

**A PRE-ATTENTIVE VISION MODEL
FOR AUTOMATED DATA EXPLORATION
APPLIED TO SPACE PLASMA REMOTE SENSING DATA**

BY

IVAN ANDREEVICH GALKIN
B.S., Irkutsk State University (1983)
M.S., University of Massachusetts Lowell (1995)

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF DOCTOR OF SCIENCE
IN COMPUTER SCIENCE
UNIVERSITY OF MASSACHUSETTS LOWELL

Signature of Author: _____ Date: _____

Signatures of Dissertation Directors:

_____ Date: _____

_____ Date: _____

Signature of Dissertation Committee Member:

_____ Date: _____

**A PRE-ATTENTIVE VISION MODEL
FOR AUTOMATED DATA EXPLORATION
APPLIED TO SPACE PLASMA REMOTE SENSING DATA**

BY

IVAN ANDREEVICH GALKIN

ABSTRACT OF A DISSERTATION SUBMITTED TO THE FACULTY OF THE
DEPARTMENT OF SOMPUTER SCIENCE
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF DOCTOR OF SCIENCE
IN COMPUTER SCIENCE
UNIVERSITY OF MASSACHUSETTS LOWELL
2004

Dissertation Supervisors:

Bodo W. Reinisch, Ph.D.
Professor, Department of Environmental, Earth, and Atmospheric Sciences

Georges G. Grinstein, Ph.D.
Professor, Department of Computer Science

ABSTRACT

As research instruments of large information capacities become a reality, automated systems for intelligent data analysis become a necessity. Scientific archives containing huge volumes of data preclude manual manipulation or intervention and require automated exploration and mining that can at least pre-classify information in categories. The large dataset from the radio plasma imager (RPI) instrument onboard the IMAGE satellite shows a critical need for such exploration in order to identify and archive features of interest in the volumes of visual information. In this research we have developed such a pre-classifier through a model of pre-attentive vision capable of detecting and extracting traces of echoes from the RPI plasmagrams.

The overall design of our model complies with Marr's paradigm of vision where elements of increasing perceptual strength are built bottom up under the Gestalt constraints of good continuation and smoothness. The specifics of the RPI data, however, demanded extension of this paradigm to achieve greater robustness for signature analysis. Our pre-attentive model now employs a feedback neural network that refines alignment of the oriented edge elements (edgels) detected in the plasmagram image by subjecting them to collective global-scale optimization. The level of interaction between the oriented edgels is determined by their distance and mutual orientation in accordance with the Yen and Finkel model of the striate cortex that encompasses findings in psychophysical studies of human vision. An additional effort has been made to pre-process the raw image

data to eliminate noise and to detect, interpret and remove the resonance signatures that resulted from the local response of the plasma to the RPI signal.

The developed models have been implemented in an operational system “CORPRAL” (Cognitive Online Rpi Plasmagram Ranking Algorithm) that currently scans daily submissions of the RPI plasmagrams for the presence of echo traces. Qualifying plasmagrams are tagged in the mission database, making them available for a variety of queries. An analysis of the performance of CORPRAL is given and directions for future research are outlined.

ACKNOWLEDGEMENTS

I dedicate this work to my parents whom I had to leave behind in pursuit of my research.

To my father, a Physicist, an Antarctic expeditionist, a man of penetrating kindness and a big, tragically fragile heart.

To my mother, for bringing me to light in many ways and for making those hard decisions that shaped my future.

My work would not come to fruition if it were not for the people that I thank here most:

My wife, Elena, for her love, care and patience over the many years of my studies.

Together with our lovely daughter, Katerina, we have built the foundation for this work.

Prof. Bodo Reinisch, my preceptor and advisor for 11 years at the Center for Atmospheric Research, for opening doors to the worlds I have never known and then leading the way.

Prof. Georges Grinstein, my thesis adviser at Department of Computer Science, for deep insights in what I was trying to accomplish, for much needed encouragement, and thoughtful suggestions.

Prof. Haim Levkowitz, the Dissertation Committee member, for interest, support, and contributions to the project.

Grigori Khmyrov and Alexander Kozlov, my fellow *Russian Troika* members of the software development team at UMLCAR, for countless contributions to the project.

Prof. Gennadii Ososkov (JINR, Dubna) and Elena Zaznobina (Ivanovo State University), for sparking my interest in neurodynamic systems and important contributions to the ANNA algorithm.

Prof. Gary Sales (UMLCAR), for continuing interest, encouragement, and proofreading of the manuscript.

Prof. Steven Neshyba (University of Puget Sound) and Prof. Robert Gamache (UML), for lending their helping heart and hand during my first years in the USA.

My research in the intelligent systems at University of Massachusetts Lowell was funded in part by NASA under subcontract 83822 to the Southwest Research Institute and through the IS Grant NAG5-13387. I gratefully acknowledge valuable discussions and suggestions from the RPI team members, especially Drs. Robert Benson (NASA), Don Carpenter (Stanford University), and Shing Fung (NASA). I would like to thank Prof. Leif Finkel (UPenn) and Prof. Mark Matlin (Sarah Lawrence College) for discussions, demonstrations, and technical assistance with the Yen and Finkel early vision model.

I am grateful to Prof. Charles Steele for his word of support when I needed it most, and to Prof. George Cheney and Prof. James Canning for the inspiration. I would also like to thank the radio propagation group at Siberian Institute for Solar-Terrestrial Physics for the unforgettable early years of my research in Irkutsk and my fellow students, Urska Cvek and Marjan Trutschl, for their help and continuing interest.

TABLE OF CONTENTS

ABSTRACT	ii
Acknowledgements	iv
TABLE OF CONTENTS.....	vii
TABLE OF FIGURES	ix
CHAPTER 1. INTRODUCTION.....	1
1.1. Scope and Goals.....	3
1.2. Plan	4
CHAPTER 2. RADIO PLASMA IMAGER DATASET	7
2.1. Radio Plasma Imager.....	7
2.2. RPI Plasmagrams.....	8
2.2.1. Plasmagram Traces	10
2.2.2. Plasmagram Resonances.....	12
2.3. RPI Data Avalanche	14
CHAPTER 3. PREVIOUS WORK.....	17
3.1. Marr's Paradigm of Vision.....	17
3.2. Transition from Raw Image to Edgels	19
3.3. Transition from Edgels to Contours	21
3.3.1. Local grouping techniques.....	22
3.3.2. Global grouping techniques.....	27
3.4. Outlook.....	45
CHAPTER 4. TRACE recognition	46
4.1. Uniqueness of Plasmagram Signatures	46
4.1.1. Variable Signal-to-Noise Ratios.....	47
4.1.2. Range Jitter	50
4.1.3. Closely Spaced Traces	52
4.1.4. Range and Frequency Resolution Artifacts	52
4.2. Approach	53
4.3. Transition from Raw Image to Edgels	56
4.4. Transition from Edgels to Rotors	61
4.5. Transition from Rotors to Saliency Map.....	62
4.5.1. Co-circular Rotor Model Neurodynamics	63
4.5.2. Improved Rotor Interaction Pattern.....	67
4.6. Transition from Saliency Map to Contour Segments	72
4.7. Transition from Contour Segments to Contours.....	73
CHAPTER 5. PROCESSING OF RESONANCE SIGNATURES	75
5.1. Resonance Detection in Plasma Sounding Data: Previous Work.....	75
5.2. Fitting composite resonance model	80
5.3. Detecting and Matching Plasmagram Resonances.....	83

5.3.1.	Resonance Detection	83
5.3.2.	Matching the Gyrofrequency Resonances	90
5.3.3.	Matching the Plasma Frequency Resonance	92
5.4.	Discussion of Results	94
5.4.1.	Errors in prediction of medium gradients	95
5.4.2.	False decisions by the signature detector	95
5.4.3.	Key resonance signatures missing or obscured	96
5.5.	Unmatched Resonance Study	97
5.6.	Removal of Resonance Signatures	97
5.7.	Conclusion.....	98
CHAPTER 6.	IMPLEMENTATION AND PERFORMANCE STUDY	99
6.1.	Performance Study	99
6.1.1.	ANNA Performance on Synthesized Patterns	101
6.1.2.	ANNA Performance on Plasmagram Data	107
6.2.	Plasmagram Processing Results.....	107
CHAPTER 7.	CONCLUSIONS AND OUTLOOK.....	110
REFERENCES	114

TABLE OF FIGURES

Figure 2.1.	Earth’s magnetosphere and IMAGE orbit	7
Figure 2.2.	Echo detection concept in RPI.....	9
Figure 2.3.	Plasmagram representation of the RPI echoes with traces and resonance signatures as indicated	10
Figure 2.4.	RPI mission planning scenario: individual RPI programs (left bottom) are assembled into RPI Schedules (right bottom) to provide appropriate coverage of diverse science goals on orbit, depending on the spacecraft location and time (top) 15	
Figure 3.1.	Marr’s paradigm of visual perception: elements of increasing perceptual strength, built bottom up.....	18
Figure 3.2.	Example of direct grouping of edgels into contours. (a) edgels produced by detection, (b) edgel subsets forming contours.	22
Figure 3.3.	An example of the “microscopic” grouping technique considering all possible line tracing decisions in 3x3 context window. From [Liow, 1991].	24
Figure 3.4.	Fitting templates to edgels and resolving conflicts.	26
Figure 3.5.	Linear Hough transform fitting straight lines to a pattern containing four signal edgels (green) and one noise edgel (red)	29
Figure 3.6.	Use of Gestalt principles of perception for grouping (adapted from [Wersing et al., 2001]).	31
Figure 3.7.	The interaction pattern of the oriented cortex cells [Yen and Finkel, 1998] in the model of pre-attentive human vision. The post-synaptic cell in the center of the pattern receives facilitating inputs from the pre-synaptic cells in the shaded areas. Degree of facilitation depends on the distance between cells and their mutual orientation.....	35
Figure 3.8.	Hopfield model of artificial neuron,.....	36
Figure 3.9.	Segment model for contour extraction by [Peterson, 1989] and [Derby, 1988]. 38	
Figure 3.10.	Edgel grouping with Denby-Peterson segment model, an example from [Peterson, 1989]	39
Figure 3.11.	Rotor model network evolves, aligning rotors along the contour.....	40
Figure 3.12.	Linear trace model for rotor NN [Peterson, 1990]	40
Figure 3.13.	Circular model for rotor network by Baginyan et al., [1994].....	41
Figure 3.14.	Effects of λ parameter in the contour energy: (a) smooth line that does not capture corners, (b) good capture of corners but high sensitivity to noise, (c) compromise. From [Lai, 1994].	43
Figure 4.1.	Plasmagram taken on Jun 29, 2001 03:13 UT containing two traces with a low signal-to-noise ratio.	47
Figure 4.2.	Plasmagram taken on Jun 29, 2001 03:00 UT showing a trace of variable signal-to-noise ratio across the frequencies.....	48

Figure 4.3.	Plasmagram taken on Jul 9, 2003 18:24 UT during high radio emission activity in the frequency band between 210 and 260 kHz (shown in detail in panel (a)). The emissions cause multiple false positives of the echo detector (b).....	49
Figure 4.4.	Hough transform applied to low (left) and high (right) range jitter patterns	50
Figure 4.5.	Simulation of range jitter causing misalignment of the edgel interaction pattern and low salience.....	51
Figure 4.6.	Extended Marr's pyramid of perception for plasmagram processing....	54
Figure 4.7.	Edgel labeling for an example plasmagram (a) taken on March 2, 2002, 04:46 UT. Direct edgel detection (b) produces a large number of false edgels even with the thresholding of the amplitude gradient. (c) Application of the smoothing 3x3 median filter prior to edgel detection reduces the noise edgels but damages thin traces (d). The chosen approach (e) uses the adaptive local thresholding to detect echoes and then labels the leading edges (f).	57
Figure 4.8.	Echo detection by adaptive amplitude thresholding. (a) One frequency scan with intervals taken by echoes, (b) amplitude histogram used to determine the threshold level.	59
Figure 4.9.	Example processing with AvTrend echo detection algorithm. Sliding window of $N=7$ points is placed at the range bin shown in black. Window average amplitude is calculated over 6 points and the threshold is set D units higher. The range bin in this example is classified as echo because its amplitude exceeds the threshold.	60
Figure 4.10.	Use of angular histogramming algorithm [Bagynyan et al, 1994] to evaluate edgel orientation (example). (a) Angular histogram is built from elevation angle of all trial straight lines going through the base edgel and all nearby edgels within the circle of radius ζ . (b) The angle of highest occurrence in the histogram is selected as the edgel orientation.....	61
Figure 4.11.	Artificial neuron for rotor interaction.	62
Figure 4.12.	Processing of a synthesized edgel pattern by a circular model of rotor interaction: (a) edgel pattern, (b) initial rotor placement by angular histogramming, (c) results of ANNA optimization.....	63
Figure 4.13.	Circular model for rotor network by Bagynyan et al., [1994].....	64
Figure 4.14.	Rotor configuration corresponding to the strong false minimum of the Bagynyan rotor model energy.....	65
Figure 4.15.	Sigmoid function of the MFT neuron with region of optimal performance as indicated.	67
Figure 4.16.	Close-range co-circular interaction of rotors in range jitter conditions (simulated example). Due to the overestimated range of edgel j , contribution of rotor V_j to rotor V_i is too small.	68
Figure 4.17.	Saliency calculations for a weak contour in vicinity of a stronger contour (simulated pattern). (a) Simulated image of two traces. (b) Edgel pattern. (c) Initial rotor orientation. (d,e) Optimized rotor orientations and resulting trace segments obtained with a conventional circular model. Weak contour rotors are attracted to the stronger contour. (f,g) ANNA model processing with added dead zones in the interaction pattern.	69

Figure 4.18.	Placement of the dead zones of rotor interaction. The interaction pattern is oriented co-axial with (a) post-synaptic rotor V_i (conventional), or (b) pre-synaptic rotor V_j . In the latter case, error in initial evaluation of V_i orientation will be corrected during evolving of the neural network.	71
Figure 4.19.	Evaluation of the α_{ij} and β_{ij} angles for the connection smoothness term in grouping score [Galkin et al., 2004].	74
Figure 5.1	Response of the cumulative median filter to synthesized signals with falling (a) and rising (b) envelopes. The filter elevates the falling envelope that the resonance signatures are expected to display and reduces the rising slope, thus improving the signature contrast.	84
Figure 5.2	Response of the cumulative median filter to synthesized signals with falling (a) and rising (b) envelopes, with the addition of 30% noise. The filter smoothes the noise and enhances the contrast of the resonance signatures.	85
Figure 5.3	Cumulative median filter applied to RPI data: (a) a frequency containing resonance, (b) a frequency without resonance and a single echo.	86
Figure 5.4	Resonance detection in RPI plasmagrams. (a) Raw plasmagram, (b) Processing with the cumulative median filter. (c) Summary amplitude function (white bars), labeled frequencies containing resonances (gray bars) and detected peaks (black bars). (d) Localized resonance signatures shown in magenta.	87
Figure 5.5	Use of the differential histogramming technique to estimate the local detection threshold for the resonance signature.	90
Figure 5.6	Choice of the actual plasmagram frequencies next to the trial frequency. If the closest frequency is not tagged as containing a resonance, the other frequency contributes to the fit quality with a reduced weight.	91
Figure 6.1.	Simulated pattern containing a single trace with gaps and noise.	102
Figure 6.2.	Processing of the same pattern as in Figure 6.1 at smaller (a-c) and larger (d-f) analysis scales.	104
Figure 6.3.	Processing of the same pattern as in Figure 6.2(d) with a more appropriate annealing regime (stop temperature increased to prevent premature saturation due to a larger scale of analysis).	105
Figure 6.4.	Sample pattern containing two closely spaced traces with some minor gaps and noise.	106
Figure 6.5.	Processing of weak short contours in vicinity of a strong contour.	106
Figure 6.6.	Example of processing stages for RPI plasmagram taken on March 01, 2002 00:02:58 UT. The raw plasmagram (a) is thresholded to obtain echoes (b), which are then reduced to the edgels (c). The edgel orientations are obtained and optimized to derive the saliency map (d) of the image. The saliency measures are analyzed to obtain trace segments (e) that are then combined together to form traces (f).	108
Figure 6.7.	Some of July 2001 plasmagrams selected by CORPRAL as containing 6 traces or more.	109

CHAPTER 1. INTRODUCTION

With the growing detailed knowledge comes an increased sophistication of our exploration techniques. We deploy scientific experiments of greater complexity and build research instruments of higher information output, all in the name of gaining further insight into the universe. Often, however, the information explosion we ignite turns into a tedious, unfathomable data avalanche. In many applications, ranging from a multi-terabyte astronomical sky survey to the ever-growing imagery archive produced by the high energy particle accelerators, the incoming flow of scientific data would demand humanly impossible effort to explore and comprehend if it were not for assistance of the Intelligent Systems (IS).

Many information-rich scientific projects spawn the IS applications that establish an automated clearinghouse for dispersed and disorganized data. The computer plays a powerful and enabling role in those projects where the size of dataset precludes manual processing. The IS applications that deal with images are perhaps the most sophisticated ones because of the need to model the intricate process of visual information processing in the brain. While a wide range of disciplines such as biology, physics, mathematics, computer science, and psychology have offered a great variety of concepts to our understanding of visual perception, a large gap still exists between the lower level of vision as described by the neurobiology of the retina cells and brain cortex, and the upper level of such brain functions as memory and learning. Lack of this understanding became the stumbling point for a real-world IS project that was started in order to automatically

classify and sort images from the Radio Plasma Imager (RPI) [Reinisch *et al.*, 2000], a radar instrument onboard NASA's IMAGE spacecraft [Burch, 2000]. The RPI makes snapshots of plasma conditions in the magnetosphere of the Earth at a rate of ~600 images per day. Exploration of an early-vision model capable of detecting signatures in the RPI images is the topic of this thesis.

The early vision extracts salient cues from images without willful concentration of attention on the image elements, and for that reason is often called *pre-attentive*. Pre-attentive vision effectively “pops up” contours of objects in the field of view. Existing model descriptions of the pre-attentive vision are based on the concept of reductionism that explains behavior of a system in terms of the simpler subsystems comprising it. Known since the late 1970s as the Marr's paradigm of vision [Marr and Nishihara, 1978], this reductionist approach builds a pyramid of perception in a bottom-up fashion, where each new level holds visual elements of higher perceptual strength and lower levels are unaware of considerations happening at the higher levels. Advancement from level to level in the Marr's pyramid is often referred to as *perceptual grouping*, which is governed by Gestalt principles of perception [Rock and Palmer, 1990]. Though largely simplified, the pyramid model is in a good agreement with *in vitro* biophysical studies of the eye retina and brain cortex cells responding to simple shape stimuli such as line segments. However, *in vivo* studies have shown that visual system cells display seemingly chaotic and highly irregular behavior [Holt *et al.*, 1996]. It is the goal here to investigate this process in a greater detail by placing the early-vision model in the framework of neurodynamic organization, where collective interaction of neurons optimizes segmentation of the contours.

1.1. Scope and Goals

The goal of this thesis is to contribute to the understanding of pre-attentive vision by investigating an artificial neural network algorithm (ANNA) for contour segmentation. ANNA is a feedback neural network with a sigmoid transfer function of the neurons and network evolving procedure based on the mean field theory [Peterson and Anderson, 1987] that employs a Gestalt-compliant interaction pattern similar to the Yen and Finkel [1998] model of the brain cortex. Although compliance with the Gestalt law of perception is a target for many models of perceptual grouping, its implementations within the neurodynamic framework to the real-world images are rare because of a high computational demand. The ANNA model is explored by identifying its performance issues on synthetic and actual RPI data and heuristically shaping its energy function and evolving process for optimal performance. Ultimately, ANNA is the central component of a Cognitive Online Rpi Plasmagram Rating Algorithm (CORPRAL) [Galkin *et al.*, 2004b] for signature extraction from the RPI images. Ratings provided by the CORPRAL are instrumental to efficient data exploration by quering the RPI mission database currently holding nearly 1,000,000 images. Other important issues of building the Marr's pyramid layers of the CORPRAL that make our approach feasible are discussed. These include principles of feature selection for perceptual grouping, as well as an independent topic of detection and matching of the plasma resonances in RPI image data [Galkin *et al.*, 2004a], which is done prior to the contour segmentation to simplify the task. The last part of the thesis presents the overall CORPRAL performance results and outlines future work.

1.2. Thesis Contributions

This thesis takes advantage of a cross-disciplinary approach to the difficult, scientifically significant problem of feature extraction from the imagery data of highly variable content and quality. Most parts of the developed pre-attentive vision model discussed earlier were previously known from other studies in related areas. They were brought together to create an innovative perceptual system that draws its strengths from established concepts in psychology, biology, mathematics, remote sensing, plasma physics, and computer science. The thesis research expands the classic Marr's classic vision paradigm by adding a previously-known concept, rotors, as a new perceptual layer in the pyramid. In the rotor layer, the oriented edge elements (edgels) detected in the image are subjected to a collective global-scale optimization to correct imperfections of their orientation originally determined by a local-scale process. Letting oriented edgels (rotors) change their orientation under collective influence of surrounding rotors is a well-known concept that can be traced back to the magnetic spin systems in physics. The principles of rotor interaction, however, remained rather simple over the years; enhancing them with the modern understanding of cell interaction in the brain cortex is another important contribution of the thesis.

The benefits and viability of the proposed techniques are shown using the operational system for plasmagram analysis, CORPRAL, which applies the pre-attentive model to the RPI plasmagrams. Special attention is paid to pre-processing of plasmagrams based on the principles of signal detection and resonance interpretation. The thesis contributed notably to the task of the automated matching of the plasma resonance signatures to their

theoretical counterparts. A new cumulative-median filter was developed for enhancing the contrast of resonance signatures, and a resonance detector was developed to register and evaluate the signatures prior to matching them to the theoretical model. The resonance matching algorithm is implemented in the flagship data visualization system for RPI data, BinBrowser, to aid the process of plasmagram interpretation.

The final contribution of the thesis is the data exploration results obtained by the CORPRAL system that works online with the RPI database holding over 800,000 plasmagram images. The subset of plasmagrams identified by CORPRAL as containing traces is made available for remote queries from the BinBrowser workstations.

1.3. Plan

The first chapter discusses the motivation for development of intelligent systems that help to organize and explore large datasets collected by scientific applications. It then describes the essence of the thesis research and defines its scope, goals, and contributions in relation to the concept of automated data exploration and the need for better understanding how the early vision functions.

The second chapter gives the description of the Radio Plasma Imager and specifics of its dataset that motivated this research.

Chapter 3 gives an overview of the existing contour extraction approaches based on the Marr paradigm of vision. The review distinguishes between local and global techniques and summarizes the main results of previous research that serves as the

foundation for this work: the Gestalt-compliant interaction of directed edge elements and feedback neural network optimizing the alignment of rotors.

In Chapter 4 the artificial neural network algorithm (ANNA) is introduced as a feedback neural network with sigmoid transfer function and the mean field evolving process that aligns rotors along the contours in the image. The ANNA uses Gestalt laws of perception in the design of the rotor interaction pattern similar to many other models of perceptual grouping. The modifications to the interaction pattern are then discussed that are necessary to provide robustness of the alignment algorithm to the specifics of the RPI image data. This also works to avoid the local minima of the energy function during the evolving of the network. The need for pre-processing of features is then discussed in the context of the computational requirements imposed by the ANNA operations. The analysis of ANNA performance is then made on synthetic and actual RPI images.

The fifth chapter introduces a related topic of the detection and matching of the plasma resonance signatures in the RPI image data, an operation done prior to contour segmentation to simplify the ANNA processing. The resonance processing algorithm includes a signature detector and an automatching algorithm based on the existing relations between several resonance frequencies.

Chapter 6 presents CORPRAL, a software tool for signature extraction in the RPI images based on the ANNA model, including implementation, first results and performance issues of the CORPRAL.

Chapter 7 summarizes the main results and discusses future research directions.

CHAPTER 2. RADIO PLASMA IMAGER DATASET

Our study was motivated and influenced by the need to build an intelligent system for the automated exploration of imagery data collected by the Radio Plasma Imager (RPI) designed to locally and remotely probe the plasma in the magnetosphere surrounding the Earth. The chapter briefly describes the RPI instrument and highlights specifics of its operations and data. Chapter 4 discusses in further detail the unique features of the RPI dataset that affected design of the pre-attentive vision model for signature extraction in the RPI images.

2.1. Radio Plasma Imager

The Radio Plasma Imager aboard the IMAGE spacecraft is a radar with the direction finding capability, designed to probe the plasma of the Earth's magnetosphere at near and far ranges (Figure 2.1).




Figure 2.1. Earth's magnetosphere and IMAGE orbit

The fundamental objective of RPI is to detect responses of surrounding plasma to the transmitted radio signals and to infer characteristics of the magnetospheric plasma environment from the detected signatures. RPI can be thought of as a point source of radio waves that illuminates the space around it using short pulse and then detects echoes arriving from various regions of ionization in the magnetosphere, plasmasphere and ionosphere, as well as plasma irregularities of various scale sizes. For the radio signal to reflect at a remote location and return to its origin, specific conditions have to be met. Not only must the frequency of the signal match the local plasma frequency in the reflection area, but also the plasma density gradients have to satisfy certain geometrical criteria. It is important to recognize that RPI, in its active sounding modes, is the radar of opportunity that “sees” targets only sporadically on any particular orbit. The need to search for features in the RPI dataset has been the rationale for intensive mission planning, data mining and signature characterization effort.

2.2. RPI Plasmagrams

To remotely sense conditions of the surrounding plasma, RPI employs a stepped frequency radio sounding concept that has been successfully practiced in a variety of applications, in particular the ionospheric sounder [Reinisch, 1996]. To search for the reflected signals, the RPI probes a specified number of the sounding frequencies by transmitting a pulse and sampling, for each frequency, the amplitude of receiver voltage at a specified number of delay times after the pulse transmission (Figure 2.2).

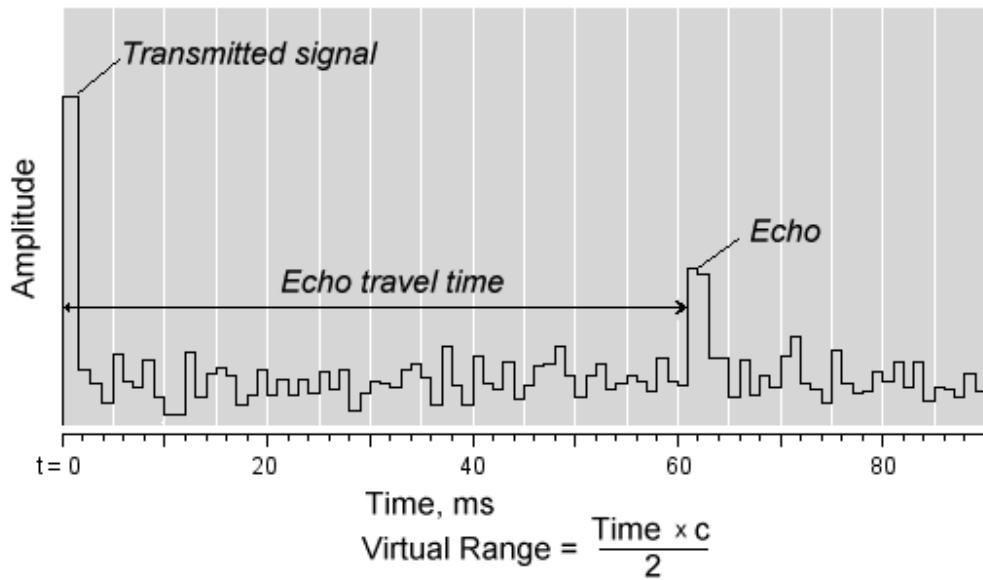


Figure 2.2. *Echo detection concept in RPI*

The echo travel time is usually converted to the so-called virtual range to the reflector using the assumption that the signal travels at the speed of light. For brevity, the virtual range is often referred to as “range” in this text. Measurements of the amplitude at each of the sampled virtual ranges present the first, simplest layer of the RPI data roster. Figure 2.3 shows an example plot of the amplitudes in the frame of sounding frequency versus virtual range, a technique that has been in practice since the first experiments with the swept frequency sounding of the ionosphere in the 1930s.

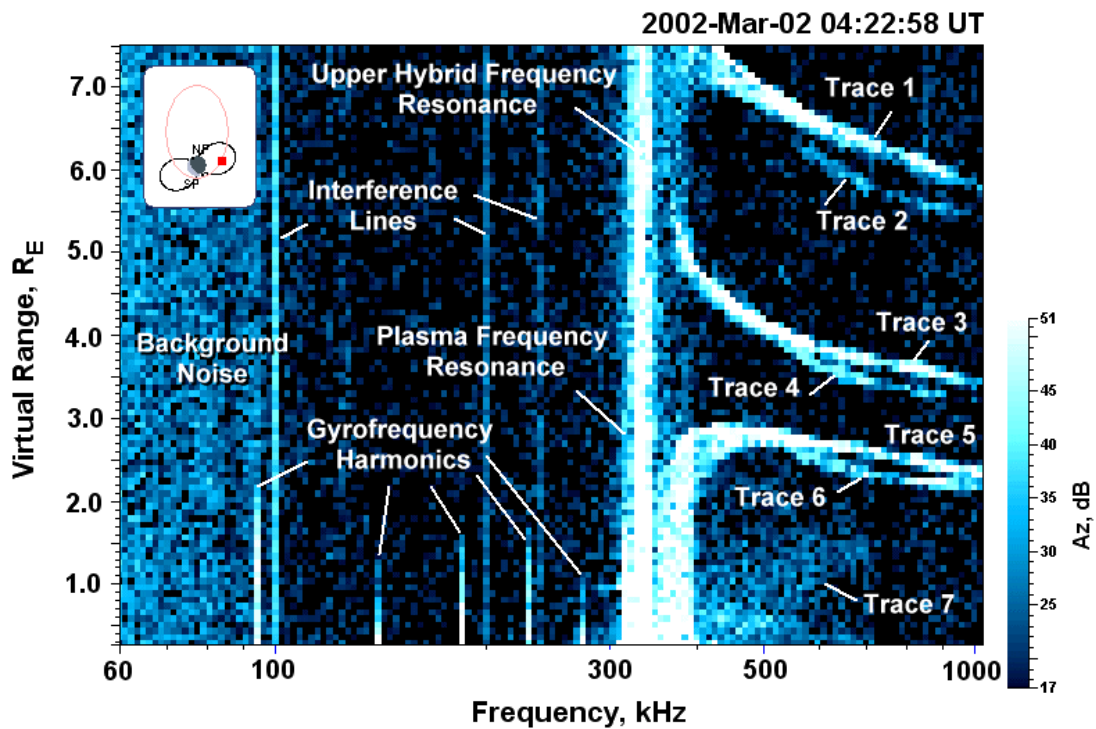


Figure 2.3. Plasmagram representation of the RPI echoes with traces and resonance signatures as indicated

Plots of RPI amplitudes in the frequency-range frame are called *plasmagrams*. The example plasmagram shown in Figure 2.4 is plotted using the BinBrowser visualization tool [Galkin *et al.*, 2001]. As can be seen in this particular plasmagram taken on March 20, 2002, 04:28 UT, the plasma responses to the RPI transmitter signal in two basic ways, (a) reflecting the signal from a remote location back to its origin, and (b) sustaining short-range plasma waves at characteristic frequencies, resonances.

2.2.1. Plasmagram Traces

Individual echoes, when viewed in the plasmagram representation, typically form curves of various shapes and lengths, called *traces*. Each individual plasmagram trace is

x-monotonic, i.e., for each frequency only one range value exists. The traces are not y-monotonic, though, and the same range value may appear twice (or more) in a trace.

Interpretation of the traces found in the RPI plasmagrams still presents a considerable difficulty because of the experiment's novelty. Preliminary investigation showed that the traces can be classified in at least three basic categories:

1. Wave polarization ^a
 - a. O
 - b. X
 - c. Z

2. Wave propagation mechanism ^b
 - a. Direct
 - b. Field-aligned propagation (FAP)

3. Reflector ^c

a For RPI electromagnetic waves, the space plasma in the Earth's magnetic field is an anisotropic medium that causes splitting of the RPI signal into the left- and right-circular polarized waves traveling with different phase and group velocities in the medium and yielding different traces on the plasmagram. The left circular and right circular waves (with respect to the geomagnetic field) are referred to as ordinary (O), and extraordinary (X) polarizations, whereas their combined mode involving mode transformations is called Z-trace.

b Propagation of the electromagnetic waves in the magnetized plasma of the Earth has been studied extensively to identify two major propagation mechanisms of special interest to RPI, (1) direct (involving refraction) and (2) field aligned (following the curve of Earth's magnetic field lines [Reinisch *et al.*, 2001]). The physical mechanism causing waves to deviate from the straight line is explained in terms of guided propagation.

c The conditions for RPI signal to reflect and return to the spacecraft location may be satisfied in a variety of plasma structures of global scale. Those are typically either abrupt gradients of plasma density (plasmopause, magnetopause, cusp), or plasma of smoothly increasing density (plasmosphere, ionosphere), where the local plasma frequency eventually matches the sounding frequency causing the refractive index to become zero and wave to reflect.

- a. Plasmasphere
- b. Plasmopause
- c. Ionosphere
- d. Polar Cap area
- e. Cusp
- f. Magnetopause

Finding plasmagrams with traces and interpretation of the found traces constitute a significant part of the RPI science data analysis. This process usually involves manual screening of the plasmagrams and extraction of the traces, a tedious effort often requiring hours to process one day of raw data.

2.2.2. Plasmagram Resonances

A radio transmitter immersed in the plasma is capable of stimulating short-range plasma-wave echoes and plasma emissions when its sounding frequency matches one of the characteristic frequencies of the plasma, often called resonance frequencies, or simply resonances. Typical resonance signatures have been observed by a number of space missions carrying a topside ionosonde or a relaxation sounder (see Muldrew [1972] for a representative review). Detection of stimulated resonances and wave cutoffs in the radio sounding data provides a measurement of local plasma density and magnetic field intensity that has demonstrated an accuracy and diagnostic potential superior to what conventional magnetometers and density probes are able to achieve.

The resonance signatures appear in plasmagrams as the vertical line segments extending upward from zero virtual range (see Figure 2.3), each corresponding to a

particular resonant response of the local plasma to the transmitter excitation. Five major types of resonances in the RPI plasmagrams have been identified, all pertaining to oscillations of the plasma electrons, observed at the following frequencies:

- a) electron cyclotron frequency f_{ce} , and its harmonics, $n:f_{ce}$, ($n = 1, 2, \dots$)
- b) electron plasma frequency f_{pe} ,
- c) upper-hybrid frequency f_{uh} ,
- d) Qn resonances f_{Qn} (also known as Bernstein-mode resonances)
- e) Dn resonances f_{Dn}

A number of relations exist between the plasma resonance frequencies [Stix, 1992]. The upper-hybrid frequency, f_{uh} , for example, is given in terms of the plasma and gyro frequencies:

$$f_{uh} = \sqrt{f_{pe}^2 + f_{ce}^2} \quad (2.1)$$

The Qn resonances can be related to the plasma and gyro frequencies by an approximate expression [Warren and Hagg, 1968]:

$$f_{Qn} \approx f_{ce} \left[n + \frac{0.46}{n^2} \frac{f_{pe}^2}{f_{ce}^2} \right], \quad (2.2)$$

which is a good approximation when f_{Qn}/f_{ce} is near an integer value. Benson *et al.* [2001] provided curves for f_{Qn}/f_{ce} based on electrostatic dispersion equation solutions for zero-group velocity plasma waves perpendicular to the Earth's magnetic field. Finally, the sequence of Dn resonances is described by the following expressions [Osherovich and Benson, 1991]:

$$\begin{aligned}
f_{Dn} &= 0.95\sqrt{f_p f_{ce}} \sqrt{n} \quad n = 1, 2, \dots \\
f_{Dn}^+ &= \sqrt{f_{Dn}^2 + f_{ce}^2} \\
f_{Dn}^- &= \sqrt{f_{Dn}^2 - f_{ce}^2}
\end{aligned} \tag{2.3}$$

Equations (2.1-2.3) can be used to build a composite model of the resonance signatures driven only by the electron gyrofrequency f_{ce} and the electron plasma frequency f_{pe} . In addition to the known resonance frequencies, plasmagram scalars commonly utilize, for the resonance interpretation, the so-called X cutoff and Z cutoff frequencies, corresponding to the left cusp end points of the Trace 5 and 6 shown in Figure 2.4:

$$f_x = \frac{f_{ce}}{2} \left[1 + \sqrt{1 + 4 \frac{f_{pe}^2}{f_{ce}^2}} \right] \tag{2.4}$$

$$f_z = f_x - f_{ce} \tag{2.5}$$

2.3. RPI Data Avalanche

In addition to the remote active sensing mode discussed in Section 2.2, where RPI receives reflections of its transmitter signal from remote regular and irregular plasma structures, and stimulates and observes the local plasma resonances, RPI also monitors the natural electromagnetic noise environment and participates in the joint campaigns with other radio spectrum instruments in space and on the ground as both transmitter and receiver. Furthermore, even within the same observation mode, RPI cannot keep its measurement program parameters constant as it orbits the Earth. Because of the highly elliptical orbit of the IMAGE spacecraft, RPI observes a dramatic range of the plasma

and magnetic field conditions and needs to accommodate this variability by appropriate selection of its plasmagram frequency and range intervals.

Figure 2.4 illustrates the elaborate mission planning concept for RPI operations, originally described by Reinisch *et al.*, [2000].

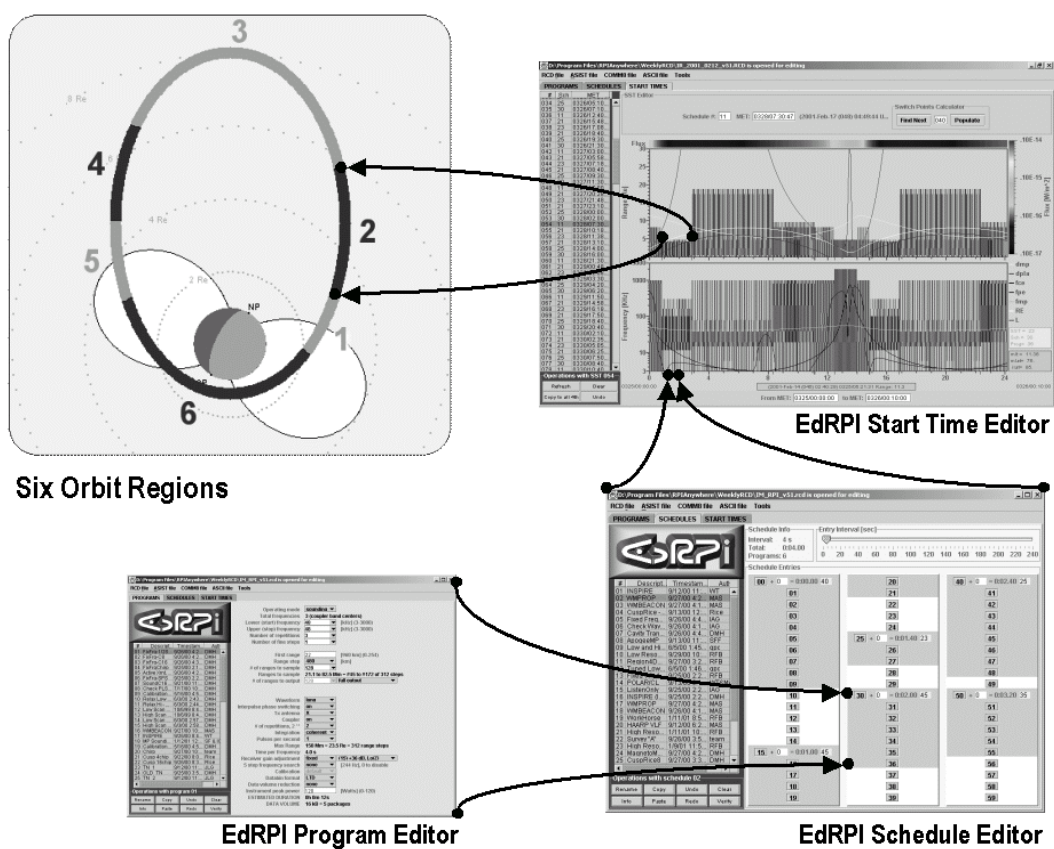


Figure 2.4. RPI mission planning scenario: individual RPI programs (left bottom) are assembled into RPI Schedules (right bottom) to provide appropriate coverage of diverse science goals on orbit, depending on the spacecraft location and time (top)

The spacecraft orbit is sectioned into six regions whose boundaries are calculated individually for each 14.5 hour orbit, depending on the location of the intersection points

of the orbit with the predicted location of the Earth's plasmopause^d. Within each orbit region a particular repetitive sequence of measurements is run, called a *schedule*. Figure 2.4 includes the screen captures from the RPI mission planning tool, *EdRPI*. Individual RPI programs are designed with the Program Editor, which are then assembled into schedules using the Schedule Editor, and then finally assigned to the appropriate orbit regions using the Start Time Editor.

As individual measurements may or may not contain useful information, the RPI data requires a substantial exploration and classification effort. Unfortunately, because of the irregular content and format of the RPI data stream, it is impossible to present its images in a constant frame so that many frames can be assembled in movies for a fast visual analysis. Data analysts have to browse through the RPI archive image by image to inspect plasmagrams for useful information. In July 2002, the measurement count reached 1,000,000 records. With another three years of anticipated mission life, the total number of data records is expected to reach 3,000,000, with projected total of 2,000,000 plasmagrams. With average 5 sec response time of the database to the retrieval request, it will take about 2,700 man-hours to simply glance at the data. This condition can be classified as the "data avalanche".

^d Plasmopause is the outer boundary of the Earth's plasmasphere. Plasmasphere comprises the higher density plasma controlled by the Earth's magnetic field that co-rotates with the Earth.

CHAPTER 3. PREVIOUS WORK

In this chapter we put the task of finding traces in the RPI plasmagrams in perspective of existing approaches to contour detection in images.

3.1. Marr's Paradigm of Vision

We constrict this review to the contour extraction techniques that comply with Marr's paradigm of computer vision [Marr, 1976], where the elements of increasing perceptual strength are built successively bottom-up using such operations as selecting and grouping. The Marr's paradigm has been widely adopted for its ability to break the complex task of visual analysis into smaller, independent and manageable components (layers). There is ample evidence of modularity in the mammal vision system supporting this approach. On the other hand, independence of the layers in the Marr's pyramid constitutes a fundamental flaw: lower stages of analysis are unaware of the model considerations on the higher levels, and errors flow irreversibly from lower stages of analysis upward to the next. Nonetheless, the Marr's paradigm is commonly accepted as the founding principle of the pre-attentive vision, and its extensions are sought along various avenues.

Figure 3.1 shows the Marr's pyramid of perception that is used here to review existing work on contour extraction. The pyramid consists of the following layers:

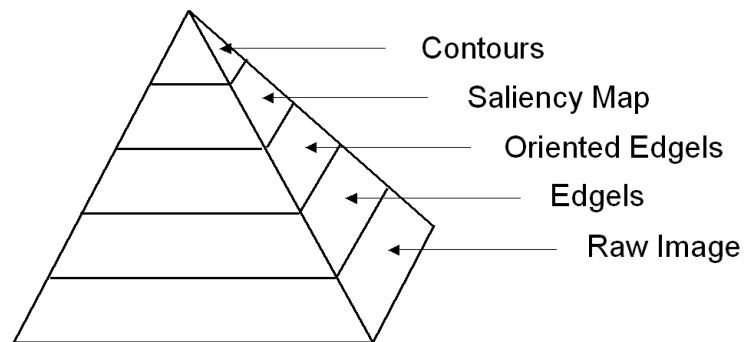


Figure 3.1. *Marr's paradigm of visual perception: elements of increasing perceptual strength, built bottom up.*

Raw Image. Raw image is an ($N \times M$) matrix of pixels, each pixel representing the image intensity.

Edgels. Edgel is an “edge element”, a term introduced in visual scene analysis where objects are analyzed by first locating their boundaries (edges). The edgel corresponds to a sharp gradient in the image intensity typical of a boundary. Individual edgels form the contour.

Oriented Edgels. Evaluation of local orientation of the contour at the edgel location adds an additional perceptual value to edgels. The edgel orientation is commonly obtained by seeking gradients of image intensity within a local context window around the edgel.

Saliency map. Local context techniques do not yield reliable edgel orientations in various situations requiring a larger context area to correctly identify the trace. Long-range, collective analysis of edgel data results in evaluation of perceptually stronger quantity, saliency. The saliency measure reflects how likely an edgel is part of a

contour. The saliency measure calculated for all edgels constitute the saliency map of the image.

Contours. The saliency map is analyzed to find subsets of edgels belonging to the same contour. Found subsets may correspond to the contour segments that need to be completed into the full contours, but this operation is outside the scope of the pre-attentive vision.

A wide range of disciplines such as biology, physics, mathematics, computer science, and psychology have offered a great variety of concepts that contribute to the understanding how perception progresses from one level to the next in the Marr's pyramid. The following section reviews the major results of these studies.

3.2. Transition from Raw Image to Edgels

In the seminal work on the edge detection, Marr and Hildreth [1980] laid the mathematical foundation that tied together several physiological mechanisms previously discovered in experiments with the retinal cells responding to a point light stimulus. It appears that the retina is not simply a photoreceptor matrix, but a rather complicated parallel bio-processor of the image. It partitions the image at multiple scales (channels) using cells of various receptive field sizes. Within each channel, a smoothing filter is applied, second derivative of the image intensity is taken, the short segments across which the second derivative changes its sign (zero crossings) are detected, and the magnitude of the change gradient is evaluated. Then the zero-crossing segments found in each scale channel are combined into a map describing all gradients of the image

intensity. The map is a symbolic early visual code presenting the original raw image to the higher levels of the perception.

The merit of biophysical plausibility goes to the edge detection computational techniques that employ, similar to the retina,

- a differentiation operator sensitive to the sharp gradients of the image intensity corresponding to an object boundary or a line,
- a smoothing filter providing protection of the differentiation operator from the false positives due to noise, and
- a multi-scale analysis ensuring that gradients of various scales are localized.

A suite of differentiation operators and filters has been devised for detection of edges and ridges in images [Ziou and Tabbone, 1998]. The fundamental dilemma is to find the right balance between noise reduction and inevitable loss of information due to the damage that smoothing imposes on the fine structures in an image. This trade-off problem has motivated development of a class of detection algorithms that manipulate the analysis scale to better balance the tasks of noise protection and edge detailing. This class includes the classic Canny's "feature synthesis" detector [Canny, 1986] that implements the fine-to-coarse scale combination strategy, "edge focusing" filter by Bergholm [1987] that applies the coarse-to-fine scale processing, and a later "local scale control" detector by Elder and Zucker [1996] that selects the most appropriate scale for each edgel.

Incomplete presentation of the raw image by the edge detector to higher stages of the processing is a serious source of recognition errors. Securing completeness of the representation code is a non-trivial task, considering the need to balance the completeness requirement with compactness, robustness and precision of the obtained code [Elder, 1999]. Although many multi-scale representations can be made complete or near-complete (e.g., the zero-crossing operator at multiple scales used by Marr and Hildreth is complete), they are definitely not compact. Wavelets-based codes [Mallat, 1989] secure mathematical completeness of the representation but infer little understanding of the explicit image structure and therefore offer just another way to compress the image. The variable scale method where each edgel is represented not by multiple scales, but rather by a uniquely selected scale [Elder and Zucker, 1996, Lindenberg, 1996] comes very close to fulfilling the list of requirements, but the completeness of the output code is achieved by loading the edgel with additional information (such as blur scale and asymptotic intensities) that are yet to be included in the models at the higher layers of the pre-attentive vision pyramid. These models commonly operate on the multiple scales independently and in order to select a feature across the scales implement a combination of “winner take all” (WTA) and “inhibition of return” strategies [Koch and Ullman, 1985], discussed in greater detail in Section 3.3.2.5 below.

3.3. Transition from Edgels to Contours

After the edgels are detected, a solution to the task of combining them into contours can be sought immediately, without going through the intermediate layers of the Marr’s

pyramid shown in Figure 3.1. Figure 3.2 presents a realistic example illustrating several facets of the task.




Figure 3.2. Example of direct grouping of edgels into contours. (a) edgels produced by detection, (b) edgel subsets forming contours.

The edgels rarely line up in distinct, continuous sequences. In most real-world applications, the set of detected edgels is prone to various degrees of false positive edgels (noise), false negative edgels (gaps), and localization errors (jitter). However, a simple concept of seeking neighbors in vicinity of each edgel can readily make a working algorithm for edgel grouping. Early edgel grouping algorithms worked quite efficiently in a number of applications where problems of noise, gaps and jitter were not severe.

3.3.1. Local grouping techniques

Early edgel grouping methods worked in two steps, seeding and tracing. The seeding algorithm sought an edgel or a group of edgels that could be used as the start of a contour, and the tracing attempted to find the rest of the edgels constituting the contour by testing edgels in the vicinity. Grouping algorithms based on this simple principle are

found in the literature as far back as the early 1960s (e.g., [Galkin, 1962]). These algorithms employed a rudimentary logic that allowed them to bridge narrow gaps in the contour and avoid false edgels that did not lie along the contour line.

Rosenfeld *et al.*, [1976] pioneered an approach that not only made the early methods more robust against noise and gaps, but in fact established a new field of the research in perceptual grouping. It started with the idea to analyze the close vicinity of edgels before linking them, trying to identify and reverse false positive and false negative decisions (i.e., noise and gaps). Because the authors used the term “labeling” for the operation of edge detection, the revising procedure was named “relabeling”. False negative pixels were considered for relabeling via a “relaxation” procedure that simply weakened the labeling criterion if the trial pixel fitted well with the neighboring edgels (for more details, refer to Hancock and Kittler [1990]). Similarly, the false positive edgels were removed if they were not supported by either strong or weak edgels. Finally, an important decision was made to make the relaxation procedure iterative and run it until no further relabeling changes were observed. The iterative nature of the relaxation algorithm had an intriguing impact on its further development. As one individual application of the algorithm affects only a few immediate neighboring context windows, the relaxation labeling can be performed in the parallel fashion and can be therefore thought as dynamic, parallel evolving of a network of locally interacting operators. This observation linked the simple heuristic edgel relaxation concept to mechanisms of the pre-attentive vision and resulted in an intensive study of statistical and neural aspects of the *relaxation labeling networks* [Kittler and Illingworth, 1985; Pelillo and Faneli, 1997]. Eventually the

relaxation labeling network concept was recasted into the framework of evolutionary autonomic agents [Liu *et al.*, 1997] and took its place among other combinatorial optimization methods discussed later in Section 3.3.2.

The iterative relaxation procedure analyzes the context window around the trial pixel many times. We call a local grouping technique *microscopic* if its context window is small and the relabeling outcome or edgel linking decisions are pre-computed following certain design rules (cellular automata) and are often stored in a lookup table for better speed. A popular example of such microscopic technique is so-called extended border tracing by Liow [1991] where all configurations of the labeled pixels in 3x3 context window were reduced to the set of 12 situations, each determining the next step of the tracer (see Figure 3.3).

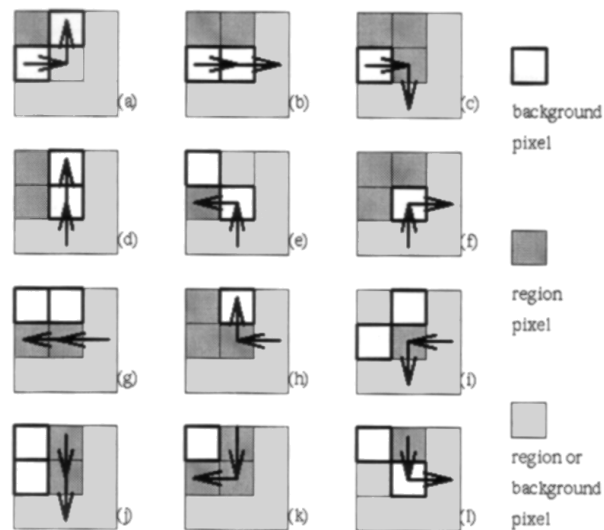


Figure 3.3. An example of the “microscopic” grouping technique considering all possible line tracing decisions in 3x3 context window. From [Liow, 1991].

The fundamental problem with the microscopic local grouping techniques is insufficient protection against noise, gaps and jitter due to inherent “locality” problem (see, for example, analysis by Nelson, [1994]) that is directly related to their inability to abstract from precise edgel locations within the small context window. The locality problem manifests itself as failure of the algorithm to universally link together longer line segments and at the same time be robust against the noise and edgel jitter. The process of enlarging the scale of collective analysis of edgels constituting the lines leads to development of local, *macroscopic* grouping techniques.

Simple enlargement of the context window so that wider gaps can be bridged by the grouping algorithm causes aggravation of the locality problem. A popular class of techniques strengthens the line smoothness requirement by constraining the search of edgels through the gaps to the sectors whose location is predicted by extrapolation of an already found line segment. The predicting filter can be thought of as a variant of the well-known Kalman filter used to predict the next state of systems in time. To use the predicting filter for edgel grouping, a line segment needs to be found first to setup the filter (equivalent to the system history in Kalman filtering). To find the seeding segment in the edgel data, certain assumptions have to be used on its quality in terms of noise, gaps, and jitter, which is a drawback of the algorithm.

Numerous implementations of the predicting filter grouping technique can be found in the literature since 1960s (for a review of the early work on line tracing in high energy physics refer to [Strand, 1972], a good example of a predicting filter for ionogram autoscaling is given by Fox and Blundel [1989]). With time, these techniques were

improved for better noise robustness during both seeding and tracing phases (e.g., an optimizing technique by Nelson [1994]). However, the predicting filter remains prone to locality problems because each grouping decision involves only a single edge.

Techniques that manipulate with longer line primitives instead of individual edgels or small context windows may have better robustness to noise and jitter. In one approach by the author [Galkin, 1992], a set of predetermined linear templates is applied to each edgel to test if there are neighboring edgels that fit together within one template. Because two or more templates can claim the same edgel, a recursive conflict resolution algorithm is run to determine one-to-one assignment of edgels to lines (Figure 3.4).




Figure 3.4. *Fitting templates to edgels and resolving conflicts.*

Although fairly successful in the case of smooth traces, the algorithm still appeared to be sensitive to the jitter observed during disturbed plasma conditions, and its gap bridging

characteristics were rather modest. Although widening of the template could improve robustness to the jitter, it also creates excessive conflict rate and deteriorates protection against noise.

A generic bottom-up clusterization with a distance metric for the agglomeration criterion will not produce elongated clusters for it is supporting the proximity principle instead of the continuity (see Figure 3.1). Murtagh and Raftery [1984] were able to generalize the distance metric to include a constraining parameter that scaled down the criterion across the short dimension of the clusters. The constrained clusterization is able to group edgels in long and smooth lines, but use of the distance metric makes it overly sensitive to noise and jitter.

To extend the context area of grouping analysis, it became necessary to analyze all edgels in the context neighborhood collectively. Collective analysis naturally led to a suite of algorithms featuring synaptic interactions that weigh contributions from individual edgels. The weights were determined based on mutual placement of edgels and other characteristics. The key principle of the weighted interaction is to facilitate contributions from the edgels belonging to the same contour and to suppress all others.

3.3.2. Global grouping techniques

We use the term "global" to describe the next class of edgel grouping techniques to highlight the difference to the local techniques that manipulate with individual edgels separately. In contrast, the global grouping method involves collective analysis of many edgels to make grouping decisions. Global methods employ a model, though simple at

times, that describes the class of lines to look for and then extracts the lines by bringing the model into optimal agreement with the curvilinear features in the image.

The Hough transform [Hough, 1962], together with its many derivatives and generalizations, has been known in the literature as the rigid contour approach because it makes direct assumptions on the geometric shape of the sought contours and exhaustively fits the contour model to the image, looking for the best fit(s). Although it is very rare that contours have perfect geometric shape, the rigid processing is still commonly used to find segments for further grouping.

3.3.2.1. *Rigid Contours*

The Hough transform is a classic example of the global model fitting technique for searching rigid contours in edgel patterns. In its classic form, the Hough transform attempts to fit a rigid model to all possible subsets of given points, thus building a histogram of estimated model parameters. Figure 3.5 illustrates this idea with an example of the linear Hough transform applied to a point pattern containing four signal points and one noise point. The four signal points on Figure 3.5 are consistent with the linear model and therefore all parameters estimated from the six signal pairs contribute to the same bin of the histogram. The pairs with the noise point produce four different parameter sets. After the exhaustive fitting is done, the histogram can be searched for the local maxima, corresponding to the fitting solutions. The solutions are optimal in terms of detection error.




Figure 3.5. Linear Hough transform fitting straight lines to a pattern containing four signal edgels (green) and one noise edgel (red)

The advantages of Hough transform include robustness of the fitting against noise and gaps without any assumptions on the number or location of the curves in the image. The classic Hough transform (CHT) is not particularly “intelligent”. Drawing analogy with the chess game strategies, it corresponds to the algorithm that takes the time to generate all possible moves on the board. For this simple reason, the CHT becomes computationally inefficient when the number of model parameters is greater than 2. Because of that restriction, the models that work well with CHT are very simple, and corresponding solutions lack locality. The lines restored from the found solutions span the whole image, and another step is required to locate the actual edgels in the image that belong to the solution and thus identify the start and end points of the line. For this reason, the HT fails to process busy images with short segments.

A number of ideas were suggested to improve CHT computation time. One popular class of enhanced HT is called Probabilistic Hough Transform (PHT) (a good review of PHT can be found in [Kalviainen *et al.*, 1995]), where instead of the exhaustive test of all edgel combinations, a repetitive, random sampling of the distribution is organized, and when a distinct maximum in the histogram appears, the corresponding line segment is immediately identified and the edgels constituting the segment are removed from the edgel pattern, thus conforming to another principle of human vision, inhibition of return [Posner and Cohen, 1984; Itti *et al.*, 1999]. The Random Hough Transform (RHT) [Xu *et al.*, 1990] is a version of PHT that samples the edgel distribution itself. A notable version of HT, called the Importance Sampling Hough transform, ISHT, was suggested by Walsh and Raftery [2001], where the model parameter space is repetitively sampled to verify for a randomly selected subset of model lines the corresponding quality-of-fit (called “importance” by the authors). The histogram of importances is thus obtained and searched for the peak corresponding to the best solution. As in all PHTs, the solution is then identified in the edgel pattern and the corresponding edgels are deleted from the image. The importance sampling operation is repeated in a loop until a stop condition is met. The ISHT is reported to produce smaller errors in the resulting model fit comparing to other PHTs, because it samples the parameter distribution instead of the edgel distribution and therefore is not susceptible to the errors in the edgel localization.

The key problem of Hough Transform method remains unchanged in its improved versions: it does not have a means of adjusting itself to deviations from the rigid model

that it was based on. Another class of techniques exists that employs *deformable* models that conform themselves to the image linear features with greater flexibility.

3.3.2.2. *Deformable Contours*

In contrast to the rigid approach, interaction of the edgels in deformable models is governed not by a particular contour model, but rather by a set of generic perceptual restrictions. These restrictions are known in the literature since the 1930s as the Gestalt principles of perception [Rock and Palmer, 1990, and the references therein]. Continuity and proximity are two principles frequently mentioned in the literature; Figure 3.6, adopted from [Wersing *et al.*, 2001], is a good illustration of two typical grouping tasks performed under the guidance of proximity and continuity constraints.




Figure 3.6. Use of Gestalt principles of perception for grouping (adapted from [Wersing *et al.*, 2001]).

The Gestalt continuity interaction in the lower panel of Figure 3.6 uses an additional characteristic of the edgel, its orientation. We shall briefly discuss the process of inferring edgel orientation for the purpose of the continuity interaction.

3.3.2.3. *Transition from Edgels to Oriented Edgels*

Orientation of the contour at the edgel location is relatively easy to determine with a local context filter. Remarkably, such filter exists in the mammal visual system beyond the retina, in a receptive field of cells in the brain called “primary visual cortex”, or striate cortex, in agreement with Marr’s concept of separating the perception process in the pyramid layers. Individual cells of the striate cortex are sensitive to segments of particular orientation and therefore can be thought of as “bar detectors”. Among models of the cortical cells, the Gabor function is considered the closest equivalent [MacLennan, 1991]. Local estimates of the edgel orientation can also be obtained by other, computationally lighter filters such as Sobel x and y operators [Gonzalez and Wintz, 1987].

3.3.2.4. *Transition from Oriented Edgels to Saliency maps*

First introduced by Sha’ashua and Ullman [1988], the concepts of saliency measure, saliency map, and saliency network attracted considerable attention in the field. The saliency measure is a score calculated for any contour in the given image using principles based on the Gestalt laws of perception:

Continuity:

- the score is higher for longer curves,
- the score is penalized for gaps

Constancy of curvature (*Prägnanz*):

- the score is penalized for changes of curvature

Proximity:

- individual scores from pixels along the curve are added, so that contributions from distant pixels decay with accumulating gaps and curvature changes along the curve.

Effectively, the saliency measure gives a higher score to the long, continuous and smooth lines over the short, discontinuous and wiggly lines. It is a stronger perceptual entity inferred from previously used local measures of the edgel perceptual strength, such as position, gradient orientation, etc. Using the saliency measure, Sha'ashua and Ullman defined a saliency map of an image as another image where each pixel's intensity is proportional to the most salient curve emanating from that pixel. Once the saliency map is available, it is possible to find, for example, the most salient curve in the image that corresponds to the element of the map with the highest saliency value. The most intelligent part of the proposed approach is the *saliency network*, a tool to efficiently calculate saliency maps for images. The saliency network is a dynamic system consisting of locally interacting elements (edgels). Similar to the local relaxation labeling process described previously in this chapter, the saliency network optimizes its state iteratively by a relaxation procedure that is capable of gradually building the contours that conform to a family of "extensible" shapes (such as the circle). The concept of iterative optimization is very important for our approach described in Chapter 4.

Guy and Medioni [1996] reviewed a number of saliency measures, algorithms for their evaluation, and saliency map analysis techniques that were developed since the pioneer work by Sha'ashua and Ullman in 1988. They also introduced their own, iteration-less approach to calculation of the saliency maps, where the saliency score is calculated once as a weighted sum of all contributions from the neighboring edges.

Major contribution to the perceptual saliency approach came from the psychophysical research of pre-attentive vision. Yen and Finkel [1998] built a model of the striate cortex that encompassed many of the findings in this research. Their model is based on interaction between oriented edges represented by the cortical cells. Discovery and analysis of such interactions presents another strong argument in favor of modularity of the vision captured by the Marr's paradigm. According to the model of Yen and Finkel, each cortical cell receives a collective facilitation from a network of other cells in its vicinity. The degree of facilitation is determined individually for each pair of interacting cells using the argument of Gestalt continuity. Figure 3.7 shows the typical interaction pattern for the post-synaptic cell (in the center of the pattern) receiving inputs from the pre-synaptic cells in the shaded areas.

The co-axial areas of the patterns implement the co-circularity constraint that frequently appears in other models for perceptual grouping [Parent and Zucker, 1989, Guy and Medoni, 1996, Galkin *et al.*, 1996]. The co-circularity constraint is in agreement with the Prägnanz Gestalt principle of curvature constancy. Biophysical studies show that co-circularity is enforced only in the long-range co-axial pattern, whereas the trans-axial pattern features simpler requirement of parallelism.

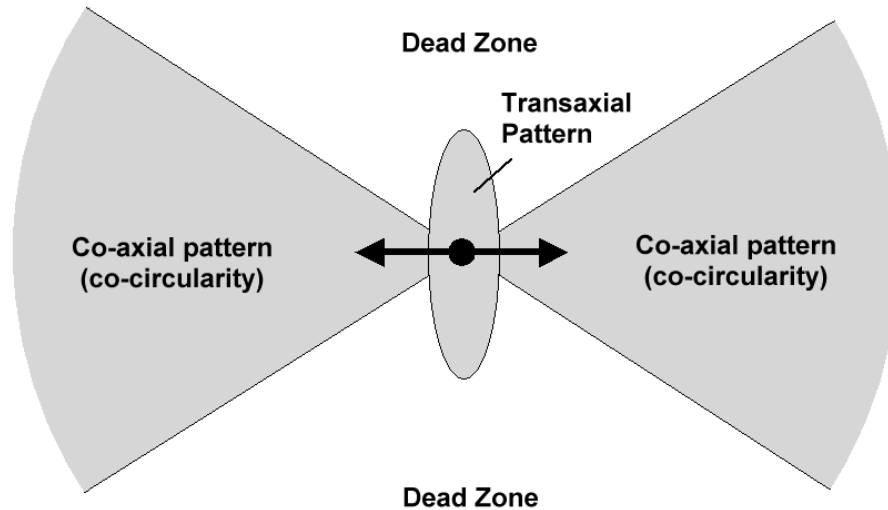


Figure 3.7. *The interaction pattern of the oriented cortex cells [Yen and Finkel, 1998] in the model of pre-attentive human vision. The post-synaptic cell in the center of the pattern receives facilitating inputs from the pre-synaptic cells in the shaded areas. Degree of facilitation depends on the distance between cells and their mutual orientation.*

The saliency measure is typically based on the local perceptual strengths of the edgels such as their orientation evaluated with the help of a local-context steerable filter. Alter and Basri [1998] demonstrated high sensitivity of the saliency evaluation to discretization effects, which makes this approach highly susceptible to the edgel jitter that causes incorrect estimates of the local orientation. Indeed, errors in evaluation of the post-synaptic edgel orientation (Figure 3.8) cause misalignment of the whole long-range interaction pattern. These considerations warrant attention to those techniques in which the edgel orientation is an integral part of the global combinatorial optimization scheme, so that a greater robustness to the discretization problems can be achieved. The task of

perceptual grouping was approached in terms of an evolving dynamic system, where constituting parts interact under certain restrictions, an artificial neural network (ANN).

Dynamic interaction of the neurons in the neural network is similar to the iterative relaxation optimization in the saliency networks. However, the ANN models are better studied and have a good biological counterpart, the brain. Each neuron in the network has connections to many others contributing their outputs for the neuron's analysis (Figure 3.8).




Figure 3.8. Hopfield model of artificial neuron,

The artificial neuron is a simple element that constantly evaluates the weighted sum of inputs I_j from other neurons to decide what output it shall assume [Hopfield, 1982]:

$$x = \sum_{j=1}^N W_j I_j, \quad (3.6)$$

where W_j are the synaptic weights that scale inputs I_j from other neurons to produce the summary output x . This feature of the ANN is in a very good agreement with the concept

of establishing a cumulative context in support of the grouping decisions. The synaptic weights capture the restrictions under which the edgels interact (in our case, the Gestalt principles). The system of neurons starts interacting until they reach a stable state. A major contribution to understanding of the dynamic processes in ANNs is due to Hopfield who described the evolving process in terms of an energy function that always decreases. The network design process can be then thought of as “energy engineering”, where the energy function is written down so that it becomes minimal for the desired outcome of the optimization. The energy function is then used to derive the rules governing the evolving of the neural network.

The ANN model shown in Figure 3.8 employs an additional mechanism in the evolving rules for the neurons that allows the network to avoid local minima of the energy function on its way to the global minimum. Using a statistical treatment of neuron states known as the “mean field theory”, or MFT [Peterson and Anderson, 1987], it became possible to induce regulated amounts of thermal noise into the evolving rule so that it escapes the local minima. In the Hopfield network the neuron output, O , is obtained by processing the weighted sum of inputs, x , with a sigmoid function:

$$O = \tanh\left(\frac{x}{T}\right), \quad (3.7)$$

where T is the noise temperature. Larger temperatures correspond to larger amount of noise, and the evolving procedures benefit from various rules of the simulated annealing, where the temperature decreases as the network approaches its global minimum of the energy function.

One of the first energy-minimization techniques to be applied to the contour extraction problems was the segment model introduced independently by Peterson [1989] and Denby [1988]. The segment model is based on a “spin” network where neurons accept only a discrete set of states. In this segment model, the spin variable s_{ij} denotes whether two edgels i and j are connected by a contour segment of length r_{ij} . The energy function describing this simple idea is (see Figure 3.9):

$$E = -\frac{1}{2} \sum_{ijkl} \delta_{jk} \frac{\cos^m \theta_{ijl}}{r_{ij} r_{jl}} + \frac{\alpha}{2} \left[\sum_{ik} s_{ik} s_{kl} + \sum_{jl} s_{ij} s_{jl} \right] \quad (3.8)$$




Figure 3.9. Segment model for contour extraction by [Peterson, 1989] and [Derby, 1988].

The first term of the energy function imposes continuity and smoothness constraints on the way connected segments are combined into lines. It favors adjacent segments s_{ij} and s_{jl} that are short (small r_{ij} and r_{jl}) and aligned (small angle θ_{ijl} between r_{ij} and r_{jl}). The second term punishes bifurcated lines. Figure 3.10 shows the results of stringing by Peterson’s segment model for a family of circular traces with no noise, where N_{sweep} denotes the number of neural network iterations.

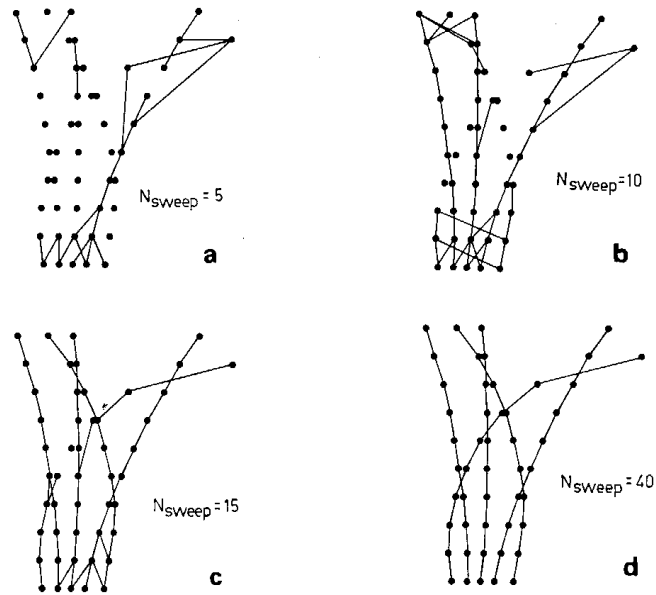


Figure 3.10. Edgel grouping with Denby-Peterson segment model, an example from [Peterson, 1989]

This segment model could not produce a reliably converging network configuration because of the cumbersome formulation of its energy. Another approach was suggested by Peterson [1990], where a “rotor model” was introduced as shown in Figure 3.11. In this approach, the network has control over the length and orientation of rotors placed on top of each edgel. The energy function of Petersons’ rotor model is minimal when rotors are aligned with a trace line (see Figure 3.12):

$$E = - \sum_{i \neq j} \frac{1}{\rho_{ij}^m} V_i V_j (\cos \alpha + k \cos \beta) \quad (3.9)$$




Figure 3.11. Rotor model network evolves, aligning rotors along the contour.

where ρ_{ij} is distance between points i and j , α is angle between rotors \mathbf{V}_i and \mathbf{V}_j , β is angle between rotor \mathbf{V}_i and the line connecting two points, and m, k are algorithm constants.




Figure 3.12. Linear trace model for rotor NN [Peterson, 1990]

The first term of the equation 3.9 forces the rotors to turn parallel to each other and the second term aligns them along the trace segment. The energy function favors the rotors pointing in the same direction as the straight line between them, and therefore has a tendency to favor straight lines. Its efficiency degrades when traces have significant

curvature. Peterson's model was successfully generalized to circular traces by Baginyan *et al.*, [1994] (Figure 3.13 and Eq. (3.10)):




Figure 3.13. Circular model for rotor network by Baginyan *et al.*, [1994]

$$E = - \sum_{i \neq j} \frac{1}{C_{ij}^m} V_i V_j \cos \alpha_{ij} \quad (3.10)$$

Here α_{ij} is angle between rotors \mathbf{V}_i and \mathbf{V}_j' , where \mathbf{V}_j' is obtained by flipping \mathbf{V}_j around the chord \mathbf{C}_{ij} . The circular model of rotor interaction turns rotors tangential to the trace curve.

The concept of energy minimization provides a better foundation to the deformable contour approach. The deformable contour can now be thought of as a dynamic system described by its energy function with three terms:

$$E = E_{\text{int}} + E_{\text{img}} + E_{\text{ext}} \quad (3.11)$$

where internal energy E_{int} imposes smoothness and continuity constraints on the contour, image energy E_{img} attracts the contour to the edgels, and external energy E_{ext} moves and

stretches/shrinks the contour. This dynamic system seeks optimal balance between the three terms.

A great variety of energy functions have been engineered for the purpose of contour extraction, appearing in the literature under such names as “snake”, “elastic arm”, and “active contour”. These algorithms can be thought of as dynamic optimization systems seeking a balance between fit of the contours to the image features while preserving their smoothness and continuity. In the absence of external force, the contour balances the force that pulls it to the edgels and the elastic force that enforces local smoothness of the contour and global conformation to the contour model, if it is available. The balance can be biased to either side by using regularization coefficients $0 < \lambda_i < 1$ (Lai, [1994]):

$$E_V = \sum_{i=1}^n \lambda_i E_{\text{int}}(v_i) + (1 - \lambda_i) E_{\text{img}}(v_i) \quad (3.12)$$

Here the contour is a set of edgels $\{v_i\}$, $i = 1..n$. Figure 3.14 illustrates how choice of λ changes the shape of the contour, where higher values of λ results in a smooth line that cannot capture sharp corners, and lower λ make the line follow the sharp turns but is too sensitive to noise.

The contour without external force that can change its location, orientation or length needs to start with a good initial configuration to optimize it in accordance to the energy formulation. The initial configuration can be obtained by some other technique, including manual specification of the seed edgels. The common automated solution to the task of finding the number of lines in the image and initialize the “snakes” is to use rigid

contours extracted via the Hough transform and then evolve the deformable contour around each of them.




Figure 3.14. Effects of λ parameter in the contour energy: (a) smooth line that does not capture corners, (b) good capture of corners but high sensitivity to noise, (c) compromise. From [Lai, 1994].

A deformable contour that does not contain the internal energy term but allows more freedom in terms of the number of lines and their expected shape was suggested by Ohlsson [1993] to solve the task of grouping N edgels in M lines:

$$E = \sum_{i=1}^N \sum_{a=1}^M S_{ia} M_{ia} + \lambda \sum_{i=1}^N \left(\sum_{a=1}^M S_{ia} - 1 \right)^2 \quad (3.13)$$

Here the S_{ia} is the binary decision term which is 1 if edgel i is assigned to line a , and 0 otherwise. The first term of energy E in equation 3.13 specifies the image force acting on the contour by calculating the cumulative deviation of edgels from the line (M_{ia} is the squared Euclidean distance between point i and trace a). The second term is external force imposing a penalty for edgels not included in any line (i.e., the edgels with

$\sum_{a=1}^M S_{is} = 0$). The deformable contour described by Eq. (3.5) was found to have difficulty reliably evolving into the global minimum of energy in presence of noise and gaps. Muresan *et al.*, [1997] improved the convergence by setting the contours' initial configurations to the rigid contours obtained by the Hough transform.

3.3.2.5. *From saliency map to contours*

In the absence of a top-down supervision, analysis of the saliency map is driven by the map features called “activation spots” that stand out against the background. A model architecture for this process was proposed by Koch and Ullman [1985] based on a study of human visual psychophysics by Posner and Cohen [1984]. To extract more than one activation spot that may be present in the saliency map, an iterative procedure attends the spots sequentially and then blocks them from further search. Itti *et al.* [1999] built a visual attention system that implements this concept using the “winner take all” (WTA) neural network for selection of most salient features in the saliency map. The WTA neuron is a dynamic “integrate-and-fire” unit that constantly sums its synaptic input and compares to a threshold value. The WTA network is placed on top of the saliency map and starts the integration process; the neuron at the map location of the highest saliency fires first, while all others remain suppressed. This event triggers a switch of the focus of attention (FOA) to the winning location and activation of the “inhibition of return” mechanism demonstrated in human visual psychodynamics [Posner and Cohen, 1984] that resets WTA network and suppresses the saliency map at the winning location from subsequent analysis.

3.4. Outlook

A range of early vision techniques exists with good potential for inferring traces from the RPI plasmagrams. In reviewing them, a greater attention was paid to approaches of biological plausibility, primarily because of the fact that human vision solves the task. There is a class of models that comply with Marr's paradigm, a bottom-up archetype that builds the visual target out of simpler elements of increasing perceptual strength. There is an abundance of evidence that mammal pre-attentive vision conforms to the Marr's paradigm. In the next chapter we will analyze how the task of plasmagram trace recognition fits into the frame of pre-attentive vision models and discuss the measures taken to improve the robustness of the processing.

CHAPTER 4. TRACE RECOGNITION

An increasing number of the real-world applications can be delegated to an artificial vision system. Searching for the traces in the RPI plasmagrams is one such application presenting particular importance to the RPI mission team and space research community at large. This chapter discusses the unique features of the plasmagram images and describes the pre-attentive vision model for their processing.

4.1. Uniqueness of Plasmagram Signatures

The task of plasmagram trace recognition presents a unique challenge to existing methods because of a combination of factors. Fundamentally, there is little *a priori* information available on the possible occurrence and number of the traces, and there is no simple function that models their shape. Thus, even though more knowledge is collected about plasmagram traces, the trace extraction algorithm is still unable to use top-down considerations at this time. In the bottom-up, pre-attentive approach to the recognition task, the vision model regards the plasmagrams as non-specific images with lines. Furthermore, the plasmagrams are collected in an unknown, tremendously variable environment using a low-power probing signal that often arrives near or below the background noise level associated with natural radio emissions. The need for a trace extraction algorithm to be adaptive to a range of signal-to-noise ratios (SNR) makes the recognition task unique and difficult.

4.1.1. Variable Signal-to-Noise Ratios

The echo trace is often weak, as in Figure 4.1 showing an example plasmagram taken on June 29, 2001 03:13 UT where a faint trace can barely be seen spanning frequencies from 65 to 120 kHz. Besides the trace being globally weak, it can display various levels of SNR within a single plasmagram due to differences between noise characteristics on different frequencies and changing conditions for the signal propagation depending on orientation of the RPI antenna system that slowly rotates in space.

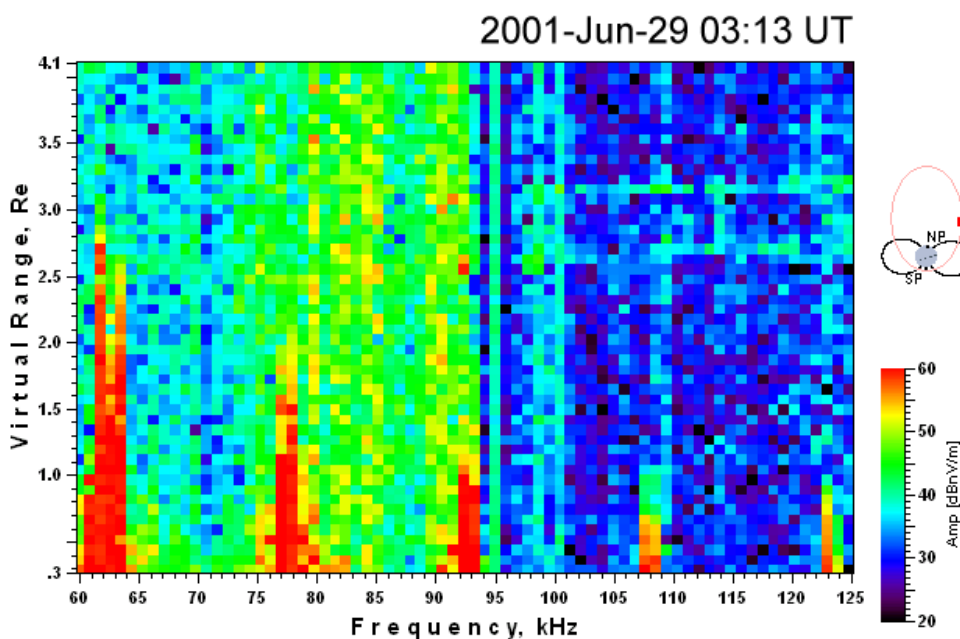


Figure 4.1. Plasmagram taken on Jun 29, 2001 03:13 UT containing two traces with a low signal-to-noise ratio.

Figure 4.2 shows a plasmagram taken on June 29, 2001 03:00 UT, where some parts of the trace are missing whereas others are well-defined.




Figure 4.2. Plasmaprogram taken on Jun 29, 2001 03:00 UT showing a trace of variable signal-to-noise ratio across the frequencies.

Noise characteristics vary substantially not only from frequency to frequency, but also with spacecraft location and time. Global events in the solar-terrestrial system cause frequent increases in the natural radio emissions in space (e.g., auroral kilometric radiation, solar radio bursts). These emissions often have highly irregular structure causing numerous false positives of the differentiation operator in the edge detector. Figure 4.3 presents a plasmaprogram taken during one of such event on Jul 9, 2003 18:24 UT with an increased level of noise in the frequency band between 210 and 260 kHz. Each frequency within the band contains multiple pseudo-echo signatures resulting from

the highly variable structure of the noise as shown in the Panel 4.3(a) in detail; the echo detector produces a large number of false positives for this case (b).

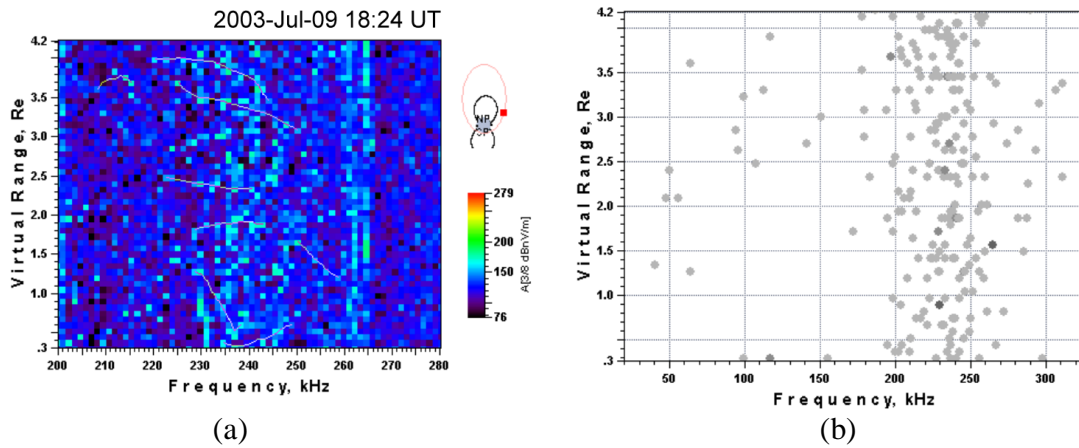


Figure 4.3. *Plasmagram taken on Jul 9, 2003 18:24 UT during high radio emission activity in the frequency band between 210 and 260 kHz (shown in detail in panel (a)). The emissions cause multiple false positives of the echo detector (b).*

Low SNR results in false positive and negative rates of the edge detector. The problem of trace gaps is at times severe, and faint traces can mix with stronger traces in the same plasmagram because certain propagation modes experience signal attenuation because of increased absorption in plasma or unfavorable orientation of antenna. Inconsistent trace quality makes local grouping techniques inapplicable because of the need to process edgels over a large context scale for proper identification.

4.1.2. Range Jitter

In contrast to the airport radar whose echoes are coming from well-defined targets, the RPI echo is a result of reflection from a relatively large area of a complex structure exhibiting dynamic movements and irregularity of various scales. A single echo detected by RPI may be a combination of multiple echoes arriving with similar propagation times, where the overlapping causes distortions of the echo envelope. In this case it is difficult to determine its leading edge precisely, and it is not unusual for the resulting trace line to display a substantial range jitter (deviation up or down from the trace line).

This range jitter causes errors in the rigid contour approaches and grouping techniques that employ the Gestalt principle of smoothness. Figure 4.4 illustrates performance of the linear Hough transform applied to an edgel pattern consisting of four echoes comprising the straight line and one noise spike (left panel) and to the same pattern affected by a range jitter (right panel). Because edgel positions are used for fitting directly, the corresponding histogram peaks are washed out.




Figure 4.4. Hough transform applied to low (left) and high (right) range jitter patterns

Muresan *et al.* [1997] studied the washout problem to conclude that if edgels deviate from the trace line, the HT is applicable only as a pre-processing algorithm that derives an approximate number of trace segments and a rough evaluation of their position.

Range jitter causes errors in evaluation of the local contour orientation for the edgels, subsequently resulting in misalignment of the long-range interaction pattern. Figure 4.5 shows a simulated edgel pattern illustrating this effect.




Figure 4.5. Simulation of range jitter causing misalignment of the edgel interaction pattern and low salience

In this example, two edgels in the middle of the pattern are slightly misplaced from the trace line and data gaps are inserted on both sides. Local context algorithm produces wrong orientation of contour at both edgel locations. The saliency measure is then made very low for both edgels, a direct consequence of the range jitter and data gaps.

4.1.3. Closely Spaced Traces

Another problem for conventional saliency measure approaches is the need to resolve plasmagram traces that lie close to each other. If one of the traces is short, its saliency measure will not be comparable to its longer neighbor's, resulting in the complete loss of this shorter trace in the map as the point of attention attends the strongest saliencies in the image. In plasmagrams, such a shorter trace typically corresponds to the O-polarization waves that accompany the stronger X-polarization trace.

4.1.4. Range and Frequency Resolution Artifacts

Most of the time, the frequency and range resolution of plasmagram images are insufficient to capture traces adequately. The RPI operates at very low pulse repetition rate because of the need to record echoes arriving at great distances. Thus there is an upper limit on the number of probed frequencies to keep total plasmagram measurement time reasonable to satisfy a requirement of high cadence of measurements as the spacecraft orbits the Earth. With the number of frequencies kept under control, the need for the RPI to accommodate a broad range of possible scenarios in the environment means that wider frequency coverage is always chosen over better frequency resolution. Another important requirement limits the telemetry data volume, which decreases the number of range bins and makes the range resolution coarser. Thus the traces are commonly thin and often are just one pixel wide. Conventional smoothing filters that protect the edge detector from noise can damage thin trace signatures.

4.2. Approach

Range jitter, data gaps, false echoes, and nearby traces are the reasons for frequent problems with calculation of local edgel orientations and subsequent saliency analysis techniques that are based on the static oriented patterns of edgel interaction [Guy and Medioni, 1996, Yen and Finkel, 1998]. The saliency evaluation has to allow edgel orientations to be modified by the global scale process, because the problems with local estimates can only be identified on the global scale. The saliency calculation then becomes an optimization process that refines the local orientations. An additional layer is therefore introduced in the Marr's pyramid for the oriented edgels that can change their orientation, *rotors* (as in Peterson's rotor model [1989]). In contrast to the oriented edgels, the rotors have freedom to rotate depending on facilitation from other rotors in the vicinity. As every change of a rotor modifies the degree of facilitation, the whole optimization process is made iterative and thus is best described in terms of the energy minimization (as discussed in the Chapter 3). The interactions between the rotors are not forced to comply with any specific trace model, but rather are governed by the general Gestalt laws of perception, plus some other considerations as discussed later in the chapter.

Optimizing the rotor pattern can be computationally demanding, considering that number of single rotor interactions to evaluate in a fully connected network grows with the number of rotors as $O(N^2)$. This is the reason why Gestalt-based saliency calculations based on the fully connected iterative neural networks are rather exceptional. A number of measures were devised to decrease the computational difficulty. In particular, we use

the fact that RPI echoes are recognized only if they are above the noise level, and so we apply echo detection techniques to suppress irrelevant information from entering the saliency calculations. These questions are further discussed in the Section 4.3.

Figure 4.6 shows extended version of the Marr's pyramid for trace extraction in the plasmagrams. Below we describe the perceptual elements on all levels of the pyramid in relation to RPI plasmagrams.

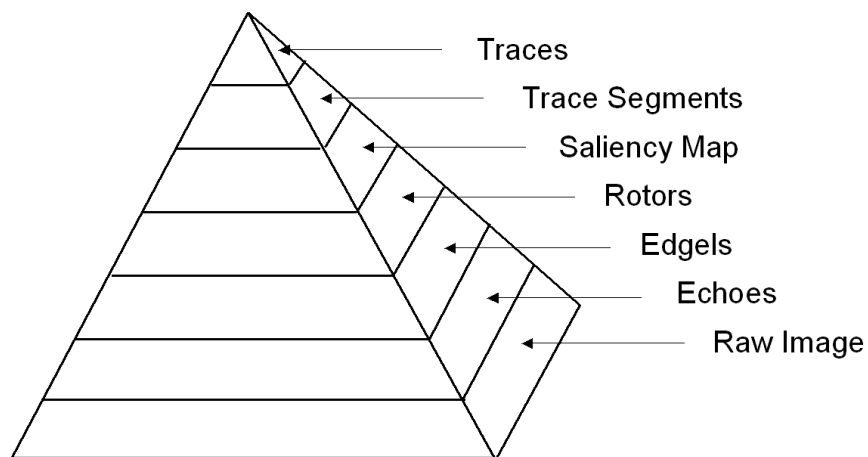


Figure 4.6. *Extended Marr's pyramid of perception for plasmagram processing.*

Raw Image. Raw image is an ($N \times M$) matrix of pixels, each pixel representing the RPI antenna voltage.

Echoes. The first stage of processing labels the pixels of the plasmagram image that potentially belong to the RPI signal. The labeling is accomplished using an adaptive thresholding algorithm, a frequency-scanning radar echo detection technique that

evaluates the threshold level for each frequency individually using an estimate of background noise and then labels those pixels that exceed the threshold levels as echoes.

Edgels. The range interval taken by an echo is reduced to a single point corresponding to its time of arrival (TOA). This operation is very similar to selection of the edgels, but using only the leading edges of the echoes. Later in the text the TOA points are referred to as edgels (edge elements) to indicate that they constitute the leading edge of a trace.

Rotors. Classic early vision model for contour extraction suggests *oriented edgels* as the next degree of perceptual strength in the pyramid (see Section 3.1) The edgel orientation is commonly obtained by seeking gradients of image intensity within a local context window around the edgel. Local context techniques, however, do not yield reliable edgel orientations in various situations requiring a larger context area to correctly identify the trace. Orientation of edgels is modified by a global-scale optimization process. Rotors are oriented edgels that can rotate under facilitation from neighboring rotors.

Saliency map. Long-range, collective analysis of rotor data results in evaluation of perceptually stronger quantity, *saliency*. The saliency measure reflects how likely an edgel is part of a trace. Saliency measures calculated for all edgels constitute the saliency map of the image.

Segments. The saliency map is analyzed to find subsets of edgels belonging to the same trace. This is a bottom-up procedure typical of the pre-attentive vision, and the

found subsets may correspond to the segments of traces to be further grouped into complete traces.

Traces. Perceptual grouping of segments into traces corresponds to the attention-driven recognition, where certain assumptions are made about the model of the trace shape. The grouping algorithm switches attention between found pre-attentive vision cues to evaluate their support of the model.

The next section discusses the questions of reducing the number of edgels before the optimization of the rotor alignment starts.

4.3. Transition from Raw Image to Edgels

Direct application of the conventional leading edge detectors to the plasmagrams would cause a severe false positive rate because of their sensitivity to every gradient of intensity in the image. Robust edgel detectors explore various possibilities to distinguish signal from noise. Common approaches to the problem (e.g., [Ziou and Tabbone, 1998]) involve smoothing to remove the noise variability, thresholding to suppress weaker noise, and local tests of the signal integrity across adjacent pixels. Most of these commonly used measures are damaging, to varying degree, on the signal in plasmagrams. In contrast to the typical scenarios of object identification in visual scenes, the plasmagram traces are thin and faint signatures in the noise background of irregular and varying nature.

Figure 4.7 presents an example of the leading edge detection in the plasmagram taken on March 2, 2002, 04:46 UT.

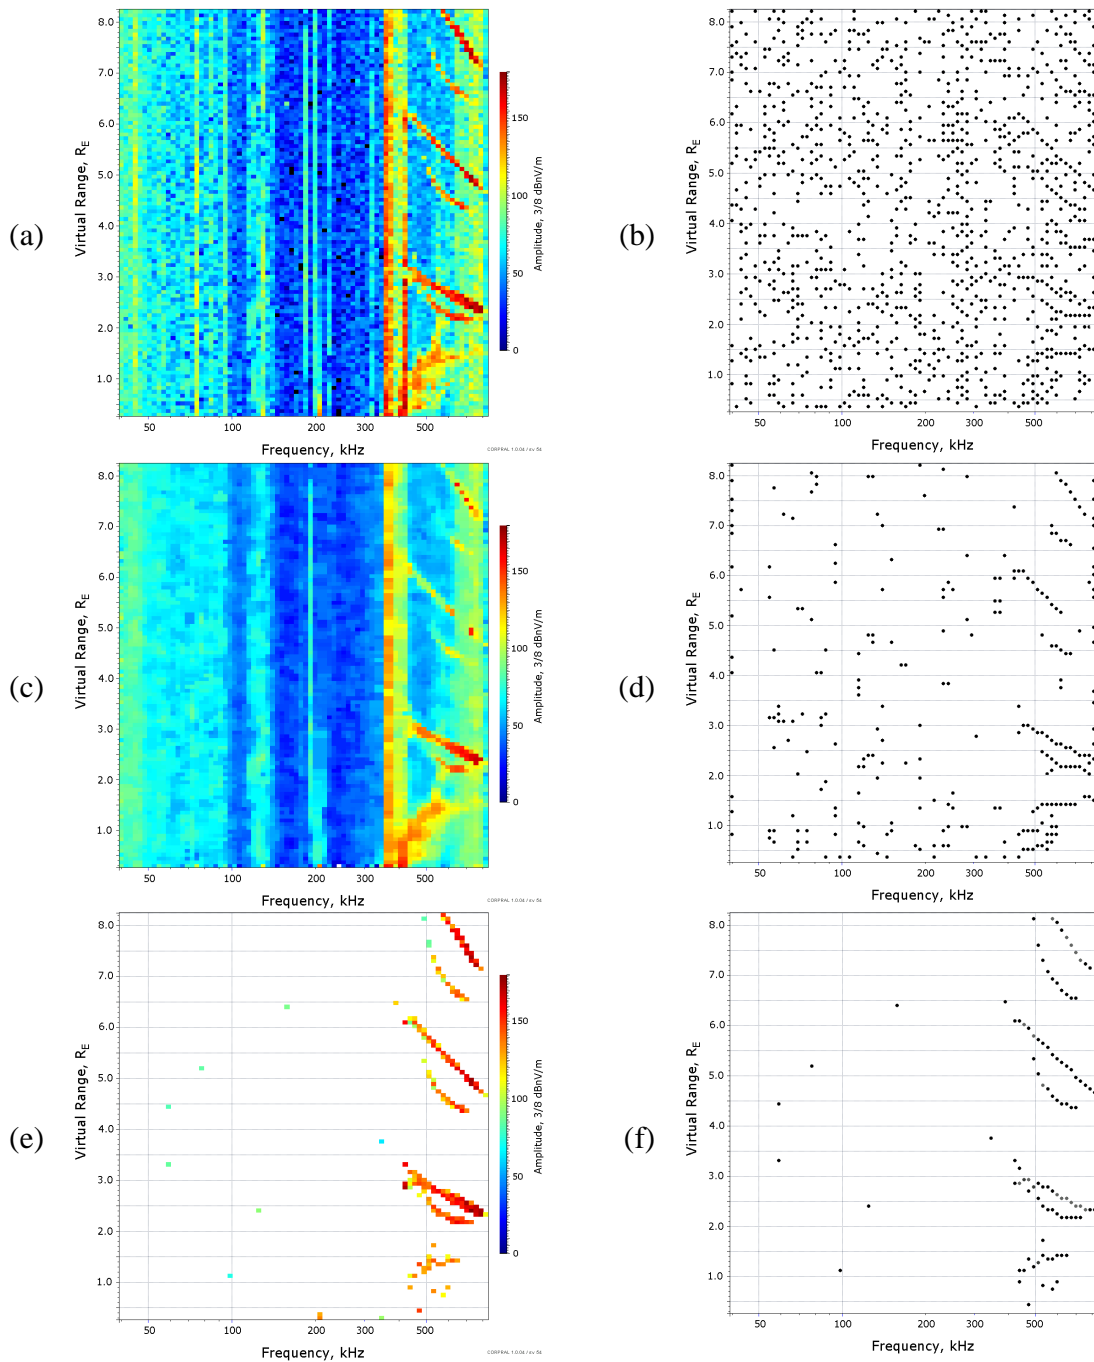


Figure 4.7. Edgel labeling for an example plasmagram (a) taken on March 2, 2002, 04:46 UT. Direct edgel detection (b) produces a large number of false edges even with the thresholding of the amplitude gradient. (c) Application of the smoothing 3×3 median filter prior to edgel detection reduces the noise edges but damages thin traces (d). The chosen approach (e) uses the adaptive local thresholding to detect echoes and then labels the leading edges (f).

Panel 4.7(a) shows the raw plasmagram image and the direct detection results obtained by the “zero-crossing” double differentiation operator that seeks the point of largest intensity gradient, Panel 4.7(b). Even though found amplitude gradients that are smaller than 6 dB were suppressed, the background noise captured by the RPI receivers still produce an overwhelming number of false edgels in Panel 4.7(b). The second row of the Figure 4.7 illustrates common effects of a smoothing filter that protects the differentiation operator from noise. The plasmagram image is smoothed with a 3x3 median filter (Panel 4.7(c)), and then the edgel detection is done with the same algorithm as before. The results shown in Panel 4.7(d) indicate that the amount of smoothing is still not sufficient to suppress enough noise, whereas the thin traces in the upper frequency band are already damaged.

The specific nature of the remote sensing data opens an opportunity for noise protection without smoothing, known in radar literature as *echo detection*. In remote sensing, the useful information comes from the signal returns (echoes) that are overlaid with the background noise. Remote sensing systems are designed to ensure that the echoes are detectable in the variable noise environment. A suite of adaptive (robust) methods has been developed to dynamically estimate the characteristics of noise background and apply *thresholding* to appropriately suppress the noise, leaving the echoes intact (see, for example, Schleher [1980]). The adaptive echo detection techniques relies on a statistical approach, where the probability distribution function of noise is often assumed known and its parameters are estimated from the data, for each frequency individually. The class of detection techniques that adjust their threshold level as noise

characteristics vary is often called “constant false alarm rate” (CFAR), from the early days when they were successful in relieving radars from excessive false alarms during jamming or periods of higher interference. Figure 4.8 gives an example of the adaptive echo detection algorithm used in Digisonde [Reinisch, 1996], which calculates the most probable value of the amplitude distribution to estimate the noise level.

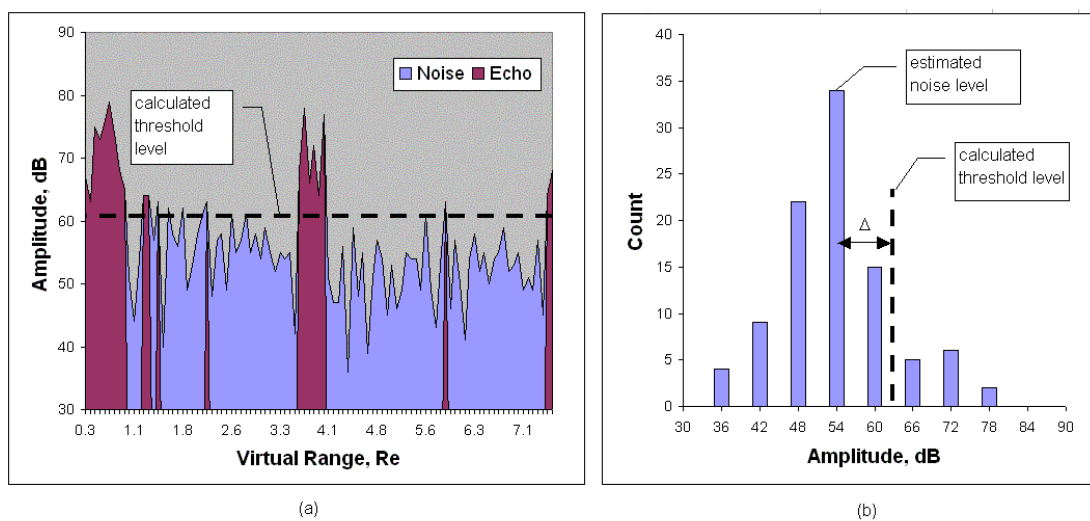


Figure 4.8. Echo detection by adaptive amplitude thresholding. (a) One frequency scan with intervals taken by echoes, (b) amplitude histogram used to determine the threshold level.

The echo detection algorithm that selects the mode of the amplitude distribution as the noise level assumes that this noise level is the same for all ranges, which is not necessarily true in case of the RPI sounding. We use an adaptive detection algorithm “AvTrend” [Galkin *et al.*, 2004] that uses a short-length surround window that slides over the ranges (see Figure 4.9). The AvTrend algorithm labels a pixel i as echo if its amplitude A_i exceeds a threshold value T_i calculated over the surround window of size N placed at the tested pixel. The threshold value is set to the average amplitude within the

window (excluding the pixel i itself) plus a fixed value D . The algorithm is designed to detect signals above the noise level that are narrow enough to fit within the surround window and leave room for the background amplitudes that produce the threshold value. The one-dimensional window is placed on the plasmagrams vertically to avoid influence of the neighboring frequencies that may have significantly different levels of noise and signal. The third row of Figure 4.7 shows the plasmagram processed by this algorithm that eliminates background noise without the artifacts of smoothing.

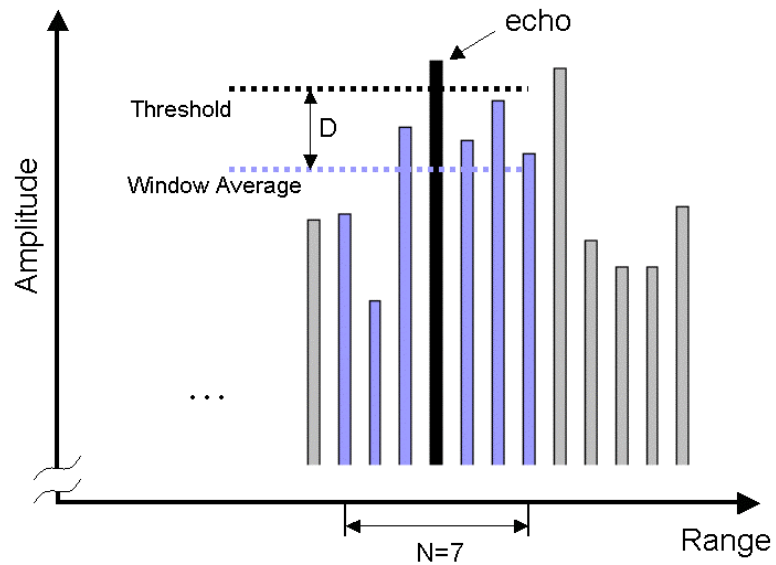


Figure 4.9. Example processing with AvTrend echo detection algorithm. Sliding window of $N=7$ points is placed at the range bin shown in black. Window average amplitude is calculated over 6 points and the threshold is set D units higher. The range bin in this example is classified as echo because its amplitude exceeds the threshold.

4.4. Transition from Edgels to Rotors

The orientation of edgels is obtained using a steerable narrow sector placed on top of the edgel. After a number of sector orientations is tested, the edgel orientation is set along the direction of maximum number of other edges found within the sector. The actual steering algorithm implements the so-called angular histogramming technique [Bagynyan *et al.*, 1994], which is a version of the linear Hough Transform (HT) reworked for a higher computing efficiency. The angular histogramming algorithm samples the edgel distribution itself instead of the parameter space as the classic HT does by fitting trial straight lines through all pairs of the base edgel with the neighboring edges (Figure 4.10). The elevation angle of the trial lines is binned to obtain a histogram of angles and select the angle of highest occurrence.

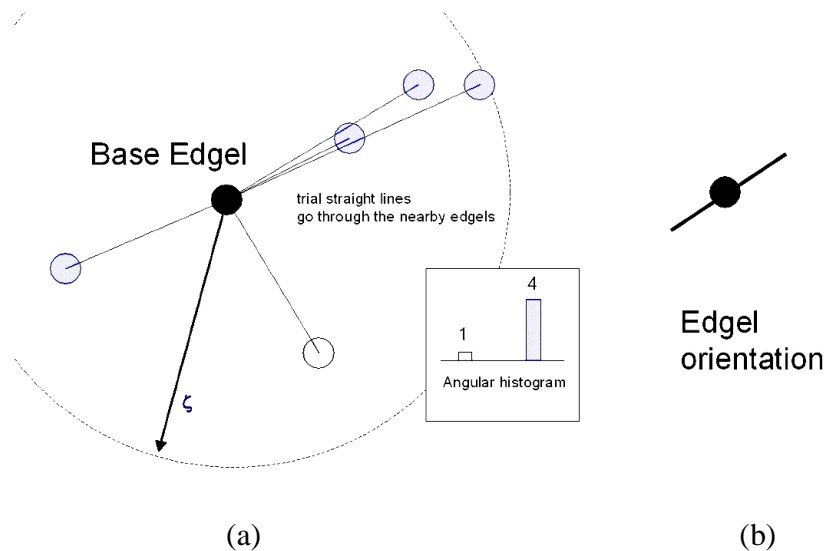


Figure 4.10. Use of angular histogramming algorithm [Bagynyan *et al.*, 1994] to evaluate edgel orientation (example). (a) Angular histogram is built from elevation angle of all trial straight lines going through the base edgel and all nearby edgels within the circle of radius ζ . (b) The angle of highest occurrence in the histogram is selected as the edgel orientation.

4.5. Transition from Rotors to Saliency Map

Calculation of the saliency score for an edgel (see discussion in Section 3.3.2.4) involves counting contributions from many other edgels in its vicinity. The score calculation is commonly done for the static oriented edgels that do not change their orientation during the process. This scheme can benefit from the iterative optimization approach of the rotor models [Peterson, 1990], where the rotors have freedom to rotate seeking the best saliency score. Figure 4.11 shows how a single neuron in the neural network receives facilitating inputs from other neurons in its vicinity. The synaptic weights W_i are not constant; instead, they are calculated dynamically depending on distance and mutual orientation of two interacting rotors. The optimization process corrects errors of local estimates of the initial rotor orientations by extending the context area to a larger scale where these errors are visible.




Figure 4.11. Artificial neuron for rotor interaction.

Previous implementations of the rotor model optimizers in high-energy physics [Bagynian et al., 1994] and ionospheric sounding [Galkin et al., 1996] were able to demonstrate successful operation on the imagery data of rather consistent quality and content. We first studied performance of the conventional Peterson model with the co-circularity criterion introduced by Bagynian *et al.* [1994] on the RPI plasmagrams and then discuss improvements to the optimizer design to better handle specifics of the task. We will refer to the improved optimizer for RPI plasmagram processing as ANNA (Artificial Neural Network Algorithm).

4.5.1. Co-circular Rotor Model Neurodynamics

The neural optimizer by Bagynian et al. [1994] discussed in Section 3.3.2.4 reaches the global minimum in its energy function when the rotors are aligned tangential to the circle going through the interacting rotors. This is equivalent to the Gestalt principle of co-circularity, so we will refer to this rotor model as “co-circular”. Figure 4.12 illustrates this concept with a simulated edgel pattern where 8 edgels are placed on a circle.




Figure 4.12. Processing of a synthesized edgel pattern by a circular model of rotor interaction: (a) edgel pattern, (b) initial rotor placement by angular histogramming, (c) results of ANNA optimization.

Left panel of the figure shows the pattern with randomly chosen initial orientations of the rotors, and right panel present the results of rotor optimization by the co-circular model, obtained after 25 iterations with the constant temperature of 2.

For convenience, we repeat here Figure 3.13 illustrating the algorithm of calculating facilitation contributions in the co-circular rotor model:




Figure 4.13. Circular model for rotor network by Baginyan et al., [1994]

The strength of interaction is maximal when both edges are tangential to the circle connecting edges i and j . The circle is determined uniquely for each pair of edges i and j using the location of edges and orientation of the post-synaptic rotor V_i . Facilitating contribution of pre-synaptic V_j on post-synaptic V_i is:

$$L_{ij} = \frac{V'_j (\cos \alpha_{ij})^a}{|C_{ij}|^b} \quad (4.1)$$

Here α_{ij} is angle between rotors \mathbf{V}_i and \mathbf{V}_j' , where \mathbf{V}_j' is obtained by flipping \mathbf{V}_j around the chord \mathbf{C}_{ij} . Use of the V_j' factor assures that longer vectors give higher contributions, whereas factor $(\cos \alpha_{ij})^a$ enforces co-circularity of the vectors. Finally, $\frac{1}{|C_{ij}|^b}$ factor attenuates facilitation of more distant edgels. Parameters a and b regulate how fast the contribution falls with misalignment and distance.

Study of the co-circular optimizer performance on the random initial orientation of the rotors identifies a stable false minimum of its energy function where the rotors are aligned normal to the circle instead of tangential. As can be seen from the Figure 4.14, the flipped vector \mathbf{V}_j' aligns with the post-synaptic vector \mathbf{V}_i perfectly, so that rotors deviating from this orientation will be pulled back.




Figure 4.14. Rotor configuration corresponding to the strong false minimum of the Baginyan rotor model energy.

Because of the strong false minimum in the energy function of the Baginyan rotor model, it is important for the initial rotor orientation to be as close to the optimal configuration as

possible; otherwise the rotor network may evolve into the false minimum. It is also important to properly design the annealing scheme so that the mean-field theory (MFT) mechanism makes it possible for the network evolving to tunnel through the barriers on the way to the global minimum of the energy.

It was not possible to use a fixed temperature or a static annealing scheme for processing of plasmagram data due to the fact that number of interacting rotors varies substantially from plasmagram to plasmagram. In the MFT evolving of the neural network, the sum of weighted inputs in the neuron goes through the sigmoid function (see also Figure 3.8):

$$O = \tanh \left(\frac{\sum_{all\ i} W_i I_i}{T} \right) \quad (4.2)$$

where O is the post-synaptic neuron output, I_i are the pre-synaptic neuron inputs, W_i are synaptic weights, and T is annealing temperature. Figure 4.15 shows the sigmoid function chart illustrating how inappropriate choice of temperature T may lead to saturation if the argument of hyperbolic tangent is greater than 1 or smaller than -1. If neurons go into saturation too early in the evolving process (starting temperature too small), the network loses its ability to tunnel through the barriers of the energy function, in which case it is likely to stop in a local minimum of the energy function. If the starting temperature is set too high, neuron outputs become insignificant at the first step of evolving and may remain small during the course of annealing, so that the final saliency measure is too small for the segment extraction algorithm to consider it.




Figure 4.15. Sigmoid function of the MFT neuron with region of optimal performance as indicated.

To keep the neural network in the optimal operating range, ANNA annealing scheme uses estimates of the edgel density to calculate starting and ending temperature of annealing individually for each plasmagram.

4.5.2. Improved Rotor Interaction Pattern

The co-circular rotor interaction model (Eq. 4.1) does not work well at short distances due to the range jitter in the edgel data. Figure 4.16 shows a simulated example of co-circular processing for a set of 5 edgels belonging to a straight horizontal line segment. The range of edgel j was overestimated because of distortions in the echo shape, which resulted in a weak contribution of \mathbf{V}_j to \mathbf{V}_i as the angle α_{ij} between flipped rotor \mathbf{V}_j and \mathbf{V}_i is too large.




Figure 4.16. Close-range co-circular interaction of rotors in range jitter conditions (simulated example). Due to the overestimated range of edgel j , contribution of rotor V_j to rotor V_i is too small.

Poor performance of co-circular interaction model at short distances warrants introduction of a different model of interaction for rotors that are close to each other. The same idea can be found in other vision models (see Section 3.3.2.4) where the short distance interaction pattern is different from the long-distance pattern. For example, the Yen and Finkel [1998] model (Figure 3.7) contains a transaxial pattern for short-range interaction where simple parallelism constraint is used instead of the co-circularity. Similarly, ANNA uses the parallelism constraint for facilitation at short distances:

$$S_{ij} = \frac{V_j}{|C_{ij}|^b} \quad (4.3)$$

A recognized problem of saliency calculating algorithms is analysis of weak contours in the vicinity of strong contours. Figure 4.17 (a-b) show a simulated pattern of two sets of edgels belonging to two different contours, one set having a weaker saliency.

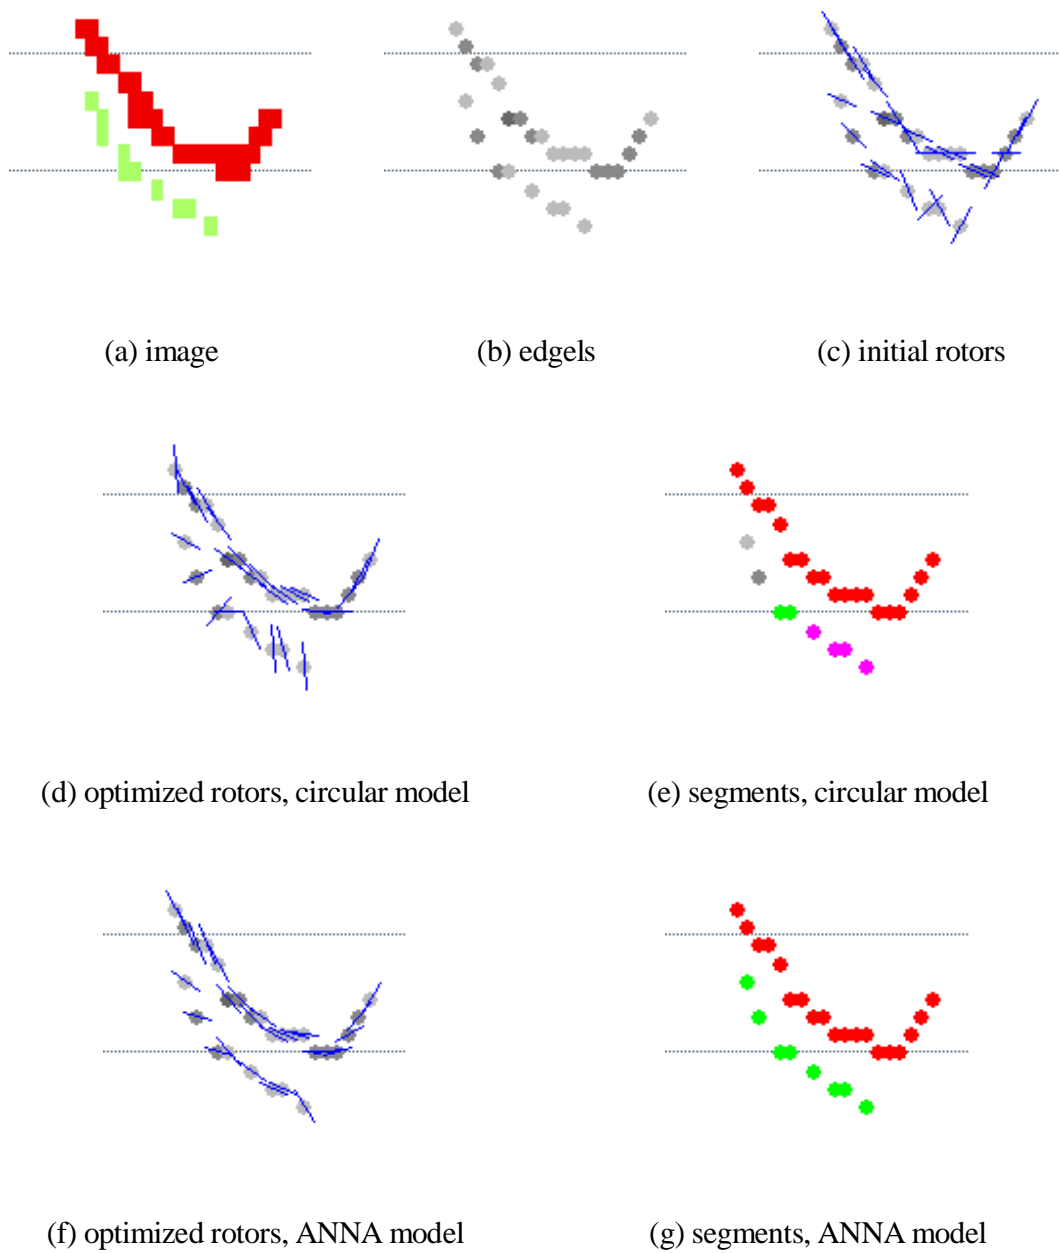


Figure 4.17. Saliency calculations for a weak contour in vicinity of a stronger contour (simulated pattern). (a) Simulated image of two traces. (b) Edgel pattern. (c) Initial rotor orientation. (d,e) Optimized rotor orientations and resulting trace segments obtained with a conventional circular model. Weak contour rotors are attracted to the stronger contour. (f,g) ANNA model processing with added dead zones in the interaction pattern.

In spite of attenuation of contributions from the rotors belonging to the stronger contour (because they are not aligned perfectly with the weak contour), they still dominate the rotors of the weak contour and rotate them in a wrong direction. This problem happens because the summary facilitating input from the strong contour still exceeds total input from the member edgels of the weak contour, thus resulting in attraction of the post-synaptic rotors to the wrong contour (Figure 4.17 (d-e)).

Psychophysical studies of the human vision show existence of two dead zones outside the co-axial pattern sector [Yen and Finkel, 1998, Guy and Medioni, 1996] (see Figure 3.7). Typically the dead zones are added by testing the elevation angle θ_{ij} between the chord C_{ij} and post-synaptic rotor V_i , so that the overall facilitation pattern becomes

$$F_{ij} = \begin{cases} S_{ij}, & C_{ij} < \rho \\ L_{ij}, & C_{ij} > \rho \text{ and } |\theta_{ij}| < \Theta/2 \\ 0, & C_{ij} > \rho \text{ and } |\theta_{ij}| > \Theta/2 \end{cases} \quad (4.4)$$

where ρ is the radius of short-term interaction, and Θ is the central angle of the long-term interaction sector.

Orienting the dead zones on the direction of the post-synaptic rotor V_i has disadvantage of being susceptible to mistakes of calculating initial orientation of the edgel V_i . Instead, ANNA calculates the elevation angles of the dead zone θ_{ij} between the chord C_{ij} and pre-synaptic rotors V_j , not V_i . In Figure 4.18, the advantage of this approach is illustrated by a simulated example where it becomes possible to correct the local estimate of V_i orientation if the dead zones are not placed co-axial with the wrongly

determined V_i . Rotor V_i in Panel 4.18(a) is not influenced by the fellow rotors of the same contour because they fall in the dead zones of the pattern whose orientation was determined by rotor V_i and used for all pairs of V_i and V_j . In Panel 4.18(b) the dead zones are placed individually, depending on orientation of pre-synaptic rotors in the vicinity of V_i , thus providing better reflection of the majority of rotors aligned to the contours in the plasmagram image.




Figure 4.18. Placement of the dead zones of rotor interaction. The interaction pattern is oriented co-axial with (a) post-synaptic rotor V_i (conventional), or (b) pre-synaptic rotor V_j . In the latter case, error in initial evaluation of V_i orientation will be corrected during evolving of the neural network.

Figure 4.17 (f-g) shows improvement of the rotor optimization on an edgel pattern with introduction of the dead zones. It also serves as a good illustration of advantages of our approach to trace extraction. Although some of the local estimates of edgel

orientations (Figure 4.17(c)) turn out to be wrong, the rotor optimization process refines them to produce a better saliency map and ultimately better trace segments.

Finally, the saliency measure for edgel i is calculated as vector sum of all contributions F_{ij} from the neighboring edgels on edgel V_i :

$$Y_i = \sum_{j \neq i} F_{ij} \quad (4.5)$$

Thus, the saliency measure is a vector quantity reflecting both perceptual strength and trace orientation of the edgels.

4.6. Transition from Saliency Map to Contour Segments

Biophysical studies of contour extraction in mammals' vision system show that areas of brain cortex stimulated by high-saliency edgels of the same contour exhibit synchronous excitation [Yen *et al.*, 1999]. Common model for synchronization of individual saliency measures in the saliency map suggests evaluation of both alignment and distance between cortex cells and binding (synchronizing) pairs of cells in agreement. In our case, this process is best accomplished by bottom-up clustering driven by the rotor alignment. The clustering starts with the highest degree of alignment and continues until the stop value is reached. Additional measures are taken to avoid bifurcation during the process of joining segments together.

4.7. Transition from Contour Segments to Contours

Additional grouping may be required to combine segments produced by the clustering algorithm. All possible segment combinations are searched exhaustively, and a grouping score Q_{ij} is determined for each pair of segments i and j similarly to the reconnection criterion of Lee and Talbot [1995] as:

$$Q_{ij} = \lambda A_{ij} + (1 - \lambda) G_{ij} = \lambda \sqrt{\frac{\alpha_{ij} + \beta_{ij}}{\pi}} + (1 - \lambda) \frac{d_{ij}}{D} \quad (4.6)$$

The score Q_{ij} includes a connection smoothness term A_{ij} and a connection gap term G_{ij} . Coefficient λ regulates relative contribution of terms A_{ij} and G_{ij} to the overall score. The smoothness of connection, A_{ij} , is characterized by angles α_{ij} and β_{ij} between the straight line connecting two segments and the straight lines fitted through the connection area to the segment tips. Figure 4.19 illustrates evaluation of A_{ij} with an example where two segments i and j , each containing 7 points, are considered for connection. The gap connector line is first built, joining the end points of the segments (shown in red) with a straight line. Then two fit windows are placed at the end points of the segments to determine orientation of the segment tips in the gap area. The fit window in this example is 7 point wide. A straight line is fitted to the points within the window using a least-square fit technique (shown in blue). Inclusion of the gap connector points in the fit window accomplishes a greater robustness to the range resolution jitter. Angles α_{ij} and β_{ij} are taken between the gap connector line and the lines fitted to the segment tips.




Figure 4.19. Evaluation of the α_{ij} and β_{ij} angles for the connection smoothness term in grouping score [Galkin et al., 2004].

The connection gap term G_{ij} in the grouping score governs contribution of the gap between segments, d_{ij} , to the grouping score. D is the scale parameter keeping G_{ij} within the interval of (0,1). At each exhaustive trial of segment pairs for reconnection, only two segments with the best connection score are combined together, provided that their grouping score is below a fixed threshold value. The segment grouping is continued until no segments can be joined.

Chapter 6 discussed performance of the trace extraction algorithms in further detail.

CHAPTER 5. PROCESSING OF RESONANCE SIGNATURES

The RPI plasmagrams contain not only a variety of traces pertaining to the remote plasma structures and boundaries, but also a suite of the local plasma resonances stimulated by the RPI radio transmissions. Detection and interpretation of the resonance signatures is a valuable diagnostic tool providing the actual electron density and magnetic field strength at the spacecraft location, which are needed for the accurate processing of the remote sensing information in the plasmagrams. Resonance matching is performed on the plasmagrams prior to trace extraction to remove identified resonance signatures from the image thereby simplifying further processing.

5.1. Resonance Detection in Plasma Sounding Data: Previous Work

Resonance signatures have been observed by a number of space missions carrying a topside ionosonde or a relaxation sounder. Table 5.1 and 5.2 list all missions that included radio transmission in plasma and collected data that could be analyzed for resonance signatures. The missions are roughly divided in two sections for the sounders capable of stepped-frequency remote sensing of the distant plasma structures and the relaxation sounders whose low power is enough only to stimulate resonances of the local plasma.

Table 5.1. Remote Sensing Sounders

Mission	Nation	Starting Year	Automatic data processing
Alouette 1, 2	USA	1962, 1965	“TOPIST”
ISIS 1, 2	USA	1965, 1971	[Huang <i>et al.</i> , 2002]
ISS B	Japan	1978	[Igi <i>et al.</i> , 1982]
InterKosmos-19	USSR	1979	[Voevudsky <i>et al.</i> , 1981]
Kosmos 1809	USSR	1986	
EXOS B, C, D	Japan	1981, 1984, 1989	[Obara <i>et al.</i> , 1990]
CORONAS-I	Russia	1994	
MIR Priroda	Russia	1999	
IMAGE RPI	USA	2000	CORPRAL [Galkin <i>et al.</i> , 2004]

Table 5.2. Relaxation Sounders

Mission	Nation	Starting Year	Operation regions	Automatic processing
GEOS 1, 2	USA	1965, 1968	Topside ionosphere	
ISEE 1	USA	1977	Magnetosphere, magnetotail, Solar wind	[Trotignon <i>et al.</i> , 1986]
EXOS B	Japan	1978	Ionosphere, magnetosphere	
VIKING	Sweden	1986	Magnetosphere	
Ulysses	International	1990	Interplanetary	
Interball-2 POLRAD	International	1996	Magnetosphere	
Cluster 2 Whisper	Europe	2000	Various	[Trotignon <i>et al.</i> , 2001]

Interpretation of resonances requires matching of all observed signatures to model values (see Chapter 2, Section 2.2.2 for theoretical background). The matching process involves a certain amount of simple calculations that are time consuming to carry out manually. Automatic resonance matching algorithms have been developed since the early 1980s, all using the known relations between resonance frequencies to build a resonance

model and fit it to the data by testing candidate combinations of the gyrofrequency and plasma frequency f_{ce} and f_{pe} that drive the model. Huang and Reinisch [1982] and Huang *et al.*, [2002] developed an automated resonance detection algorithm that is now a part of the TOPIST system for autoscaling topside ionograms from the Alouette and ISIS sounders [<http://nssdc.gsfc.nasa.gov/space/isis/isis-status.html>]. This algorithm seeks the “best” combination of the X-mode cutoff frequency and the gyrofrequency that maximizes the amplitude sum over range bins extending to 500 km on 5 frequencies: (1) O-mode cutoff or local plasma frequency f_{pe} , (2) X-mode cutoff frequency f_x , (3) upper hybrid frequency f_{uh} , and (4) and (5) two neighboring gyrofrequency harmonics nf_{ce} and $(n+1)f_{ce}$, selected to be close to the other resonances within the ionogram image. Igi *et al.* [1982] developed a similar algorithm for the topside sounder on the ISS-B spacecraft using two separate fits, one for the gyrofrequency f_{ce} , and another for the triplet f_{pe}, f_{uh}, f_x .

Trotignon *et al.* [1986] reported successful algorithmic solutions to the task of extracting resonance signatures from the ISEE-1 relaxation sounder data. The ISEE-1 satellite had a highly elliptic orbit with an apogee of $\sim 23 R_E$ (Earth’s radius = 6,375 km) and a perigee of 280 km and therefore observed a wide range of plasma densities in the solar wind, magnetosheath, magnetosphere, and magnetospheric tail. The resonance signatures in the magnetosphere were found to exhibit the greatest variety and present the greatest challenge to automated recognition. As a definite advantage for the automated data analysis, the ISEE-1 had an onboard magnetometer to measure the value of the gyrofrequency with an accuracy of $\sim 1\%$. The magnetometer’s estimate of the gyrofrequency was further improved by fine fitting of the f_{ce} harmonics to the sounding

data. The candidate plasma frequency values were tested in a separate fit. In contrast to the topside sounders, the best fit was sought for the series of Q-type resonances, f_{Qn} , using a weighted sum to favor the lowest orders of f_{Qn} . Matching of f_{Qn} was found to be critically sensitive to the accuracy of the gyrofrequency determination, with errors of only 0.7% still causing misalignment of the fit. A similar matching concept was used in the resonance interpretation algorithm by Trotignon *et al.* [2001] for the Whisper relaxation sounder aboard the Cluster-2 spacecraft. An amplitude envelope technique developed by Trotignon *et al.* [1986] had limited success, producing unacceptably high rates of false resonance recognitions.

Considering the results of previous efforts, automated detection and identification of the resonance signatures in the RPI plasmagrams presents a great challenge because of the large range of plasma densities probed by RPI in the magnetosphere and the often less than optimal resolution of the frequency scans. The diversity of the scientific goals that the RPI targets along the orbit often results in a frequency range and resolution unfavorable for automatic detection of the key resonance signatures. Finally, no onboard magnetometer measurements were available to help constrain the f_{ce} fitting procedure.

Two avenues can be pursued regarding the recognition of the resonances, (1) detection of the individual resonance signatures and their subsequent classification, or (2) search of the optimal fitting of the composite resonance model to the entire plasmagram image. Both approaches are detailed and compared below.

5.2. Fitting composite resonance model

Because of specifics of the RPI measurement scenario, the automated fitting algorithm cannot rely on the guaranteed presence of any particular resonance on each plasmagram. RPI is a versatile instrument that adjusts its sounding frequencies along the spacecraft orbit to match changes in the plasma density and to meet a variety of scientific goals. Because of dynamic allocation of the frequencies from the operating band of 3 kHz to 3 MHz with either constant or logarithmically changing frequency steps ranging in size from 300 Hz to tens of kilohertz, the RPI measurement parameters are often far from optimal to properly study the resonances. Compromises in the operating frequency band and resolution are commonly made to balance requirements for a high sounding cadence and long range coverage. Thus, plasmagrams may contain none of the nf_{ce} resonances or more than 20 of them, and the frequency resolution may be fine enough to cover the resonance amplitude peak with several frequency steps, or be as coarse as to completely miss the resonance. The fitting scheme had to be adjusted to work for a varying number of resonances/harmonics that may be present on plasmagrams. The coarse frequency resolution effects were considered by determining whether the tested resonance is too far from the closest plasmagram frequency. As a result, the number of frequencies whose summary amplitudes contribute to the total fit quality becomes different for different sets of the driving f_{ce} and f_{pe} . With a varying number of contributors at each fitting step, the fit quality cannot be calculated as the total sum, and therefore the “average fit” was used as the fit quality criterion:

$$Q_{ij} = \frac{\sum_{\{f\}_{ij}} S_{1R_E}(f)}{N}, \quad (4)$$

where

$\{f\}_{ij}$ - subset of plasmagram frequencies f corresponding to the model resonance frequencies calculated for the trial gyrofrequency $f_{ce}^{(i)}$ and the trial plasma frequency $f_{pe}^{(j)}$, such that each model frequency of the subset falls within the interval of $(f-\Delta_r, f+\Delta_r)$ of a plasmagram frequency f , where Δ_r is the half-width of the expected frequency band taken by the resonance signature,

N - size of $\{f\}_{ij}$,

$S_{1R_E}(f)$ - summary amplitude calculated over all ranges $r < 1 R_E$ at the plasmagram frequency f , i.e.,

$$S_{1R_E}(f) = \sum_{r < 1R_E} A_r(f) \quad (5)$$

where $A_r(f)$ is linear signal amplitude of the plasmagram bin (r,f) .

For this fitting method to work, the true values of f_{ce} and f_{pe} have to be in the set of tested combinations, and the fit has to be robust to other data features such as traces and interference lines. The set of trial frequencies $f_{ce}^{(i)}$ and $f_{pe}^{(j)}$ is obtained by selecting low and high boundaries around the predicted model values of gyrofrequency $(f_{ce})_p$ and

plasma frequency $(f_{pe})_p$. Then the frequency steps, Δ_{ce} and Δ_{pe} , are selected that determine how many trial frequencies are selected from the interval around the model values.

The Geopack Model of the Earth's magnetic field [Tsyganenko, 1990] is used to obtain $(f_{ce})_p$. Most of the time, the interval of $\pm 5\%$ around $(f_{ce})_p$ is sufficient, except for periods of high magnetic activity at large radial distances from the Earth, when the interval is increased to $\pm 40\%$. The prediction of the local plasma frequency $(f_{pe})_p$ is even more uncertain, especially when the spacecraft is in the vicinity of the modeled plasmopause or magnetopause. The actual value may be one or two orders of magnitude off the modeled value. To make computational matters worse, the composite model appears to be sensitive to even a fraction of one per cent deviations of the input f_{ce} and f_{pe} from the true values, so that practically all plasmagrams require sub-pixel accuracy of the resonance signature localization (i.e., better than the frequency resolution of the plasmagram), and the trial frequency steps, Δ_{ce} and Δ_{pe} , are forced to be as small as 0.1 kHz. The resulting computing time of the composite model fit becomes unacceptable in the framework of the interactive data analysis with the BinBrowser tool [Galkin *et al.*, 2001], where the scaler starts the automated resonance matching and waits for its results.

Splitting the composite model in two separate fits, as in Igi *et al.*, [1986], reduces the computing time but leads to unsatisfactory false recognition rates, primarily because of the sensitivity of the gyrofrequency fitting to the presence of other resonances, noise, and natural emission bands. Introduction of heuristics to improve noise resistance was attempted, but did not solve the problem. Even if there were no resonance signatures on

the plasmagram, the model fit approach would still give a best-fit solution. Therefore a method was developed that is able to locate and match the resonance signatures when they are present as well as identify the lack of resonance signatures when they are absent.

5.3. Detecting and Matching Plasmagram Resonances

Instead of fitting the composite model of resonances directly to a plasmagram image, the image is first analyzed to identify frequencies that are potential resonances, and then a matching algorithm is applied to classify these candidates as either a particular resonance type or a noise (interference) line.

5.3.1. Resonance Detection

5.3.1.1. SMOOTHING. A great variety of noise reduction and signal enhancement techniques have been devised to improve the quality of signature detection in images. Ziou and Tabbone [1998] give a review of conventional image smoothing filters and arguments for their use. For this approach a novel filter was designed to improve the quality of resonance detection in plasmagrams, since the analysis of applicable classic 2D ridge detection filters [e.g., Subirana-Vilanova, 1992] showed that they tend to wash out and displace the subtle resonance signatures. The new 1D filter replaces each amplitude of the scan, A_r , with the median calculated over the amplitudes with smaller virtual range,

$$F(A_r) = \text{median}\{A_j\}, j = [1, r]. \quad (5.1)$$

This concept is further referred to in the text as the “cumulative median”. The cumulative median filter has a smoothing scale that gradually increases with range, so that it applies less smoothing at the lower ranges, preserving the resonance signature shape, and eventually arrives at a robust estimate of the average background level for that frequency. Besides providing protection of further processing stages from random noise and jitter, the cumulative median filter enhances a particular type of signature whose amplitude decreases with range and suppresses all others. It is effectively matched to the expected shape of the resonance envelope that should have a falling slope due to eventual loss of the wave power. Figure 5.1 illustrates the response of the cumulative median filter to synthesized signals with rising and falling envelopes.




Figure 5.1 Response of the cumulative median filter to synthesized signals with falling (a) and rising (b) envelopes. The filter elevates the falling envelope that the resonance signatures are expected to display and reduces the rising slope, thus improving the signature contrast.

This filter enhances the falling shape (Figure 5.1a) and reduces the rising slope (Figure 5.1b), thus improving the contrast of resonances in the plasmagram. Figure 5.2 demonstrates the noise-suppression performance of the filter applied to the same signal shapes with the addition of 30% white noise. In both cases the filter successfully removes the jitter and improves the contrast of the resonance.

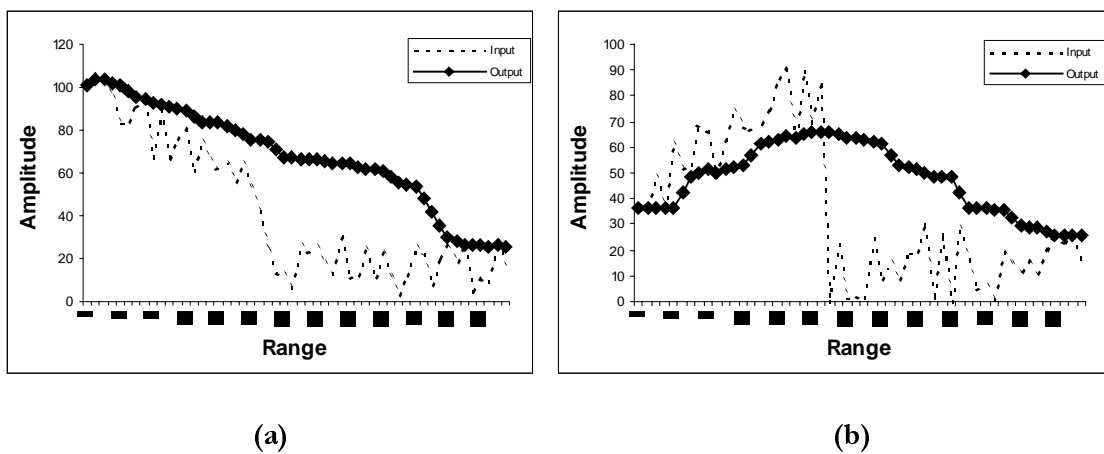


Figure 5.2 Response of the cumulative median filter to synthesized signals with falling (a) and rising (b) envelopes, with the addition of 30% noise. The filter smooths the noise and enhances the contrast of the resonance signatures.

Figure 5.3 illustrates the filter performance on two samples of RPI data collected (a) on a frequency containing a resonance, and (b) on another frequency without resonance but with an echo. The cumulative median filter enhances the resonance signature, smooths out noise jitter, and removes the echo signature. Figure 5.4 shows a sample RPI plasmagram obtained on June 28, 2001 at 23:58 UT (a) before and (b) after the filtering

operation. Smoothing out the noise jitter and the remote echoes visually simplifies the resonance detection without compromising the accuracy of the frequency registration.

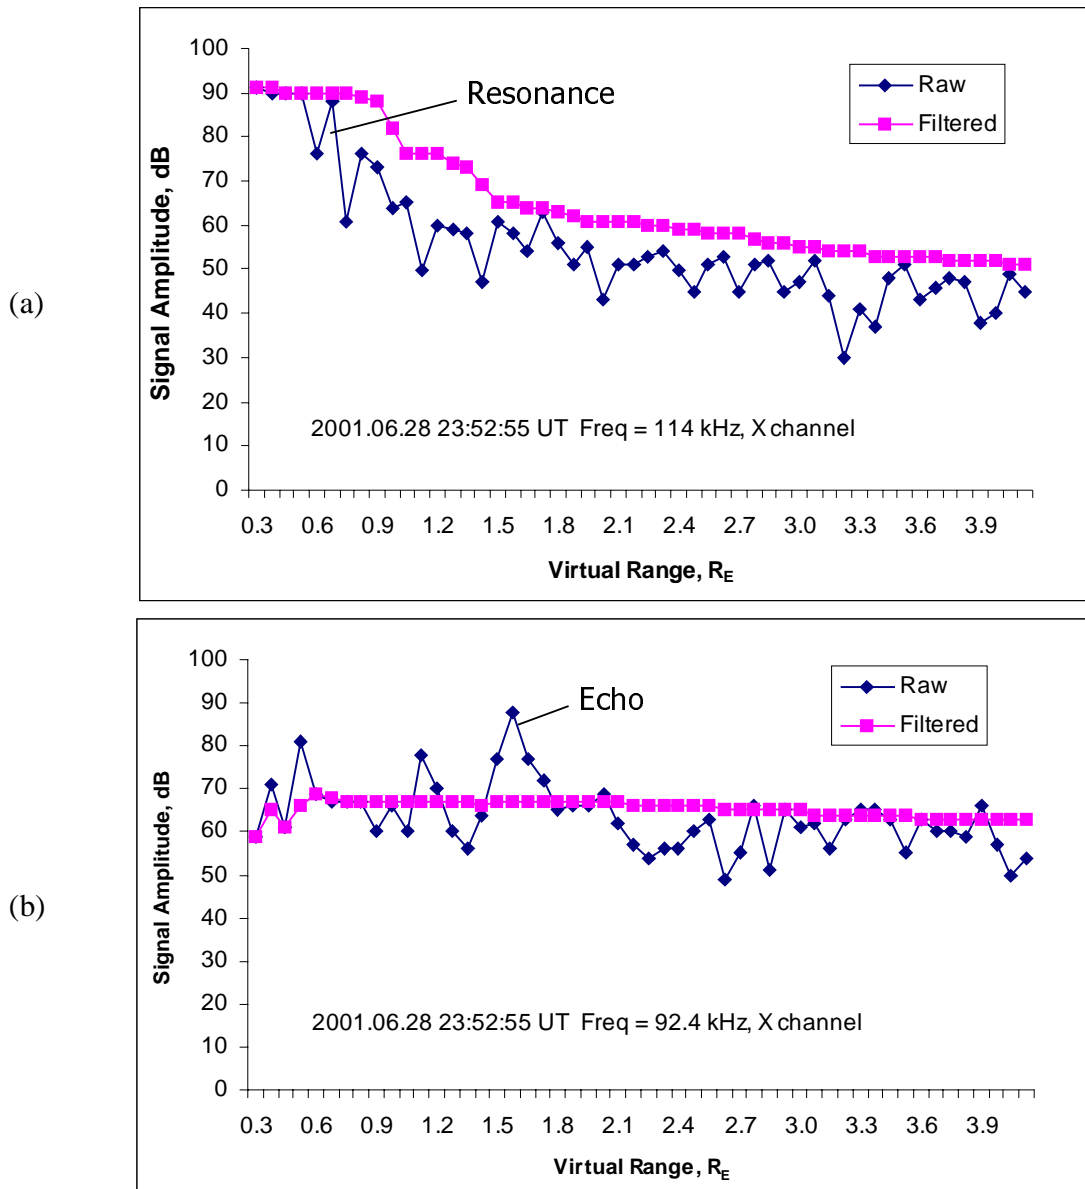


Figure 5.3 Cumulative median filter applied to RPI data: (a) a frequency containing resonance, (b) a frequency without resonance and a single echo.

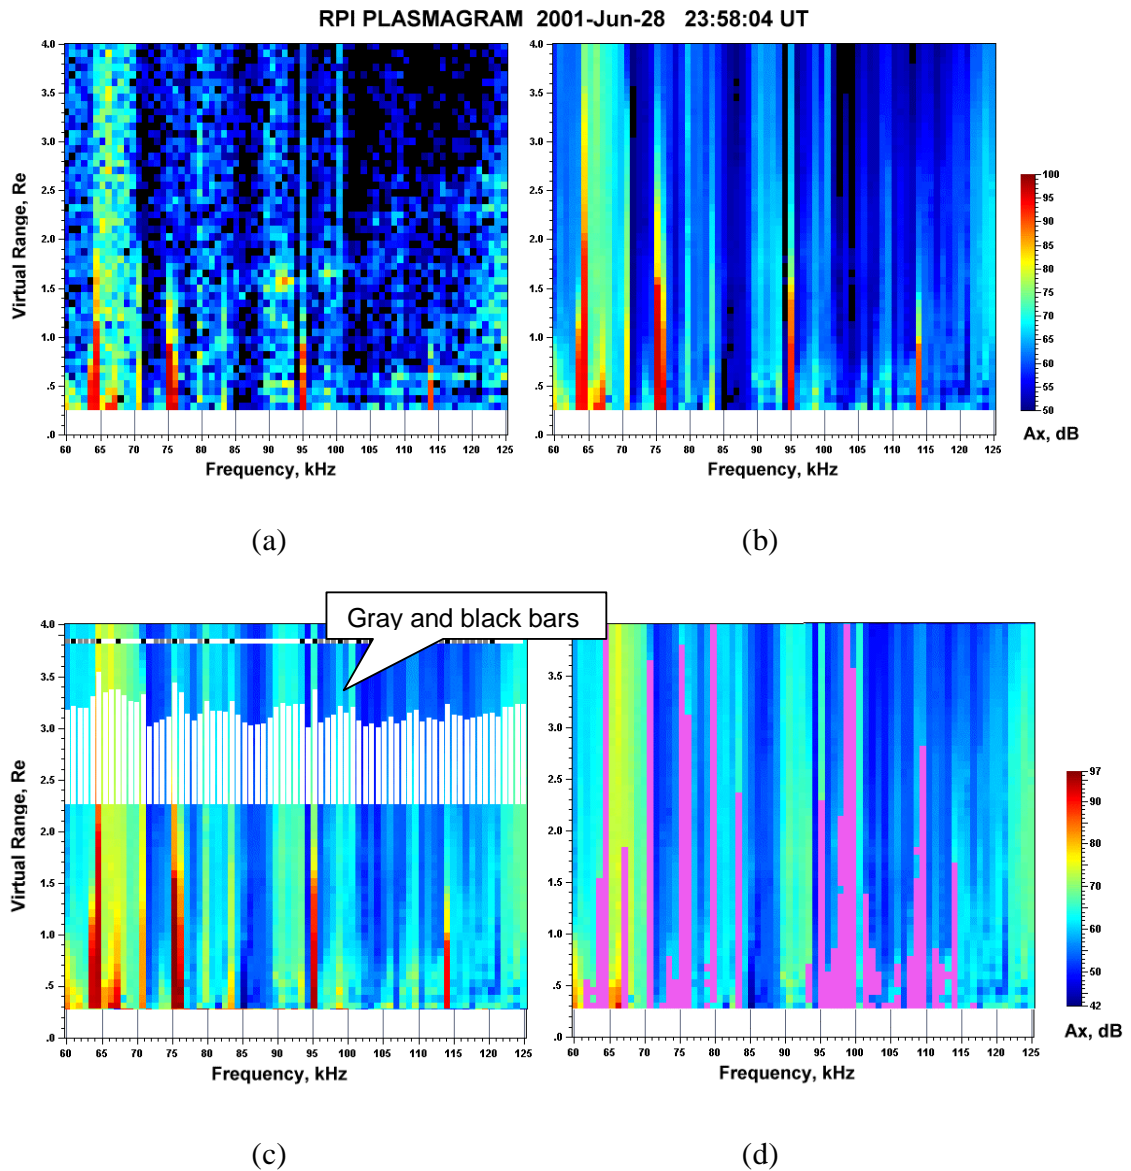


Figure 5.4 Resonance detection in RPI plasmagrams. (a) Raw plasmagram, (b) Processing with the cumulative median filter. (c) Summary amplitude function (white bars), labeled frequencies containing resonances (gray bars) and detected peaks (black bars). (d) Localized resonance signatures shown in magenta.

5.3.1.2. LABELING. After applying the cumulative median filter, the next processing step is the resonance signature labeling that tags the frequencies containing resonances using pattern recognition methods. A two-step labeling procedure is used, where frequencies that contain no resonances are first excluded, and then frequency groups are identified that belong to the same resonance. A frequency is excluded from the resonance analysis if its amplitude envelope does not display the appropriate decay with time. This is determined by evaluating the slope of the straight line fitted to the envelope using the least-square fitting technique. This procedure is used together with receiver saturation criteria that prevent elimination of strong resonances that do not show the expected amplitude decay. Saturation artifacts are occasionally observed for measurements with high receiver gains. At any particular fixed base gain, the system dynamic range is determined by the bit resolution of the voltage digitizer (12 bits = 72 dB) and choice of 4, 8, and 16-chip phase code waveforms that can add up to 24 dB of dynamic range after pulse compression. The total dynamic range then varies from 72 dB (plain pulse waveform, fixed gain) to 126 dB (16 chip waveform, 30 dB range of the autogain adjustment). Commonly the resonance study plasmagrams are made at a smaller dynamic range setting: (a) pulse compression is not used as it suppresses the plasma waves that do not have appropriate phase code, and (b) the autogain evaluation is disabled as it requires an additional 200 ms per frequency that becomes a considerable overhead for the resonance study measurements. To avoid losing resonances due to the saturation we do not exclude any frequencies where 90% of the amplitudes are above the saturation threshold.

After exclusion of all frequencies without resonance signatures, the frequency intervals belonging to the same resonance signatures are determined by collapsing the image intensities along the vertical axis and seeking the peaks of the resulting summary amplitude function, $S(f)$. The status line on top of the plasmagram in Figure 5.4(c) shows valid frequencies in gray, and the peaks identified in black.

5.3.1.3. CONTRAST EVALUATION. To further improve the robustness of the matching algorithm, a measure of contrast is calculated for each signature identified, so that stronger signatures contribute more to the fit quality. To determine the contrast, the average amplitude of the signature is compared to the average background amplitude. A statistical technique of differential histogramming was used to distinguish resonances from the background. In Figure 5.5, the lower amplitude histogram is built using all the amplitudes from all the range bins in each frequency step contained in a narrow ± 1 kHz band around the resonance peak frequency, and the upper histogram is calculated over a wider ± 2.5 kHz band that is guaranteed to include both resonance and background noise. Subtracting histogram 1 from histogram 2 leaves only background amplitudes, and their upper boundary defines the amplitude threshold for the resonance detection. Figure 5.3(d) shows in magenta the amplitudes exceeding the threshold. After the signature is localized, its contrast against the background is calculated by averaging the gradients between the tagged amplitudes and their immediate neighbors.

The resonance detection procedure labels a subset of plasmagram frequencies as potentially belonging to resonances, and determines tentative location and contrast of the resonance signatures. This information enters the interpretation algorithm that matches

the frequencies to their theoretical counterparts. The matching starts with identification of the resonances produced at the gyrofrequency and its harmonics. Without knowledge of the actual value of f_{ce} , further interpretation of the plasmagram resonances is impossible.

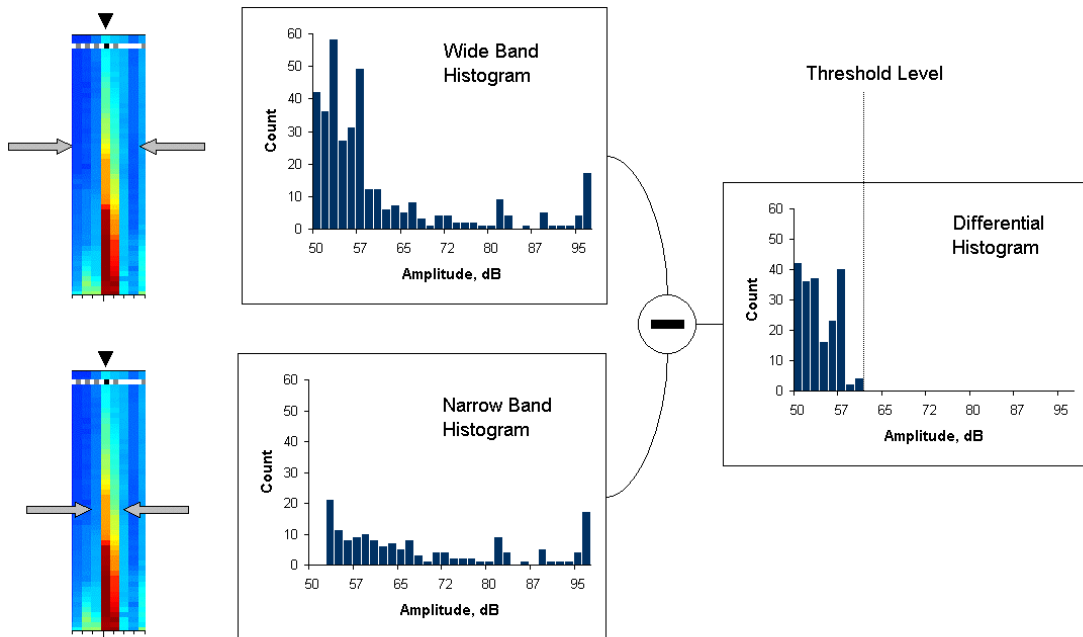


Figure 5.5 Use of the differential histogramming technique to estimate the local detection threshold for the resonance signature.

5.3.2. Matching the Gyrofrequency Resonances

At this point of the analysis, a list of detected resonance-like signatures and a predicted value of the gyrofrequency, $(f_{ce})_p$, exist. The task then is to find $n f_{ce}$ resonances in the list and thus determine the actual value of f_{ce} . A number of trial $f_{ce}^{(i)}$ values around the predicted value can be examined to determine which one produces the best match. This fit is different from the previously considered composite or separate model fitting tasks in that the pattern recognition analysis has provided the actually detected resonance

signatures. Only those frequencies contribute to the fit that are labeled as containing a resonance, and the resonances of higher contrast contribute more to the fit quality.

Analysis of the RPI plasmagrams for the gyrofrequency harmonics showed that they do not always match perfectly to multiples of f_{ce} . One of the causes of this mismatch is the change of spacecraft location (and therefore f_{ce}) during the plasmagram measurement, which is typically 1-3 minutes. The fitting scheme was modified to correct the higher orders of nf_{ce} for the expected change in the magnetic field strength as the spacecraft changes its location. Another consideration was given to possible localization errors and to insufficient frequency resolution that prevents separation of neighboring resonances. Figure 5.6 illustrates the general principle of selecting matching signatures next to a trial frequency, where first preference is given to the closer actual plasmagram frequency.




Figure 5.6 Choice of the actual plasmagram frequencies next to the trial frequency. If the closest frequency is not tagged as containing a resonance, the other frequency contributes to the fit quality with a reduced weight.

If, however, the closest frequency is not identified as valid (i.e., containing a resonance), the other frequency, if valid, is allowed to contribute to the fit with a lower weight. Both

plasmagram frequencies are tested if they are still within the expected frequency band occupied by the resonance through analysis of their deviations from the trial frequency, Δ_+ and Δ_- .

5.3.3. Matching the Plasma Frequency Resonance

Fitting of the plasma frequency is affected by the same problems discussed above for the gyrofrequency fit: (a) insufficient frequency coverage and resolution, (b) changes in the medium during the plasmagram measurement time, (c) need for sub-pixel accuracy of the frequency determination, (d) prediction errors. In addition to these common issues, the model equations describing dependencies of the Q- and D-type resonances on the plasma frequency, are only approximate. For example, equation (2.2) for the Q-type resonance frequencies does not always apply to the plasma in the magnetosphere that may contain a hot component in addition to the (dominant) cold population. In a plasma that is not described by a Maxwellian distribution, the best match of Q type resonances based on equation (2.2) will yield f_{pe} and f_{uh} values that are not likely to match any of the resonances in the plasmagram [Benson *et al.*, 2003a]. Our present algorithm does not include a match of resonances to the Dn frequencies given by (2.3).

Figure 5.7 shows examples of fully automated resonance processing with the algorithm developed here.

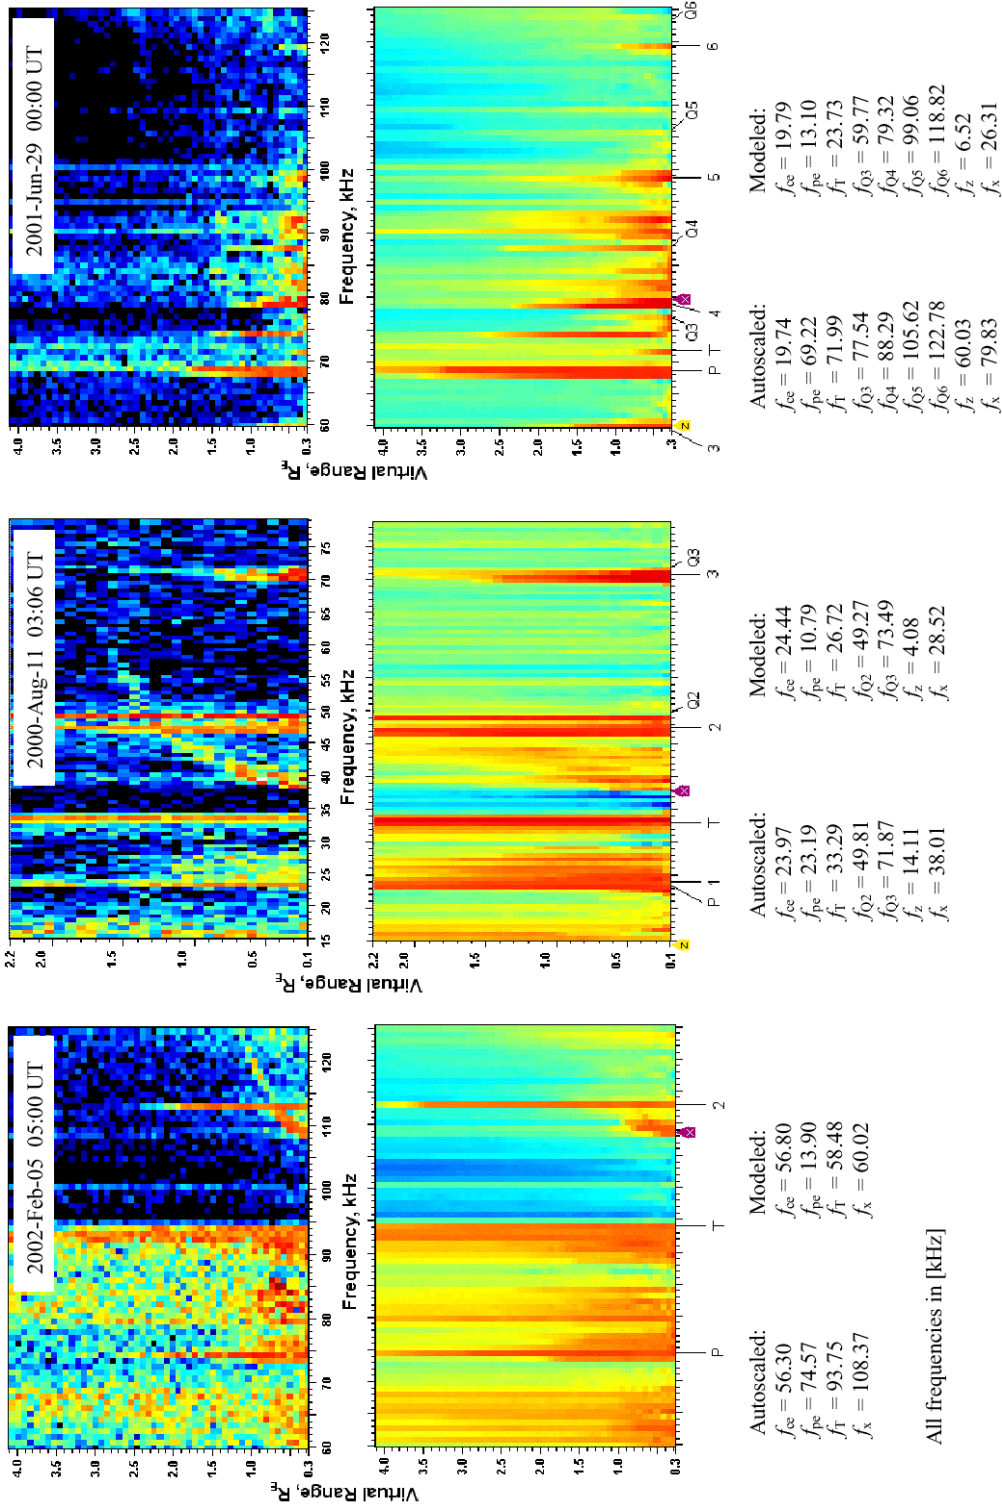


Figure 5.7 Examples of automated resonance matching. The upper panel shows raw plasmagrams; the middle row displays results of accumulative median filtering operation and places markers at automatically scaled resonance and cutoff frequencies. (P = f_{pe} , T = f_{tub} , x = f_x (yellow arrow), z = f_z (purple arrow), 1, 2, ... = nf_{ce} , Qn = f_{Qn}). Lower panel gives autoscaled frequencies in comparison with the predicted values.

The scaled values are compared with predicted f_{ce} values based on the Tsy 96-1 magnetic-field model [Tsyganenko, 1996] and the ad-hoc f_{pe} model [Reinisch *et al.*, 2001]. For the three examples, all resonance signatures were recognized and matched correctly. The modeled f_{ce} values agreed closely with the measured values, whereas for f_{pe} , large differences between the predicted and observed values were observed, signifying that the plasma conditions were significantly different from the model. As a result, f_{uh} , f_z , and f_x were also very different from the model. Use of the high frequency resolution in the plasmagram in the central panel allowed to observe an approximate 1% mismatch of the model Qn resonances, which is comparable to the observational uncertainty but may also be caused by deviations of the plasma distribution from the Maxwellian assumed in the expressions used for Qn.

5.4. Discussion of Results

The three major features that led to successful automated interpretation of the resonances are (a) use of pattern recognition techniques to detect and evaluate resonance signatures prior to matching, (b) accounting for changes in the medium during the time required to acquire a plasmagram, and (c) use of the cumulative median filter for enhancement of the resonance signatures. The new algorithm now correctly and reliably interprets a wide range of scenarios found in RPI plasmagrams. Whereas the overall quality of gyrofrequency automatching is satisfactory, it is often difficult to correctly identify the plasma frequency resonance. In the case of the former, there is a harmonic sequence that aids the autodetection algorithm and there is a fairly reliable model to

constrain the search criteria. In the case of the plasma frequency, neither is available. While the observed magnetic-field strength can be higher than the quiet-condition model by tens of percent on disturbed days, the deduced electron density can differ from the model by factors of ten [Benson *et al.*, 2003b].

Mismatches fall into three general categories, (1) no match can be found because the medium gradients are different from predicted ones, (2) key signature(s) are either outside the plasmagram coverage or they are obscured, (3) a false match is selected due to errors by the resonance signature detector.

5.4.1. Errors in prediction of medium gradients

Noticeable mistakes in predicting the general gradient of gyrofrequency, Δf_{ce} , are very rare. It is more common to observe natural fluctuations in the magnetic field that cause occasional mismatches of f_{ce} harmonics in plasmagrams taken with high frequency resolution. The frequency bin selection technique discussed in Section 3.2 allows mismatched harmonics to still contribute to the quality of the fit if their deviation does not exceed one frequency step. The gradient of the plasma frequency, Δf_{pe} , is frequently predicted incorrectly in the vicinity of the plasmopause because of the difficulty in accurately modeling the location and gradient at this boundary.

5.4.2. False decisions by the signature detector

While selecting plasmagram frequencies as resonances based on the amplitude decay pattern (Section 3.1.2) discriminates against interference lines, it can also eliminate

resonances that do not display any significant amplitude decay over the listening time interval in the plasmagram. It is difficult to select the tagging criterion optimally so as to minimize both false positive and false negative decisions; the current choice is to allow false “valid” frequencies to enter the matching phase rather than to remove good signatures. Although the contribution of such false frequencies to the fit is typically small because of their low contrast against the background, they still increase the likelihood of wrong matches.

5.4.3. *Key resonance signatures missing or obscured*

The automatching algorithm works best when there are a few distinct resonance signatures present in the plasmagram so that the fit quality maximizes when the model and the measurements match. If the frequency coverage is not sufficient to include enough signatures in the plasmagram, the best fit may not correspond to the correct answer. Also, when the nf_{ce} and f_{Qn} separation is not significant (for low ratios of f_{pe}/f_{ce}) the matching may not be optimal because it relies completely on the proper identification of f_{pe} and f_{uh} which can fall outside the frequency coverage or can be overlapped by other resonances. Proper interpretation of such cases often requires analysis of the previous and following plasmagrams to confirm changing patterns as plasma conditions change. Such an analysis, as well as the inclusion of the Dn frequencies in the matching routine, remains to be done.

5.5. Unmatched Resonance Study

After the resonance matching is complete, there still may be detected resonance signatures left unmatched. Unmatched signatures fall into four categories:

1. True signature unmatched to a known resonance frequency because of deviations of plasma from assumed Maxwellian distribution causing observed difference from the model.
2. True signature unmatched to a known resonance frequency because of wrong prediction of f_{ce} or f_{pe} gradients within the time of plasmagram measurement.
3. False positive decision by the signature detector.
4. True signature of unknown nature.

The unmatched resonances are stored in the mission database for further studies.

5.6. Removal of Resonance Signatures

All identified resonance signatures are removed from the plasmagram image to simplify analysis of traces. The pixels belonging to a signature were previously identified as a part of contrast evaluation (Section 5.3.1.3), and now they are set to a “missing data” value to exclude them from further analysis.

5.7. Conclusion

Detection and interpretation of the resonances in RPI plasmagrams stands as an important diagnostic tool and a necessary step towards further interpretation of the remote sensing information contained in these records. Automation of this procedure relieves the RPI data analyst from tedious calculations. The method described in this thesis presents a novel approach to automated resonance detection and interpretation in a plasma environment with large variability like the one encountered by the RPI instrument on IMAGE. Implementation of this approach into the BinBrowser data visualization and analysis tool [Galkin *et al.*, 2001] has made it possible to automatically scale all of the approximately 600 plasmagrams recorded every day by RPI. The algorithm could also become a useful tool for on-board processing of future active plasma-wave instruments.

CHAPTER 6. IMPLEMENTATION AND PERFORMANCE STUDY

The trace extraction and resonance automatching algorithms have been implemented in a Cognitive Online Rpi Plasmagram Ranking ALgorithm (CORPRAL) [Galkin *et al.*, 2004]. The CORPRAL system operates in a research mode where its performance can be evaluated step-wise on both synthesized patterns and actual plasmagram data, as well as in a batch mode online with the RPI mission database to process newly acquired plasmagrams arriving daily from Goddard Space Flight Center (NASA) and submit extracted traces and resonances to the database where they are available for queries. This chapter discusses various aspects of CORPRAL performance.

6.1. Performance Study

Table 6.1 summarizes the parameters that need to be determined for CORPRAL operations. For each listed parameter, Table 6.1 provides a brief explanation how the parameter changes the algorithm performance and a general concept of selecting the optimal value. Sections of the Chapter 4 containing a more detailed description of the algorithms are also given. The CORPRAL parameters are divided into four groups, (1) AvTrend echo detection, (2) oriented edgel evaluation, (3) ANNA segment extraction, and (4) perceptual grouping of segments. A suite of synthesized patterns was designed to test robustness of the algorithms, study parameter sensitivity and determine optimal settings for the CORPRAL.

Table 6.1. CORPRAL parameters.

Name	Symbol	Definition	Effect	Considerations
AvTrend Echo Detection				
1. Detection threshold	D	Fig. 4.9	Higher D results in less signal	
2. Detection window length	N	Fig. 4.9	Smaller N increases sensitivity to small features in the image	Determines the scale of analysis
Oriented Edgel Evaluation				
3. Context area radius	ζ	Fig. 4.10	Higher ζ increases scale of analysis	Smaller scale increases sensitivity to jitter; larger scale loses detail
ANNA Segment Extraction				
4. Rotor misalignment factor for long-distance interaction	a	4.1	Higher a causes smaller contribution from poorly aligned rotors	An estimate can be found in human vision research data
5. Attenuation with distance	b	4.1, 4.3	Higher b causes distant rotors to contribute less	R and b determine scale of analysis. (R is introduced for computational efficiency)
6. Long-distance cutoff radius	R	-	Smaller R decreases size of co-axial pattern	
7. Short-distance zone radius	ρ	4.4	Smaller ρ makes transaxial pattern smaller	Needs to be large enough to accommodate range jitter
8. Coaxial pattern central angle	Θ	4.4	Larger Θ makes dead zone smaller	An estimate can be found in human vision research data
9. Annealing regime temperatures	T_{start} T_{stop}	4.2 , Fig.4.15	Smaller T brings MFT neuron closer to saturation areas of sigmoid function	Annealing scheme needs to be adaptive to avoid premature saturation or fading (Section 4.5.1)
10. Clustering threshold	L	-	Lower L causes merging of clusters of less alignment	Merging segments is done better at perceptual grouping stage
Perceptual Grouping of Segments				
11. Relative contribution coefficient	λ	4.6	Higher λ weighs smoothness of the connection over the gap length	
12. Gap scale	D	4.6	Higher D attenuates connection gap term	
13. Score threshold	Q	-	Higher Q rejects more connections	

Robust performance of the algorithms on the imagery data of highly variable nature is one of the greatest challenges of the computer vision models. Devising an algorithm that

would be able to process a particular plasmagram correctly is an attainable task; using the same algorithm to process successfully any other plasmagram remains a challenging task. Robustness studies indicate that the greatest difficulty lies in adapting algorithms to the variable scale of signatures. Though slow compared to computers, the natural vision processor analyses the same input on multiple scales in parallel, selecting the highest saliency among all scales. The current version of the CORPRAL is designed for deployment on consumer-grade computing platforms that are not optimal for parallel multiple scale analysis. Although single scale analysis is prone to errors, the adaptive and optimizing qualities of CORPRAL help to accommodate typical range of the plasmagram signature scales. Extension of our current approach to the multiple scale treatment of the problem does not pose any difficulty in principle. Related to the issue of non-optimal scale selection is the problem of insufficient resolution, corresponding to the case of algorithm scale being larger than the feature scale. Robustness to the deviation from the minimal scale and artifacts of low resolution were studied on a suite of synthesized patterns.

6.1.1. ANNA Performance on Synthesized Patterns

Figure 6.1 shows a typical test pattern we use to study gap bridging and the noise protection features of the ANNA segment extractor and perceptual grouping algorithm. This pattern presents certain difficulty to the saliency calculation techniques, as it does not comply with the constant curvature constraint of the Gestalt perception. Figure 6.1(a) shows pattern of echoes (as resulted from the echo detection) created manually using the interactive pattern editor of CORPRAL.

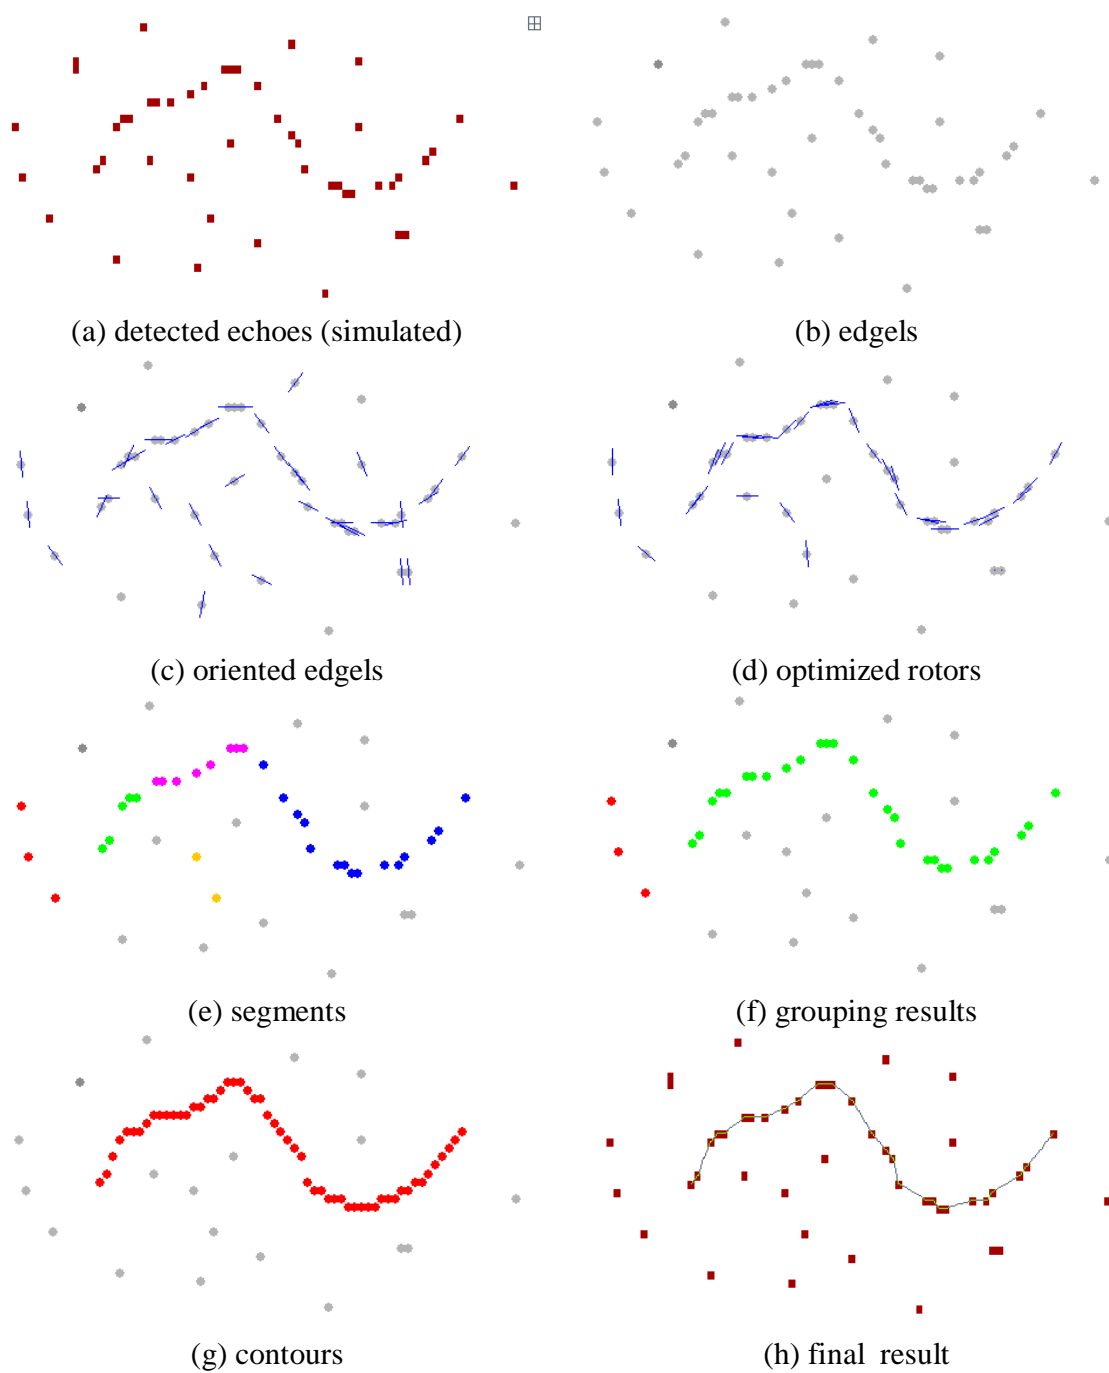


Figure 6.1. *Simulated pattern containing a single trace with gaps and noise.*

Pixel resolution in x (frequency) and y (range) axes are shown in the upper right insert of Figure 6.1(a). Using the simulated echo pattern, the edgels are obtained by selecting their leading edge (lower pixel) as shown in Figure 6.1(b). Figures 6.1(c) and (d) illustrate the rotor optimization process. Initial orientation of the edgels shows a few errors due to the locality of the edgel processing, subsequently improved by the ANNA algorithm. Using the optimized rotor pattern, contour segments are extracted as shown in Figure 6.1(e), and the perceptual grouping algorithm is run to attempt their connection. Thus the contours obtained, shown in Figure 6.1(f), pass final quality control that evaluates smoothness, length, and percentage of gaps in each contour to reduce false positives (Figure 6.1(g)). This result is shown in Figure 6.1(h) superimposed on the original echo pattern.

Smaller scale of analysis can be set by decreasing context area radius ζ of the angular histogramming and adjusting coefficients b and R regulating the attenuation of contributions with distance in ANNA (see Table 6.1). Figure 6.2 illustrates how changes of the analysis scale influence initial placement and optimization of rotors for the same pattern as in Figure 6.1(c) and (d). Smaller scale affects ability to bridge gaps in larger scale contours; larger scale results in engagement of all rotors, including those that are not part of any contour. Subsequent processing compensates in part both problems: perceptual grouping can bridge gaps in the smaller scale results, and segment clustering avoids connecting false edgels in the larger scale data. Also, rotor optimization results shown in Figure 6.1(d) and Figure 6.2(e) are obtained using an identical fixed annealing regime with linearly decreasing temperature T from 12 to 0.1 in 50 steps.

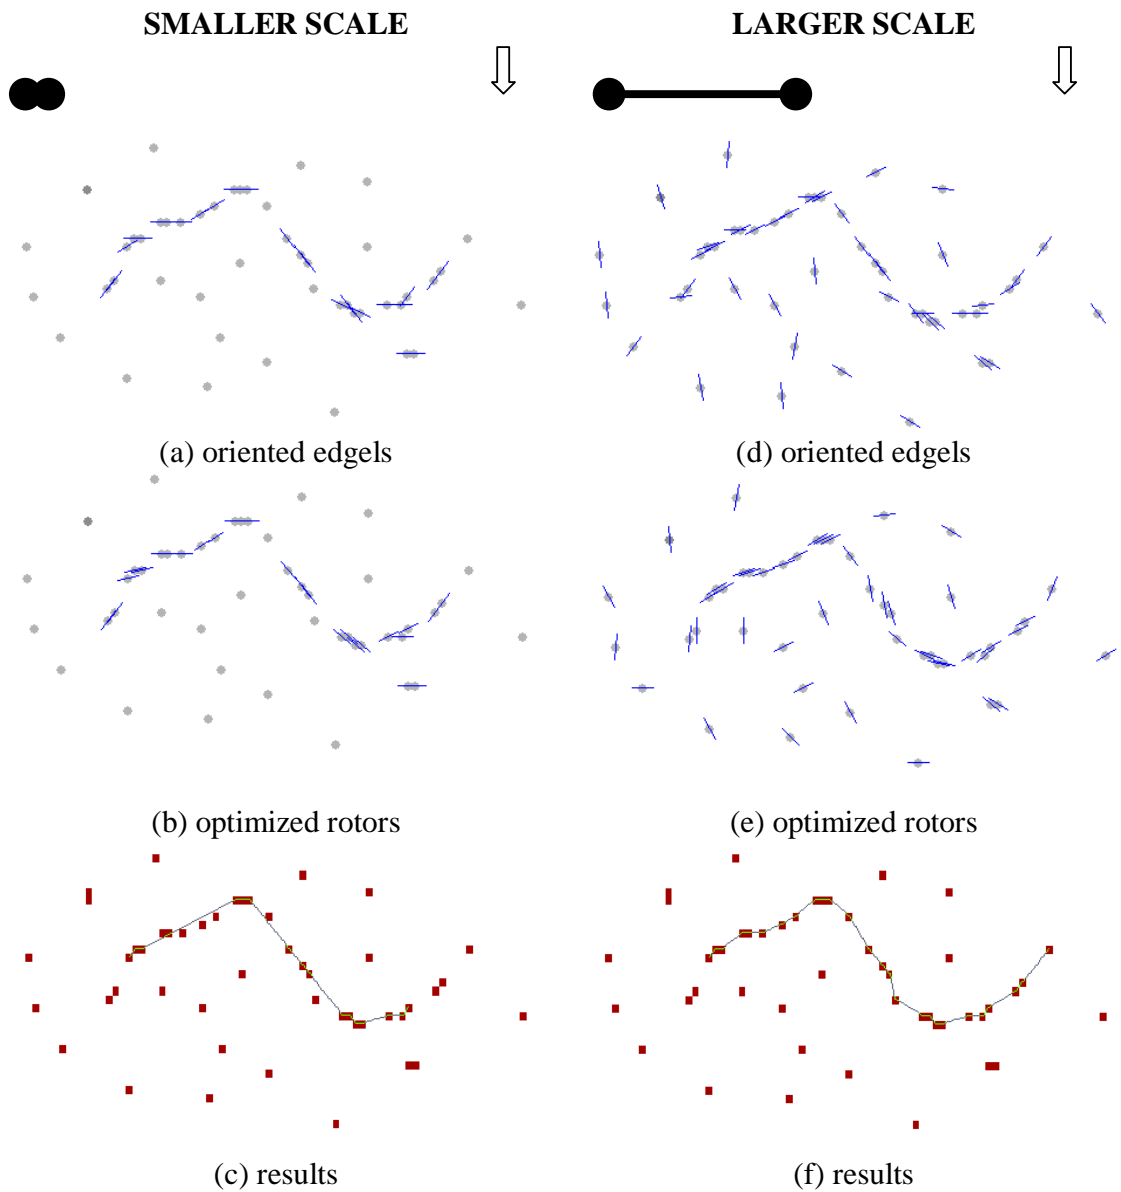


Figure 6.2. Processing of the same pattern as in Figure 6.1 at smaller (a-c) and larger (d-f) analysis scales.

Since the scale of optimization is larger in the latter case, a higher stopping temperature can be used to avoid fast saturation of the neurons. Figure 6.3 illustrates the advantage of selecting a more appropriate annealing regime with the stopping temperature of 4 instead

of 0.1. Implementation of an algorithm for selecting the appropriate annealing regime based on the choice of analysis scale and estimates of density of echoes allows good performance at larger scales and low sensitivity to the selection of the scale parameters ζ , b and R .




Figure 6.3. Processing of the same pattern as in Figure 6.2(d) with a more appropriate annealing regime (stop temperature increased to prevent premature saturation due to a larger scale of analysis).

Figure 6.4 shows another typical test pattern with two closely spaced traces that are difficult to separate for the saliency algorithms that are based on the static oriented edgels, especially if the locality scale is not optimal. It can be seen from Figure 6.4.(c) that orientations of many edgels are initially obtained incorrectly due to a strong influence of the nearby contour, which would result in poor saliency calculations if it were not for the rotor optimization processing, Figure 6.4(d). Analysis of another test pattern containing a short weak contour in vicinity of a longer contour is illustrated in Figure 6.5.

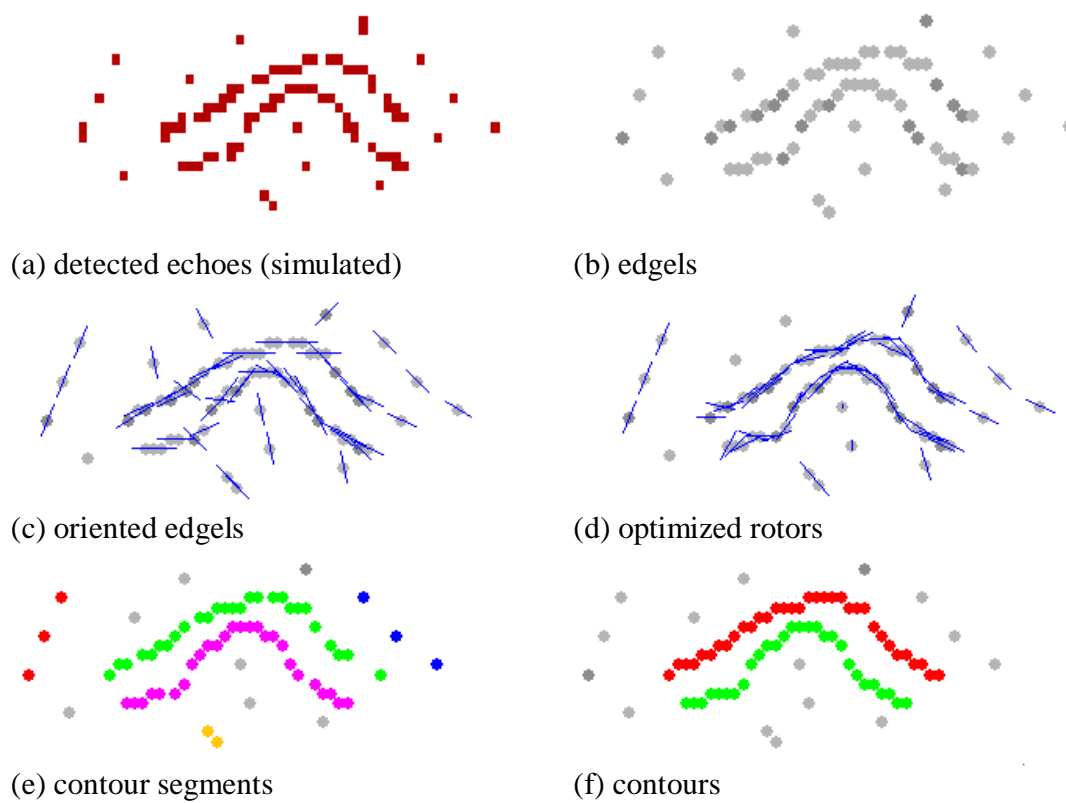


Figure 6.4. Sample pattern containing two closely spaced traces with some minor gaps and noise.




Figure 6.5. Processing of weak short contours in vicinity of a strong contour.

The test in Figure 6.5 was used to study the ability of ANNA to separate low-saliency features in the presence of nearby high saliency features, which has been a known difficulty of the static pattern saliency algorithms.

6.1.2. ANNA Performance on Plasmagram Data

Figure 6.6 illustrates the processing stages on the RPI plasmagram taken on March 1, 2002 00:02:58 UT. Raw plasmagram image is processed by the AvTrend adaptive echo detection algorithm to select all pixels belonging to echoes arriving from remote plasma structures (Figure 6.6(b)). Using detected echo data, edgel locations are evaluated, as shown in Figure 6.6(c), and rotor initialization and optimization process is applied to build the saliency map shown in Figure 6.6(d). Bottom-up clusterization is then used to derive contour segments from the saliency map (Figure 6.6(e)). The perceptual grouping algorithm is applied to chain qualifying segments together to form the traces shown in Figure 6.6(f). All six found traces are then submitted to the RPI mission database as expert knowledge data derived by CORPRAL.

6.2. Plasmagram Processing Results

The implemented version of CORPRAL works online with the RPI database of telemetry, derived, and expert knowledge data [Galkin *et al.*, 2001], delivering traces and expert ratings from the ~600 plasmagram images arriving daily. The CORPRAL results can be queried remotely from online BinBrowser workstations over the Internet to bring up subsets of plasmagrams that contain traces. By March 1, 2004 CORPRAL has tagged 138,074 out of 842,674 plasmagrams.

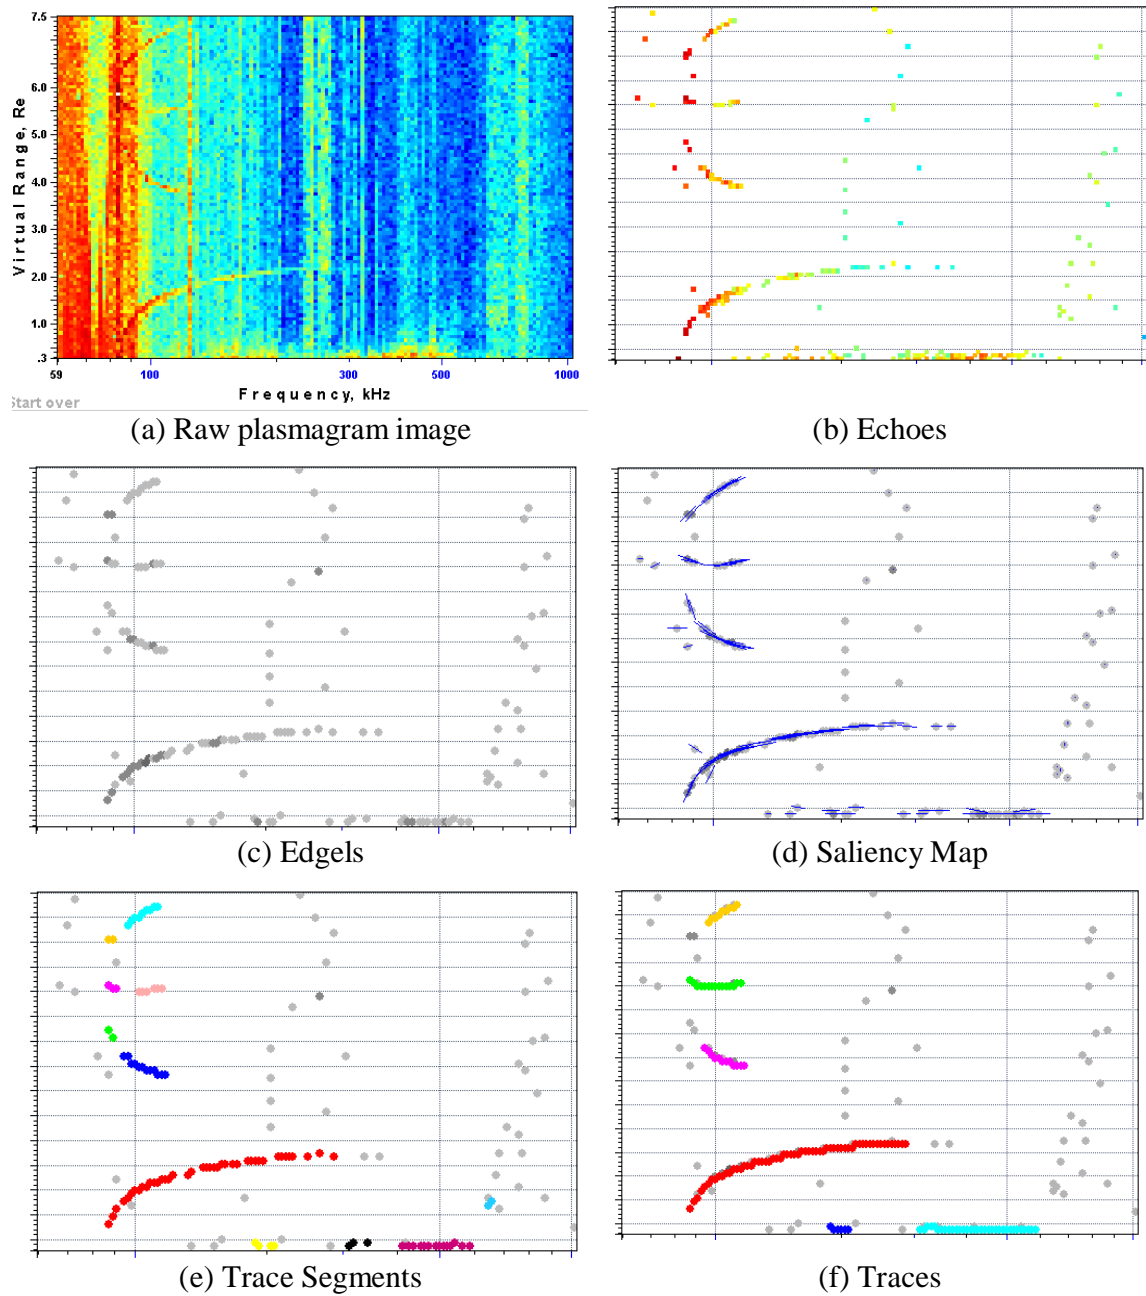


Figure 6.6. Example of processing stages for RPI plasmagram taken on March 01, 2002 00:02:58 UT. The raw plasmagram (a) is thresholded to obtain echoes (b), which are then reduced to the edgels (c). The edgel orientations are obtained and optimized to derive the saliency map (d) of the image. The saliency measures are analyzed to obtain trace segments (e) that are then combined together to form traces (f).

Further insight into the subset of CORPRAL-selected plasmagrams can be achieved by using the number of found traces as the query parameter. Figure 6.7 shows a sample subset of plasmagrams taken in July 2001 that contain 6 traces or more. Querying one month of ratings takes a few seconds to complete.




Figure 6.7. Some of July 2001 plasmagrams selected by CORPRAL as containing 6 traces or more.

CHAPTER 7. CONCLUSIONS AND OUTLOOK

In this chapter we summarize the main results of the thesis and discuss possible directions for future research.

This work was originally inspired by the practical need of locating scientifically significant data records in the large archive of RPI plasmagrams. The research eventually sharpened its focus on robust performance of the pre-attentive vision models presented with imagery data of highly variable content and quality. Based on previous results of trace recognition in ionograms and early experiments with the RPI plasmagrams, the choice of algorithmic approaches to the trace recognition was narrowed down to the techniques that analyze edge elements (edgels) detected in the image *collectively*, looking for subsets of edgels that form salient contours. Much effort was taken in this thesis to identify weaknesses of saliency calculation approach that became apparent as we applied it to the real-world data, and then to improve the method.

Low signal to noise ratio, range jitter, and discretization artifacts due to insufficient resolution of the plasmagram images cause problems with the local estimate of edgel orientation that is required for the saliency calculations. Errors in the edgel orientations propagate to the saliency measures that are directly based on relative orientation of the edgels and therefore are sensitive to imperfections of direction calculation. The major improvement of the collective edgel analysis was achieved by allowing edgel orientations to change in the process of iterative optimization of the edgel alignment. Edgels that can

change orientation were named *rotors*, in recognition of previously known studies of rotor models that optimize mutual rotor alignments under the co-circularity restriction governing the process of rotor interaction. The co-circularity restriction of the conventional rotor models was then revised to reflect modern understanding of the pre-attentive vision concepts that suggests existence of trans-axial pattern of interaction based on simple parallelism and dead zones of interaction outside the co-axial pattern cone. These modifications resulted in a remarkable enhancement of the algorithm performance compared to both static edgel saliency calculators and classic rotor models. Further studies suggested revision of the trans-axial pattern to better handle the range jitter in plasmagrams and placement of the dead zones on the pre-synaptic rotors instead of the central post-synaptic rotor to allow optimization process to correct errors in local estimation of edgel orientation. Other enhancements were also suggested to avoid false minima of the energy function of the optimizing neural network and to adapt simulating annealing scheme to the image features.

Introduction of the rotor optimization layer in the perceptual analysis of salient contours involves a significant increase of computations. Direct implementations of the combinatorial optimization schemes that handle all detected edgels are rather exceptional for the reason of high computational demand. An important part of this thesis was to find a pre-processing scenario that would reduce the number of detected edgels in the image with a minimal trade-off in terms of associated false negative and positive decisions. Investigations of the available approaches to noise reduction showed that conventional smoothing is inapplicable to plasmagrams because of its damaging effect on the traces

that are commonly thin and faint, and an adaptive thresholding was chosen instead to preserve only those features in plasmagram that stand out against the background.

Special attention in this thesis was given to the task of locating, identifying and then the removal of the resonance signatures as a part of the plasmagram pre-processing. Conventionally, the resonance processing employs a best-fit matching of the resonance frequency model to the raw plasmagram image by testing a range of trial electron plasma and cyclotron frequencies. The conventional approach failed to produce reliable results because of specifics of the RPI scientific mission that exposes the instrument to highly variable plasma and magnetic field conditions. A greater degree of robustness was achieved by implementing a resonance signature detector to eliminate contributions from the frequencies that do not display the anticipated decay of the resonance wave energy. Additional performance enhancement was accomplished by setting contributions to the fit proportional to the contrast of resonance signatures against the background.

All developed algorithms were implemented in a CORPRAL system that processes RPI plasmagrams on a daily basis as they arrive from the IMAGE operations center at Goddard SFC. The results of CORPRAL analysis are stored in the mission database that can be queried with a variety of search criteria to select plasmagrams with signatures.

Pre-attentive vision has been drawing attention of researchers since the 1950s. Understanding of visual information processing in the retina and brain cortex, a Nobel Prize result, has spawned a generation of biologically plausible models replicating the key components of this analysis. We plan to seek further enhancements of the CORPRAL by bringing its design closer to the solutions developed over the million years

of evolution. In particular, we are fascinated by the natural system of parallel analysis of the same image at multiple scales with subsequent fusion of the results, a concept that could provide robust detection of both spread and thin traces in plasmagrams. Noise protection can be improved by complementing the facilitating interaction of the rotors with the inhibitive counterpart that apparently has been discovered in the brain cortex. Edgel analysis can be improved by implementing 2D techniques that analyze data integrity across frequencies. It must be beneficial to preserve a greater amount of image information in rotors (such as edge gradient, color, scale width) so that their interaction can be made more efficient. Other optimization vehicles, such as the linear threshold neural networks, should be attempted to make use of the expanded set of edgel characteristics. Alternative combinatorial techniques can be studied for the task of compilation of aligned rotors into contour segments. Finally, the segment grouping algorithms can be improved by introducing higher order interpolating functions to evaluate the segment connection score. These measures shall prepare us for the next important step in our research, modeling of the attention-driven recognition of contours in images.

REFERENCES

- Alter, T.D., and R. Basri (1998), Extracting salient curves from images: an analysis of the saliency network, *IJCV*, **27**, 51-69.
- Baginyan, S., A. Glazov, I. Kisel, E. Konotopskaya, V. Neskromnyi, and G. Ososkov (1994) Tracking by a modified rotor model of neural network, *Comp. Phys. Comm.*, **79**, 165-178.
- Benson, R.F., V. A. Osherovich, J. Fainberg, A. F.-Vinas, and D. R. Ruppert (2001), An interpretation of banded magnetospheric radio emissions, *J. Geophys. Res.*, **106**, 13,179-13,190.
- Burch, J.L. (Ed.) (2000) *The IMAGE Mission*, Kluwer Acad.Pub., Dordrecht, 506 pp.
- Canny, J.F. (1986), A computational approach to edge detection, *IEEE Pattern Analysis Machine Intelligence*, **8**, 679-698.
- Cook, G.W., and E.G. Delp (1995), Multiresolution sequential edge linking, *Proc. IEEE Int. Conf. Image Proc*, Oct 23-26, 1995, Washington, 41-44.
- Denby, B. (1988), Neural networks and cellular automata in experimental high energy physics, *Comp. Phys. Comm.*, **49**, 429-448.
- Durbin, R. and D. Willshaw (1987), An analog approach to the traveling salesman problem using an elastic net method. *Nature*, **326**, 689.
- Elder, J.H. (1999), Are edges incomplete? *Int. J. Comp. Vision*, **34**, 97-122.
- Fox, M.W., and C. Blundell (1989), Automatic scaling of digital ionograms. *Radio Science*, **24**, 747-761.
- Freeman, W.T., and E.H. Adelson (1991), The design and use of steerable filters, *IEEE Trans.*, **PAMI-13**, 891-906.
- Galkin, A.I. (1962) Some aspects of automatic processing of vertical incidence sounding data (in Russian), *Geomagn. Aeron.*, **2**, 782-790.
- Galkin, I.A. (1987) Software development in support of automatic processing of vertical incidence ionograms. I. Ionogram preprocessing., *SibIZMIR Comm. #20-87*, USSR Acad. of Sci., Sib. Division, Irkutsk. 17pp. (in Russian)

- Galkin, I.A. (1992) On the problem of automatic scaling of vertical incidence and backscatter ionograms, *Proc. ISAP-92*, Sapporo, 1197-1200.
- Galkin, I.A. (1995) Neural networks for recognition of ionogram traces, M.S. Thesis, 51 pp, University of Massachusetts, Lowell, MA.
- Galkin, I.A., B.W. Reinisch, G.A. Ososkov, E.G. Zaznobina, and S.P. Neshyba (1996) Feedback neural networks for ARTIST ionogram processing, *Radio Science*, **31**, 1119-1129.
- Galkin, I.A., G. M. Khmyrov, A. Kozlov, B.W. Reinisch, X. Huang, and G. Sales (2001), New tools for analysis of space-borne sounding data, Proc. 2001 USNC/URSI Nat. Radio Sci. Meeting, Boston, MA, 304.
- Galkin, I.A., B.W. Reinisch, X. Huang, R. F. Benson, and S. F. Fung (2004a), Automated diagnostics for resonance signature recognition in RPI plasmagrams, (*to appear in Radio Science*).
- Galkin, I.A., B.W. Reinisch, G. Grinstein, G. M. Khmyrov, A. Kozlov, X. Huang, and S. F. Fung (2004b), Automated Exploration of the Radio Plasma Imager Data, (*submitted to Journal of Geophysical Research*).
- Gonzalez, R.C., and P. A. Wintz (1987), *Digital image processing*, 2nd edition, Addison-Wesley, Reading, MA.
- Guy, G., and G. Medioni (1996), Inferring global perceptual contours from local features. *IJCV*, **20**, 113-133.
- Hancock, E. and J.Kittler (1990), Discrete Relaxation, *Pattern Recognition*, **23**, 711-733.
- Hough, P.V.C. (1962), A method and means for recognizing complex patterns. US Patent 3,069,654.
- Hopfield, J.J. (1982), Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. U.S.A.*, **79**, 2554-2558.
- Huang, X. and B.W. Reinisch (1982), Automatic calculation of electron density profiles from digital ionograms. 1. Automatic O and X trace identification for topside ionograms, *Radio Science*, **17**, 421-434.
- Huang, X., B.W. Reinisch, D. Bilitza, and R.F. Benson (2002), Electron density profiles of the topside ionosphere. *Annals of Geophysics*, **45**, 125-130.

- Igi, S., K. Aikyo, R. Nishizaki (1982), Efficient coding and resonance spike identification for topside ionogram processing, *Memoirs of National Institute of Polar Research. Special Issue*, **22**, 140-154.
- Igi, S., and K. Aikyo (1986), Automatic identification of resonance spikes for ionograms obtained by ionospheric sounding satellite-b (ISS-b), *J. Radio Res. Lab.*, **33**, 169-175.
- Itti, L., C. Koch, and E. Niebur (1999), A model of saliency-based visual attention for rapid scene analysis, *IEEE Tran. PAMI*, **20**, 1254-1259.
- Kalviainen, H., P. Hirvonen, L. Xu, and E. Oja (1995), Probabilistic and non-probabilistic Hough transforms – overview and comparisons. *Image and Vision Computing*, **13**, 239-252.
- Kittler, J., and J. Illingworth (1985), Relaxation labeling algorithms – a review, *Image and Vision Computing*, **3**, 206-216.
- Lai, K.F., Deformable contours: modeling, extraction, detection, and classification (1994), PhD dissertation, University of Wisconsin-Madison.
- Laptev, I., H. Mayer, T. Lindeberg, W. Eckstein, C. Steger, and A. Baumgartner (2000), Automatic extraction of roads from aerial images based on scale-space and snakes. *Machine Vision and Applications*, **12**, 23-31.
- Lee, T.C.M. and H. Talbot (1995), A fast method for detecting and matching linear features in images. *Proc. DICTA-95*, 649-654.
- Mallat, S (1989), Multifrequency channel decompositions of images and wavelet models. *IEEE Trans. Inform. Theory*, **37**, 2091-2110.
- Marr, D. (1976), Early processing of visual information. *Phil. Trans. R. Soc. Lond.*, **B-275**, 483-524.
- Marr, D., and H.K. Nishihara (1978), Visual information processing: Artificial intelligence and the sensorium of light, *Technology review*, **81**, 2-23.
- Marr, D., and E. Hildreth (1980), Theory of edge detection, *Proc. R. Soc. London*, **B-207**, 187-217.
- MacLennan, B. J. (1991), *Gabor Representations of Spatiotemporal Visual Images*, Technical Report CS-91-144 (Knoxville: University of Tennessee Computer Science Department); accessible via URL <http://www.cs.utk.edu/~mclennan>.
- Morgan, M. and W. Evans (1951), Synthesis and analysis of elliptic polarization loci in terms of space-quadrature sinusoidal components. *Proc. IRE*, **39**, 552-556.

- Muresan, L., R. Muresan, and G.A. Ososkov (1997), Deformable templates for circle recognition. *Proc. SPIE Symp. Aerospace/Defense Sensing, Simulation and Control*, 20-25 April 1997, Florida, **3007**, 723-733.
- Nelson, R.C. (1994), Finding line segments by Stick Growing, *IEEE Trans. Pattern Analysis and Machine Intelligence*, **16**, 519-523.
- Obara., T., K. Syutoh, T. Kato, *et al.*, A new system for operation and data handling of Akebono (EXOS-D) satellite (1990), *J.Geomag. Geoelectr.*, **42**, 565-577.
- Ohlsson, M., C. Peterson, A. Yille (1992), Track finding with deformable templates – the elastic arm approach. *Computer Physics Communications*, **71**, 77.
- Osherovich, V.A., and R.F. Benson (1991), The lower subsidiary diffuse plasma resonances and the classification of radio emissions below the plasma frequency. *J.Geophys.Res.*, **96**, 19,331-19,341.
- Pelillo, M., and A.M. Faneli (1997), Autoassociative learning in relaxation labeling networks, *Pattern Recognition Letters*, **18**, 3-12.
- Peterson, C. (1989), Track finding with neural networks, *Nucl.Instr.Meth.Phys.Res.*, **A279**, 537-545.
- Peterson, C. and J.R. Anderson (1987), A mean field theory learning algorithm for neural networks, *Complex Systems*, **1**, 995-1019.
- Posner, M.I., and Y. Cohen (1984), Components of visual orienting, *Attention and Performance X*, Lawrence Erlbaum Associates Publishers, London, 531-556.
- Pulinets, S.A., and R.F. Benson (1999), Radio frequency sounders in space. *Review of Radio Science*, Ed. W. Ross Stone, Oxford University Press, 711-733.
- Rangarajan, A. (1997), Self-annealing: unifying deterministic annealing and relaxation labeling, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, M.Pelillo and E.R.Hancock, ed., Springer-Verlog, Berlin, 229-244.
- Reinisch, B.W. (1996), Modern Ionosondes, in: *Modern Ionospheric Science*, European Geophysical Society, 440-458.
- Reinisch, B.W., Haines, D.M., Bibl, K., Cheney, G., Galkin, I.A., Huang, X., Myers, S.H., Sales, G.S., Benson, R.F., Fung, S.F., Green, J.L., Boardsen, S., Taylor, W.W.L., Bougeret, J.-L., Manning, R., Meyer-Vernet, N., Moncuquet, M., Carpenter, D.L., Gallagher, D.L., and Reiff, P. (2000), The Radio Plasma Imager investigation on the IMAGE spacecraft, *Space Science Reviews*, **91**, 319-359.

Reinisch, B.W., G.S. Sales, D.M. Haines, S.F. Fung, and W.W.L. Taylor (1999), Radio wave active Doppler imaging of space plasma structures: Arrival angle, wave polarization, and Faraday rotation measurements with the radio plasma imager. *Radio Science*, **34**, 1513-1524.

Reinisch, B.W., X. Huang, P. Song, G.S. Sales, S.F. Fung, J.L. Green, D.L. Gallagher, and V.M. Vasyliunas (2001), Plasma Density Distribution Along the Magnetospheric Field: RPI Observations from IMAGE, *Geophys. Res. Ltrs.*, **28**, 24, 4521-4524.

Rock, I., and S. Palmer, The legacy of Gestalt psychology, *Sci. American*, SCA9012, 84-90, 1990.

Rosenfeld, A., R.A. Hummel, S.W. Zucker (1976), Scene labeling by relaxation operations, *IEEE Systems, Man, Cybernetics*, **6**, 420-433.

Schleher, D.C. (1980), *Automatic Detection and Radar Data Processing*, Artech House, Dedham, MA, 636 pp.

Shashua, A., and S. Ullman (1988), Structural saliency: the detection of globally salient structures using a locally connected network, *2nd ICCV*, 321-327.

Strand, R.C. (1972), Optical-image recognition for experiments in the track chambers of high-energy physics, *Proceedings of IEEE*, **60**, 1122-1137.

Subirana-Vilanova, J.B., and K.K. Sung (1992), Multi-scale vector-ridge-detection for perceptual organization without edges. *A.I.Memo No,1318, Artificial Intelligence Laboratory*, MIT, Cambridge, MA, 48 pp.

Trotignon, J.G., J. Etcheto, and J.P. Thouvenin (1986), Automatic determination of the electron density measured by relaxation sounder on board ISEE 1. *J. Geophys. Res.*, **91**, 4302-4320.

Trotignon, J.G., P.M.E. Decreau, J.L.Rauch, *et al.* (2001), How to determine the thermal electron density and the magnetic field strength from the Cluster/Whisper observations around the Earth. *Annales Geophysicae*, **19**, 1711-1720.

Walsh, D. and A.E. Raftery (2001), Accurate and efficient curve detection in images: the Importance Sampling Hough Transform, Tech. Rep. No. 388, Dept. Statistics, University of Washington, Seattle, WA. 24 pp.

Warren, E.S., and E.L. Hagg (1968), Observation of electrostatic resonances of the ionospheric plasma, *Nature*, **220**, 466-468.

Wersing, H., J.J.Steil, and H. Ritter (2001), A competitive layer model for feature binding and sensory segmentation of features. *Neural Computation*, **13**, 357-387.

- Wolfram. S.(ed.) (1986), *Theory and Applications of Cellular Automata*. World Scientific. pp.
- Xu, L., E. Oja, and P. Kultanen (1990), A new curve detection method: Randomized Hough Transform (RHT). *Pattern Recognition Letters*, **11**, 331-338.
- Yen, S.C. and L.H. Finkel, Extraction of perceptually salient contours by striate cortical networks. *Vision Research*, *38*, 719-741, 1998.
- Yen, S.C., E.D. Menschik, and L.H. Finkel, Perceptual grouping in striate cortical networks mediated by synchronization and desynchronization. *Neurocomputing*, *26-27*, 609-616, 1999.
- Yosida M. (1989), Automatic scaling of ionospheric parameters, *Rev. Comm. Res. Lab.*, **35**, 33-40. (in Japanese)
- Yuille, A.L. (1990), Generalized deformable models, statistical physics, and matching problems. *Neural Computation*, **2**, 1.
- Ziou, D., and S. Tabbone (1998), Edge detection techniques – an overview, *Int. J. Pattern Rec. & Image Analysis*, **8**, 537-559.