are interested in reconstructing curves, or more specifically sets of curves, we will need a few local properties of curves. To simplify our work we assume that $c(t)$ is arclength parameterized.

Theorem 2 *If c is closed then $\kappa(c) \geq 2\pi$.*

Proof: Richardson, theorem 4.1

This next lemma will be used throughout the paper. It basically states that absolutely continuous curves of finite length and finite total absolute curvature have one-sided tangents everywhere.

Lemma 1 *The following hold for all t in the interior of the domain of c, and the corresponding one-sided versions of (i) and (ii) hold at the endpoints of non-closed curves.*

(i) $lim_{\epsilon\to 0^+}\kappa(c|_{[t,t+\epsilon]}) = lim_{\epsilon\to 0^+}\kappa(c|_{[t-\epsilon,t]}) = 0.$

(ii) T_c has right and left limits at t, $T_c(t+)$ and $T_c(t-)$, respectively.

(iii) If $t \in (a,b)$, then $\kappa(c|_{[a,b]}) = \kappa(c|_{[a,t]}) + |\angle(T_c(t-), T_c(t+))| + \kappa(c|_{[t,b]})$

Proof: Richardson, lemma 4.3.

Line (ii) states that the one-sided tangents exist everywhere. Line (iii) states that the total arc curvature is equal to the sum of the curvature upto the point and the curvature after the point and any curvature at the point. That is $\angle(T_c(t-), T_c(t+))$ denotes the the angle through which the tangent moves at the point $c(t)$.

In general we will allow $T_c(t)$ to be any element in $\angle(T_c(t-), T_c(t+))$. (I.e. it is a generalized tangent.) Let $\mathcal{J}_c := \{t : T_c(t+) \neq T_c(t-)\}$. The *tangent hull* is $\tilde{T}_c(t) := \angle(T_c(t-), T_c(t+))$. Also we will find it convenient later to talk about $\pm\tilde{T}_c(t) := \tilde{T}_c(t) \cup -\tilde{T}_c(t)$. We will in general be interested in the unoriented properties of c. This is because the line intersection function n_C does not tell us anything about the direction of parameterization of the curves in C.

Lemma 2 *Assume $T_c(t+) \neq -T_c(t-)$. If $t_i' \uparrow t, t_i'' \downarrow t$, and $\theta_i \to \theta$, where $\theta_i = \frac{c(t_i'')-c(t_i')}{|c(t_i'')-c(t_i')|}$, then $\theta \in \tilde{T}_c(t)$. Conversely, if $\theta \in \tilde{T}_c(t)$, then such sequences exist.*

Proof: Richardson, lemma 4.5

Note that when $T_c(t+) = -T_c(t-)$ it is not clear which way the tangent will turn.

2.3 Extended Indicatrix

As usual let $c(t)$ be an arclength parameterized curve. This time we insist it be closed. In this section we will show how the curve c induces a curve on $\mathcal{G}$ called the *extended indicatrix*. The indicatrix will be closed, parameterized by arc curvature, and have period $\kappa(c)$.

We will use the variable s to denote the arc curvature parameter. Intuitively $s(t) = \int_0^t |\kappa(t')|dt' + \sum$ (angles at all corners upto time t).

The mapping $s(t)$ is not invertible in general so we define:

$$\tau_c(s) := inf\{t \in [0, \infty) : \kappa(c|_{[0,t]}) \geq s\}.$$

Lines have no curvature so they do not contribute anything to the total absolute curvature. We say an interval $[a, b]$ is a *line segment* in c if there exists a θ such that $c(t) = c(a) + (t - a)\theta$ for $t \in [a, b]$. An interval is a *maximal line segment* if it is not a proper subinterval of another line segment. Let $L_c := \cup\{(a, b) : [a, b]$ is a maximal line segment in $c\}$.

Note that if $\tau_c(s+) > \tau_c(s-)$, then $[\tau_c(s+), \tau_c(s-)]$ is a maximal line segment in c. (Richardson, lemma 5.1)

Definition 1 *If c is a closed rectifiable curve of finite total absolute curvature in R^2, then we define its extended indicatrix $\psi_c : R \to \mathcal{G}$ by*

$$\psi_c(s) := (\rho_c(s), \theta_c(s)),$$

where $\theta_c(s) \in \tilde{T}_c(\tau_c(s))$ is defined by the condition $\kappa(c|_{[0,\tau_c(s)]}) + |\angle(\theta_c(s), T_c(\tau_c(s)-))| = s$. And $\rho_c(s) := \langle c(\tau_c(s)), *\theta_c(s)\rangle$.

Lemma 3 *If $\theta \in \tilde{T}_c(t)$, then there exists a unique $s \in R$ such that $t \in [\tau_c(s-), \tau_c(s+)]$*

and $\psi_c(s) = (\langle c(t), *\theta\rangle, \theta)$.

Proof: Richardson, lemma 5.2.

This lemma shows that if θ is a tangent or a generalized tangent to c at $c(t)$ then $(\langle c(t), *\theta\rangle, \theta)$ must lie on the extended indicatrix.

Theorem 3 *The function $\psi_c(s)$ is a closed rectifiable curve of period $\kappa(c)$, and the function $\theta_c(s)$ is an arclength parameterized closed rectifiable curve.*

Proof: Richarson, theorem 5.3.

An important part of the reconstruction of c depends on the tangents to ψ_c on $\mathcal{G}$.

Lemma 4 *For any $s \in R$, there exists sequences $s_j \uparrow s$ and $s_k \downarrow s$ such that*

$$(i) lim_{j\to\infty} \frac{\rho_c(s_j) - \rho_c(s)}{\mathcal{R}(\theta_c(s_j), \theta_c(s))} = -\langle c(\tau_c(s-)), \theta_c(s)\rangle.$$

$$(ii) lim_{k\to\infty} \frac{\rho_c(s_k) - \rho_c(s)}{\mathcal{R}(\theta_c(s_k), \theta_c(s))} = -\langle c(\tau_c(s+)), \theta_c(s)\rangle.$$

Proof: Richardson, lemma 5.4.

A point $x \in R^2$ can be decomposed into $x = \langle x, *\theta\rangle * \theta + \langle x, \theta\rangle\theta$. Where θ and $*\theta$ are taken to be unit vectors as described in section 2.1.2. We know $\rho = \langle x, *\theta\rangle$ and by lemma 4 we know $\langle x, \theta\rangle$.

When (i) and (ii) converge to the same limit for all sequences $s_j \uparrow s$ and $s_k \downarrow s$ we see that ψ_c has one-sided tangents at $\psi_c(s)$. They are equal only when $\tau_c(s-) = \tau_c(s+)$. When $\tau_c(s-) = \tau_c(s+)$ and the limits are unique then there exists a tangent at $\psi_c(s)$ and therefore there is a locally unique preimage. Richardson defines the following map $\mathcal{F} : \mathcal{T}(\mathcal{G})\backslash\{\xi : d\theta(\xi) = 0\} \to R^2$ by

$$\mathcal{F}(\xi) := \rho * \theta - \frac{d\rho(\xi)}{d\theta(\xi)}\theta$$

where $(\rho, \theta) = \pi(\xi)$ and $\mathcal{T}(\mathcal{G})$ is the tangent bundle to $\mathcal{G}$. This mapping allows us to map a given tangent $G \in tr\psi_c$ to a point $x \in trc$ where G is tangent to c at x. In some sense this is akin to the inverse function theorem. Since ψ_c is absolutely continuous we know that tangents exist almost everywhere.

Theorem 4 *A closed rectifiable curve of finite total absolute curvature is uniquely represented by its extended indicatrix.*

Proof: Richardson, theorem 5.5.

If $G = (\rho, \theta)$ then let $-G = (-\rho, -\theta)$. Define $g\_tr \pm \psi_c := g\_tr\psi_c + g\_tr - \psi_c$ and $tr \pm \psi_c := tr\psi_c \cup tr - \psi_c$ This allows us to remove any explicit orientation on the curves.

Note when $T_c(t+) = -T_c(t-)$ then $\pm\theta(s(t))$ is the whole space S^1. That is we can turn both clockwise and counterclockwise π degrees at the point $c(t)$.

Theorem 5 *If ξ is tangent to $tr \pm \psi_c$ and $\pm\pi_{\mathcal{G}}(\xi) = \pm\psi_c(s)$, then $c(\tau_c(s)) = \mathcal{F}(\xi)$.*

Proof: Richardson, theorem 5.6. If ξ is tangent to $tr \pm \psi_c$ at $\psi_c(s)$, then, for every sequence $G_i \to \pi_{\mathcal{G}}(\xi), \{G_i\} \subset tr \pm \psi_c$, we have $\frac{d\rho(\xi)}{d\theta(\xi)} = lim_{i\to\infty}\frac{\rho_i - \rho}{\mathcal{R}(\theta_i, \theta)}$. By lemma 4, $\tau_c(s-) = \tau_c(s+)$ and $-\langle c(\tau_c(s)), \theta\rangle = \frac{d\rho(\xi)}{d\theta(\xi)}$, hence $c(\tau_c(s)) = \mathcal{F}(\xi)$. If $\pi_{\mathcal{G}}(\xi) = -\psi_c(s)$, then $-\langle c(\tau_c(s)), \theta\rangle = \frac{d\rho(\xi)}{d\theta(\xi)}$ still holds since $\mathcal{R}(\theta_1, \theta_2) = \mathcal{R}(-\theta_1, -\theta_2)$, and the theorem follows. □

We assume our image space is bounded by a circle of radius R centered at the origin. That is $A \subset \{x \in R^2 : |x| \leq R\}$.

Lemma 5 *If the image space is bounded by a circle of radius R then both ρ and $\frac{d\rho(\xi)}{d\theta(\xi)}$ are bounded by R.*

Proof: We know $x = \mathcal{F}(\xi) := \rho * \theta - \frac{d\rho(\xi)}{d\theta(\xi)}\theta$. Since $|x|$ is bounded by R. We see that ρ and $\frac{d\rho(\xi)}{d\theta(\xi)}$ are bounded by R. This is true for all $x \in A$ except on a set of measure zero. But by continuity the result follows on this set of measure zero. □

2.4 $\mathcal{K}$-sets and Regularity

Until now we have been discussing properties of one curve c. But Richardson has shown that we can generalize this to $\mathcal{K}$-sets. We will define the generalized indicatrix to be $g\_tr \pm \psi_C := \sum_{c \in C} g\_tr \pm \psi_c$ and $tr \pm \psi_C := \cup_{c \in C} tr \pm \psi_c$.

We define an *arc-segment* of A to be any subset $E \subset \underline{A}$ such that $E = trc|_{(a,b)}$ for some simple curve c define on $[a, b]$ and $A(x)$ is constant on E.

Lemma 6 *If A is an exact $\mathcal{K}$-set and if $x \in \underline{A}$, then there exist arbitrarily small r and $\epsilon = \epsilon(r,x)$ such that $\underline{A} \cap (B_{r+\epsilon}(x) \backslash B_{r-\epsilon}(x))$ is a union of disjoint arc-segments of A.*

Proof: Richardson, lemma 6.1.

A corollary of this lemma is:

Corollary 1 *If C is equivalent to C' then $g\_trC = g\_trC'$.*

This states that there is only one exact $\mathcal{K}$-set for a $\mathcal{K}$-set A but that there are many realizations of it in terms of sets of curves C.

We can now generalize theorem 5.

Theorem 6 *If ξ is tangent to $tr \pm \psi_C$ and $\pm\pi_{\mathcal{G}}(\xi) = \pm\psi_c(s)$, then $c(\tau_c(s)) = \mathcal{F}(\xi)$.*

Proof: Richardson, theorem 6.2. We note that either ξ is tangent to one of the extended indicatrices ψ_c or that $tr \pm \psi_C$ has a tangent line at $\pi_{\mathcal{G}}(\xi)$ but none of the extended indicatrices $tr \pm \psi_c$ do. In the latter case it then follows that either ξ or $-\xi$ is a one-sided tangent to $tr \pm \psi_c$. $\square$

The exceptional case described in theorem 6 occurs when there exist curves $c_1, c_2 \in C$ which have inflectional points with the same tangent at the same point $x \in R^2$. See Figure 2-1.

To understand this better note that $\psi_c(s)$ has a tangent everywhere except at points where $lim_{j\to\infty} \frac{\rho_c(s_j)-\rho_c(s)}{\mathcal{R}(\theta_c(s_j),\theta_c(s))} = -\langle c(\tau_c(s-)), \theta_c(s)\rangle$ or $lim_{k\to\infty} \frac{\rho_c(s_k)-\rho_c(s)}{\mathcal{R}(\theta_c(s_k),\theta_c(s))} = -\langle c(\tau_c(s+)), \theta_c(s)\rangle$ fail to exist or whenever $\tau_c(s-) \neq \tau_c(s+)$. And $\psi_c(s)$ has a tangent line everywhere it has a tangent except when $\frac{d\theta_c(s-)}{ds} \neq \frac{d\theta_c(s+)}{ds}$. Which of course represents an inflection point.

Lemma 7 *If $n_C(G) = \infty$ then $G \in tr \pm \psi_C$, but $tr \pm \psi_C$ does not have a tangent at G.*

Proof: Richardson, lemma 6.3.

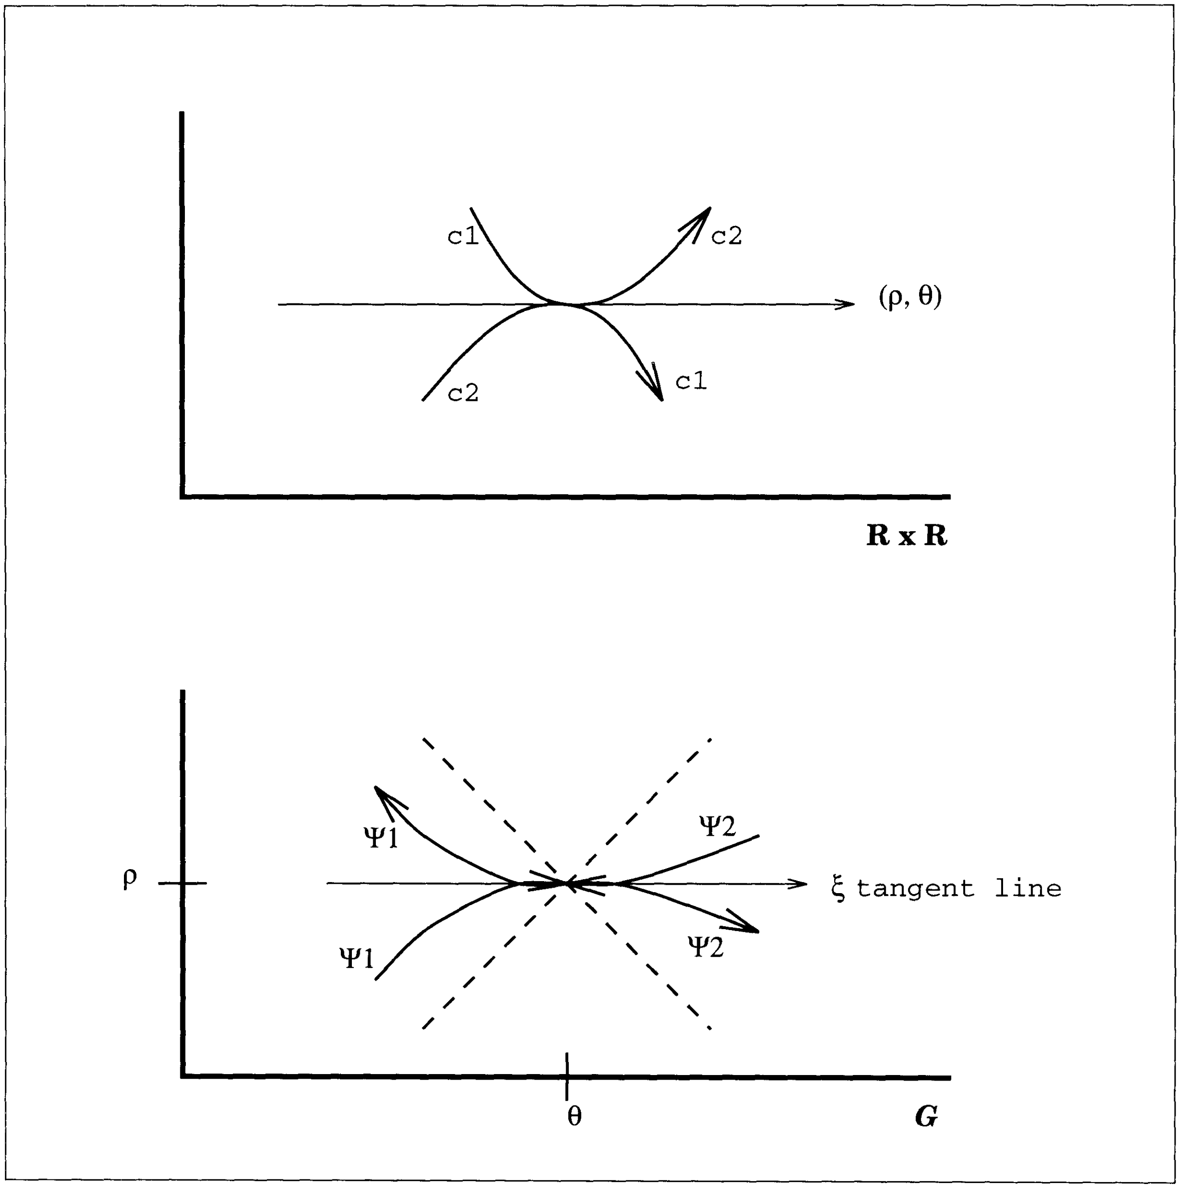

Figure 2-1: Example for Theorem 6

The only points of the extended indicatrix that do not have tangents are those $G \in tr \pm \psi_C$ that represent lines in the curves C, multi-tangents, and those points that represent the limit point of an infinite oscillation.

It turns out that if $n_C(G)$ is uncountably infinite then G must lie on a line segment of c. If $n_C(G)$ is countably infinite then c must oscillate around the line G an infinite number of times. We can find curves with infinite oscillation that still have finite turn and finite length. For example take the function $e^{-\frac{1}{t^2}} sin(\frac{1}{t})$ for $t \in [0, 1]$. We will discuss this again in chapter 3.

Next we show why the jumpset of n_C can be used to find the extended indicatrix.

Lemma 8 *If $tr \pm \psi_C$ has a tangent at $G = (\rho, \theta)$ and $\pm\psi_c(s) = \pm G$, then*

$$lim_{\delta \to 0^+} lim_{\epsilon \to 0^+} |c^{-1}(\underline{G}_\epsilon) \cap (\tau_c(s) - \delta, \tau_c(s) + \delta)| \in \{0, 1, 2\}$$

$$lim_{\delta \to 0^+} lim_{\epsilon \to 0^+} |c^{-1}(\underline{G}_\epsilon) \cap (\tau_c(s) - \delta, \tau_c(s) + \delta)| + |c^{-1}(\underline{G}_{-\epsilon}) \cap (\tau_c(s) - \delta, \tau_c(s) + \delta)| = 2$$

Proof: Richardson, lemma 6.5.

Where $\underline{G}_\epsilon = (\rho + \epsilon, \theta)$ and $\underline{G}_{-\epsilon} = (\rho - \epsilon, \theta)$. The value one occurs in the first equation of lemma 8 when $c(\tau_c(s))$ is an inflection point on the curve c.

See Figure 2-2.

2.5 Total Curvature Minimization

For any exact $\mathcal{K}$-set A corollary 1 tells us that all exact representations have the same generalized traces; but they may have different total absolute curvature. As we will see in the next section the function n_C gives us information only on the curvature minimal representation. This section describes some properties of curvature minimal sets. We will also show the surprising fact that even though curvature minimality is a global property it can be described completely locally.

Theorem 7 *Every $\mathcal{K} -$ set has a curvature minimal representation.*

Proof: Richardson, theorem 6.6.

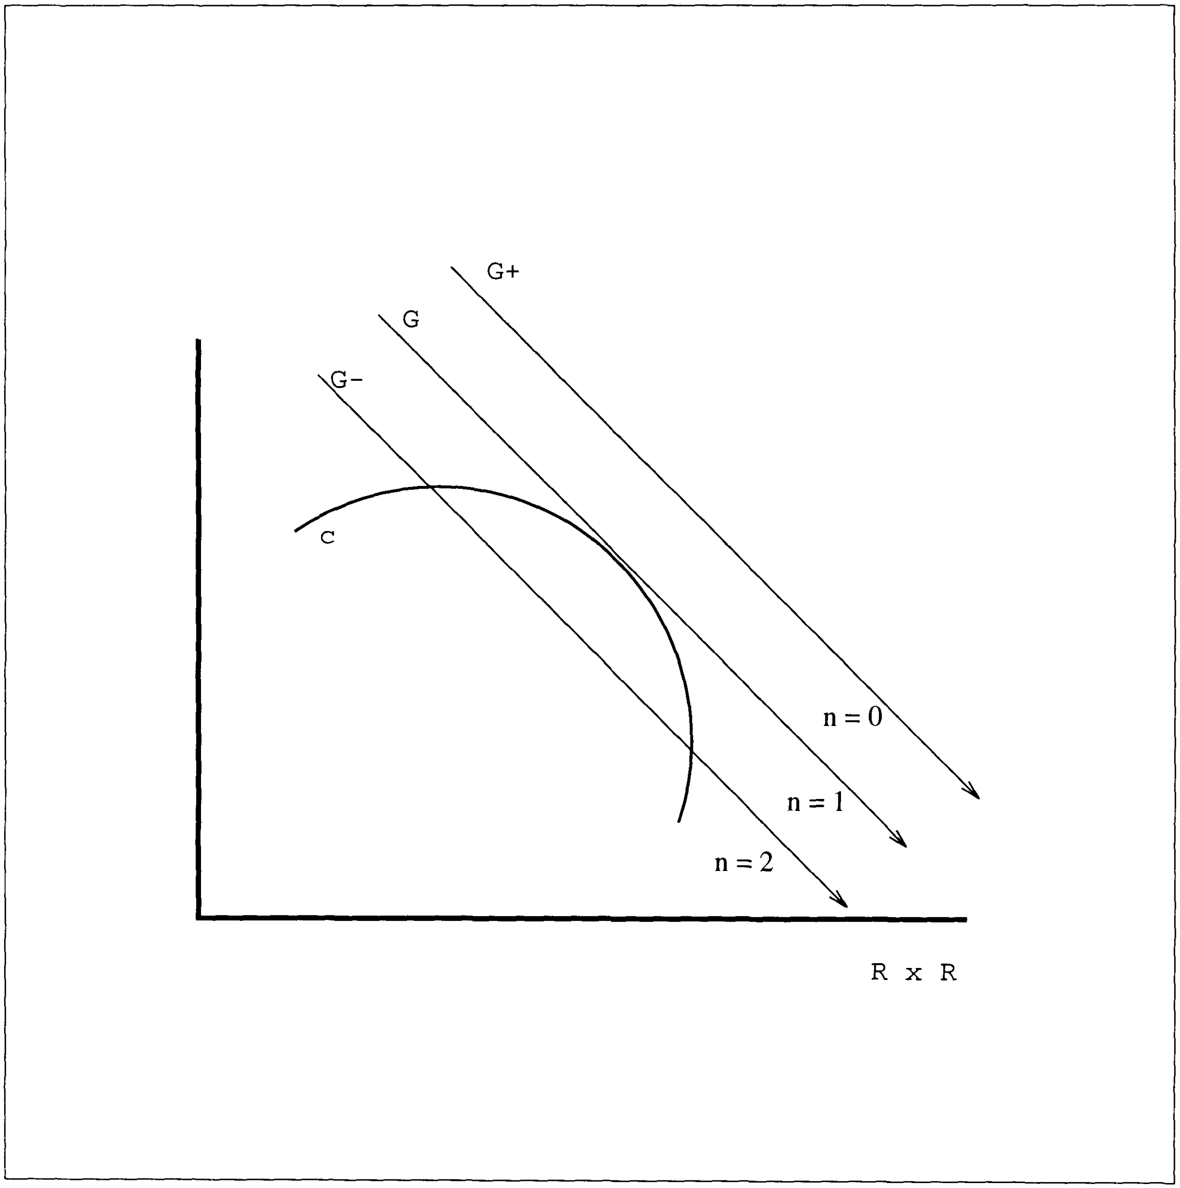

Figure 2-2: Tangents to c and the jumpset of n_c

To measure total absolute curvature we need a way to measure the curvature at every point $x \in R^2$. Define:

$$g\_TanC(x)(\theta) := \sum_{c \in C} \{ |\{t \in c^{-1}(x) : T_c(t+) = \theta\} mod \mathcal{L}(c)|$$
$$+ |\{t \in c^{-1}(x) : -T_c(t-) = \theta\} mod \mathcal{L}(c)| \}$$

That is $g\_TanC(x)(\theta)$ equals the number of arc-segments that leave the point x in the θ direction.

Let $g\_TanC(x) = \{\theta_1, ...\theta_m\}$ equal the generalized set of angles with associated multiplicity. It contains all the tangential angles of all the arc-segments leaving the point x. Our job is to match the different arc-segments in a curvature minimal way. We call a matching that matches $T_C(t+)$ to $-T_C(t-)$ at x, $\omega_C(x)$. We call any such matching K-optimal if it is a curvature minimal matching. That is for every possible matching $\omega(x)$ over $g\_TanC(x)$ we have to minimize the function $k(x) = \sum_{i=1}^{m} |\angle(\theta_i, -\omega(\theta_i))|$.

Lemma 9 *If C is curvature minimal , then the matching $\omega_C(x)$ is K-optimal on $g\_TanC(x)$ for all $x \in R^2$.*

Proof: Richardson, lemma 6.7.

The converse is also true. We can describe curvature minimality in a completely local way. The intuition is as follows: given two equivalent sets C and C' their generalized traces must agree everywhere. Thus C and C' can differ only in the way they connect at crossings. Specifically they can only differ at points of crossings where there is a choice of how to match paths.

Lemma 10 *If $\omega_C(x)$ is K-optimal for all $x \in R^2$, then C is curvature minimal.*

Proof: We prove that if C is not curvature minimal then it is not K-optimal everywhere. Let C be any non-minimal curvature representation of A. Let $\tilde{C}$ be a curvature minimal representation. By corollary 1 $g\_trC = g\_tr\tilde{C}$. Now $\tilde{k}(x) \leq k(x)$ for all $x \in R^2$. If there was one point x where $\tilde{k}(x) > k(x)$ then $\tilde{C}$ would not be a curvature minimal representation since we could change the matching of $\tilde{C}$ at x to that of the matching C uses at x.

Now there must be at least one point x where $\tilde{k}(x) < k(x)$ is a strict inequality. Otherwise $\kappa^*(C)$ and $\kappa^*(\tilde{C})$ would be the same. But this states that C is not K-optimal at the point x. □

An obvious question to ask is whether the curvature minimal representation is unique. As we will see in the next section the answer is no.

2.6 Relation of n_C to ψ_C

As stated before the jumpset of the function n_C is closely related to the $tr \pm \psi_C$. Lemma 8 gives us some intuition as to why this is the case. See Figure 2-2.

Lemma 11 *If $G \notin tr \pm \psi_C$, then n_C is constant on $B_\rho(G)$ for some $\rho > 0$.*

Proof: Richardson, lemma 7.1.

Note this implies that $S_{n_C} \subset tr \pm \psi_C$. By lemma 8 we see that inflection points in C are not counted in S_{n_C} so the subset is strict.

The determination of $tr \pm \psi_C$ is not enough to reconstruct C. We need the multiplicities of the ψ_c functions: $g\_tr \pm \psi_C$.

Theorem 8 *If $tr \pm \psi_C$ has a tangent line ξ at $G = (\rho, \theta)$, then $n_C^+(G) - n_C^-(G) \leq g\_tr \pm \psi_C(G)$. If furthermore C is curvature minimal then the equality holds for $\mathcal{H}^1$-almost all $G \in tr \pm \psi_C$.*

Proof: We will give a pseudo-proof of this result. A more rigourous proof can be found in Richardson, theorem 7.3. First we note that the slope of ξ is bounded. Now since ξ is a tangent line to $tr \pm \psi_C$ we know that $G = \pi_{\mathcal{G}}(\xi)$ is not tangent to an isolated inflection point of some curve $c \in C$. We say isolated because it is possible for G to be a tangent line to separate inflectional points as described in theorem 6. This tells us that the curve $tr \pm \psi_C$ does not make a π degree turn at G. Therefore there exists a ball $B_\rho(G)$ that is divided into two regions, an upper and a lower region, by $tr \pm \psi_C$. n_C will be constant over the two connected components, by lemma 11 and lemma 6. Thus $n_C^+(G) - n_C^-(G)$ can be replaced with $n_C(G_+) - n_C(G_-)$. Thus by lemma 8 the jump in n_C at G divided by two is equal to the $g\_tr + \psi_C(G)$. The factor of two comes into play because $g\_tr + \psi_C(G) = g\_tr - \psi_C(G)$. □

2.7 Richardson's Reconstruction

In this section we describe Richardson's approach to reconstructing a $\mathcal{K}$-set A from n_A. In the next chapter we will propose a similar approach with some changes that aid in algorithmic constructions later.

We define two functions: $\Gamma_A^{curve}(x,\theta), \Gamma_A^{linear}(x,\theta) : R^2 \times S^1 \to R$.

Let $\xi \in \mathcal{T}(\mathcal{G})$ where $\pi_{\mathcal{G}}(\xi) = G = (\rho, \theta)$. If ξ is tangent to S_{n_A}, set

$$\Gamma_A^{curve}(x,\theta) := n_A^+(G) - n_A^-(G).$$

If no such ξ exists for a given (x,θ) pair then set $\Gamma_A^{curve}(x,\theta) = 0$. This captures all the curved parts of the set A.

Now let

$$\Gamma_A^{linear}(x,\theta) := 2\lfloor liminf_{r\to 0^+} \frac{\int_{\underline{G}\cap B_r(x)} A d\mathcal{H}^1}{2r} \rfloor.$$

Where $\lfloor\cdot\rfloor$ is the floor function. This captures all the linear parts of the set A.

Now let us define the generalized set Γ_A in R^2 by

$$\Gamma_A(x) := sup_{\theta\in S^1}(\Gamma_A^{linear} + \Gamma_A^{curve})(x,\theta).$$

This recovers most of the original set A except points of inflection, endpoints, and multi-tangents.

Lemma 12 *Let A be an exact K-set and C any curvature minimal representation of A. Then $\Gamma_A \leq \frac{1}{2} g\_trC$ and the equality holds on a dense subset of $\underline{A}$.*

Proof: Richardson, lemma 8.3.

Finally we can reconstruct all of A by the following two operations:

$$A = \odot cl \Gamma_A.$$

Where $\odot A(x) := \frac{1}{2} \operatorname{limsup}_{r \to 0^+} \int_{\delta B_r(x)} A d\mathcal{H}^0$ and $clA(x) := \operatorname{limsup}_{r \to 0} sup_{y \in B_r(x)} A(y)$. The closure operation allows us to pick up the endpoints and inflectional points that were missed because they do not have tangent lines in their respective extended indicatrix. The circledot function adjusts the multiplicity of A for crossings.

In summary Richardson's reconstruction requires two scans over R^2. One scan looks for curved parts the other looks for linear parts. For every point $x \in R^2$ we look to see what contributes more the curved or the linear parts of the image. Then finally we "clean-up" the measure zero error by using the circledot and closure functions.

Chapter 3

Reconstruction

We are interested in implementing a reconstruction algorithm. To this end it will be necessary to enforce some regularity conditions on the types of sets we would like to reconstruct.

Richardson's reconstruction algorithm, section 2.7, requires two passes of n_A. The first pass looks for the the curved parts of A and the second pass looks for the linear parts of A. One of Richardson's open problems concerns whether A can be recovered just from Γ_A^{curve}. Another way to state this question is are both parses of R^2 necessary? By restricting the class curves to be reconstructed we can use Γ_A^{curve} to reconstruct the linear parts of A.

In what follows we will show that one can reconstruct the linear parts of A by looking at the so-called singularities of $tr \pm \psi_C$.

Let C be a curvature minimal representation of the exact $\mathcal{K}$-set A.

Definition 2 *A singularity of $tr \pm \psi_C$ is any point $G \in tr \pm \psi_C$ that does not have a tangent line and/or is not a regular point.*

We will now motivate why looking at singularities can be helpful in reconstructing the linear parts of A. Take for example a curve c containing a line segment and its corresponding indicatrix ψ_c. We know that ψ_c will have a corner at a point $G \in \mathcal{G}$ where the line G represents the tangent to the line segment. To determine the length and position of the line segment, though, we need to know the slopes of $tr\psi_c$ coming into and leaving the point G. But once we know these, and by using the fact that the curve c is continuous, we can reconstruct the whole line segment.

We will show that one can classify the singularities and thus reconstruct all the linear segments. In general there are a countable number of singularities. We need to know both where the singularities are and also what the slopes of ψ_C are coming and leaving the singularity G.

That is we need to ensure that the left and right tangents of ψ_C exist at every singularity. We first give an example of a curve where the right and left tangents of ψ_C do not exist around a singularity. Then we will show how to restrict the class of curves so that this problem will not occur.

Specifically consider the curve $c(t) = (t, e^{-\frac{1}{t^2}} sin(\frac{1}{t}))$ over the interval $t \in [-1, 1]$. It can be shown that this curve has both finite length and finite turn. But note that the line equal to the x-axis crosses this curve an infinite number of times. Thus by lemma 7 the right and left-sided tangents at $t = 0$ of ψ_c do not exist.

To get around these problems we will assume that our curves only have a finite number of inflection points. This avoids the problems of rapid oscillation. This condition allows us to partition C into a finite number of concave and convex pieces. Thus we are able to do away with the differentiability problems, such as the one discussed in the example above.

The next section will show that the set of curves with finite inflection is still very large and expressive.

3.1 Curves With Finite Number of Inflection Points

Definition 3 *Let $D_1 = (s_0 - \delta, s_0)$ and $D_2 = (s_0, s_0 + \delta)$. Then x is an inflection point if there exists a $\delta > 0$ such that $\theta(s)$ is increasing (decreasing) on D_1 and $\theta(s)$ is decreasing (increasing) on D_2.*

Lemma 13 *If $n_c(G) = \infty$ and $\tau_c(s_0-) = \tau_c(s_0+)$ then the curve c has an infinite number of inflection points.*

Proof: By Richardson, lemma 6.3, there exists $t_i \uparrow t$ such that $c(t_i) \in \underline{G}$ and there exist $\tilde{t}_i \in (t_i, t_{i+1})$ such that $\langle c(t') - c(t), *\theta \rangle$ attains a non-zero extrema at $t' = \tilde{t}_i$. This in turn implies $\theta \in T_c(\tilde{t}_i)$. WLOG assume $\tilde{t}_i$ is a maximal extrema. Then there must exist a minimal extrema $\tilde{t}_j$, where $j > i$. This implies $\theta_c(t')$ is decreasing for t' close to but greater than $\tilde{t}_i$ and $\theta_c(t')$ is increasing for t' close to but less than $\tilde{t}_j$. By continuity of $\theta_c(t')$ there must exist a $\hat{t} \in (\tilde{t}_i, \tilde{t}_j)$ such that $c(\hat{t})$ is an inflection point. Thus there are an infinite number of inflection points in c. □

If our curves C have finite number of inflection points then we will not have the rapid oscillation problem.

We now show that curves with finite number of inflection points are $\epsilon-$dense in the space of curves of finite length and turn. We will do this by approximating curves by piecewise linear curves. Let $C_{K,L}$ equal the space of curves with length less than L and turn less than K and let $C_{K,L,n}$ be those curves in $C_{K,L}$ with $\leq n$ inflection points.

A piecewise linear curve of n segments has at most $n - 3$ inflection points. The worst case of $n - 3$ occurs when the piecewise linear curve has a sawtooth behavior.

Lemma 14 *Let $c \in C_{K,L}$. Then we can find a piecewise linear curve $\hat{c}$ with n segments such that $d_I(c, \hat{c}) < \epsilon = \frac{K^2 L}{8n}$.*

Proof: Kulkarni, lemma 6, implies that one can get $d_I(c, \hat{c}) < \epsilon$ with a curve $\hat{c}$ with at most $\frac{K^2 L}{8\epsilon}$ linear segments. □

Theorem 9 *The space $C_{K,L,n}$ is $\epsilon-$dense with respect to d_I in $C_{K,L}$ where $\epsilon(n) = \frac{K^2 L}{8(n+3)}$.*

Proof: Use lemma 14 and fact that piecewise linear curves with n segments have at most $n - 3$ inflection points. □

In chapter 5 we will be using the Hausdorff metric to approximate curves. $C_{K,L,n}$ is still $\epsilon-$dense in $C_{K,L}$ with respect to this metric.

Lemma 15 *Let $\hat{c}$ be the line segment that connects the endpoints of c. Then $d_H(c, \hat{c}) \leq \mathcal{L}(c)/2$.*

Divide the curve c and c' into segments c_i and c_i' such that $trc = \cup trc_i$ and $trc' = \cup trc_i'$.

Lemma 16 $d_H(c, c') \leq max\ d_H(c_i, c_i')$.

Lemma 17 *Let $c \in C_{K,L}$. Then we can find a piecewise linear curve $\hat{c}$ with n segments such that $d_H(c, \hat{c}) \leq \epsilon = L/2n$.*

Proof: Divide the curve c into n segments c_i each of length $\mathcal{L}(c)/n$. Then let $\hat{c}_i$ be the line segment connecting the endpoints of c_i. Then $d_H(c_i, \hat{c}_i) \leq \mathcal{L}(c)/2n \leq L/2n$. And $d_H(c, \hat{c}) \leq L/2n$. □

Theorem 10 *The space $C_{K,L,n}$ is $\epsilon-$dense with respect to d_H in $C_{K,L}$ where $\epsilon(n) = \frac{L}{2(n+3)}$.*

Proof: Use lemma 17 and fact that piecewise linear curves with n segments have at most $n - 3$ inflection points. □

Thus we have shown that as $n \to \infty$ we can get arbitrarily close approximations in both the d_I and d_H sense.

3.2 Reconstruction Theorem

We provide a reconstruction algorithm for all exact $\mathcal{K}-$sets A represented by C where the curves $c_i \in C$ have a finite number of inflection points.

Define $\Delta(G) := n_C(G_+) - n_C(G_-)$. Then $g\_tr \pm \psi_C(G) = \odot cl|\Delta(G)|$ for all points G where $g\_tr \pm \psi_C(G)$ has finite value. We do not "catch" those points G with multiplicity of infinity. See lemma 7. But it does not matter in our reconstruction. The closure fills in cusps in the ψ_c curve caused by inflection points in the respective c curve. And the circledot fills in crossings of the ψ_c curves caused by multi-tangency and π degree turns in the curves c.

Then it is easy to see that $\Gamma_A^{curve}(x, \theta) = \{x = \mathcal{F}(\xi) :$ where ξ is a tangent line at a regular point of $g\_tr \pm \psi_C\}$.

We know ψ_C is rectifiable by theorem 3. Because it has a finite number of inflection points it also has finite turn. To see this note that each inflection point gives at most a turn of π. But there are only a finite number of them. There may be a countable number of corners. But the sum of the angles in each corner must sum to a finite value. This is because the size of the angle measures the length of the line segment. And C is rectifiable.

Now by the local Euclidean nature of $\mathcal{G}$ and the fact that ψ_C has finite length and turn we can extend lemma 6 to the extended indicatrix.

Lemma 18 *If $G \in g\_tr \pm \psi_C$ and C has a finite number of inflection points then there exists arbitrarily small r and $\epsilon = \epsilon(r, G) > 0$ such that $g\_tr \pm \psi_C \cap (B_{r+\epsilon}(G) \backslash B_{r-\epsilon}(G))$ is a union of disjoint arc-segments of $g\_tr \pm \psi_C$.*

This allows us to identify the singularities by the slopes of the incoming arc-segments.

Theorem 11 *Let A be an exact $\mathcal{K}$-set with curvature minimal representation C. Let the curves $\{c_i\} = C$ all have finite number of inflection points. Then one can reconstruct A from $\Delta(G)$.*

Proof: From $\Delta(G)$ we can reconstruct $g\_tr \pm \psi_C$. In turn we can reconstruct $\Gamma_A^{curve}(x, \theta)$. Which in turn we can reconstruct Γ_A^{curve}. To reconstruct the linear parts we can use the height information of $\Delta(G)$. Specifically we will show in the next chapter that every singularity can be uniquely identified and mapped back to the curves in R^2 that created it. $\square$

Chapter 4

Singularities

In this chapter we classify all the kinds of behavior that one can expect of the function n_C (more specifically its jump set S_{n_C}), ψ_c for all $c \in C$, and $g\_tr \pm \psi_C$. As was shown in chapter two all points of trC that have nonnegative curvature can be reconstructed from the function $\mathcal{F}$ (except for a countable set of points representing inflection points and multi-tangents, and points of infinite oscillation). Unfortunately line segments do not map under $\mathcal{F}$. By assumption, though, the curves c are closed and continuous and therefore their extended indicatrices ψ_c are closed and continuous. By using this continuity condition we will show how to reconstruct those parts of $c \in C$ that have zero curvature.

Just as for points in R^2 there will be points $G \in \mathcal{G}$ where multiple arc-segments of $tr \pm \psi_C$ meet. For R^2 the matching problem consisted of matching the incoming arc-segments in a curvature minimal way. In $\mathcal{G}$ the condition for matching is different but we can use a similar notation. And as in R^2 the matching will not in general be unique. This is because curvature minimal representations of a $\mathcal{K}$-set A are not necessarily unique.

First we describe the singularities of $tr \pm \psi_C$. We prove some continuity conditions on n_C. We then describe the basic curve structures found in R^2, provide a matching algorithm, and prove its uniqueness in reconstructing the original curve.

4.1 Types of Singularities

Let C be a curvature minimal representation of the exact $\mathcal{K}$-set A. From now on assume the curves in C have a finite number of inflection points. For every $c \in C$ let ψ_c be its extended indicatrix.

We know from section 2.3 that ψ_c will have a tangent when $\psi_c(s)$ is differentiable at s and $\tau_c(s-) = \tau_c(s+)$. It will have a tangent line everywhere it has a tangent except when $\frac{d\theta_c(s-)}{ds} \neq \frac{d\theta_c(s+)}{ds}$.

If $\tau_c(s-) \neq \tau_c(s+)$, then by lemma 4, we see that $\frac{d\rho_c}{d\theta_c}|_{s-} \neq \frac{d\rho_c}{d\theta_c}|_{s+}$. That is the incoming and out going slopes are different. We see ψ_c has a corner at $\psi_c(s)$ which represents a line segment in R^2.

When $\tau_c(s-) = \tau_c(s+)$ but $\frac{d\theta_c(s-)}{ds} \neq \frac{d\theta_c(s+)}{ds}$ then ψ_c has a cusp at $\psi_c(s)$ which represents an inflection point in R^2.

ψ_c has a countable number of corners because each corner represents a line segment in c and c has finite length. ψ_c has a finite number of cusps because each cusp represents an inflection point in c and $c \in C_{K,L,n}$.

Note that locally $\mathcal{G} = R^1 \times S^1 = R^1 \times R^1$. We can thus define things like tangents on $\mathcal{G}$ in terms of the parameterized curves that lie on $\mathcal{G}$. Specifically let $T_{\psi_c(s+)} := lim_{s_i \uparrow s} \frac{\psi_c(s_i) - \psi_c(s)}{|\psi_c(s_i) - \psi_c(s)|}$ and $T_{\psi_c}(s-) = lim_{s_j \downarrow s} \frac{\psi_c(s_j) - \psi_c(s)}{|\psi_c(s_j) - \psi_c(s)|}$. Note if ξ is tangent to $tr \pm \psi_c$ at $\pm\psi_c(s)$ then $\pm T_{\psi_c} = \pm \frac{\xi}{|\xi|}$.

By an extension of lemma 1 we get:

Lemma 19 *For all s in the interior of the domain of ψ_c we have T_{ψ_c} has right and left limits at s, $T_{\psi_c}(s-)$ and $T_{\psi_c}(s+)$ respectively.*

We have shown that $tr \pm \psi_c$ will have a tangent line everywhere except at points $\psi_c(s)$ which represent inflection points, line segments, and multi-tangents in R^2. The same happens when we generalize to $tr \pm \psi_C$. It has tangent lines everywhere except at points which represent inflection points, line segments, and multi-tangents. The only new event that can happen occurs when two cusps in R^2 meet at a point. Like in theorem 6. See Figure 2-1.

Lemma 20 *There are a countable number of singularities in $tr \pm \psi_C$.*

4.2 The Matching Problem

Remember that $n_C(G_+) = \mathrm{limsup}_{\epsilon\downarrow 0} n_C(G_\epsilon)$ and $n_C(G_-) = \mathrm{limsup}_{\epsilon\downarrow 0} n_C(G_{-\epsilon})$. Define, as before, $\Delta(G) = n_C(G_+) - n_C(G_-)$.

Following Richardson's lead:

Define

$$g\_Tan\psi_C(\alpha) := \sum_{c\in C}\{|\{s \in \psi_c^{-1}(G) : T_{\psi_c}(s+) = \alpha\} mod \mathcal{L}(\psi_c(s))|$$

$$+|\{s \in \psi_c^{-1}(G) : T_{\psi_c}(s-) = -\alpha\} mod \mathcal{L}(\psi_c(s))|\}$$

Note that $g\_Tan\psi_C(G)$ is a generalized set of angles. If $\alpha \in g\_Tan\psi_C(G)$ then $\alpha = T_{\psi_c}(s+)$ or $\alpha = -T_{\psi_c}(s-)$ for some $c \in C$ and $s \in (0, \kappa(c))$. Define $\Delta(\alpha) = n_C((\psi_c(s+))_+) - n_C((\psi_c(s+))_-)$ or $\Delta(\alpha) = n_C((\psi_c(s-))_+) - n_C((\psi_c(s-))_-)$ respectively. $\Delta(\alpha)$ measures the sign and the magnitude of the n_C jump on the arc-segment represented by α. Note for closed curves c, Δ will always be even and that $\frac{1}{2}\Delta$ equals the multiplicity of the underlying arc-segment of ψ_C.

Since the tangents ξ are bounded we see that the angles α are confined to two separate intervals of S^1. Note when ξ is tangent to ψ_c at $\psi_c(s)$ and $T_{\psi_c}(s) = \alpha$ then $\frac{\langle\alpha,\hat{\rho}\rangle}{\langle\alpha,\hat{\theta}\rangle} = \frac{d\rho(\xi)}{d\theta(\xi)}$ where $\hat{\theta}, \hat{\rho}$ are the elementary basis vectors on $\mathcal{G}$.

Let us order the elements of $g\_Tan\psi_C(G) = \{\alpha_1, ...\alpha_m\}$ by slope. Specifically for $i < j$ we have $\frac{\langle\alpha_i,\hat{\rho}\rangle}{\langle\alpha_i,\hat{\theta}\rangle} \geq \frac{\langle\alpha_j,\hat{\rho}\rangle}{\langle\alpha_j,\hat{\theta}\rangle}$.

The matching we choose on $g\_Tan\psi_C(G)$ has no effect on the total absolute curvature of C. But it does effect the lengths of line segments in C. Specifically let ψ_C have a corner at $G = (\rho, \theta)$ with incoming and outgoing angles α_1 and α_2. Then the line segment represented by this corner has length $|\rho * \theta - \frac{\langle\alpha_1,\hat{\rho}\rangle}{\langle\alpha_1,\hat{\theta}\rangle}\theta - (\rho * \theta - \frac{\langle\alpha_2,\hat{\rho}\rangle}{\langle\alpha_2,\hat{\theta}\rangle}\theta)| = |\frac{\langle\alpha_1,\hat{\rho}\rangle}{\langle\alpha_1,\hat{\theta}\rangle} - \frac{\langle\alpha_2,\hat{\rho}\rangle}{\langle\alpha_2,\hat{\theta}\rangle}|$.

We will show in section 4.6 how to appropriately match the angles α.

Lemma 21 *Let $c_1, c_2 \in C$. Then $n_{c_1\cup c_2} = n_{c_1} + n_{c_2}$ and $g\_Tan\psi_{c_1\cup c_2}(G) = g\_Tan\psi_{c_1}(G) + g\_Tan\psi_{c_2}(G)$ and $\Delta_{c_1\cup c_2}(\alpha) = \Delta_{c_1}(\alpha) + \Delta_{c_2}(\alpha)$.*

Remember $n_c = n_{g\_trc}$ and $n_{c_1 \cup c_2} := n_{g\_trc_1 + g\_trc_2}$. This last lemma is important for it allows us to "stack" elementary curves on top of each other to create more complex behaviors in $\mathcal{G}$.

4.3 Types of Curve Behavior

Assume C is curvature minimal with a finite number of inflection points. There are nine basic types of curve behavior for $c|_{(t-\delta,t+\delta)}$ around the point $x = c(t) \in R^2$ for finite inflection curves.

Let $x \in trC$. Then there exists a $c \in C$ such that $x = c(t)$ for some $t \in [0, \mathcal{L}(c))$. Let $\theta \in \tilde{T}_c(t)$. Note this is a directed tangent. By lemma 3 there exists a uniquely determined $s_0 \in R$ such that $t \in [\tau_c(s_0-), \tau_c(s_0+)]$ and $\psi_c(s_0) = (\langle c(t), *\theta \rangle), \theta)$.

We say x is an *endpoint* if $+T_c(t+) = -T_c(t-)$.

We list the cases:

Case A: No endpoints near x.

$$+T_c(\tau_c(s_0-)+) \neq -T_c(\tau_c(s_0-)-)$$

$$+T_c(\tau_c(s_0+)+) \neq -T_c(\tau_c(s_0+)-)$$

A1 A point $x = c(t)$ is *concave up with respect to* θ if there exist δ_1, δ_2 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle > 0$ and for all $t' \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t') - c(\tau_c(s_o-)), *\theta \rangle > 0$. See Figure 4-1.

A2 A point $x = c(t)$ is *concave down with respect to* θ if there exist δ_1, δ_2 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle < 0$ and for all $t' \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t') - c(\tau_c(s_o-)), *\theta \rangle < 0$. See Figure 4-2.

A3 A point $x = c(t)$ is *up-down inflectional with respect to* θ if there exist δ_1, δ_2 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle > 0$ and for all $t' \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t') - c(\tau_c(s_o-)), *\theta \rangle < 0$. See Figure 4-3.

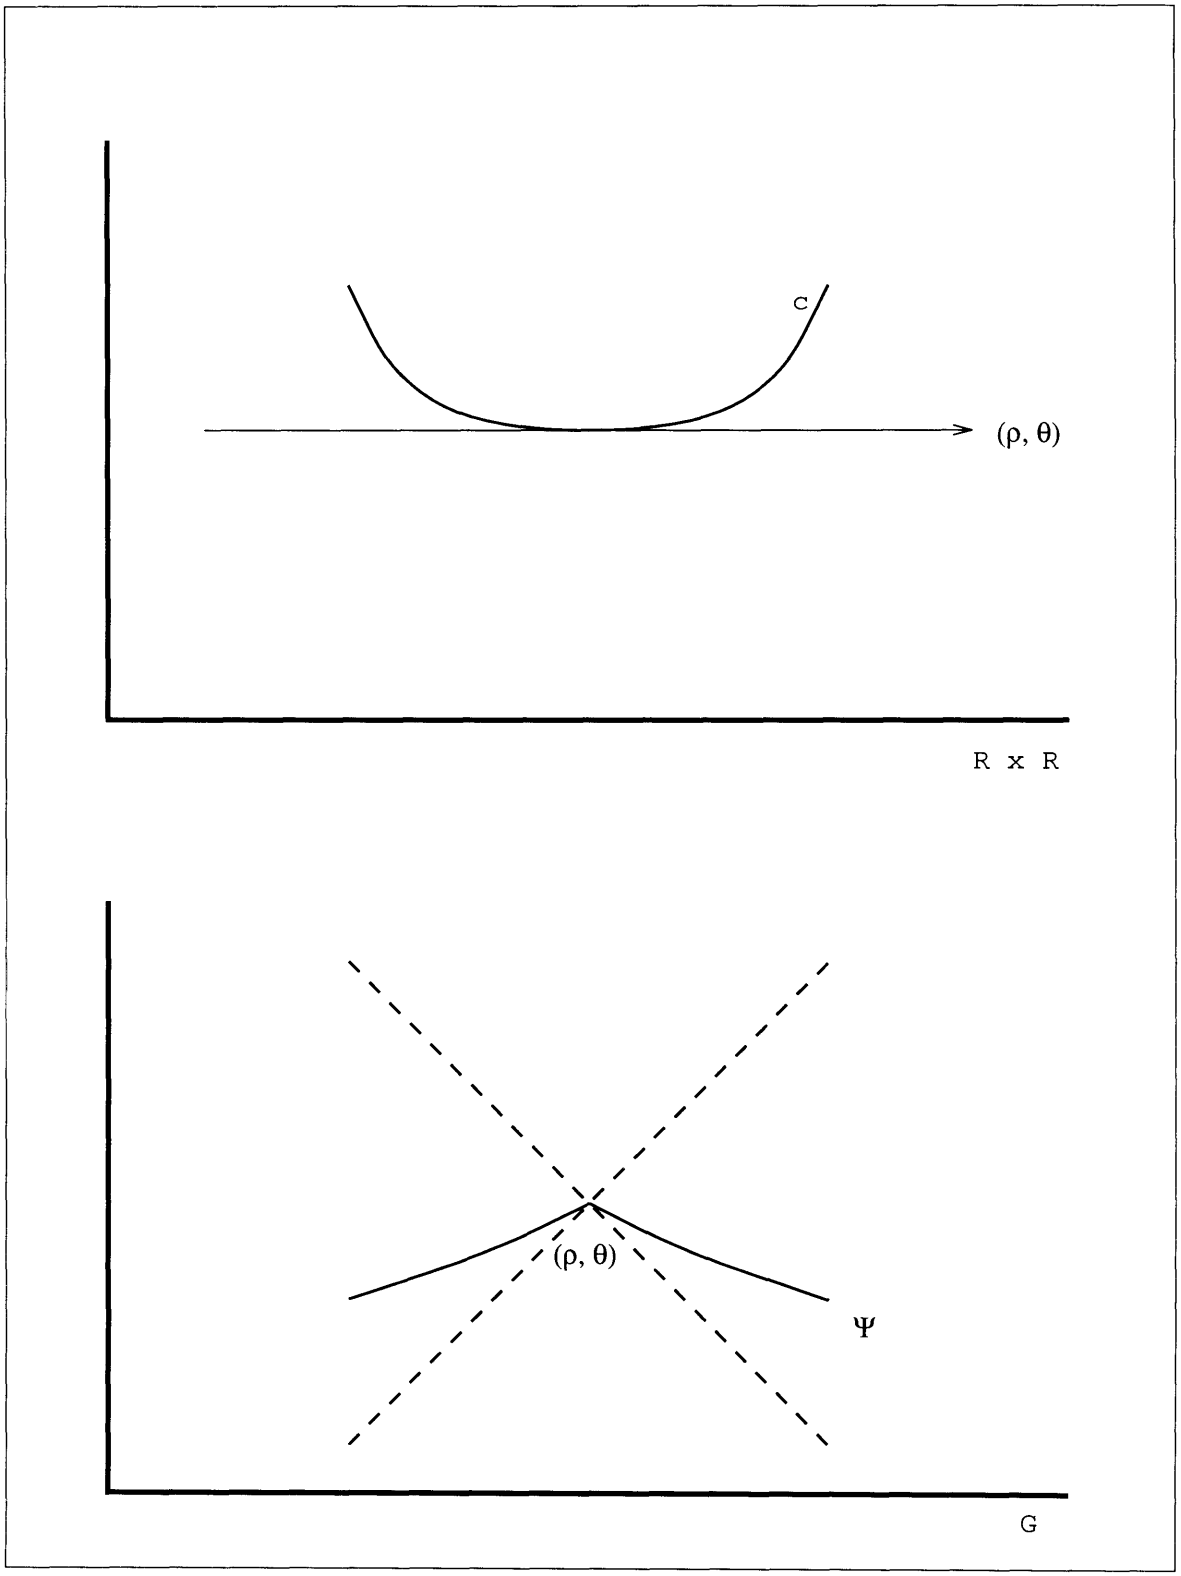

Figure 4-1: Type A1

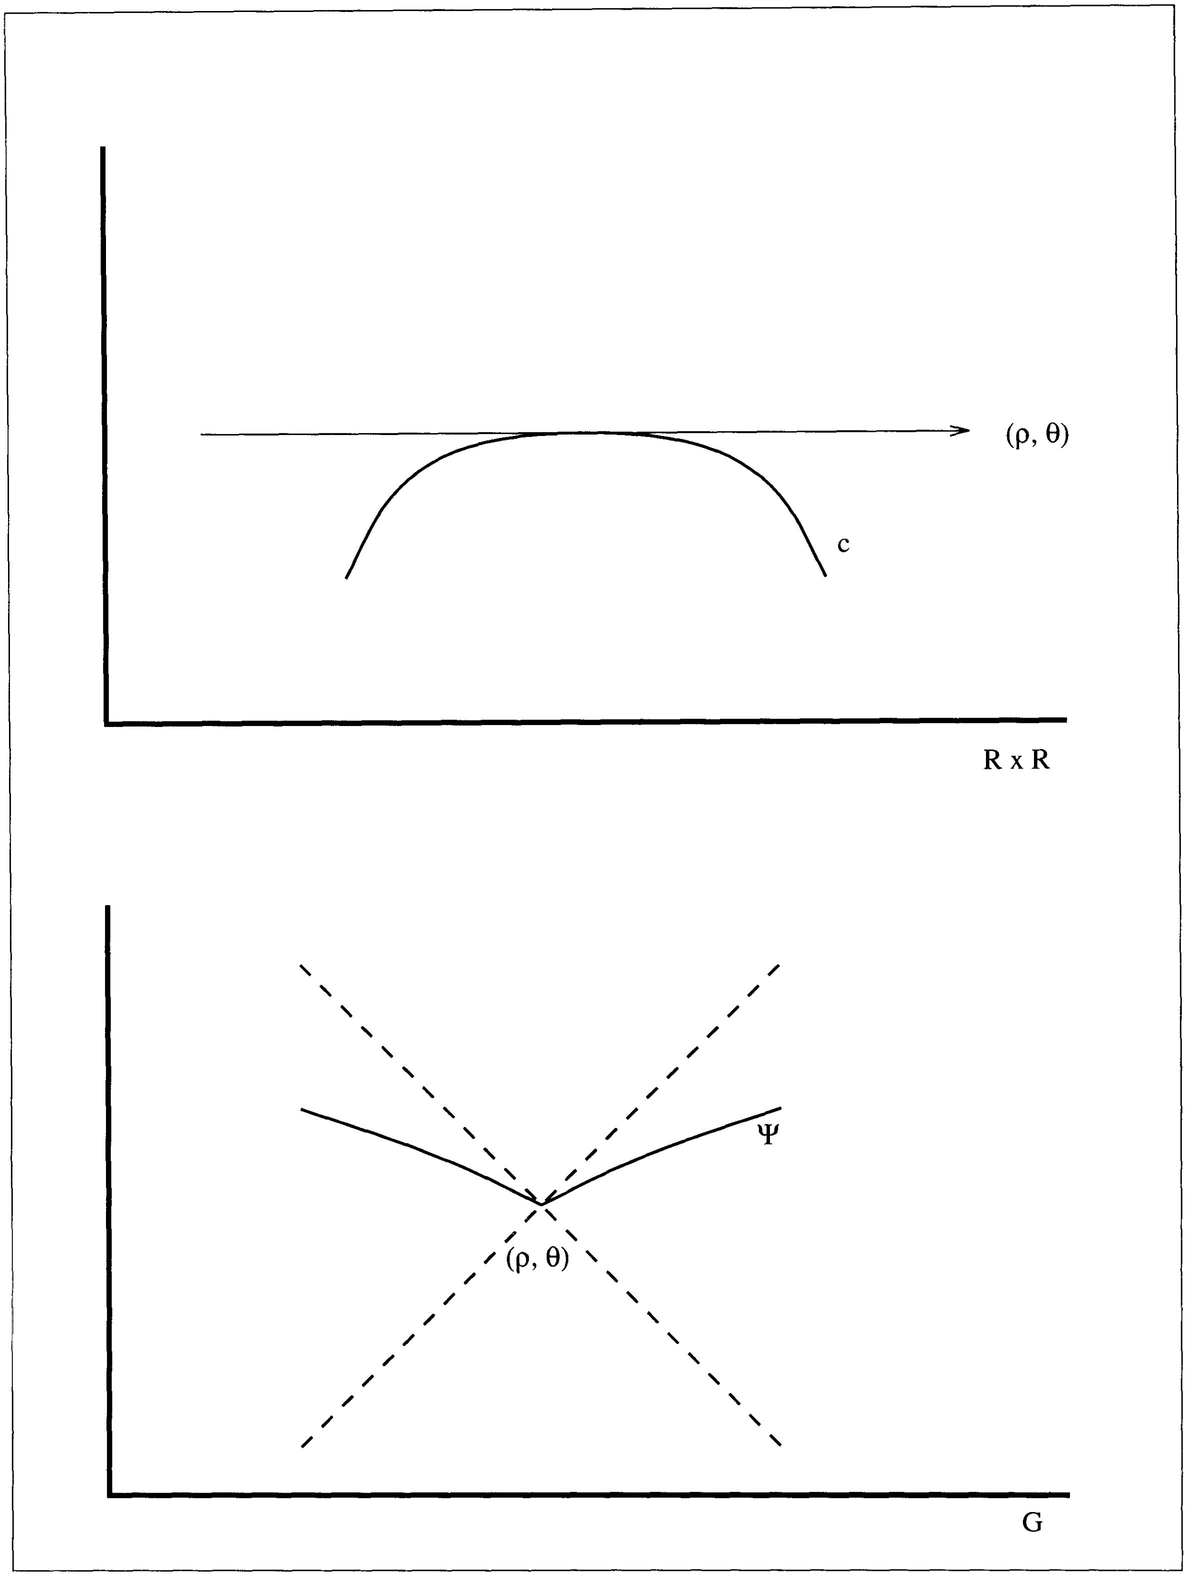

Figure 4-2: Type A2

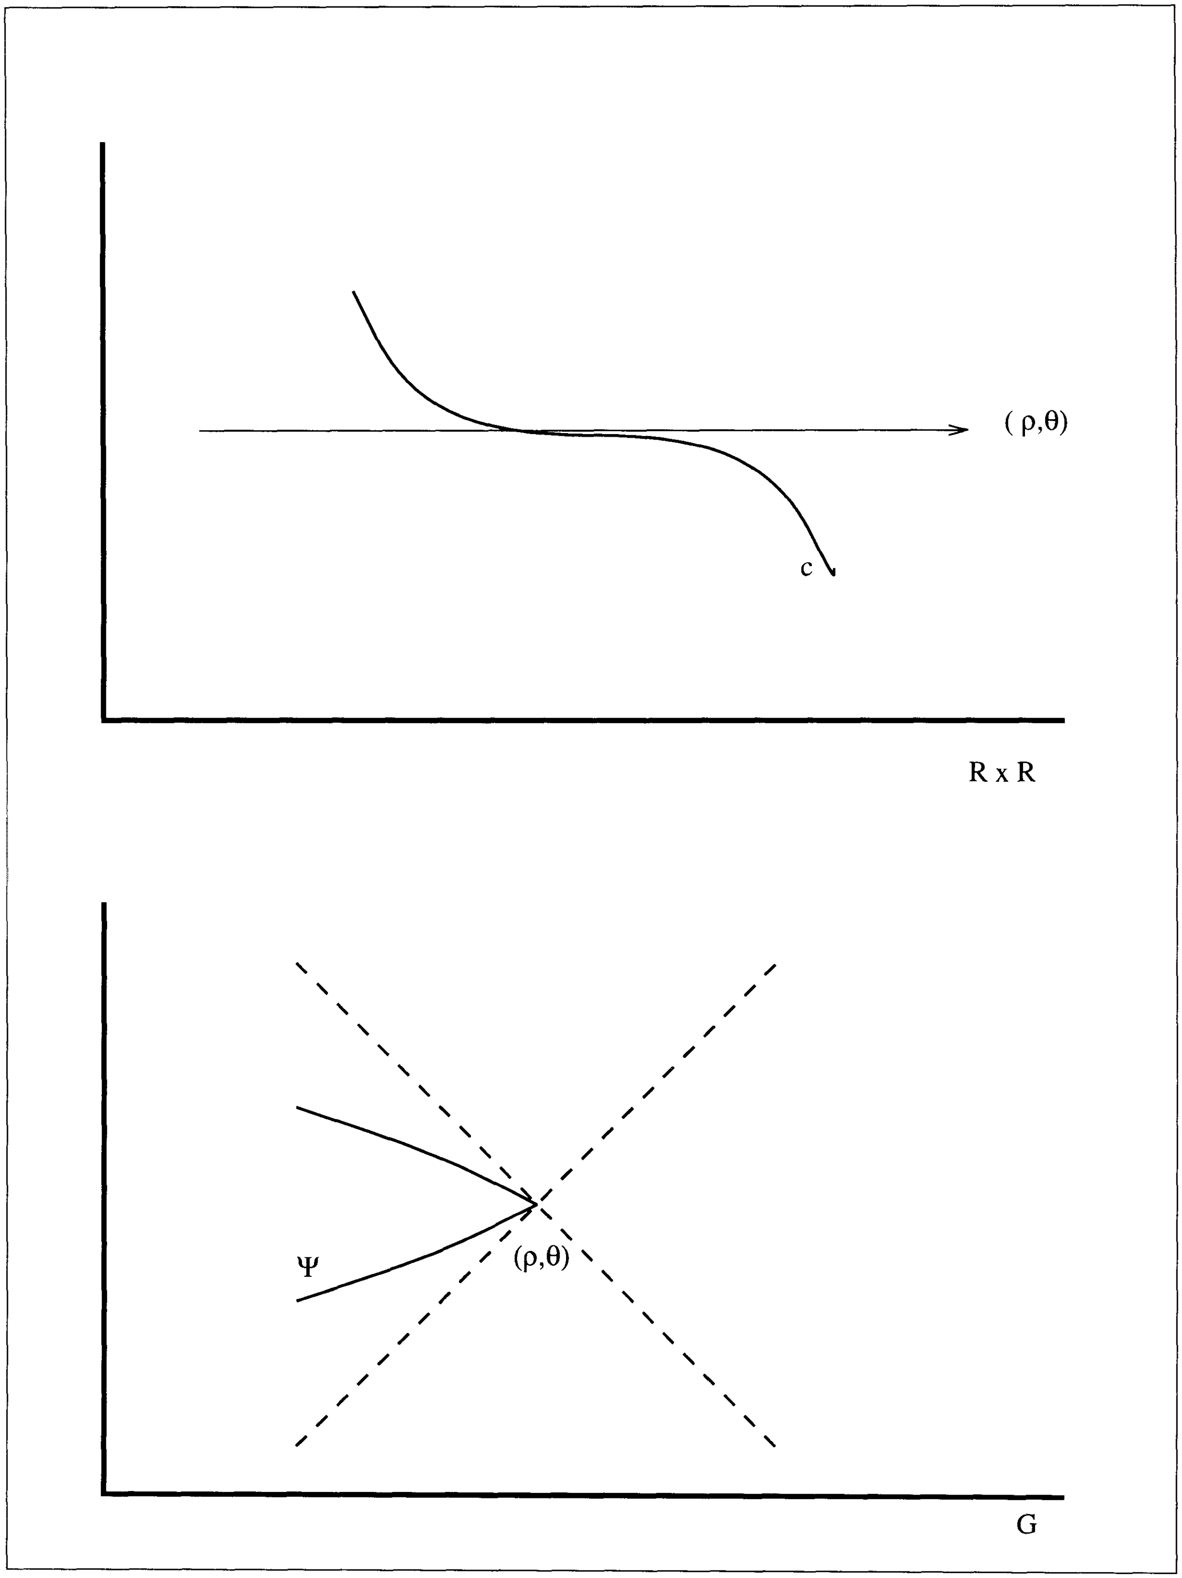

Figure 4-3: Type A3

A4 A point $x = c(t)$ is *down-up inflectional with respect to* θ if there exist δ_1, δ_2 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle < 0$ and for all $t' \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t') - c(\tau_c(s_o-)), *\theta \rangle > 0$. See Figure 4-4.

Case B Right endpoint near x.

$$+T_c(\tau_c(s_0-)+) \neq -T_c(\tau_c(s_0-)-)$$

$$+T_c(\tau_c(s_0+)+) = -T_c(\tau_c(s_0+)-)$$

B1 A point $x = c(t)$ is *concave up with a right endpoint with respect to* θ if there exists a δ such that for all $t \in (\tau_c(s_0-) - \delta, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle > 0$. See Figure 4-5.

B2 A point $x = c(t)$ is *concave down with a right endpoint with respect to* θ if there exists a δ such that for all $t \in (\tau_c(s_0-) - \delta, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_o-)), *\theta \rangle < 0$. See Figure 4-6.

Left endpoint near x.

$$+T_c(\tau_c(s_0-)+) = -T_c(\tau_c(s_0-)-)$$

$$+T_c(\tau_c(s_0+)+) \neq -T_c(\tau_c(s_0+)-)$$

B3 A point $x = c(t)$ is *concave up with a left endpoint with respect to* θ if there exists a δ such that for all $t \in (\tau_c(s_0+), \tau_c(s_0+)+\delta)$ we have $\langle c(t) - c(\tau_c(s_o+)), *\theta \rangle > 0$. See Figure 4-7.

B4 A point $x = c(t)$ is *concave down with a left endpoint with respect to* θ if there exists a δ such that for all $t \in (\tau_c(s_0+), \tau_c(s_0+)+\delta)$ we have $\langle c(t) - c(\tau_c(s_o+)), *\theta \rangle < 0$. See Figure 4-8.

Case C Two endpoints near x.

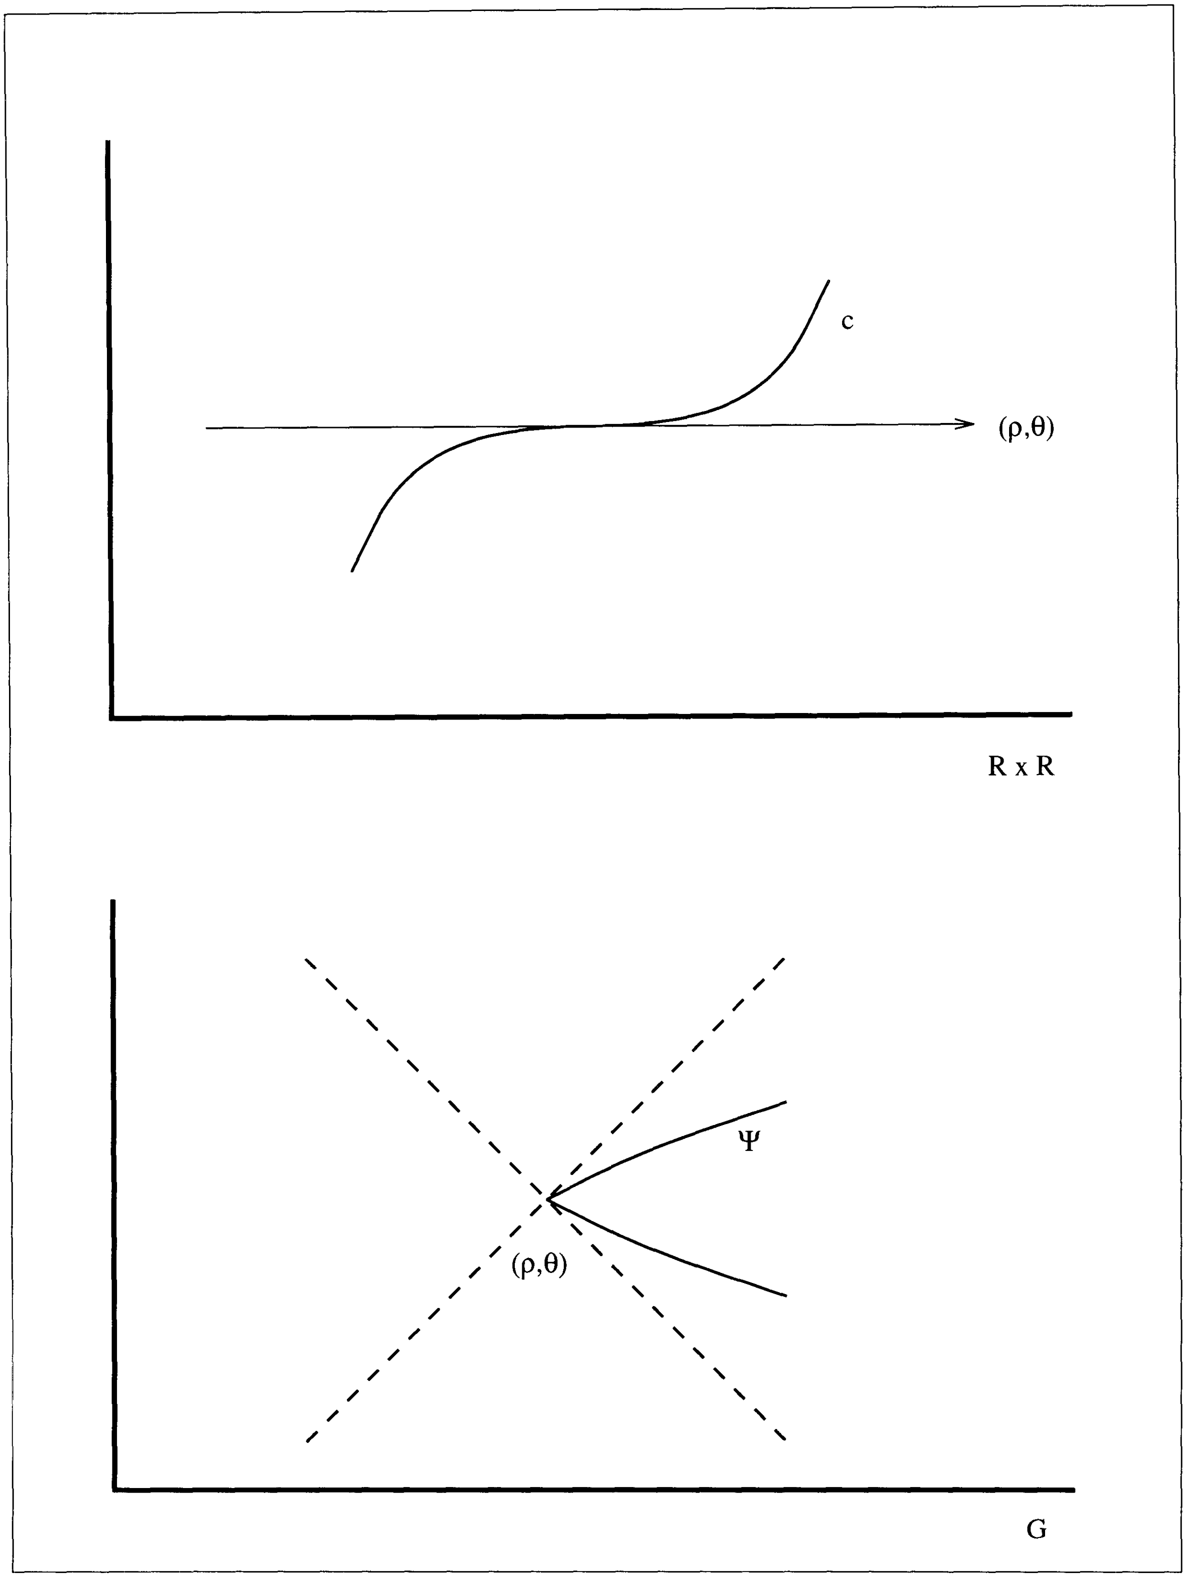

Figure 4-4: Type A4

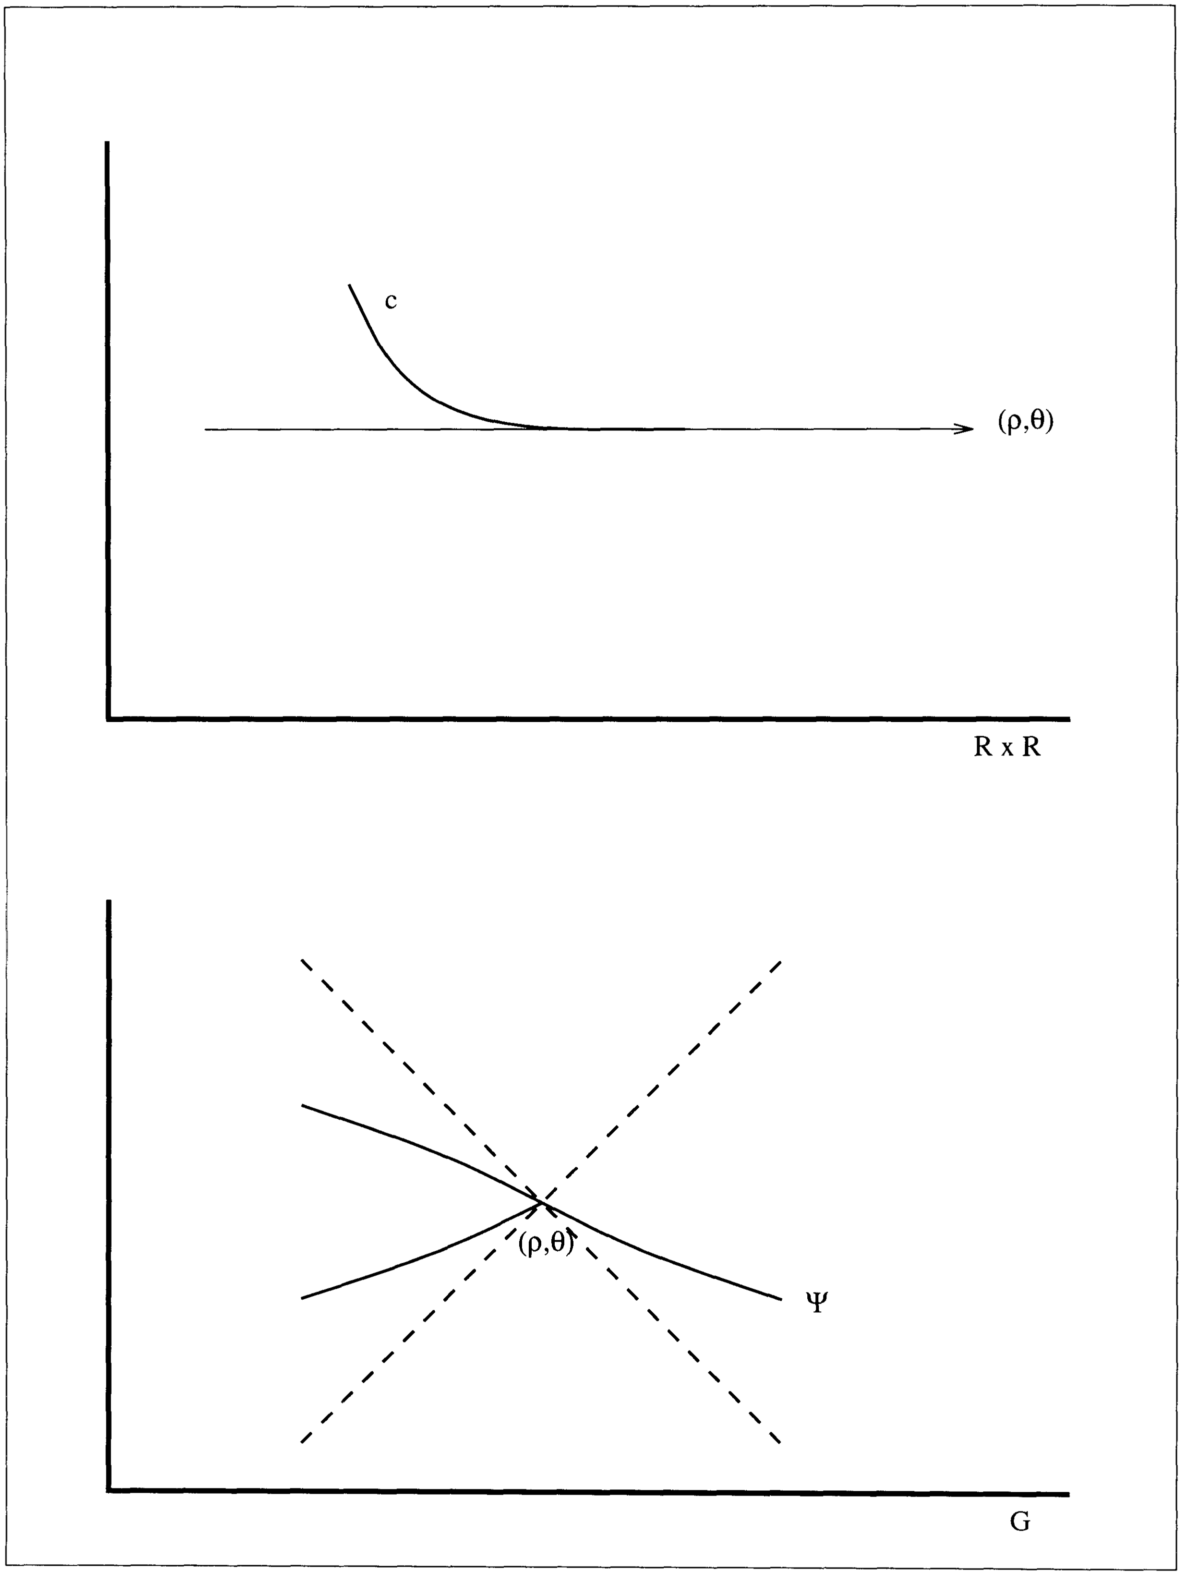

Figure 4-5: Type B1

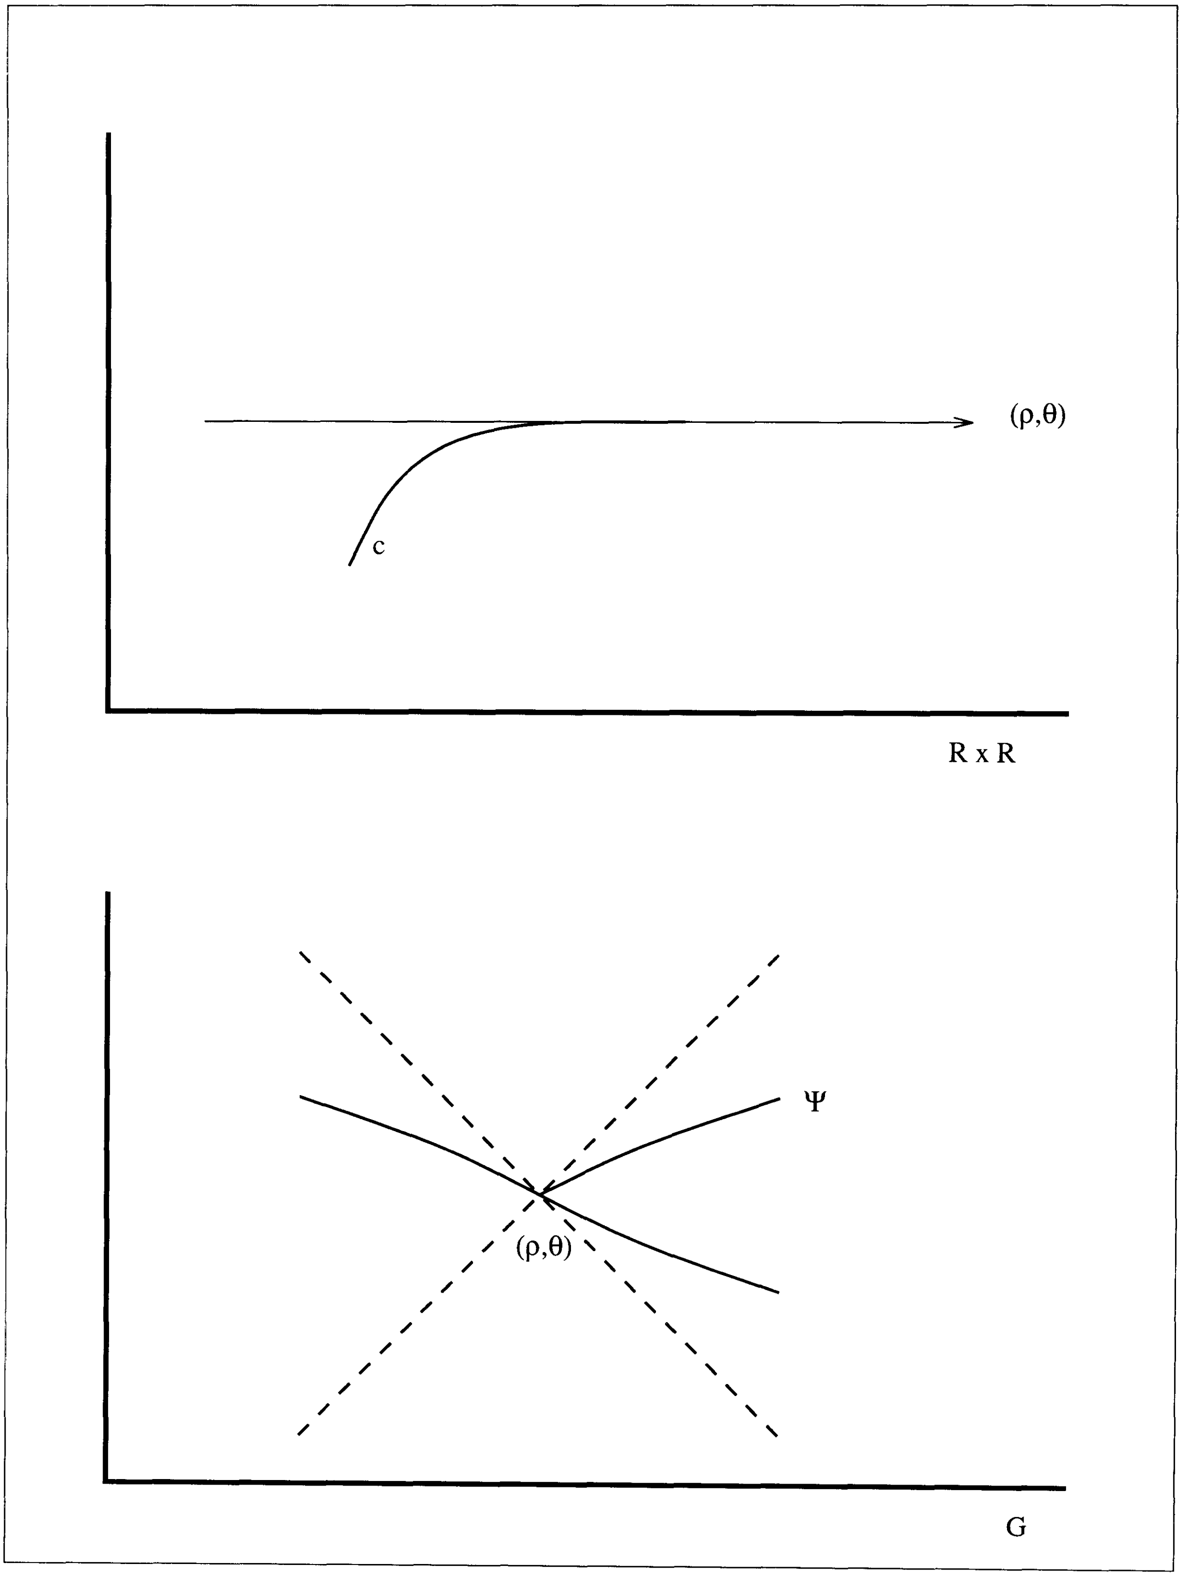

Figure 4-6: Type B2

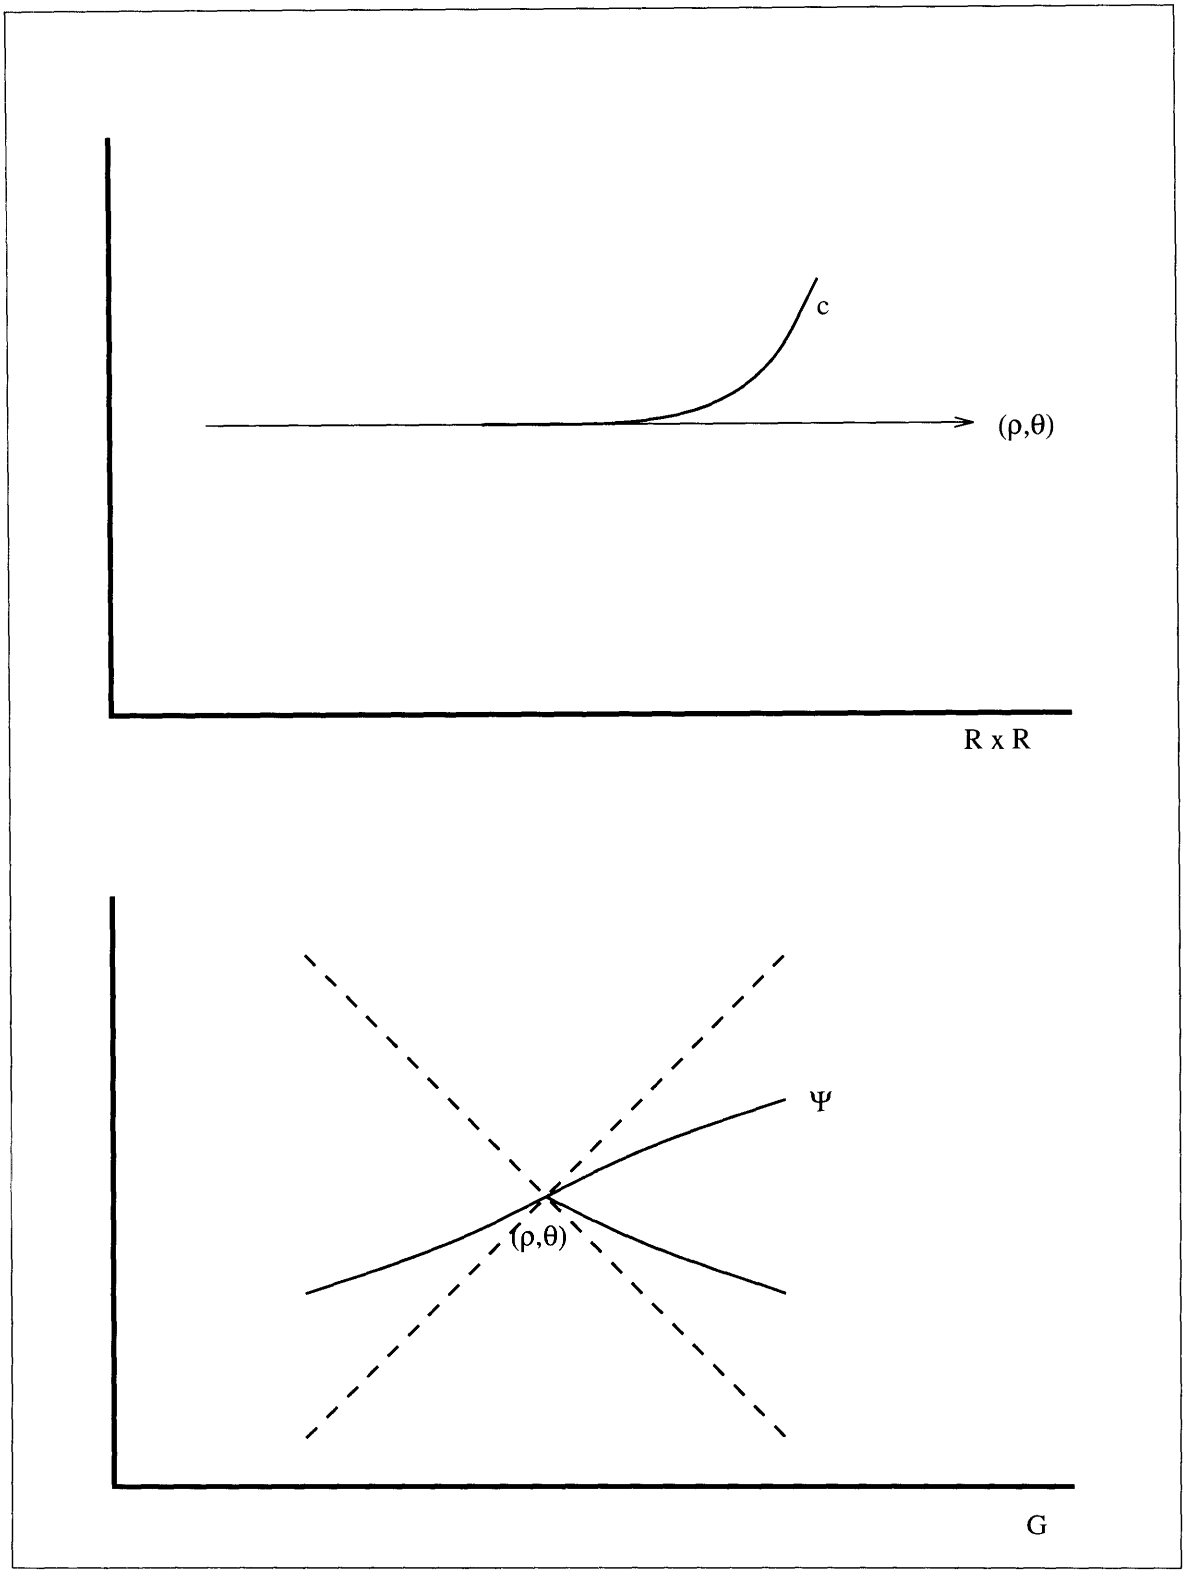

Figure 4-7: B3

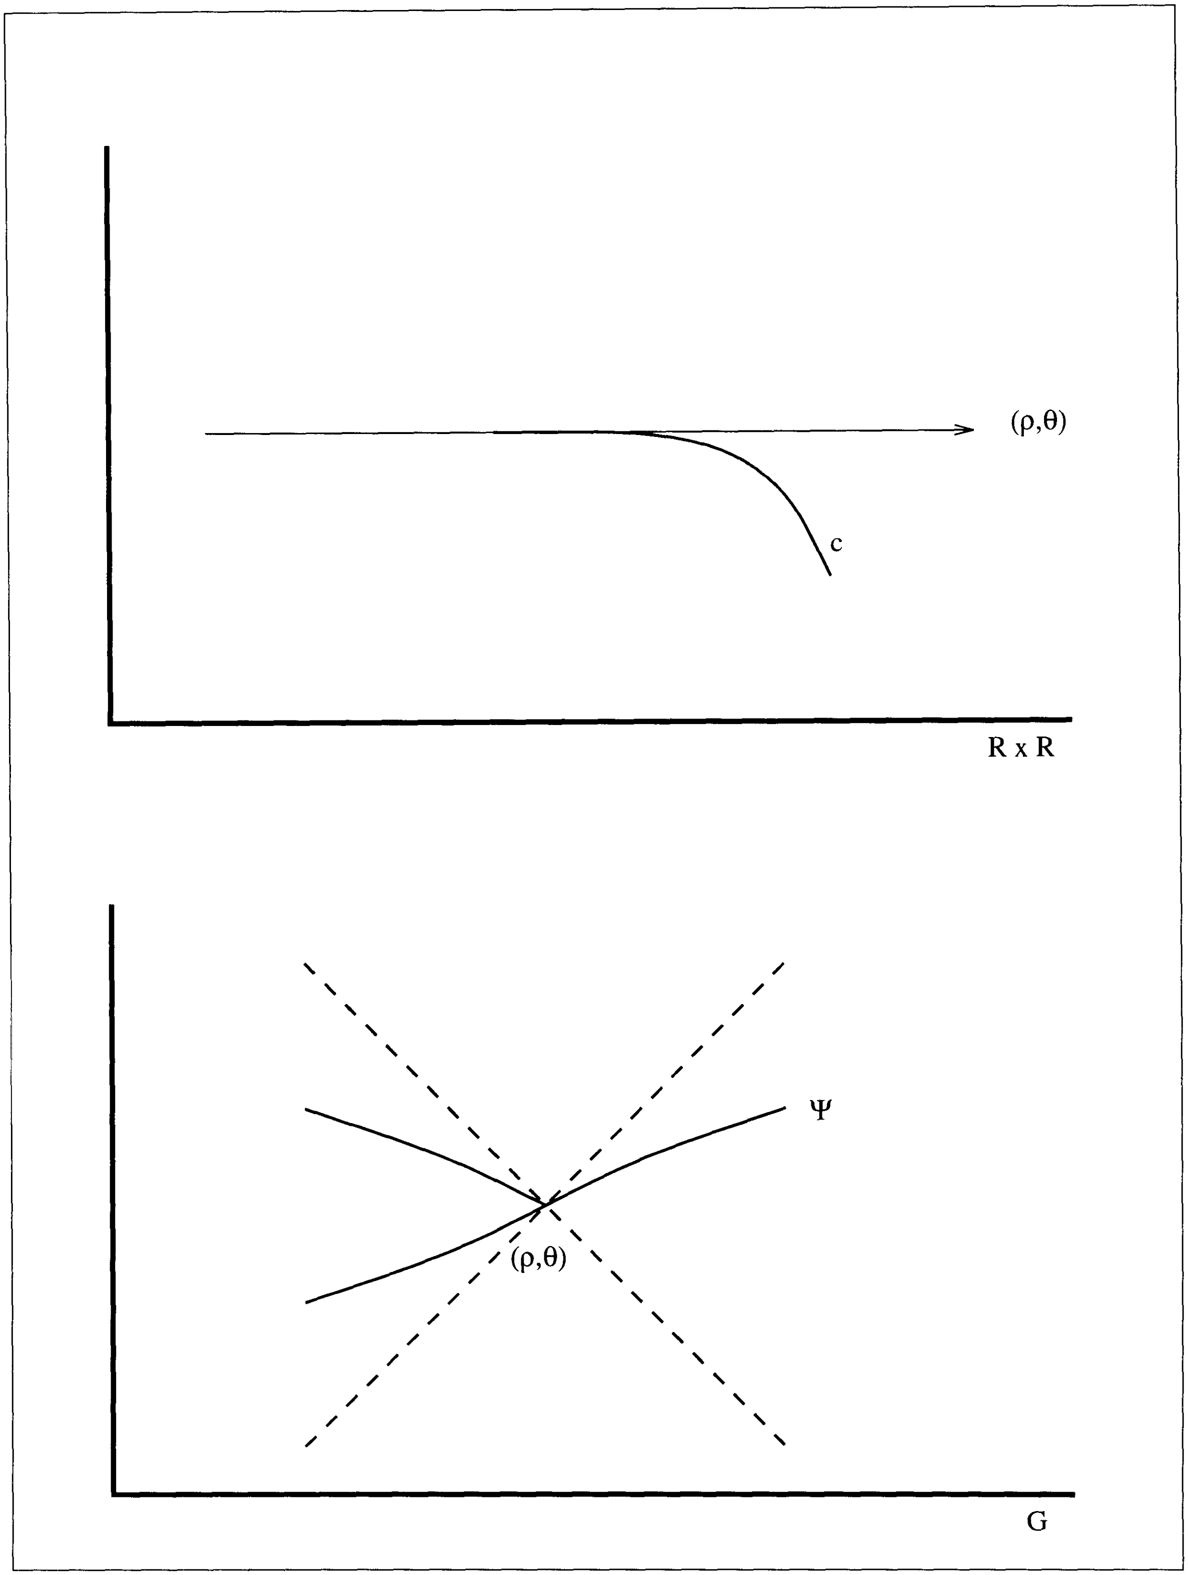

Figure 4-8: Type B4

$$+T_c(\tau_c(s_0-)+) = -T_c(\tau_c(s_0-)-)$$

$$+T_c(\tau_c(s_0+)+) = -T_c(\tau_c(s_0+)-)$$

C A point $x = c(t)$ is a *line segment with respect to* θ if the above conditions hold. See Figure 4-9.

Clearly we have exhausted all the possible curve structures in a neighborhood of $x = c(t)$ with respect to some interval $(t - \delta, t + \delta)$.

Becuase we only have a finite number inflection points we do not need to worry about the cases where there exist t_1, t_2 arbitrarily close to but less than $\tau_c(s_o-)$ such that $\langle c(t_1) - c(\tau_c(s_o-)), *\theta\rangle > 0$ and $\langle c(t_2) - c(\tau_c(s_o-)), *\theta\rangle < 0$. Similarly we could have t_1, t_2 arbitrarily close to but greater than $\tau_c(s_o+)$ such that $\langle c(t_1) - c(\tau_c(s_o+)), *\theta\rangle > 0$ and $\langle c(t_2) - c(\tau_c(s_o+)), *\theta\rangle < 0$. That is we do not have the problems of rapid oscillation.

Note for the above nine cases we were able to incorporate line segments. We are interested in what $\psi_c|_{(s_0-\epsilon, s_0+\epsilon)}$ looks like in $\mathcal{G}$. If we understand what the neighborhoods looks like for each of the above elementary curve arc-segments we can reconstruct more complicated phenomena by stacking them.

In order to study neighborhoods of $G \in \mathcal{G}$ we need some continuity and regularity results for $tr \pm \psi_C$ and S_{n_C}.

4.4 Continuity Results

Assume C is curvature minimal with a finite number of inflection points.

Theorem 12 *Let G be a regular point of $tr \pm \psi_C$ with a tangent line ξ. Then*

1. *$\mathcal{F}(\xi)$ is concave up with respect to θ if and only if $\Delta(G) > 0$.*

2. *$\mathcal{F}(\xi)$ is concave down with respect to θ if and only if $\Delta(G) < 0$.*

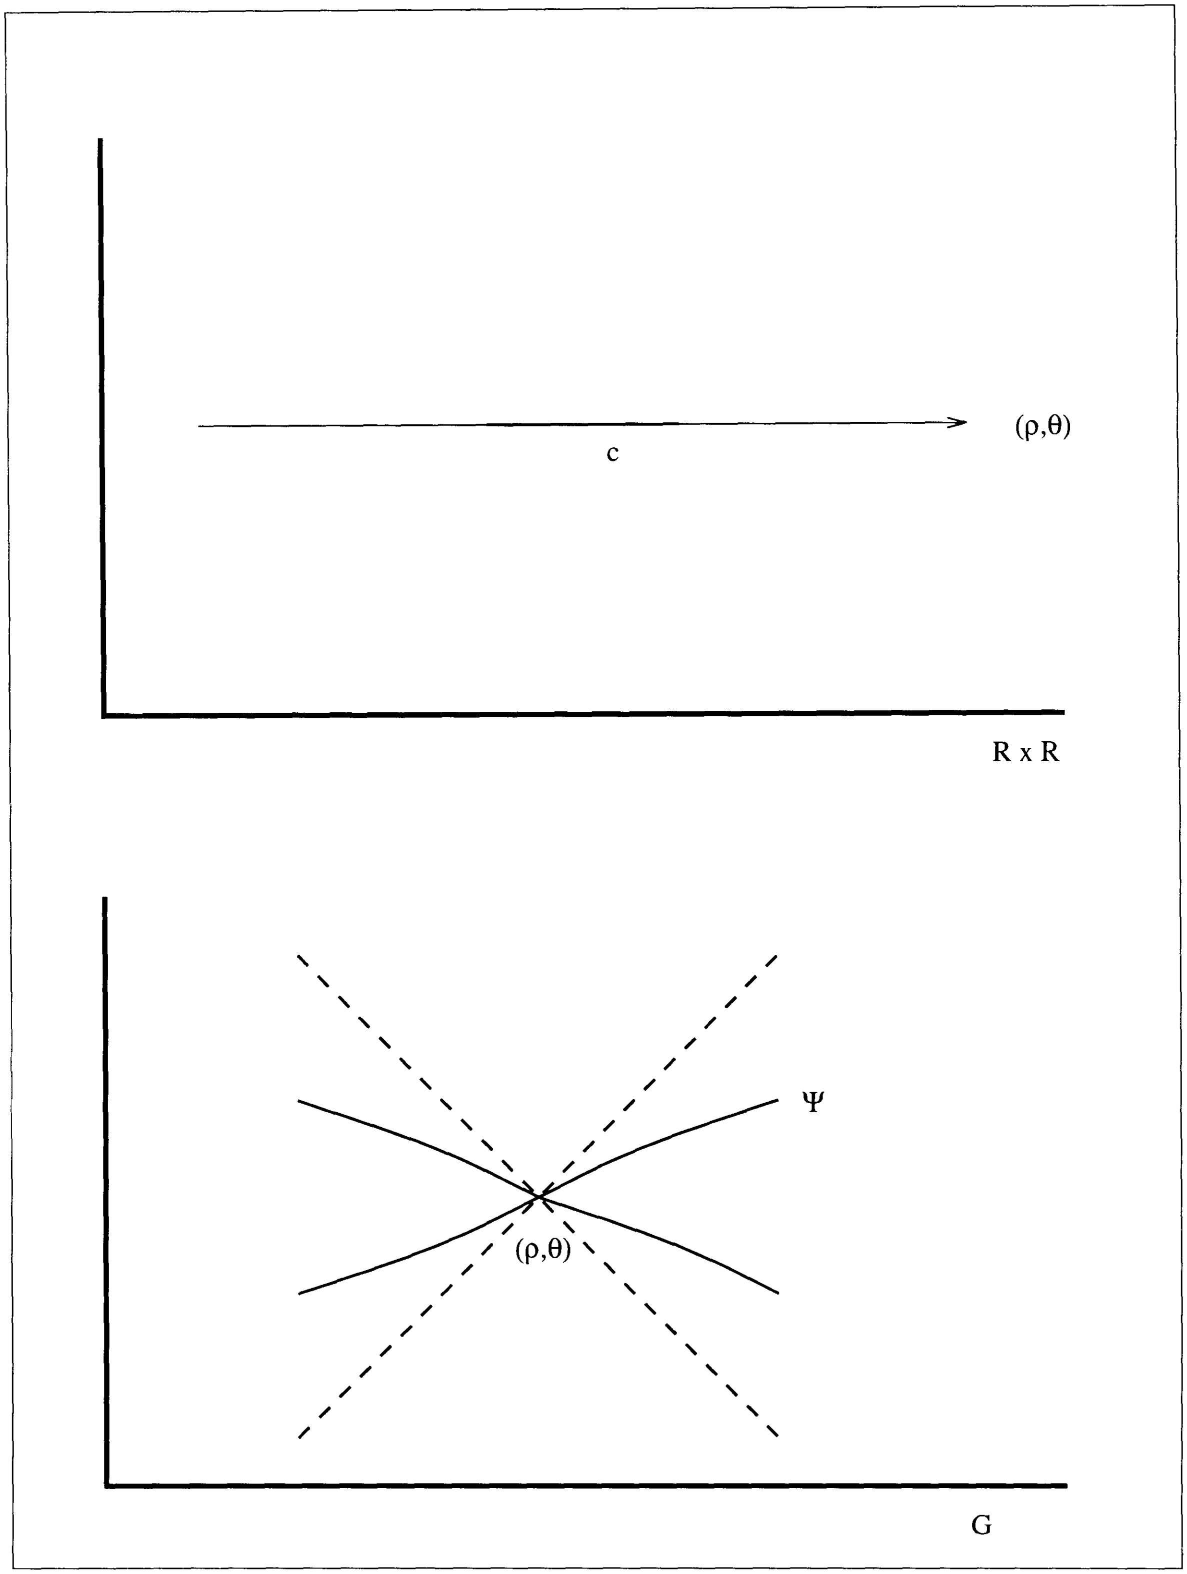

Figure 4-9: Type C

Proof: We will prove statement one. The other follows analogously. ($\Rightarrow$) Let $\mathcal{F}(\xi) = x = c(t)$ be concave up with respect to θ. Then there exist δ_1, δ_2 such that for all $t' \in (t - \delta_1, t)$ and $t'' \in (t, t + \delta)$ we have $\langle c(t') - c(t)), *\theta \rangle > 0$ and $\langle c(t'') - c(t), *\theta \rangle > 0$. By continuity of c there must exist some $\delta > 0$ such that for all $\gamma \in (t - \delta, t + \delta) \backslash \{t\}$ we have $\langle c(\gamma) - c(t), *\theta \rangle > 0$. Therefore for ϵ small enough we can find a $t_1 < t$ and a $t_2 > t$ such that $c^{-1}(G_\epsilon) = \{t_1, t_2\}$. That is $\Delta(G) > 0$.

($\Leftarrow$) $n_C(G_+) - n_C(G_-) > 0$ implies there exists an $\epsilon > 0$ such that $n_C(G_\epsilon) - n_C(G_{-\epsilon}) > 0$. Therefore there exists $\{t', t''\} = c^{-1}(G_\epsilon)$ where $t' < t$ and $t'' > t$. Which implies $\langle c(t') - c(t), *\theta \rangle > 0$ and $\langle c(t'') - c(t), *\theta \rangle > 0$ □

Theorem 13 *If G is a regular point of $tr \pm \psi_C$ with tangent line then there exists a neighborhood $B_\epsilon(G)$ such that*

1. *all points $G' \in B_\epsilon(G) \cap tr \pm \psi_C$ are regular with tangent lines.*

2. *$\Delta(G') = \Delta(G)$ for all $G' \in B_\epsilon(G) \cap tr \pm \psi_C$*

Proof: The first statement is just a result of lemma 18 and lemma 19. Assume a fixed $c \in C$. For r and ϵ_1 small enough we see by lemma 6 that $tr \pm \psi_C \cap (B_{r+\epsilon_1}(G) \backslash B_{r-\epsilon_1}(G))$ is a union of disjoint arc-segments of $tr \pm \psi_C$. Since $tr \pm \psi_C$ has a countable number of singularities we see that each arc-segment must be regular and have a tangent line over some ϵ_2 ball. We can relate this ϵ_2 to a neighborhood of the parameterization variable s. Specifically for each $c \in C$ and each $s_0 \in \psi_c^{-1}(G)$ there exists an ϵ_3 such that $tr\psi_c|_{(s_0-\epsilon_3, s_0+\epsilon_3)}$ is regular and has tangent lines.

To prove the second item we show there exists an $\epsilon < \epsilon_3$ such that for all $s \in (s_o - \epsilon, s_0 + \epsilon)$ we have $n_c(\psi_c(s)_+) - n_c(\psi_c(s)_-) = \Delta(\psi_c(s_0))$ for each $c \in C$.

Without loss of generality assume a given $c \in C$ and assume $\Delta(\psi_c(s_0)) > 0$ then we have by theorem 10 $c(\tau_c(s_0))$ is concave up with respect to θ. Therefore there exists a $\delta > 0$ such that for $t \in (\tau_c(s_0) - \delta, \tau_c(s_0) + \delta) \backslash \{\tau_c(s_0)\}$ we have $\langle c(t) - c(\tau_c(s_0)), *\theta \rangle > 0$

We will show the existence of ϵ such that for all $s \in (s_o - \epsilon, s_0)$ we have $n_c(\psi_c(s)_+) - n_c(\psi_c(s)_-) = \Delta(\psi_c(s_0))$ The other side, $(s_0, s_0 + \epsilon)$, follow analogously.

Since $c(\tau_c(s_0))$ is not an inflection point nor on a line segment we see there are two cases:

I) If $T_c(\tau_c(s_0)-) \neq T_c(\tau_c(s_0)+)$ then we have a corner at $c(\tau_c(s_0))$. Choose ϵ such that $T_c(\tau_c(s_0)-) < \theta_c(s_0-\epsilon) < \theta_c(s_0)$. It is clear that for $t \in (\tau_c(s_0), \tau_c(s_0)+\delta)$ we have $\langle c(t) - c(\tau_c(s_0)), *\theta_c(s_0 - \epsilon)\rangle > 0$. Now $T_c(\tau_c(s_0)-) = \lim_{t_i \uparrow \tau_c(s_0)} \frac{c(t_i)-c(\tau_c(s_0))}{t_i - \tau_c(s_0)}$. We also have $\langle T_c(\tau_c(s_0)-), *\theta_c(s_0 - \epsilon)\rangle < 0$. Thus for t_i sufficiently close to $\tau_c(s_0)$ we have by continuity that $\langle \frac{c(t_i)-c(\tau_c(s_0))}{t_i-\tau_c(s_0)}, *\theta(s_0 - \epsilon)\rangle < 0$. Which implies $\langle c(t_i) - c(\tau_c(s_0)), *\theta(s_0 - \epsilon)\rangle > 0$.

II) If $T_c(\tau_c(s_0)-) = T_c(\tau_c(s_0)+)$ then choose ϵ such that $\tau_c(s_0) - \delta < \tau_c(s_0 - \epsilon)$. Let $D_1 = (\tau_c(s_0) - \delta, \tau_c(s_0 - \epsilon))$ and $D_2 = (\tau_c(s_0 - \epsilon), \tau_c(s_0))$. Note by continuity and the fact that there are no inflection points over $(\tau_c(s_0) - \delta, \tau_c(s_0))$, the function $\theta_c(s)$ is monotonically increasing over $(\tau_c(s_0) - \delta, \tau_c(s_0))$. Assume towards a contradiction that there exists a $t_1 \in D_1$ and a $t_2 \in D_2$ such that $\langle c(t_1) - c(\tau_c(s_0 - \epsilon)), *\theta(s_0 - \epsilon)\rangle < 0$ and $\langle c(t_2) - c(\tau_c(s_0 - \epsilon)), *\theta(s_0 - \epsilon)\rangle < 0$. By continuity this would imply that there were $t_1' \in D_1$ and $t_2' \in D_2$ such that $\theta_c(t_1') < \theta_c(s_0 - \epsilon)$ and $\theta_c(t_2') > \theta_c(s_0 - \epsilon)$ which contradicts the fact that θ_c is monotonically increasing. □

By theorem 13 Δ changes sign or magnitude only when we reach a singularity. Thus if $tr\psi_c|_{(a,b)}$ is regular and has tangent lines then Δ is constant over it.

4.5 Types of Indicatrix Behavior

We are now ready to describe the associated indicatrix behavior to each of the nine curve behaviors.

Let $c \in C$ such that $x = c(t)$ for some $t \in [0, \mathcal{L}(c))$. Let $\theta \in \tilde{T}_c(t))$. By lemma 3 there exists a uniquely determined $s_0 \in R^1$ such that $t \in [\tau_c(s_0-), \tau_c(s_0+)]$ and $\psi_c(s_0) = (\langle c(t), *\theta\rangle), \theta)$.

Let $D_1 = (s_0 - \epsilon, s_0)$ and $D_2 = (s_0, s_0 + \epsilon)$.

If there exists a δ_1 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_0-)), *\theta\rangle > 0$ then $\theta_c(s)$ increases over D_1. Similarly if there exists a δ_1 such that for all $t \in (\tau_c(s_0-) - \delta_1, \tau_c(s_0-))$ we have $\langle c(t) - c(\tau_c(s_0-)), *\theta\rangle < 0$ then $\theta_c(s)$ decreases over D_1.

If there exists a δ_2 such that for all $t \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t) -$

$c(\tau_c(s_0+)), *\theta\rangle > 0$ then $\theta_c(s)$ increases over D_2. Similarly if there exists a δ_2 such that for all $t \in (\tau_c(s_0+), \tau_c(s_0+) + \delta_2)$ we have $\langle c(t) - c(\tau_c(s_0+)), *\theta\rangle < 0$ then $\theta_c(s)$ decreases over D_2.

A1 Point $x = c(t)$ is concave up with respect to θ. Thus we see that $\theta_c(s)$ is increasing over D_1 and D_2. Now by theorem 13 we see that $tr\psi_c|_{D_1}$ and $tr\psi_c|_{D_2}$ are regular and have tangent lines so Δ is constant over them. Since x is not an inflection point we see that Δ is constant over $(s_0 - \epsilon, s_0 + \epsilon)$ and since x is concave up we see by theorem 12 that $\Delta = 2$. See Figure 4-1.

A2 Point $x = c(t)$ is concave down with respect to θ. Thus we see that $\theta_c(s)$ is decreasing over D_1 and D_2. Now by theorem 13 we see that $tr\psi_c|_{D_1}$ and $tr\psi_c|_{D_2}$ are regular and have tangent lines so Δ is constant over them. Since x is not an inflection point we see that Δ is constant over $(s_0 - \epsilon, s_0 + \epsilon)$ and since x is concave down we see by theorem 12 that $\Delta = -2$. See Figure 4-2.

A3 Point $x = c(t)$ is up-down inflectional with respect to θ. Thus we see that $\theta_c(s)$ is increasing over D_1 and decreasing over D_2. Now by theorem 13 we see that $tr\psi_c|_{D_1}$ and $tr\psi_c|_{D_2}$ are regular and have tangent lines so Δ is constant over them. Since x is an inflection point we see that Δ changes at x. That is $\Delta = +2$ over D_1 and $\Delta = -2$ over D_2. It will always be the case that $tr\psi_c|_{D_1}$ will lie under $tr\psi_c|_{D_2}$. See Figure 4-3.

A4 Point $x = c(t)$ is down-up inflectional with respect to θ. Thus we see that $\theta_c(s)$ is decreasing over D_1 and increasing over D_2. Now by theorem 13 we see that $tr\psi_c|_{D_1}$ and $tr\psi_c|_{D_2}$ are regular and have tangent lines so Δ is constant over them. Since x is an inflection point we see that Δ changes at x. That is $\Delta = -2$ over D_1 and $\Delta = +2$ over D_2. It will always be the case that $tr\psi_c|_{D_1}$ will lie over $tr\psi_c|_{D_2}$. See Figure 4-4.

Note that when $T_c(t+) = -T_c(t-)$ we are at an endpoint. Thus the tangent T_c has two choices on which direction to turn. Either clockwise or counter-clockwise at $c(t)$. Note that clockwise behavior implies $\theta_c(s)$ is decreasing over the turn and counter-clockwise behavior implies that $\theta_c(s)$ is increasing over the turn.

Because c is closed and we have an endpoint in a neighborhood of t, c will trace the image of trc near t twice.

B1 Point $x = c(t)$ is concave up with right endpoint with respect to θ. Thus we see that θ_c is increasing over D_1. $\psi_c(s)$ can move in one of two directions over D_2. Thus we see that $\Delta = 4$ over D_1 and $\Delta = \pm 2$ over D_2 depending on whether we turn counter-clockwise or clockwise. See Figure 4-5.

B2 Point $x = c(t)$ is concave down with right endpoint with respect to θ. Thus we see that θ_c is decreasing over D_1. $\psi_c(s)$ can move in one of two directions over D_2. Thus we see that $\Delta = -4$ over D_1 and $\Delta = \pm 2$ over D_2 depending on whether we turn counter-clockwise or clockwise. See Figure 4-6.

B3 Point $x = c(t)$ is concave up with left endpoint with respect to θ. Thus we see that θ_c is increasing over D_2. $\psi_c(s)$ can move in one of two directions over D_1. Thus we see that $\Delta = 4$ over D_2 and $\Delta = \pm 2$ over D_1 depending on whether we turn counter-clockwise or clockwise. See Figure 14-7.

B4 Point $x = c(t)$ is concave down with left endpoint with respect to θ. Thus we see that θ_c is decreasing over D_2. $\psi_c(s)$ can move in one of two directions over D_1. Thus we see that $\Delta = -4$ over D_2 and $\Delta = \pm 2$ over D_1 depending on whether we turn counter-clockwise or clockwise. See Figure 4-8.

C Point $x = c(t)$ is a line segment with respect to θ. In this case we have a choice of turning clockwise or counter-clockwise at either endpoint. So we have four paths leaving G with Δ heights $= \pm 2$. See Figure 4-9.

Note that each indicatrix behavior can be determined uniquely from the others.

Theorem 14 *The elementary curve behaviors are uniquely determined by their elementary indicatrix behavior.*

In summary if $G \in tr \pm \psi_C$ and is not a cusp then Δ will have the same value and sign near G. If $G \in tr \pm \psi_c$ and is a cusp then Δ will have the same magnitude but different signs on the incoming and outgoing arc-segments. Specifically the arc-segment that is above the other will always have a negative Δ and the bottom one will always a positive Δ.

4.6 Matching Algorithm and Uniqueness

We are now ready to give a matching algorithm. The algorithm always constructs the shortest line segment that fits the data given so far. It always tries to match a given angle α with the one closest to it in magnitude. Remember line segments are determined by the difference between the incoming and outgoing α angles.

Let $g\_Tan(G) = \{\alpha_1, ..., \alpha_n\}$.

Algorithm:

for $i = 1$ to n

    for $j = i + 1$ to n

        if sgn $\langle \alpha_j, *\theta \rangle =$ sgn $\langle \alpha_i, *\theta \rangle$

            if $\langle \alpha_j, *\theta \rangle < 0$ (shows α_j, α_i are to the left of θ)

                if $\Delta(\alpha_j) < 0$ and $\Delta(\alpha_i) > 0$ then match α_j to α_i

                else next j

            if $\langle \alpha_j, *\theta \rangle > 0$ (shows α_j, α_i are to the right of θ)

                if $\Delta(\alpha_j) > 0$ and $\Delta(\alpha_i) < 0$ then match α_j to α_i

                else next j

        else if sgn $\langle \alpha_j, *\theta \rangle) \neq$ sgn $\langle \alpha_i, *\theta \rangle$

                if $\Delta(\alpha_j) = \Delta(\alpha_i)$ then match α_j to α_i

                else next j

next i

This is a legal matching because we are following the rules of reconstruction set in the previous section.

Theorem 15 *Let C have finite inflection. Given any behavior of $tr \pm \psi_C$ around a* *point G the above algorithm will find the unique behavior of trC determined by said behavior around G.*

Chapter 5

Reconstruction Algorithms

In this chapter we discuss some requirements for efficient reconstruction algorithms. Specifically we are given partial information about the function n_A and from this we are to reconstruct A. First we discuss what it means to be a good approximation. Then we discuss singularity detection. Next we describe a linear interpolation algorithm. Then we take up the issue of noise.

5.1 Approximation

To determine if a reconstruction is good we need a measure of accuracy. For curves people often use the following Hausdorff metric. Let c_1, c_2 be curves in R^2. Then $d_H(trc_1, trc_2)$:=inf$\{trc_1 \subset c_2^\eta$ and $trc_2 \subset c_1^\eta\}$ where c^η is the η neighborhood of trc. Specifically $c^\eta := \{x : \inf_{y \in trc}|x - y| < \eta\}$. $d_H(trc_1, trc_2) < \epsilon$ means that both curves can fit in an ϵ-wide tube of the other curve.

Unfortunately there are some problems with this measure. Specifically there is no constraint on the smoothness of the reconstructed curve. Some pathological problems that can occur are: the reconstructed curve's length and total absolute curvature may not be finite, and even if they are they may not tend smoothly to the real length and curvature as ϵ goes to zero.

The above definition was given for curves but it can be easily generalized to $\mathcal{K}$-sets. Specifically if A is an exact K-set with curvature minimal representation C and

$\tilde{A}$ is a reconstruction with representation $\tilde{C}$ we say

$d_H(A, \tilde{A}) := \{ \inf \eta : A \subset \tilde{A}^\eta \text{ and } \tilde{A} \subset A^\eta\}$ where $A^\eta := \{x : \inf_{y\in A}|x - y| < \eta\}$.

If for each $c \in C$ there is a corresponding $\tilde{c} \in \tilde{C}$ then we see that $d_H(A, \tilde{A}) < \max_{c\in C} d_H(trc, tr\tilde{c})$.

In section 5.3 we will show how we can achieve a good approximation of trC by approximating $\Delta(G)$.

I would like to contrast the d_H measure with that of $d_{\mathcal{I}}$ and $\mathcal{H}^1$ used by Richardson. Specifically Richardson showed that if A_1 and A_2 are $\mathcal{K}$-sets then $d_{\mathcal{I}}(A_1, A_2) = 0$ if and only if $\mathcal{H}^1(|A_1 - A_2|) = 0$ (Theorem 1). Both of these measures are good for dealing with outliers and multiplicities. Unfortunately they do not work well in the d_H sense. $\mathcal{H}^1$ takes into account the set difference between the two curves but does not take into account the proximity of one curve to another curve. It is possible for $\mathcal{H}^1(c_1, c_2) = \mathcal{L}(c_1) + \mathcal{L}(c_2)$ but $d_H(c_1, c_2) < \epsilon$ for arbitrarily small ϵ.

We will use d_H as our fitness measure. We will attempt to find at least one curve in the ϵ-ball around c. A better formulation of this problem would use a calculus of variations approach to select one of the curves in the ϵ-ball. Richardson has suggested that minimum total absolute curvature be the criterion for selection over this ϵ-ball.

The algorithm works by operating on the sphere bundle of $\mathcal{G}$. Specifically we are given some discretization of n_A. From this we reconstruct the jump set of n_A. From which we reconstruct $g\_tr \pm \psi_A$ and the singularity set. Then by determining the tangents to the set $g\_tr \pm \psi_A$ we can use the function $\mathcal{F}$ and the singularity matching algorithm of the last chapter to reconstruct A.

5.2 Approximation in $\mathcal{G}$

Let A be an exact $\mathcal{K}$-set with curvature minimal representation C.

We want a condition on our approximation of $g\_tr \pm \psi_C$ that will insure a good approximation of trC.

We know that a dense subset of the curved parts of trC can be recovered from $g\_tr \pm \psi_C$ by using the function $\mathcal{F}(\xi) = \rho * \theta - \frac{d\rho(\xi)}{d\theta(\xi)}\theta$ where ξ is tangent to $g\_tr \pm \psi_C$ at

$G = (\rho, \theta)$. Note that this reconstruction depends on both $g\_tr \pm \psi_C$ and its tangents. Therefore we need to approximate both $g\_tr \pm \psi_C$ and its tangents in order to get a good approximation of c.

Let $c \in C$ and let ψ_c be its respective extended indicatrix. We know $tr\psi_c$ will have a countable number of singularities. Let this set of singularities be S. Segment ψ_c into pieces ψ_c^i where the set of endpoints is S.

Define $\Xi^i(G) := \frac{d\rho(\xi(G))}{d\theta(\xi(G))}$, whenever $tr\psi_C$ has a tangent ξ at G. We are interested in the slope.

Define $\vec{\psi_c^i}(G) := (\psi_c^i(G), \Xi^i(G))$. This is like a curve in the sphere bundle except that we are interested in the slope and not the direction of the tangent.

Let $\vec{G} = (\rho, \theta, \Xi)$.

Then let $\vec{\psi_c} = \cup_i \vec{\psi_c^i}$. Note this a curve in $\vec{\mathcal{G}} := R^1 \times S^1 \times [-R, R]$ (Where R is the radius of the boundary of our image space.) Note that this curve is discontinuous at every singularity). We define a metric on $\vec{\mathcal{G}}$: $d_{\vec{\mathcal{G}}}(\vec{G_1}, \vec{G_2}) := ((\rho_1 - \rho_2)^2 + |\angle(\theta_1, \theta_2)|^2 + (\Xi_1 - \Xi_2)^2)^{\frac{1}{2}}$.

We say two curves $\vec{\psi}$ and $\vec{\tilde{\psi}}$ are *similar* if the singularity sets of ψ and $\tilde{\psi}$, $S, \tilde{S}$ are similar. That is if $s \in S$ then there exists and $\tilde{s} \in \tilde{S}$ such that $d_{\mathcal{G}}(s, \tilde{s}) < \epsilon$ and vice-versa.

Theorem 16 *If $\vec{\psi_1}$ and $\vec{\psi_2}$ are similar, and $d_{H_{\vec{\mathcal{G}}}}(\vec{\psi_{c1}}, \vec{\psi_{c2}}) < \epsilon$ then $d_{H_{R^2}}(trc_1, trc_2) < 2\epsilon(R+1)$ where R is the radius of the boundary circle.*

The sub-subscripts of the d_H measure are there to remind us of the spaces they are operating on.

Proof: We will show for every $x \in trc_1$ there exists a $y \in trc_2$ such that $|x - y| < 2\epsilon(R+1)$. There are two cases.

Case 1: x is a point of c_1 with nonzero curvature, with tangent G_1 where G_1 is not a singular point of $tr\psi_{c_1}$

Case 2: x is a point of c_1 not accounted for by case 1. Specifically x is a point of zero curvature, an endpoint, or a multi-tangency point. The point G_1 that represents it is a singular point of $tr\psi_{c_1}$.

For case 1: Since G_1 is regular it has a tangent slope Ξ_1. Since $d_H(\vec{\psi_{c_1}}, \vec{\psi_{c_2}}) < \epsilon$ there exists a $G_2 \in \psi_{c_2}$ such that G_2 is regular with tangent slope Ξ_2 and $d_{\mathcal{G}}(\vec{G_1}, \vec{G_2}) < \epsilon$.

Letting $\vec{G_1} = (\rho_1, \theta_1, \Xi_1)$ and $\vec{G_2} = (\rho_2, \theta_2, \Xi_2)$ we can define a finer partition of the distance between them. Specifically we know $d_{\mathcal{G}}(\vec{G_1}, \vec{G_2}) < \epsilon$ let $|\rho_1 - \rho_2| < \epsilon_\rho, |\angle(\theta_1, \theta_2)| < \epsilon_\theta$, and $|\Xi_1 - \Xi_2| < \epsilon_\Xi$ where $\epsilon_\rho^2 + \epsilon_\theta^2 + \epsilon_\Xi^2 < \epsilon^2$

Let $x = \mathcal{F}(\vec{G_1}) = \rho_1 * \theta_1 - \Xi_1\theta_1$ and $y = \mathcal{F}(\vec{G_2}) = \rho_2 * \theta_2 - \Xi_2\theta_2$ where we take the θ's to be vectors. Thus $|x - y| = |\rho_1 * \theta_1 - \Xi_1\theta_1 - (\rho_2 * \theta_2 - \Xi_2\theta_2)| \leq |\rho_1 * \theta_1 - \rho_2 * \theta_2| + |\Xi_1\theta_1 - \Xi_2\theta_2| \leq 2R\epsilon_\theta + \epsilon_\rho + \epsilon_\Xi + \epsilon_\theta\epsilon_\rho + \epsilon_\theta\epsilon_\Xi$. The θ terms are bounded by ± 1 and both the ρ and the Ξ terms are bounded by $\pm R$. If we ignore higher order terms we get $|x - y| < 2\epsilon(R + 1)$.

For case 2: There are three subcases depending on whether, x, is a point of inflection, of multi-tangency, or on a line. The first two are irrelevant since they have measure zero and can be reconstructed by interpolation. The last case, though, takes special consideration.

Let G be the line that the line segment containing x is on and let a_x and b_x be the endpoints of the line segment. For some $t \in (0, 1)$ we have $x = ta_x + (1-t)b_x$. We know from case 1 that there exist points points $a_y, b_y \in c_2$ such that $|a_x - a_y| < 2\epsilon(R + 1)$ and $|b_x - b_y| < 2\epsilon(R + 1)$. We need to find a point on the line connecting a_y to b_y that is $2\epsilon(R + 1)$ close to x. Let $y = ta_y + (1 - t)b_y$. Then $|x - y| = |ta_x + (1 - t)b_x - ta_y - (1 - t)b_y| \leq t|a_x - a_y| + (1 - t)|b_x - b_y| \leq 2\epsilon(R + 1)$. □

We notice the error $2R\epsilon_\theta + \epsilon_\rho + \epsilon_\Xi$ depends heavily on ϵ_θ. This is because small changes in the angle θ can have large effects on the radial components ρ and Ξ.

One way to get a better bound is to reduce the ϵ_θ error. We impose a parameterization on the two curves $\vec{\psi_{c_1}}$ and $\vec{\psi_{c_2}}$ that allow for zero error in θ. Specifically we will parameterize the curve $\vec{\psi_c}$ by its θ_c component. That is we match-up the graphs of $\vec{\psi_{c_1}}$ and $\vec{\psi_{c_2}}$ with respect to the θ_c component.

We can parameterize $\vec{\psi_c}$ in a natural way. Remember from theorem 3 that θ_c is arc-length parameterized by the function $s(t)$. Where $s(t)$ is the total absolute curvature of the curve $c|_{[0,t]}$. This says that if $\frac{d\theta}{ds}$ exists at s_0 then there exists a neighborhood around s_0 such that $\theta_c(s_1) - \theta_c(s_2) = \pm(s_1 - s_2)$ depending on whether $\frac{d\theta}{ds} = \pm 1$ for all s_1, s_2 in neighborhood of s_0.

To parameterize $\vec{\psi}_c$ pick some point $G \in tr\psi_c$. Let this be our starting point. It does not matter where we start because all the ψ_c curves are closed. Let $\theta_c(G) = \theta_c(0)$. Do the same with $\rho_c(s)$ and $\Xi_c(s)$. Note $s \in [0, \kappa(c))$.

Given a parameterization of $\vec{\psi}_{c_1}$ we can induce a parameterization on a similar curve $\vec{\psi}_{c_2}$. Given $\theta_{c_1}(0)$ find that point $G_2 = (\rho_2, \theta_2) \in tr\psi_{c_2}$ such that $\theta_{c_1}(0) = \theta_2$ and then call this G_2 the starting point for $\vec{\psi}_{c_2}$. Note this mutual parameterization is local. We cannot extend it to the whole curve $\vec{\psi}_{c_2}$ because the two curves c_1, c_2 may not have the same total absolute curvature.

The following theorem is local approximation theorem for mutually parameterized curves.

Theorem 17 $d_H(trc_1, trc_2) \leq sup_s\{|\rho_{c_1}(s) - \rho_{c_2}(s)| + |\Xi_{c_1}(s) - \Xi_{c_2}(s)|\}$.

Proof: Since we have eliminated ϵ_θ errors we see from theorem 14 that $d_H(trc_1, trc_2) \leq \epsilon_\rho + \epsilon_\Xi$. □

5.3 Singularity Detection

Any reconstruction algorithm we propose depends on accurate knowledge of $g\_tr \pm \psi_C$ and accurate knowledge of its singularities. To this end we need a way to detect singularities in our n_A data.

In the case where we know n_A completely and thus $g\_tr \pm \psi_C$ completely we can use the concept of density to determine the structure of the singularity.

Let $G = (\rho, \theta)$.

Let $S_\epsilon(G) = \{G' \in \mathcal{G} : |\theta - \theta'| < \epsilon$ and $|\rho - \rho'| < R|\theta - \theta'|\}$. All arc-segments of $tr \pm \psi_C$ that intersect G must go through this region. The area is $2\epsilon R$ where R is the radius of the boundary of the image set.

The number of arc-segments entering a point $G \in \mathcal{G}$ is just $\lim_{\epsilon \to 0^+} \frac{\mathcal{H}^1(S_\epsilon(G) \cap tr \pm \psi_C)}{2\epsilon R}$.

To get the appropriate multiplicity information just take $\lim_{\epsilon \to 0^+} \frac{\mathcal{H}^1(S_\epsilon(G) \cap g\_tr \pm \psi_C)}{2\epsilon R}$.

The number of arc-segments entering G from the left and right can be computed by taking one-sided density limits. That is take the limit over the right or left side of $S_\epsilon(G)$. Once we know the number of arc-segments entering and leaving we can start to piece the dual lattice information to construct the actual arc-segments and Δ height information.

In the case of partial information we have two choices. If we know the values of n_A over a fixed lattice then we can take a discrete density to determine the number of arc-segments entering G. For a given lattice we construct a dual lattice where the dual lattice values are just the differences between the n_A values. We know that all points on the dual lattice with nonzero differences must be close to a member on $tr \pm \psi_C$. Thus by taking the discrete density over these points we can estimate the number of arc-segments entering G.

The other approach is to use some sort of template matching. We have a set of canonical singularities. For each singularity in our set we have another set of deformable templates. These templates can be used in detection. This works well when the number of types of singularities is finite. We can achieve finiteness by assuming that there are no triple tangents in the image set.

5.4 Linear Interpolation

Assume we are able to detect singularities. Also assume we have n_A information on a rectangular lattice. We will give some conditions on the lattice spacing to ensure that $d_H(c_1, c_2) < \epsilon$. Let the lattice spacing distance be δ_θ and δ_ρ.

Given a lattice construct its dual lattice. We say a node in the dual lattice is "on" if its value is nonzero. We will use linear interpolation to connect the "on" nodes. Specifically if $G = (\rho, \theta)$ is an "on" node and does not represent an inflection point then there must be points $G' = (\rho \pm m\delta_\rho, \theta - \delta_\theta)$ and $G'' = (\rho \pm n\delta_\rho, \theta + \delta_\theta)$ that are "on", where m, n integers with absolute value less than $R\delta_\theta$. (We can put this bound on m, n because we know the slopes of ψ_C are bounded.) If G represents an inflection point then there must exist $G' = (\rho \pm m\delta_\rho, \theta - \delta_\theta)$ and $G'' = (\rho \pm n\delta_\rho, \theta - \delta_\theta)$

that are "on" or $G' = (\rho \pm m\delta_\rho, \theta + \delta_\theta)$ and $G'' = (\rho \pm n\delta_\rho, \theta + \delta_\theta)$ that are "on." Note that locally the linear interpolated curve and the original curve will be mutually parameterized.

Given some ψ let $\tilde{\psi}$ curve be our linear interpolation of it. Then by theorem 17 we have $d_H(\vec{\psi}, \vec{\tilde{\psi}}) \leq sup_s\{|\rho(s) - \tilde{\rho}(s)| + |\Xi(s) - \tilde{\Xi}(s)|\} < \delta_\rho + \frac{\delta_\rho}{\delta_\theta}$. If we want $d_H(c, \tilde{c}) < \epsilon$ we need $\delta_\rho + \frac{\delta_\rho}{\delta_\theta} < \epsilon$. One possible solution is $\delta_\rho = \frac{\epsilon^2}{2}$ and $\delta_\theta = \epsilon$ for $\epsilon < 1$. This makes intuitive sense. We need to know both ρ and the slope Ξ. Thus $\frac{\delta_\rho}{\delta_\theta} < \epsilon$. We need $O(\frac{1}{\epsilon^3})$ lattice points to get ϵ accuracy.

The problem with this construction is that our reconstructed $\tilde{\psi}$ will have many more singularities than ψ. This is because we have a corner at every "on" dual node. We know corners represent line segments in R^2. So our reconstructed curve will be very "jumpy." But as ϵ gets smaller the incoming and outgoing tangents at an "on" node will get closer and thus the length of the corresponding line segment in R^2 will get smaller.

5.5 Spline Interpolation

One way around the "jumpiness" problem in linear interpolation is to use some other kind of smooth interpolation. Splines are the obvious answers. They can be locally determined so they satisfy our local approximation theorem. It is not so clear, though, what order spline to use. We need to approximate ρ and its slope Ξ, so we need a spline that in some sense minimizes the magnitude of its slopes.

A good heuristic for setting δ_ρ and δ_θ is $\delta_\rho = \epsilon^2$ and $\delta_\theta = \epsilon$.

Note we should pick that spline algorithm that best suits our fidelity criteria. For example minimum absolute curvature.

Because splines are smooth we do not have the problem of false singularities.

Splines also allow us to do one or many step predictions. Given a curve c and a point $x = c(t)$ we know that we can approximate c around x by a circle of radius $\frac{1}{\kappa(t)}$. Circles in R^2 get mapped to sine waves in $\mathcal{G}$. Therefore given some partially reconstructed curve $\tilde{\psi}_c$ we can interpolate with respect to the underlying sine function.

We can also use this sort of sine interpolation to detect singularities. If a point is not where we have predicted it to be we have a candidate for a singularity.

5.6 Noise

Finally we come to the issue of noise. Noise is very domain specific. In the case of LIT there are two main domains. One is that of medical imaging and the other is that of computer vision. The former is continuous in some sense and the latter is discrete in some sense.

For medical imaging problems there is uncertainty in the information we are given about n_A. This noise cannot in general be modeled unless one has direct access to the equipment taking the measurements.

For computer vision problems we have the usual "discretization noise." The lines that we "drop" on our image will have a width associated with them. It is then possible that we may count more or less crossings of our "thick" line on our discretized image. We do not have sub-pixel accuracy.

These are issues that need to be studied in order to construct a robust LIT algorithm.

Chapter 6

Conclusions and Open Problems

It is clear that LIT has potential in both the fields of medical imaging and computer vision. It works well for detecting boundaries of non-occluded objects and transparent objects. In the general case it can reconstruct a large class of one-dimensional sets.

In this thesis we have given summary of Richardson's mathematical groundwork. We have described the line intersection function in complete detail. An approach to implementing an algorithm was given.

A more general approach at reconstructing the original set A might be to consider the problem as a regularization problem. We approximate the reconstructed curve by a smooth function that tapers to zero outside some compact set around A. In the limit these functions should uniformly approach A. Richardson has given some conditions for uniform convergence of curves.

We should try to formulate this reconstruction problem as a calculus of variations problem. This allows us to incorporate criteria, like smoothness or minimal total curvature, that will help us select an appropriate curve $\tilde{c}$ in the ϵ-ball around c.

There remains the problem of implementing an algorithm on "real - life" data. The issues of noise need to be looked at more closely.

Finally another open problem includes the generalization of this problem to higher dimensional problems. It is not clear if one-dimensional objects lying in R^n can be reconstructed from its projections.

Bibliography

A. D. Alexandrov and Yu. G. Reshetetnyak. General Theory of Irregular Curves, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1989.

K. J. Falconer. The Geometry of Fractal Sets, Cambridge University Press, Cambridge, 1985.

Sanjeev Kulkarni, Sanjoy Mitter, John Tsitsiklis, and Ofer Zeitouni. "PAC Learning with Generalized Samples and an Application to Stochastic Geometry," IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 15, Number 9, September 1993.

V. F. Leavers. Shape Detection in Computer Vision Using the Hough Transform, Springer Verlag, London, 1992.

Thomas J. Richardson. Total Curvature and Intersection Tomography, working paper, 1993.

Luis Santalo. Integral Geometry and Geometric Probability, Addison-Wesley Publish Company, Reading, Massachusetts, 1976.

Jean-Phillipe Thirion. "Segmentation of Tomographic Data Without Image Reconstruction," IEEE Transactions on Medical Imaging, Volume 11, Number 1, March 1992.