

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

and

CENTER FOR BIOLOGICAL INFORMATION PROCESSING
WHITAKER COLLEGE

A.I. Memo 773
C.B.I.P. Paper 001

April, 1984

ILL-POSED PROBLEMS AND REGULARIZATION ANALYSIS IN EARLY VISION

Tommaso Poggio and Vincent Torre

Abstract

One of the best definitions of early vision is that it is inverse optics — a set of computational problems that both machines and biological organisms have to solve. While in classical optics the problem is to determine the images of physical objects, vision is confronted with the inverse problem of recovering three-dimensional shape from the light distribution in the image. Most processes of early vision such as stereomatching, computation of motion and all the "structure from" processes can be regarded as solutions to inverse problems. This common characteristic of early vision can be formalized: *most early vision problems are "ill-posed problems" in the sense of Hadamard.* We will show that a mathematical theory developed for regularizing ill-posed problems leads in a natural way to the solution of early vision problems in terms of variational principles of a certain class. This is a new theoretical framework for some of the variational solutions already obtained in the analysis of early vision processes. It also shows how several other problems in early vision can be approached and solved.

© Massachusetts Institute of Technology, 1984

This report describes research done within the Artificial Intelligence Laboratory and the Center for Biological Information Processing (Whitaker College) at the Massachusetts Institute of Technology. Support for the A. I. Laboratory's research in artificial intelligence is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505. The Center's support is provided in part by the Sloan Foundation and in part by Whitaker College. Funds for collaborative travel were provided by NATO grant 237.81. Dr. Torre is at the Institute of Physics, University of Genoa, Italy.

Variational Solutions to Vision Problems

In recent years, the computational approach to vision has begun to shed some light on several specific problems. One of the recurring themes of this theoretical analysis is the identification of physical constraints that make a given computational problem determined and solvable. Some of the early and most successful examples are the analyses of stereomatching (Marr and Poggio, 1976, 1979; Grimson, 1981a,b; Mayhew and Frisby, 1981; Kass, 1984; for a review see Nishihara and Poggio, 1984) and structure from motion (Ullman, 1979). More recently, variational principles have been used to introduce specific physical constraints. For instance, visual surface interpolation can be derived from the minimization of functionals that embed a generic constraint of smoothness (Grimson, 1981b, 1982; Terzopoulos, 1983, 1984a). Computation of visual motion can be successfully performed by finding the smoothest velocity field consistent with the data (Horn and Schunck, 1981; Hildreth, 1984a,b) and shape can be recovered from shading information in terms of a variational method (Ikeuchi and Horn, 1981). The computation of subjective contours (Ullman, 1976; Brady et al., 1980; Horn, 1981), of lightness (Horn, 1974) and of shape from contours (Barrow and Tennenbaum, 1981; Brady and Yuille, 1984) can also be formulated in terms of variational principles. Terzopoulos (1984a, 1985) has recently reviewed the use of a certain class of variational principles in vision problems within a rigorous theoretical framework.

We wish to show that these variational principles follow in a natural and rigorous way from the ill-posed nature of early vision problems. We will then propose a general framework for "solving" many of the processes of early vision.

III-Posed Problems

In 1923, Hadamard defined a mathematical problem to be *well-posed* when its solution

(a) exists

(b) is unique

(c) depends continuously on the initial data (this means that the solution is robust against noise).

Most of the problems of classical physics are well-posed, and Hadamard argued that physical problems had to be well-posed. "Inverse" problems, however, are usually ill-posed. Inverse problems can usually be obtained from the direct problem by exchanging the role of solution and data. Consider, for instance,

$$y = Ax \tag{1}$$

where A is a known operator. The direct problem is to determine y from z , the inverse problem is to obtain z when y ("the data") are given. Though the direct problem is usually well-posed, the inverse problem is usually ill-posed¹.

Typical ill-posed problems are analytic continuation, back-solving the heat equation, superresolution, computer tomography, image restoration and the determination of the shape of a drum from its frequency of vibration, a problem which was made famous by Marc Kac (1966). In early vision, most problems are ill-posed because the solution is not unique (but see later the case of edge detection).²

Regularization Methods

Rigorous regularization theories for "solving" ill-posed problems have been developed during the past years (see especially Tikhonov, 1963; Tikhonov and Arsenin, 1977; and Nashed, 1974, 1976). The basic idea of regularization techniques is to restrict the space of acceptable solutions by choosing the function that minimizes an appropriate functional. The regularization of the ill-posed problem of finding z from the data y such that $Az = y$ requires the choice of norms $\|\cdot\|$ (usually quadratic) and of a *stabilizing functional* $\|Pz\|$. The choice is dictated by mathematical considerations, and, most importantly, by a physical analysis of the generic constraints on the problem. Three main methods can then be applied (see Bertero, 1982):

I) Among z that satisfy $\|Pz\| \leq C$ —where C is a constant—, find z that minimizes

$$\|Az - y\|, \quad (2)$$

II) Among z that satisfy $\|Az - y\| \leq C$, find z that minimizes

$$\|Pz\|, \quad (3)$$

III) Find z that minimizes

$$\|Az - y\|^2 + \lambda \|Pz\|^2, \quad (4)$$

where λ is a regularization parameter.

The first method consists of finding the function z that satisfies the constraint $\|Pz\| \leq C$ and best approximates the data. The second method computes the function z that is sufficiently close to the data (C depends on the estimated errors and is zero if the data are noiseless) and is most "regular". In the third method, the regularization parameter λ controls the compromise between the degree of regularization of the solution and its closeness to the data. Regularization theory provides techniques to determine the best λ (Tikhonov and Arsenin, 1977; Wahba, 1980). It also provides a large body of results about the form of the *stabilizing functional* P that ensure uniqueness of the result and convergence. For instance, it is possible to ensure uniqueness in the case of Tikhonov's stabilizing functionals (also called *stabilizers of p -th order*) defined by

$$\|Pz\|^2 = \int \sum_{r=0}^p p_r(\xi) \left(\frac{d^r z}{d\xi^r} \right)^2 d\xi. \quad (5)$$

Equation (5) can be extended in the natural way to several dimensions. If one seeks regularized solutions of eq.(1) with P given by eq. (5) in the Sobolev space W_2^p of functions that have square-integrable generalized derivatives up to p -th order, the solution can be shown to be unique (up to the null space of P), if A is linear and continuous. This is because for every p the space W_2^p is a Hilbert space and $\|Pz\|^2$ is a quadratic functional (see theorem 1, Tikhonov and Arsenin, 1977; p. 63). It turns out that most stabilizing functionals used so far in early vision are of the Tikhonov type (see also Terzopoulos, 1984a,b).³ They all correspond to either interpolating or approximating splines (for method II and method III, respectively).




Figure 1. Decomposition and ambiguity of the velocity field. a) The local velocity vector $V(s)$ in the image plane is decomposed according to eq.(6) into components perpendicular and tangent to the curve. b) Local measurements cannot measure the full velocity field: the circle undergoes pure translation: the arrows represent the perpendicular components of velocity that can be measured from the images. From Hildreth, 1984a.

Example I: Motion

Our first claim is that variational principles introduced recently in early vision for the problem of shape from shading, computation of motion, and surface interpolation are exactly equivalent to regularization techniques of the type we described. The associated uniqueness results are directly provided by regularization theory. We briefly discuss the case of motion computation in its recent formulation by Hildreth (1984a,b).

Consider the problem of determining the two-dimensional velocity field along a contour in the image. Local motion measurements along contours provide only the component of velocity in the direction perpendicular to the contour. Figure 1 shows how the local velocity vector $V(s)$ is decomposed into a perpendicular and a tangential component to the curve

$$V(s) = v^T(s)T(s) + v^\perp(s)N(s) \quad (6)$$

The component v^\perp and direction vectors $T(s)$ and $N(s)$, are given directly by the initial measurements, the "data". The component $v^T(s)$ is not and must be recovered to compute the full two-dimensional velocity field $V(s)$. Thus the "inverse" problem of recovering $V(s)$ from the data v^\perp is ill-posed because the solution is not unique. Mathematically, this arises because the operator K defined by

$$v^\perp = KV \quad (7)$$

is not injective. Equation (7) describes the imaging process as applied to the physical velocity field V which consists of the x and y components of the three-dimensional velocity field of the object.

Intuitively, the set of measurements given by $v^\perp(s)$ over an extended contour should provide considerable constraint on the motion of the contour. An additional generic constraint, however, is needed to determine this motion uniquely. For instance, rigid motion on the plane is sufficient to determine V uniquely but is very restrictive, since it does not cover the case of motion of a rigid object in space. Hildreth suggested, following Horn and Schunck (1981), that a more general constraint is to find the smoothest velocity field among the

set of possible velocity fields consistent with the measurements. The choice of the specific form of this constraint was guided by physical considerations — the real world consists of solid objects with smooth surfaces whose projected velocity field is usually smooth — and by mathematical considerations — especially uniqueness of the solution. Hildreth proposed two algorithms: in the case of exact data the functional to be minimized is a measure of the smoothness of the velocity field

$$\|P\mathbf{V}\|^2 = \int \left(\frac{\partial \mathbf{V}}{\partial s} \right)^2 ds \quad (8)$$

subject to the measurements $v^\perp(s)$. Since in general there will be error in the measurements of v^\perp , the alternative method is to find \mathbf{V} that minimizes

$$\beta \|K\mathbf{V} - v^\perp\|^2 + \int \left(\frac{\partial \mathbf{V}}{\partial s} \right)^2 ds, \quad (9)$$

where $\beta = \frac{1}{\lambda}$. It is immediately seen that these schemes correspond to the second and third regularizing method respectively. Uniqueness of the solutions (proved by Hildreth¹ for the case of equation (8)) is a direct consequence for both equations (8) and (9) of standard theorems of regularization theories. In addition, other results can be used to characterize how the correct solution converges depending on the smoothing parameter λ .

Example II: Edge Detection

We have recently applied regularization techniques to another classical problem of early vision – edge detection. Edge detection, intended as the process that attempts to detect and localize changes of intensity in the image (this definition does not encompass all the meanings of edge detection) is a problem of numerical differentiation (Torre and Poggio, 1984). Notice that differentiation is a common operation in early vision and is not restricted to edge detection. The problem is ill-posed because the solution does not depend continuously on the data.⁵ The intuitive reason for the ill-posed nature of the problem can be seen by considering a function $f(x)$ perturbed by a very small (in L_2 norm) “noise” term $\epsilon \sin \Omega x$. $f(x)$ and $f(x) + \epsilon \sin \Omega x$ can be arbitrarily close, but their derivatives may be very different if Ω is large enough.

In 1-D, numerical differentiation can be regularized in the following way. The “image” model is $y_i = f(x_i) + \epsilon_i$, where y_i is the data and ϵ_i represent errors in the measurements. We want to estimate f' . We chose a regularizing functional $\|P'f\| = \int (f''(x))^2 dx$, where f'' is the second derivative of f . The second regularizing method (no noise in the data) is equivalent then to using interpolating cubic splines for differentiation.⁶ The third regularizing method, which is more natural since it takes into account errors in the measurements, leads to the variational problem of minimizing

$$\sum (y_i - f(x_i))^2 + \lambda \int (f''(x))^2 dx. \quad (10)$$

Poggio et al. (1984) have shown (a) that the solution of this problem f can be obtained by convolving the data y_i (assumed on a regular grid) with a convolution filter R , and (b) that the filter R is a cubic spline⁷ with a shape very close to a Gaussian and a size controlled by the regularization parameter λ (see figure 2). Differentiation can then be accomplished by convolution of the data with the appropriate derivative of this filter. The optimal value of λ can be determined for instance by cross validation and other techniques. This corresponds to finding the optimal scale of the filter.⁸




Figure 2. The edge detection filter. a) The convolution filter obtained by regularizing the ill-posed problem of edge detection with method (III) (see Poggio et al., 1984). It is a cubic spline (solid line), very similar to a gaussian (dotted line). b) The first derivative of the filter for different values of the regularizing parameter λ , which effectively controls the *scale* of the filter. This one-dimensional profile can be used for two-dimensional edge detection by filtering the image with oriented filters with this transversal cross-section and choosing the orientation with maximum response (see Canny, 1983).

These results can be directly extended to two dimensions to cover both edge detection and surface interpolation and approximation. The resulting filters are very similar to two of the edge detection filters derived and extensively used in recent years (Marr and Hildreth 1980; Canny, 1983; see Torre and Poggio, 1984).

Other problems in early vision such as shape from shading (Ikeuchi and Horn, 1981) and surface interpolation (Grimson, 1981b 1982; Terzopoulos, 1983, 1984), in addition to the computation of velocity, have already been formulated and "solved" in similar ways using variational principles of the type suggested by regularization techniques (although this was not realized at the time). It is also clear that other problems such as stereo⁹ and structure from motion¹⁰ can be approached in terms of regularization analysis.

Physical Plausibility of the Solution

Uniqueness of the solution of the regularized problem—which is ensured by formulations such as equations (2)-(4) – is not the only (or even the most relevant) concern of regularization analysis. Physical plausibility of the solution is the most important criterion. The decision regarding the choice of the appropriate stabilizing functional cannot be made judiciously from purely mathematical considerations. A physical analysis of the problem and of its generic constraints have the upper hand. Regularization theory provides a framework within which one has to seek constraints that are rooted in the physics of the visual world. This is, of course, the challenge of regularization analysis. Conditions characterizing the physically correct solutions can be derived¹¹ (for the case of motion, see Yuille, 1983 and

for edge detection, see Poggio et al., 1984).

From a more biological point of view, a careful comparison of the various "regularization" solutions with human perception promises to be a very interesting area of research, as suggested by Hildreth's work. For some classes of motions and contours, the solution of equations (8) and (9) is not the physically correct velocity field. In these cases, however, the human visual system also appears to derive a similar, incorrect velocity field (Hildreth, 1984a,b).

Conclusion

The concept of ill-posed problems and the associated regularization theories seem to provide a satisfactory theoretical framework for part of early vision. This new perspective justifies the use of variational principles of a certain type for solving specific problems, and suggests how to approach other early vision problems. It provides a link between the computational (ill-posed) nature of the problems and the computational structure of the solution (as a variational principle). In a companion paper (Poggio and Koch, 1984), we will discuss computational "hardware" that is natural for solving variational problems of the type implied by regularization methods. The approach can be extended to other sensory modalities and to some motor control problems. For instance, a recently proposed solution to the problem of executing a voluntary arm trajectory (Hogan, 1984) can be recognized as an instance of our second regularization technique.¹²

Despite its attractions, this theoretical synthesis of early vision also shows the limitations that are intrinsic to the variational solutions proposed so far, and in any case to the simple forms of the regularization approach. The basic problem is the degree of smoothness required for the unknown function z that has to be recovered. If z is very smooth, then it will be robust against noise in the data, but it may be too smooth to be physically plausible. For instance, in visual surface interpolation, the degree of smoothness obtained with the thin plate model (from a specific form of equations (4)-(5)) smoothes depth discontinuities too much and often leads to unrealistic results (but see Terzopoulos, 1984).

These problems may be solved by more sophisticated regularization techniques, such as stochastic methods. The simple regularization techniques analyzed here rely on quadratic variational principles that lead to linear Euler-Lagrange equations. Thus the solution can be found by filtering the data through an appropriate linear filter. Analog electrical or chemical networks can be devised for the specific variational principles (Poggio and Koch, 1984). Again, the universe of solutions to quadratic variational principles is somewhat restricted.

Nonquadratic variational principles are, however, possible. They may arise naturally in one of the most fundamental problems in early vision, the problem of integrating different sources of information, such as stereo, motion, shape from shading, etc. This problem is ill-posed, not just because the solution is not unique (the standard case), but because the solution is usually overconstrained and may not exist. The use and extensions of tools from regularization theory to analyze the fusion of information from different sources is one of the most interesting challenges in the theory of early vision.^{13,14}

The problem is related to the deep question of the computational organization of a visual processor and its control structure. It is unlikely that variational principles alone could have enough flexibility to control and coordinate the different modules of early vision and their interaction with higher level knowledge. This also hints at the basic limitation of regularization methods that makes them, suitable only for the first stages of vision. They derive numerical representations—surfaces—from numerical representations—images. It is difficult to see how the computation of the more symbolic type of representations that are essential for a powerful vision processor can fit into this theoretical framework¹⁵.

In summary, we have outlined a new theoretical framework that from the computational nature of early vision leads to algorithms for solving them, and suggests a specific class of appropriate hardware. The common computational structure of many early vision problems is that they are ill-posed in the sense of Hadamard. Regularization analysis can be used to solve them in terms of variational principles of a certain type that enforce constraints derived from a physical analysis of the problem. Analog networks—whether electrical or chemical—are a simple and attractive way of solving this type of variational principles.

Acknowledgments: The idea of vision problems as ill-posed problems originated from a conversation with Professor Mario Bertero of the University of Genoa about the edge-detection work with V. Torre. Work on visual problems as ill-posed problems is now going on at MIT with Alan Yuille, Christof Koch, Harry Voorhees and at the University of Genoa with Alessandro Verri. The idea of formulating vision problems in terms of variational principles is due to a number of people, mostly at the A.I. Lab at MIT. Among them, we would like to mention Eric Grimson, Demetri Terzopoulos, Berthold Horn, Shimon Ullman, Mike Brady, Ellen Hildreth and Alan Yuille. We are grateful to A. Yuille, E. Hildreth, D. Terzopoulos and C. Koch for many discussions and comments and for critically reading the manuscript.

Footnotes

[1] Whether a problem is well- or ill-posed depends on the triplet (A, Z, Y) where Z and Y are the solution and the data space respectively.

[2] The reason for the lack of uniqueness is that the operator corresponding to A is usually not injective, as in the case of shape from shading, surface interpolation and computation of motion.

To clarify some of the structure of ill-posed problems, let us consider the linear operator

$$Az = y. \quad (1)$$

If z and y are finite vectors, then the inverse problem is easily solved by finding the inverse of A , or its pseudoinverse. It is well known that if A is a square matrix, A^{-1} exists if $\det|A| \neq 0$.

Now let us suppose that $z \in Z$ and $y \in Y$, where Z and Y are Hilbert spaces. The inverse problem is well-posed *iff* the three conditions of Hadamard are satisfied. In particular,

(1) condition (a) of Hadamard is satisfied *iff* the range of A is $R(A) = Y$.

(2) condition (b) of Hadamard is satisfied *iff* A is injective.

(3) condition (c) of Hadamard is satisfied *iff* $R(A)$ is a closed set.

If the operator A is compact and $R(A)$ does not have finite dimensions, $R(A)$ is open, and therefore the inverse problem is also ill-posed.

Most linear operators whose domain and co-domain are Hilbert spaces are compact operators. In fact, if E and F are measurable, bounded sets $E \in \mathcal{R}^n$ and $F \in \mathcal{R}^m$, and $k(t, s)$ is a measurable function defined on $E \times F$, then the linear operator $A: L_2(E) \mapsto L_2(F)$ defined as

$$(Az)(t) = \int k(t, s)z(s)ds \quad (2)$$

is compact and $R(A)$ has finite dimensions *iff* $k(t, s)$ is separable, i.e.,

$$k(t, s) = \sum_{k=1}^n \alpha_k(t)\beta_k(s). \quad (3)$$

Obviously if $R(A)$ has finite dimension, then $R(A)$ cannot coincide with Y , and therefore the inverse problem of an integral operator or a convolution is in general an ill-posed problem.

We can relax condition (2) and admit the case that A is not injective. The problem is then regularized by introducing an appropriate norm and finding the generalized pseudoinverse of the inverse problem (1).

When y is not in $R(A)$, it is not easy to regularize the problem without altering the essence of the problem itself.

[3] J. Canny's (1983) variational formulation can be derived from eq. (4) and a stabilizing functional of the form of eq. (5) (see Poggio et al., 1984).

[4] It is shown in Hildreth (1984a) that extremizing equation (8) yields a unique velocity field, since it corresponds to minimizing a positive definite functional on a convex set. The theorems of du Bois-Reymond state that, provided $\frac{\partial \mathbf{V}}{\partial \mathbf{x}}$ is continuous the solution of the minimization problem will be the solution of the corresponding Euler-Langrange equations.

[5] The problem is to find the solution z to

$$y = Az$$

with $(Az)(x) = \int_0^x z(s)ds$. Thus, z is the derivative of the data y . The problem is (mildly) ill-posed because if $z \in L_2[0, 1]$, the compact operator A is not closed in $L_2[0, 1]$.

[6] For data on a regular grid, it corresponds to convolving the data with the L_4 filters of Schoenberg (1946).

[7] A higher degree stabilizer may be used for higher derivatives, leading to higher order splines.

[8] Methods such as the Generalized Cross Validation method (GCV) (Wahba, 1980; see also Reinsch, 1967) may be used to find the "optimal" scale of the filter, i.e., the optimal λ . *Fingerprints* (Yuille and Poggio, 1983) may provide a method for finding the optimal value of the regularization parameter λ . This follows from the fact that the filter given by equation (10) is very similar to a Gaussian and that λ effectively controls the scale of the filter (see Poggio et al., 1984).

[9] Another clearly ill-posed problem is stereo-matching. It is not immediately obvious, however, what the correct regularizing procedure is. Berthold Horn has suggested (personal communication) a variational principle for stereo-matching similar to his scheme for computing optical flow. The norm to be minimized measures deviations from smoothness of the disparity field. Specifically, the norm of the derivative of the z component, the depth component, has to be minimized subject to the constraints given by the data. This can be regarded again as a variational principle of the type that is obtained directly using the standard regularization methods of ill-posed problems. We are presently developing regularization solutions to the stereo problem (Yuille and Poggio, in preparation).

The problem of shape-from-contours in the variational formulation of Brady and Yuille (1983) is an ill-posed problem but the solution is not of the standard regularization type.

[10] The rubbery constraint proposed by Ullman (1983) is more general than the rigidity constraint. It may be possible to reformulate it according to regularization techniques.

[11] A method for checking physical plausibility of a variational principle is, of course, computer simulation. A simple technique we suggest is to use the Euler-Lagrange equation associated with the variational problem.

In the computation of motion, Yuille (1983) has obtained the following *sufficient and necessary* condition for the solution of the variational principle equation (8), to be the correct physical solution

$$T \cdot \frac{\partial^2 V}{\partial s^2} = 0$$

where T is the tangent vector to the contour and V is the true velocity field. The equation is satisfied by uniform translation or expansion and by rotation only if the contour is polygonal. These results suggest that algorithms based on the smoothness principle will give correct results, and hence be useful for computer vision systems, when (a) motion can be approximated locally by pure translation, rotation or expansion, or (b) objects have images consisting of connected straight lines. In other situations, the smoothness principle will not yield the correct velocity field, but may yield one that is qualitatively similar and close to human perception (Hildreth, 1984a,b).

In the case of edge detection (intended as numerical differentiation), the solution is correct *if and only if* the intensity profile is a polynomial spline of odd degree greater than three (Poggio et al., 1984).

[12] The variational principle (minimization of jerk) corresponds to the second regularization method, with $P = d^3/dx^3$. The associated interpolating function is a quintic spline. Analog

networks for solving the problem can be devised (Poggio and Koch, 1984). It may be interesting to consider our third method of regularization in the context of the available data on arm trajectories.

[13] The variational principles that we have considered so far for early vision processes are quadratic and lead therefore to linear equations. The ill-posed problem of combining several different sources of surface information may easily lead to non-quadratic regularization expressions (though different "non-interacting" constraints can be combined in a convex way, see Terzopoulos, 1984). These minimization problems will in general have multiple local minima. Schemes similar to annealing (Kirkpatrick, Gelatt and Vecchi 1983; Hinton and Sejnowski, 1983; Geman and Geman, 1984) may be used to find the global minimum (see also Poggio and Koch, 1984).

[14] This is a list of open problems on which we are presently working:

- a) Regularized solution for stereo matching.
- b) Regularized solution for structure from motion.
- c) Full extension of the edge detection analysis to 2-D and application to surface approximation for computing differential properties of surfaces.
- d) Analysis and implementation of methods for finding the optimal regularization parameter λ . Use of fingerprints.
- e) Connection between the regularizing parameter λ , the iteration number in iterative regularizing methods (Nashed, 1976) and the truncation of a formal power series expansion of the regularizing operator.
- f) Use of stochastic regularization methods (see also Geman and Geman, 1984).

[15] But see Hummel and Zucker, 1980.

References and related papers

- Baker, H. H. and Binford, T. O., "Depth from Edge and Intensity Based Stereo," *Seventh International Joint Conference on Artificial Intelligence*, August 1981a, 631-636.
- Barrow, H.G. and Tennenbaum, J.M. "Interpreting line drawings as three dimensional surfaces" *Artif. Intelligence*, **17**, 75-117, 1981.
- Bertero, M. "Problemi lineari non ben posti e metodi di regolarizzazione", in: *Problem non ben posti ed inversi*, Istituto di Analisi Globale, Firenze, 1982.
- Brady, J.M., Grimson, W.E.L., and Langridge, D.J. "Shape encoding and subjective contours" *First Annual Natl. Conf. on Artificial Intelligence* **15-17**, 1980
- Brady, J.M., and Yuille, A., "An extremum principle for shape from contour", *I.E.E.E Trans. on Pattern Analysis and Machine Intelligence*, **288,301**, 1984.
- Canny, J. F., "Finding Edges and Lines in Images," Massachusetts Institute of Technology Artificial Intelligence Laboratory Technical Report TR-720, June 1983.
- Courant, R., and Hilbert, D., *Methods of Mathematical Physics*, Interscience Publishers, London, 1962.
- Geman, S. and Geman, D. "Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images", *PAMI*, in press, 1984.
- Grimson, W.E.L., "A computer implementation of a theory of human stereo vision", *Phil. Trans. R. Soc. Lond., B*, **292**, 217-253, 1981a.
- Grimson, W.E.L., *From Images to Surfaces: A computational study of the human early visual system*, MIT Press, Cambridge, Ma., 1981b.
- Grimson, W.E.L., "A computational theory of visual surface interpolation", *Phil. Trans. R. Soc. Lond., B*, **298**, 395-427, 1982.
- Hadamard, J., *Lectures on the Cauchy Problem in Linear Partial Differential Equations*, Yale University Press, New Haven, 1923.
- Hildreth, E.C., "Implementation of a theory of edge detection," S.M. Thesis, Department of Computer Science and Electrical Engineering, Massachusetts Institute of Technology, 1980. (see also MIT AI Lab Technical Report 597, 1980).
- Hildreth, E.C., *The Measurement of Visual Motion*, MIT Press, Cambridge, Massachusetts, 1984a.
- Hildreth, E.C., "Computation of the velocity field", *Proc. R. Soc. Lond. B*, **221**, 189-220, 1984b.
- Hinton, G.E., and Sejnowski, T.J., "Optimal perceptual inference", *Proc. IEEE 1983 Conf. Computer Vision and Pattern Recognition*, Washington, DC, 1983.
- Hogan, N. "An organising principle for voluntary movements". Submitted to *J. Neurosci.*, 1984.
- Horn, B.K.P., "Determining lightness from an image", *Computer Graphics and Image Processing*, **3**, 111-299, 1974.
- Horn, B.K.P., "The least energy curve", *AI Memo 612, MIT AI Lab.*, Cambridge, MA, 1981
- Horn, B.K.P., and Schunck, B.G., "Determining optical flow", *Artificial Intelligence*, **17**, 185-203, 1981.
- Hummel, R. and Zucker, S. "On the foundations of relaxation labeling processes." McGill University, Computer Vision and Graphics Laboratory, TR 80-7, 1980.

- Ikeuchi, K., and Horn, B.K.P., "Numerical shape from shading and occluding boundaries", *Artificial Intelligence*, **17**, 141-184, 1981.
- Kac, M. "Can one hear the shape of a drum?" *Am. Math. Monthly*, **73**, No. 4, Part II, 1-23, 1966.
- Kass, M., "A computational framework for the visual correspondence problem", *Proc. 8th Int. J. Conf. AI*, Karlsruhe, W. Germany, 1043-1045, 1983.
- Kass, M.H. "Computing stereo correspondance", *M. S., Mass. Institute of Technology*, 1984.
- Kirkpatrick, S., Gelatt, C.D., Jr., Vecchi, M.P., "Optimization by simulated annealing", *Science*, **220**, 671-680, 1983.
- Marr, D., and Hildreth, E.C., "Theory of edge detection", *Proc. R. Soc. Lond. B*, **207**, 187-217, 1980.
- Marr, D. and Poggio, T. "Cooperative computation of stereo disparity", *Science*, **194**, 283-287, 1976.
- Marr, D. and Poggio, T. "A theory of human stereo vision," *Proc. Roy. Soc. Lond. B* **204** (1979), 301-328. (an earlier version appeared as MIT AI Lab Memo 451, 1977).
- Mayhew, J.E.W. and Frisby, J.P. "Psychophysical and computational studies towards a theory of human stereopsis," *Artificial Intelligence*, **17**, 349-385, 1981.
- Nashed, M. Z., "Approximate regularized solutions to improperly posed linear integral and operator equations in *Constructive and Computational Methods for Differential and Integral Equations* G. Auger, ed., Akademie Verlag, Berlin, 289-296, 1974.
- Nashed, M. Z., ed. *Generalized Inverses and Applications*, Academic Press, New York, 1976.
- Nishihara, H.K. and Poggio, T., "Stereo vision for robotics", *Proceedings of the First International Symposium of Robotics Research*, in press.
- Poggio, T. and Koch, C., "Analog networks: a new approach to neuronal computation", *Artificial Intelligence Lab. Memo*, No. 783, MIT, Cambridge, MA, 1984.
- Poggio, G. and Poggio, T., *Annual Review of Neuroscience*, **7**, 379-412, 1984.
- Poggio, T., Voorhees, H. and Yuille A., "Regularizing Edge Detection", *Artificial Intelligence Lab. Memo*, No. 776, MIT, Cambridge, 1984.
- Reinsch, C.H., "Smoothing by spline functions", *Numer. Math.*, **10**, 177-183, 1967.
- Schoenberg, I. J., "Contributions to the problem of approximation of equidistant data by analytic functions." *Quart. Appl. Math.*, **4**, 45-99, 112-141, 1946.
- Schunck, B.G., and Horn, B.K.P., "Constraints on optical flow computation", *Proc. IEEE Conf. Pattern Recognition and Image Processing*, 205-210, 1981.
- Terzopoulos, D., "Multilevel computational processes for visual surface reconstruction", *Computer Vision, Graphics, and Image Processing*, **24**, 52-96, 1983.
- Terzopoulos, D., "Multiresolution computation of visible-surface representations," Ph.D. Thesis, Dept of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, January, 1984a.
- Terzopoulos, D., "Multilevel reconstruction of visual surfaces: variational principles and finite element representations " *Artificial Intelligence Lab. Memo 671*, reprinted in *Multiresolution Image Processing and Analysis*, A. Rosenfeld (ed.), 237-310, Springer-Verlag, 1984b.
- Terzopoulos, D., "Integrating visual information from multiple sources for the cooperative computation of surface shape", to appear in *From pixels to predicates: recent advances in computational and robotic vision*, ed. A. Pentland, Ablex, 1985.

- Tikhonov, A. N. "Solution of incorrectly formulated problems and the regularization method", *Soviet Math. Dokl.*, **4**, 1035-1038, 1963.
- Tikhonov, A.N. and Arsenin, V.Y., *Solutions of ill-posed problems*, Winston & Sons, Washington, D.C., 1977.
- Torre, V. and Poggio, T., "On Edge Detection", *Artificial Intelligence Lab. Memo*, **No. 768**, MIT, Cambridge, MA, March 1984.
- Ullman, S., "Filling in the gaps: The shape of subjective contours and a model for their generation" *Biological Cybernetics* **25**, 1-6, 1976.
- Ullman, S., "The interpretation of structure from motion", *Proc. R. Soc. Lond. B*, **203**, 405-426, 1979.
- Ullman, S., "Maximizing rigidity: the incremental recovery of 3-D structure from rigid and rubbery motion" *Artificial Intelligence Lab. Memo*, **No. 721**, MIT, Cambridge, MA, June 1983.
- Wahba, G., "Ill-posed problems: Numerical and statistical methods for mildly, moderately and severely ill-posed problems with noisy data", *Tech. Rep.*, **No. 595**, Univ. of Wisconsin, Madison, 1980.
- Yuille, A., "The smoothest velocity field and token matching schemes" *Artificial Intelligence Lab. Memo*, **No. 724**, MIT, Cambridge, MA, August 1983.
- Yuille, A., "Zero crossings on lines of curvature", *Artificial Intelligence Lab. Memo*, **No. 718**, MIT, Cambridge, MA, 1984.
- Yuille, A.L. and Poggio, T. "Fingerprints Theorems for Zero-crossings", *Artificial Intelligence Lab. Memo*, **No. 730**, MIT, Cambridge, MA, 1983b.
- Yuille, A.L. and Poggio, T. "A generalized Ordering Constraint for Stereo Correspondance", *Artificial Intelligence Lab. Memo*, **No. 777**, MIT, Cambridge, MA, 1984.