24. Digital Signal Processing

Academic and Research Staff

Prof. A.V. Oppenheim, Prof. J.S. Lim, Prof. B.R. Musicus, Professor A.B. Baggeroer, G. Aliberti

Graduate Students

M.M. Covell, W.P. Dove, M. Feder, D.W. Griffin, D. Izraelevitz, D.M. Martinez, E.E. Milios, C. Myers, T.N. Pappas, J.J. Rodriguez, A.J. Silva, P.L. Van Hove, M.S. Wengrovitz, A. Zakhor

Introduction

The Digital Signal Processing Group is carrying out research in the general area of signal processing. In addition to specific projects being carried out on campus, there is close interaction with Lincoln Laboratory, and with the Woods Hole Oceanographic Institution. While a major part of our activities focus on the development of new algorithms, there is a strong conviction that theoretical developments must be closely tied to applications. The application areas that we are involved with currently are speech, image, video and geophysical signal processing. We also believe that algorithm development should be closely tied to issues of implementation because the efficiency of an algorithm depends not only on how many operations it requires, but also on how suitable it is for the computer architecture it runs on. Also strongly affecting our research directions is the sense that while, historically, signal processing has principally emphasized numerical techniques, it will increasingly exploit a combination of numerical and symbolic processing, a direction that we refer to as knowledge-based signal processing.

In the area of knowledge-based signal processing, there are currently four projects underway. One involves the combination of numerical and symbolic processing in the context of pitch detection. The ideas being developed are potentially broadly applicable to event detection and marking in a variety of signal processing contexts. We are also developing knowledge-based signal processing within the framework of harmonic source detection and direction determination oriented toward acoustic sensor data. This work incorporates a rule-based system with signal processing algorithms. A third project in this area is the development of an interactive, signal-processing software environment incorporating both numerical and symbolic processing of signals and systems. In addition, we are currently working on algorithm development for detection and identification of artificial objects in infrared radar images. In this project, we hope to incorporate both signal processing and artificial intelligence techniques, and hope to exploit both the intensity and range map information.

In the area of speech processing, we have, over the past several years, worked on the development of systems for bandwidth compression of speech, parametric speech modeling, time-scale modification of speech and enhancement of degraded speech. Recently, based on a general class of algorithms involving the estimation of a signal after its short-time Fourier transform has been modified, we have developed a very high quality system for time-scale modification of speech and also a system for very reliable pitch detection. We are also exploring new techniques for speech enhancement using adaptive noise cancelling when multiple microphones are available.

In image processing, we are pursuing a number of projects on restoration, enhancement, and coding. One project is restoration of images degraded by additive noise, multiplicative noise, and convolutional noise. Our current work in this project involves development of a new approach to adaptive image restoration based on a cascade of one-dimensional adaptive filters. This approach, when applied to some existing image restoration systems, not only reduces the number of computations involved, but also improves the system performance. Another project we are currently exploring is the development of a very low bit-rate (around 20 Kbits/sec) video-conferencing system. The specific approach we are studying involves the transformation of an image into a one-bit intensity image by adaptive thresholding and then coding the one-bit intensity image by run-length coding.

In the area of geophysical signal processing, there are a variety of ongoing and new projects. Work on determining characteristics of the ocean bottom by inverting a measured acoustic field is continuing, and associated with this, we continue to explore algorithms for evaluating Hankel transforms. Also of potential importance to geophysical signal processing is our work on knowledge-based signal processing with distributed sensor nets, being carried out in collaboration with Lincoln Laboratory.

There are also a number of projects directed toward the development of new algorithms with broad potential applications. For some time we have had considerable interest in the broad question of signal reconstruction from partial information, such as Fourier transform phase or magnitude. We have shown theoretically how, under very mild conditions, signals can be reconstructed from Fourier transform phase information alone. This work has also been extended to the reconstruction of multidimensional signals from one bit of phase, and, exploiting duality, zero-crossing, and threshold crossing information. We have also developed a variety of theories and algorithms relating to signal reconstruction from Fourier transform magnitude and from partial short-time Fourier transform information. We are also exploring the application of some of these theoretical results to problems such as speech and image coding. We have also obtained significant results and continue to explore new algorithms for high-resolution multi-dimensional spectral estimation. Recently, we have also proposed a new approach to the problem of estimating multiple signal and/or parameter unknowns using incomplete and noisy data. Our Minimum Cross-Entropy Method applies an information-theoretic criterion to optimally estimate a separable probability density for the signal model. Not only does this new approach include all of the various maximum likelihood and maximum *a posteriori* methods as degenerate cases, but also it directly leads to a simple iterative method of solution in which we alternate between estimating the various unknowns, one at a time. We are now exploring applications to statistical problems, iterative signal reconstruction, short-time analysis/synthesis, and noisy pole/zero estimation. With the advent of VLSI technology, it is now possible to build customized computer systems of astonishing complexity for very low cost. Exploiting this capability, however, requires designing algorithms which not only use few operations, but also have a high degree of regularity and parallelism, or can be easily pipelined. The directions we are exploring include systematic methods for designing multi-processor arrays for signal processing, isolating signal processing primitives for hardware implementation, and searching for algorithms for multi-dimensional processing that exhibit a high degree of parallelism. We are also investigating highly parallel computer architectures for signal understanding, in which a mixture of intensive computation and symbolic reasoning must be executed in an integrated environment.

24.1 Low Bit-Rate Video Conferencing

National Science Foundation Fellowship
National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Jae S. Lim, Michele Covell

The attempt of this research is to achieve a drastic reduction in the channel capacity requirements for video transmissions. An algorithm based on local averaging and adaptive thresholding has been developed which reduces an 8-bit picture to a 1-bit binary picture without significant loss in quality. The resulting image is then smoothed by adaptive median filtering to remove random bit transitions. This is then suitable for run length coding. Compression factors of up to 60 are expected to result. Other forms of data encryption are also being examined. Currently the emphasis is only on intraframe redundancy reduction. The work is expected to lead to interframe analysis.

24.2 Knowledge-Based Pitch Detection

National Science Foundation (Grant ECS84-07285)
Sanders Associates, Inc.
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Webster Dove, Alan V. Oppenheim, Randall Davis

Knowledge-based signal processing (KBSP) is the application of the numerical algorithms of signal processing and the symbolic inference techniques of expert systems on problems where both numerical and symbolic information are present. One class of problems in this area includes those for which both numeric and symbolic input information is available. The Knowledge-Based Pitch Detection (KBPD) project is building a system to describe the excitation of a sentence of speech (voicing and periodicity) given a recorded waveform and a phonetic transcript of the sentence. Through this project we hope to learn more about the general problem of integrating numeric and symbolic processing.

This problem is interesting because pitch detection is an old problem with no completely successful solution (for high quality resynthesis). Existing programs attack the problem with a single unified algorithm based on the detection of periodicity. Typically, this will work some places and not others. For example, it is common for voiced sections at the ends of sentences to be irregular and not periodic. Algorithms based on periodicity cannot conclude voicing in such regions and therefore they make voiced-to-unvoiced errors there.

With more information available to it, the KBPD system can avoid such problems. Since it has a transcript, KBPD starts with strong information about the voicing of the sentence in the middle of the phonemes of the transcript. However, aligning the transcript precisely with the sentence so that the boundaries of the phonemes are correct requires interaction between the transcript representation and numerical measurements of the waveform. Also, the actual rendition of a given phoneme may vary enough between speakers and sentences that a normally voiced phoneme is not actually voiced. So the system must verify the transcript against the observed properties of the sentence to avoid mistakes.

By combining the results of several different methods of analysis rather that relying on a single algorithm, KBPD spreads the risk of errors. To the extent that these different methods both cooperate effectively and span the phenomena of the problem, one can expect the system to be more robust than a single algorithm approach. The key to this cooperation is the explicit representation of credibility in the results of each component.

Expert systems have capitalized on the representation of credibility for performing symbolic inference when both knowledge and problem information are uncertain. KBPD moves these techniques to the domain of numeric processing with the explicit representation in the program of both certainty and accuracy for numeric results. The extra information (beyond the basic result values) allows the system to merge results from different components without specific information about the nature of the components. This in turn allows the integration of new components into the system without major rewriting of existing code, and allows the final results to be based on the components that rated themselves most credible. While there is some question as to how reliably a given algorithm can be expected to rate its own performance, clearly this is an improvement over the conventional program assumptions that the algorithm's result "are the answer," and it has the virtue of presenting to the operator some information about the overall credibility of results.

This project is developing a program called the **Pitch Detector's Assistant** (PDA). It serves as both an assistant in the pitch detection process, and a workbench to evaluate numeric and symbolic processing techniques in pitch detection programming.

As an end product, it might be used to generate excitation information for off-line speech storage applications (like talking appliances) or for speech research into the properties of pitch as an information carrying medium. It is clearly not for real-time applications like vocoding, though it might be used for vocoding experiments.

As a workbench for pitch detection programming, the PDA allows us to assess what kind of symbolic information is useful in the pitch detection problem by experimenting with ways of using it.

As a KBSP catalyst, it has prompted the development of new data structures (for representing temporal estimates and pitch estimates). Furthermore, it has motivated the development of a signal processing environment (called KBSP) that dramatically reduces the time to implement signal processing algorithms. This software environment is also capable of embedding symbolic processing functions (like the generation of symbolic assertions) within numerical algorithms. That capacity, together with a rule-based system capable of invoking signal processing functions, permits a flexible, two-way communication between the numeric and symbolic parts of the system. Finally, this project has led to the clarification of some of the features of symbolic inference in the context of numeric information. Specifically, we have found that recognition of equivalent assertions is more difficult in a numeric context, and that there are more ways of combining numeric assertions than symbolic ones.

24.3 Iterative Algorithms for Parameter Estimation with Applications to Array Processing

National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Meir Feder, Ehud Weinstein[28]*, Alan V. Oppenheim*

In many signal processing problems one has to estimate unknown parameters. Usually this is done using the Maximum Likelihood (ML) criterion. Unfortunately, sometimes the parameters are connected to the observed signal in a complicated way and so maximizing the likelihood function

---

[28]Woods Hole Oceanographic Institute

is not an easy task.

The Estimate-Maximize (E-M) algorithm,[1] is an iterative algorithm that converges to the ML solution. We assume that there is an extension of the observations for which the ML estimate of the parameters is simple. We will call that extension "complete data." At each iteration we first estimate the sufficient statistics of the complete data given the observed signal and the previous values of the parameters (Estimate E-Step) and then get the new values of the parameters by maximizing the likelihood function of the complete data where we substitute the estimated statistics instead of the observed sufficient statistics (Maximize M-step).

Some examples of signal processing problems for which this idea is natural are:

* Parametric spectrum estimation from short data – the complete data will be a large data record.

* Signal parameter estimation when the signal is corrupted by noise – the complete data will be the signal alone and the noise alone.

* Array processing problems :

(i) multiple source problem – the complete data will be each source separately

(ii) multi-path problem – the complete data will be each path separately

(iii) time delay estimation – The complete data will be a large record for which the generalized cross-correlator is optimal.

In this research we will extend the E-M ideas, present a class of iterative algorithms, and determine their general properties, such as the convergence rate. Furthermore, we will suggest adaptive E-M algorithms and will consider the situation in which the number of the parameters is unknown. In addition, we will carry out the algorithm experimently on the array processing problems and try to establish this class of algorithms as a realistic procedure for solving those practical, but as yet unsolvable, problems.

Reference

1. A.P. Dempster, *et al.*, "Maximum Likelihood from Incomplete Data via the EM Algorithm," Ann. of Royal Stat. Soc., 1977, pp.1–38.

24.4 Speech Synthesis from Short-Time Fourier Transform Magnitude Derived from Speech Model Parameters

National Science Foundation (Grant ECS84-07285)
Sanders Associates, Inc.
U.S. Air Force - Office of Scientific Research (Contract F19628-85-K-0028)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Jae S. Lim, Daniel W. Griffin

Previous approaches to the problem of synthesis from estimated speech model parameters have primarily employed time-domain techniques. Most of these methods generate an excitation signal from the estimated pitch track. Then, this excitation signal is filtered with a time-varying filter obtained from estimates of the vocal tract response. This approach tends to produce poor quality "synthetic" sounding speech.

In this research, a new approach to speech synthesis from the same estimated speech model parameters is investigated. In this approach, a desired Modified Short-Time Fourier Transform Magnitude (MSTFTM) is derived from the estimated speech model parameters. The speech model parameters used in this approach are the pitch estimate, voiced/unvoiced decision, and the spectral envelope estimate. The desired MSTFTM is the product of a MSTFTM derived from the estimated excitation parameters and a MSTFTM derived from spectral envelope estimates. (The word "modified" is included in MSTFTM in these cases to emphasize that no signal exists, in general, with this MSTFTM.) Then, a signal with Short-Time Fourier Transform Magnitude (STFTM) close to this MSTFTM is estimated using the recently developed LSEE-MSTFTM algorithm.[1,2] Preliminary results indicate that this method is capable of synthesizing very high-quality speech, very close to the original speech.

This method has applications in a number of areas including speech coding and speech enhancement. In speech coding applications, the excitation parameters and spectral envelope can be coded separately to reduce transmission bandwidth. These coded parameters are transmitted and then decoded and recombined at the receiver. In speech enhancement applications, the excitation parameters and the spectral envelope can be separated, processed separately, and then recombined.

References

1. D.W. Griffin and J.S. Lim, "Signal Estimation from Modified Short-Time Fourier Transform," Proceedings of the 1983 International Conference on Acoustics, Speech and Signal Processing, Boston, Massachusetts, April 14-16, 1983, pp. 804-807.
2. D.W. Griffin and J.S. Lim, "Signal Estimation from Modified Short-Time Fourier Transform," IEEE Trans. Acoust. Speech Signal Process., vol. ASSP-32, no. 2, April 1984, pp. 236-243.

24.5 Reconstruction of a Two-dimensional Signal from its Fourier Transform Magnitude

National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
David Izraelevitz, Jae. S. Lim

In the case of an arbitrary multidimensional discrete signal, the phase and magnitude of the Fourier transform are independent functions of frequency. However, in many situations there is additional information regarding the signal which provides a very strong connection between the phase and magnitude. Specifically, it has been shown that almost all multidimensional signals which are non-zero only over a specified domain are uniquely specified, in a sense, by knowledge of the Fourier transform magnitude alone. Several algorithms have been developed for reconstructing such a signal from its Fourier transform magnitude; however, they all fall into either of two categories: they are heuristic algorithms which sometimes do not converge to the true reconstruction, or they are computationally too expensive for even moderate size signals.

In this research we present a new algorithm for reconstruction from Fourier transform magnitude which is a closed form solution to the problem and which has been used to reliably reconstruct signals of extent up to 20 by 20 pixels. The algorithm is based on posing the problem of Fourier transform magnitude reconstruction as requiring the explicit factorization of the z-transform of the autocorrelation sequence of the unknown function. This z-transform is easily computed from knowledge of the Fourier transform magnitude. A new procedure is developed for factoring large bivariate polynomials and this algorithm is applied to to the needed factorization of the autocorrelation z-transform.

We are presently looking at the data noise sensitivity of the algorithm and of the problem of reconstruction from Fourier transform magnitude in general. A detailed comparison of the behavior of the present algorithm with the previously proposed algorithms is also being developed.

24.6 Motion-Compensated Noise Reduction for Motion Video Scenes

Advanced Television Research Program
National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Dennis M. Martinez, Jae S. Lim

Motion pictures can be manipulated in a variety of ways by compensating for motion within the image sequence. An important component of all motion-compensated image processing systems is motion estimation. Previous approaches to motion estimation have encountered two primary problems, computational complexity and estimation accuracy. This work is concerned with the development and analysis of computationally efficient motion estimation algorithms which can determine motion trajectories very accurately.

A model-based motion estimation algorithm has been developed. This algorithm requires significantly less computation than traditional approaches. In addition it can determine velocity fields more accurately than commonly used region matching methods. The algorithm is based on a local three-dimensional signal model and a local translational velocity model. It is possible to estimate complex velocity fields encountered in real-life television images with the algorithm.

We have applied this algorithm to several problems in motion picture restoration and interpolation. Image restoration algorithms which operate on a single fromae at a time usually remove noise at the expense of pictue sharpness. However, we demonstrate that motion-compensated restoration systems can remove noise with little or no loss in picture sharpness.

The algorithm has also been used successfuly for frame interpolation. A motion-compensated frame interpolation system was developed which permits computing frames at arbitrary times. This system can be used in a variety of applications involving frame rate modification. A number of experiments have shown that this motion-compensated interpolation system produces motion pictures with better motion rendition than traditional frame repetition systems.

24.7 Interpretation-Guided Signal Processing of Acoustic Sensor Data

National Science Foundation (Grant ECS84-07285)
Sanders Associates, Inc.
U.S. Navy – Office of Naval Research (Contract N00014-81-K-0742)
Evangelos E. Milios, Alan V. Oppenheim

A signal processing system which integrates signal and symbol processing is being developed for acoustic waveform processing and interpretation. The design of the system was derived from the systematic observation (protocol collection) and subsequent analysis of human signal processing activity. The resulting system consists of: (1) a harmonic-set formation subsystem, which produces harmonic sets present in acoustic spectra and their symbolic description, and (2) a geometrical hypothesis formation and testing subsystem, which forms hypotheses about the acoustic source motion based on the data, and then performs detailed testing against the data. The system is being built using a hybrid methodology combining both procedural and declarative programming and accommodates both algorithmic and heuristic techniques in signal processing. Modules perform spectral peak analysis using rules that stem from human perceptual considerations. The rule-based approach is enhanced to allow thresholds to be set from examples through the same rules used for classification. In comparison to previous signal/symbol processing systems, which relied mostly on symbol processing, this system is based on the concept that a tight interaction of signal processing and interpretation can save a lot of symbol processing.

This work is being done in collaboration with the Machine Intelligence Technology Group at the M.I.T. Lincoln Laboratory.

24.8 Combined Numerical and Symbolic Signal Processing

Amoco Foundation Fellowship
National Science Foundation (Grant ECS84-07285)
Sanders Associates, Inc.
U.S. Navy – Office of Naval Research (Contract N00014-81-K-0742)
Cory Myers, Alan V. Oppenheim

One important aspect of our work on Knowledge Based Signal Processing (KBSP) is the development of a suitable environment for carrying out the research and for exploring mechanisms for integrating numerical and symbolic processing. As part of our KBSP work, we have developed the KBSP software package, an innovative software package for signal processing on the LISP Machine. The KBSP software package provides a uniform signal processing package that is integrated into an interactive, symbolic processing environment, the LISP Machine. The package provides computational speed, ease of development, and a close correspondence between abstract signal processing concepts and programming implementation.

As an outgrowth of the KBSP package, this research is devoted to the incorporation of symbolic manipulation facilities into the numerical KBSP package. Symbolic manipulation of signals involves the manipulation of signal representations rather than signal values. One example of this symbolic manipulation is the symbolic generation of Fourier transforms. The system understands many primitive signals, their Fourier transforms, and rules for the manipulation of Fourier transforms with respect to other signal processing operations, such as addition, multiplication, and convolution. For many signals, the system is able to parse the signal's representation and generate its Fourier transform without any reference to the numerical values of the signal.

One area in which the symbolic manipulation of signal representation is a natural one is the area of varying the signal processing according to "context." Context can include properties of the signals under consideration, properties of the hardware, or properties of the world. For example, different FFTs can be used when the signal is real-valued or when it is complex-valued, different filtering algorithms can be used to achieve different trade-offs among multipliers, adds, and memory references, and different spectral modeling techniques can be used for different speech sounds. The goal of this research is to automate the decision process used in algorithm selection.

As an initial step in automating the algorithm decision process, we have developed a system for the LPC analysis of a speech signal given a time-aligned phonetic transcript of the speech signal and a symbolic description of the recording environment. In this system, the parameters of the LPC analysis — the number of poles, the window size, and the frame rate — are varied according to the information present in the phonetic transcript. For example, a long window and a low frame rate are used in the steady-state portion of vowels and a short window and a high frame rate are used in the area of stop releases.

We are currently working on more general techniques that will allow many different sources of information to be used in choosing a signal processing algorithm. One aspect of this problem is the specification of properties of both signals and systems. We are trying to develop a symbolic signal processing language in which standard signal processing properties such as finite-durationness, symmetry, linearity, shift-invariance, etc., are easily represented and manipulated by the system.

We are also studying methods for the manipulation of signal representations so that the system can determine equivalent forms for algorithm implementation. For example, the system will be able to implement a filter in either the time-domain or the frequency-domain. This capability will be used to automatically choose different algorithms to implement the same signal processing operation for different input signals and for different trade-offs among the different factors that affect computational cost.

24.9 Estimation of Coronary Artery Boundaries in Angiograms

National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Thrasyvoulos N. Pappas, Jae. S. Lim

The precise and objective measurement of the severity of coronary obstructions from coronary angiograms is important in the treatment of patients with ischemic heart disease. An angiogram is an x-ray picture of the coronary arteries in which a contrast agent has been injected via a catheter.

Coronary artery imaging presents special problems because of the arteries location on the beating heart and their shape and size. As a result many techniques that have been used quite successfully for stenosis determination of other arteries, like the femoral and carotid, do not have satisfactory performance when applied to coronary arteries. These algorithms are quite heuristic and find the boundaries of the artery as the inflection points of a series of densitometric profiles perpendicular to the vessel image. Even though this approach is computationally simple, there is no theoretical justification for it.

We consider a different approach which more fully exploits the detailed characteristics of the signals involved. Specifically, we develop a model of the film density of the coronary angiograms and use it to estimate the diameter and cross-sectional area at each point along the vessel. Our model accounts for the structure of the vessel and background, as well as the distortions introduced by the imaging system (blurring and noise). We have developed both a one-dimensional model of the density profiles perpendicular to the vessel image, and a two-dimensional model of rectangular sections of the image. The parameters of the model include the vessel center-point locations and the radius at each point. The spatial continuity of the vessel is incorporated into the model and contributes to the accuracy of the estimation procedure. The algorithms are tested on synthetic data, on x-rays of contrast-medium-filled cylindrical phantoms obtained over a wide range of radiographic conditions, and on real coronary angiograms. Our results indicate that the 1-D algorithms have better performance than current methods, and preliminary results indicate that the 2-D algorithms have better performance than 1-D algorithms.

24.10 Speech Enhancement Using Adaptive Noise Cancellation

National Science Foundation Fellowship
National Science Foundation (Grant ECS84-07285)
U.S. Air Force – Office of Scientific Research Contract F19628-85-K-0028)
U.S. Navy – Office of Naval Research (Contract N00014-81-K-0742)
Jeffrey J. Rodriguez, Jae S. Lim

In many military environments, such as fighter jet cockpits, the increasing use of digital communication systems has created a need for robust vocoders and speech recognition systems. However, the high level of acoustic noise makes the vocoders less intelligible and makes reliable speech recognition more difficult. Therefore, we are using Widrow's Adaptive Noise Cancelling (ANC) algorithm to enhance the noise-corrupted speech.[1]

ANC is a noise-reduction method that uses multiple inputs. In the fighter jet application, we use two microphones: a primary and a reference. The primary microphone is located inside the pilot's oxygen face mask, and the reference microphone is attached to the exterior of the face mask. In this configuration, the primary microphone records the noise-corrupted speech, while the reference microphone ideally records only noise. When these two inputs are processed, the reference noise is filtered and subtracted from the primary signal. Hopefully, the resulting signal will contain the undegraded speech signal with very little noise.

Reference

1. B. Widrow, *et al.*, "Adaptive Noise Cancelling: Principles and Applications," Proc. IEEE, vol. 63, no. 12, December 1975.

24.11 Recovery of Undersampled Periodic Waveforms

National Science Foundation(Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Anthony J. Silva, Alan V. Oppenheim

Under certain circumstances it may be impossible to sample a signal at a rate greater than twice the highest spectral component present. In particular, the sampling rate might be limited by the A/D converter and associated hardware used. However, for certain classes of undersampled signals, complete recovery is still possible (at least theoretically), in spite of the well-known Nyquist criterion.

Periodic signals form one such class. Consider undersampling a periodic signal whose harmonic frequencies are mutually prime to the sampling rate. The harmonics have zero bandwidth[29] and though aliased into the baseband (-Fsamp/2 to +Fsamp/2 Hz), they do not overlap. The Nyquist criterion is stated generally for low-pass waveforms with energy presumably spread smoothly across their spectra. Because of this, it discounts the possibility of recovery after the non-destructive aliasing above.

In what follows, we assume the sampling rate is stable and known to a modest degree of accuracy. When the signal period is known, recovery is trivial. The time samples x[n] are sorted by interpreting the index "n" modulo the normalized signal period[30] then placing each sample in the appropriate place in a "composite" period. No samples will overlap if the mutual primality requirement above is met.

A far more interesting problem exists when the signal period is unknown. Rader[1] has presented an iterative recovery technique in which a multiplicity of trial signal periods is used as moduli in the sorting process above. The trial period yielding the best reconstruction is retained as the estimate of the true signal period. Results from number theory (Farey sequences, modular or

---

[29] Of course, the signal would have to be known and sampled for all time to yield truly zero-bandwidth harmonics

[30] The ratio of the signal and sampling periods.

congruential arithmetic, etc.) are exploited to make the approach practicable.

We will first search for ways to accelerate the Rader algorithm. An analogous frequency domain algorithm, also employing number theory, will then be developed. It will consist of determining the frequencies and amplitudes of the aliased harmonics, sorting them, and inverse transforming. Performance of all algorithms and their variants in the presence of noise, slightly wavering signal frequency and amplitude, and other undesirable conditions will be examined.

References

1. C.M. Rader, "Recovery of Undersampled Periodic Waveforms," IEEE Trans. Acoust. Speech, Signal Process., vol. ASSP-25, no. 3, June 1977.
2. G.H. Hardy and E.M. Wright, An Introduction to the Theory of Numbers, (Oxford University Press, 1956).

24.12 Relationships Between Opaque Objects and Silhouettes

National Science Foundation (Grant ECS84-07285)
U.S. Navy – Office of Naval Research (Contract N00014-81-K-0742)
Patrick Van Hove, Jae S. Lim

This research addresses the relationships between an object and the boundaries of its silhouettes, which are referred to as contours, corresponding to various three-dimensional (3-D) orientations of the line of sight. For this purpose, special models of objects and silhouettes are considered. The property sphere of an object is defined on a unit sphere which is related to the object by a 3-D Gaussian mapping. Similarly, the property circle is related to the contour it represents by a 2-D Gaussian mapping. In earlier computer vision work, property spheres and circles have been used independently, and their applications have been limited to the representation of scalar fields of object properties.

In a first stage, we have shown how the concepts of property spheres and circles can be usefully combined to relate the properties of an object and those of its contours. Specifically, it is proved that a slice through the property sphere of an object leads to the property circle of the contour corresponding to the line of sight perpendicular to that slice.

In a second stage, a new concept of object modeling has been introduced, where the property sphere is used as a domain for vector and tensor fields of object properties. In particular, a new representation of 3-D objects and 2-D silhouettes, referred to as the Curvature Transform (CT), maps the inverse of the curvature tensor field of the surface of an object on its property sphere, and the radius of curvature of the contour of a silhouette on its property circle. The key advantage of this representation is that a slice through the CT of an object followed by projection of the tensor field produces the CT of the contour corresponding to the line of sight perpendicular to the slice.

The study now progresses with attempts to use these new concepts in the reconstruction of object shape from silhouette information. Surface reconstruction procedures have been proposed in the field of machine vision, in the contexts of shape from photometric stereo, shading, texture, and motion. These procedures reconstruct a viewer-dependent 2 1/2-D sketch of the surfaces. The surface reconstruction procedure which is attempted here would provide a full 3-D sketch of the objects.

This work is being done in collaboration with the Distributed Sensor Systems group at the M.I.T. Lincoln Laboratory.

24.13 The Hilbert-Hankel Transform and Its Application to Shallow Water Ocean Acoustics

Hertz Foundation Fellowship
National Science Foundation (Grant ECS84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)
Michael S. Wengrovitz, Alan V. Oppenheim, George V. Frisk[31]

The problem of continuous-wave acoustic field propagation in shallow water is being investigated. In this environment, components of the field alternately reflect off both the ocean surface and the ocean bottom. In effect, the water can be considered as an acoustic waveguide, bounded by the ocean surface and the underlying ocean bottom. Several aspects of this waveguide propagation problem are being studied. The first concerns the development of an accurate numerical model to predict the magnitude and phase of the acoustic field as a function of range from the source, given the geoacoustic parameters of the water and the bottom. A technique has been developed which computes the field based on its decomposition into trapped (resonant) and continuum (non-resonant) components.

A second aspect being studied is the application of the Hilbert-Hankel transform to both the synthesis and inversion of these fields. This transform has a number of interesting and useful properties and forms the basis for a reconstruction method in which the real (imaginary) part of the complex-valued field is obtained from the imaginary (real) part.

Finally, the inherent sensitivity of extracting the bottom plane-wave reflection coefficient from measurements in the reverberant waveguide environment is being researched. Results indicate that there are several invariant waveguide parameters which do not depend on the properties of

---

[31] Woods Hole Oceanographic Institution

the underlying media. By exploiting the invariance of these parameters, it is possible to design an actual ocean experiment from which an improved reflection coefficient estimate results.

24.14 Computing the Discrete Hartley Transform

Hertz Foundation Fellowship

National Science Foundation (Grant ECS84-07285)

U.S. Navy – Office of Naval Research (Contract N00014-81-K-0742)

Avideh Zakhor, Alan. V. Oppenheim

The discrete Hartley transform (DHT) resembles the discrete Fourier transform (DFT) but is free from two characteristics of the DFT that are sometimes computationally undesirable. The inverse discrete Hartley transform is identical to the direct transform and so it is not necessary to keep track of +i and –i versions as with the discrete Fourier transform. Also, the DHT has real rather than complex values and thus does not require provision for complex arithmetic or separately managed storage for real and imaginary parts. The DFT is directly obtainable from the DHT by a simple additive operation.

We have found an efficient way of computing the Hartley transform which enables us to compute the DFT with half as many multiplications as the FFT. Wang[1] and Bracewell[2] have also found algorithms for the fast Hartley transform. In this research, we evaluated and compared these three algorithms in terms of their relative efficiencies. Statistical error properties of these algorithms were also investigated both theoretically and experimentally. This work was completed in June 1985.

References

1. Z. Wang, "Fast Algorithms for the Discrete W Transform and for the Discrete Fourier Transform," IEEE Trans. Acoust., Speech Signal Process., vol. ASSP-32, no. 4, August 1984.
2. R.N. Bracewell, "The Fast Hartley Transform," Proc. IEEE, vol. 72, no. 8, August 1984.