

VARIATIONAL SURFACE INTERPOLATION FROM SPARSE POINT AND NORMAL DATA

Jan Erik Solem¹ Henrik Aanæs² Anders Heyden¹

¹School of Technology and Society ²Informatics and Mathematical Modelling
Malmö University, Sweden Technical University of Denmark, Denmark
{jes,heyden}@ts.mah.se haa@imm.dtu.dk

Abstract

Many visual cues for surface reconstruction from known views are sparse in nature, e.g. specularities, surface silhouettes and salient features in an otherwise textureless region. Often these cues are the only information available to an observer. To allow these constraints to be used either in conjunction with dense constraints such as pixel-wise similarity, or alone, we formulate such constraints in a variational framework. We propose a sparse variational constraint in the level set framework, enforcing a surface to pass through a specific point *and* a sparse variational constraint on the surface normal along the observed viewing direction, as is the nature of e.g. specularities. These constraints are capable of reconstructing surfaces from extremely sparse data. The approach has been applied and validated on the shape from specularities problem.

Index terms: variational methods, computer vision, level set method, shape from specularities, multiple view stereo, surface interpolation.

1 Introduction

Reconstructing a surface from multiple known views is one of the main tasks in computer vision and occurs as a step in structure from motion systems or pre-calibrated camera networks. This is done using the various visual cues in the recorded images, where the standard approach is the similarity of pixel values, usually measured via correlation [3, 5]. Many other cues are typically available, and should ideally be used because they reduce the uncertainty of the estimate and often yield information of surface parts where correlation does not. One example of this is observed specularities in textureless areas.

Many of these additional cues are naturally sparse in nature, and can usually be specified as constraints on the surface position or its normal. Examples include surface silhouettes and specularities. Surface silhouettes specify the surface normal along the viewing direction of the observed surface contour (assuming smooth surfaces). Specularities specify the surface normal along the direction of the observed specularity, given a known light source e.g. the sun. The utilization of observed specularities is our original motivation for this work.

Although all these cues are sparse, a dense surface estimate is usually required, so a method for interpolating a dense surface from sparse data is needed. Here we propose a variational scheme for solving this problem. We formulate constraints for a surface passing through sparse points and having a specific normal along a line (typically the viewing direction) in a variational setting. We will use *sparse* in the sense that less than 50 constraints can reconstruct a surface as opposed to the dense points that can be obtained by a laser scanner or the dense normals obtained from photometric stereo. This distinction, among others, has implications for the formulation of the point constraint, although the proposed method also works for closely spaced constraints. The constraints are incorporated in a level set framework and along with a standard

surface prior gives a viable interpolating scheme. The generality of the variational formulation also allows the proposed constraints to be used alongside other visual cues e.g. image correlation. The proposed constraints can as such be seen as a step towards an all encompassing variational surface reconstruction scheme. Other recent attempts at this are [7, 18].

A first attempt to use implicit surfaces for shape from specularities with sparse data was made in [15] using a non-variational formulation. This was later improved in [16] with the variational setting presented here. Shape has also been recovered from reflections using different experimental setups cf. e.g. [1, 4, 11]. It is interesting to note that normal constraints also appear naturally from observations of apparent contours, cf. e.g. [13].

The main contribution of this paper is to give a method for surface reconstruction within the level set framework for sparse data consisting of a combination of both normals and points. Functionals and their corresponding level set gradient flows are proposed for solving both sparse normal alignment and surface interpolation of sparse point sets. As a result, a novel functional incorporating both of these components is explicitly derived to solve the shape from specularities problem with sparse normal data. This variational approach makes it possible to perform a rigorous analysis of shape from specularities for implicit surfaces. It should be noted that the method proposed in this paper can also be used for reflections of known 3D features cf. [11], not only light sources. The variational formulation makes it straight-forward to incorporate our method with variational stereo methods [3, 18], resulting in stereo methods that can handle specular reflections.

2 Background

2.1 Constraints from reflections and contours

In a structure and motion setting there are two common types of image data that give rise to normal data at unknown depths.

The first is from observations of specularities in images, cf. e.g. [16]. The condition for specular reflection is that the surface normal bisects the viewing direction and the incident light direction.

At the point where the back-projected ray (from the camera center through the image point of the specularity) intersects the surface the normal is

$$\mathbf{N} = \frac{\mathbf{l} - \mathbf{r}}{|\mathbf{l} - \mathbf{r}|}, \quad (1)$$

where \mathbf{r} is the directional vector for each ray, normalized so that $|\mathbf{r}| = 1$ and \mathbf{l} is the unit direction to the light source. It is important to note that the depths of the points where these constraints apply are unknown. The second type of image data is apparent contours. The tangent vector \mathbf{t} of the image curve γ corresponding to an apparent contour together with the camera center define a plane π for each point on γ which should coincide with the tangent plane at the surface.

At the intersection of the surface and a back-projected apparent contour point, the normal is

$$\mathbf{N} = \pm \frac{\mathbf{t} \times \mathbf{r}}{|\mathbf{t} \times \mathbf{r}|}, \quad (2)$$

where, as above, \mathbf{r} is the directional vector for the back-projection, cf. [13]. We will enumerate the constraints as \mathbf{N}_i , where the index i runs over all observations.

These *normal constraints* are not enough to determine the surface, in that they contain no depth information. Requiring that the surface also interpolates a set of points obtained from e.g. structure from motion resolves this ambiguity. However, the normal and point constraints do not uniquely define the surface and regularization or interpolation is needed. The goal is then to *find the smoothest surface that interpolates a set of points and has a given normal at the*

intersection with a fixed set of lines.

2.2 Level set surface representation

We will use an implicit level set surface representation. The level set method was introduced independently by [2] and [9] as a tool for capturing moving interfaces. The time dependent surface $\Gamma(t)$ is represented as the zero level set of a function $\phi(\mathbf{x}, t) : \mathbf{R}^3 \times \mathbf{R}_+ \rightarrow \mathbf{R}$ as $\Gamma(t) = \{\mathbf{x} ; \phi(\mathbf{x}, t) = 0\}$. The sets $\{\mathbf{x} : \phi(\mathbf{x}) < 0\}$ and $\{\mathbf{x} : \phi(\mathbf{x}) > 0\}$ are called the *inside* and the *outside* of Γ , respectively. Using this convention, the outward unit normal \mathbf{n} and the mean curvature κ of Γ are given by¹ (cf. [17])

$$\mathbf{n} = \frac{\nabla\phi}{|\nabla\phi|} \quad \text{and} \quad \kappa = \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|} . \quad (3)$$

One important, frequently used example is the signed distance function, where the additional requirement $|\nabla\phi| = 1$ is imposed.

To evolve the surface according to some derived velocity \mathbf{v} , a PDE of the form

$$\frac{\partial\phi}{\partial t} + \mathbf{v} \cdot \nabla\phi = 0 \quad \text{or} \quad \frac{\partial\phi}{\partial t} + v_n |\nabla\phi| = 0 , \quad (4)$$

is solved on a fixed grid in some domain. Here $v_n = \mathbf{v} \cdot \mathbf{n}$ is the velocity normal to the surface.

For more details on the level set method and dynamic implicit surfaces, cf. [10].

3 Variational Formulation

In a variational formulation, an energy functional is defined such that the minima correspond to desired solutions. The main purpose of this paper is to propose a functional that incorporates the deviation from all the sparse normal constraints and all the local point constraints in one single

¹Here $\nabla\phi$ denotes the gradient of ϕ , $\nabla\phi = (\phi_x, \phi_y, \phi_z)$, and $\nabla \cdot$ denotes the divergence.

expression. This expression then contains two data terms, one for the normals, cf. Section 3.1, and one for the points, cf. Section 3.2. Since the data is sparse, some form of regularization is also needed. This section describes the components of this novel functional. The corresponding level set motion is derived as a gradient descent of the functional, cf. [14]. This motion makes the energy decrease until equilibrium. The desired surface is then found by solving the motion PDE and minimizing this functional. The constraints, as defined in Section 2.1, are localized around points. Since all functions are only defined on a fixed grid it is necessary to increase the volume affected by each constraint since we can only “measure” at these grid points. More details can be found in [16].

3.1 Normal constraints

Given a set of normals \mathcal{N} and a set of corresponding points \mathcal{X} where each normal constraint $\mathbf{N}_i \in \mathcal{N}$ is associated with a point in space² $\mathbf{x}_i \in \mathcal{X}$, we would like to minimize the difference between the desired normal and the normal of the surface Γ . To incorporate the normal constraints, we introduce an energy functional as the L_p norm of the angular deviation of the surface normal from the desired normal

$$E_N(\Gamma, \mathcal{N}, \mathcal{X}) = \left(\int_{\Gamma} (1 - N(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x}))^p d\sigma \right)^{1/p}, \quad (5)$$

where $N(\mathbf{x})$ is a unit-length extension of the desired normals \mathbf{N}_i at points \mathbf{x}_i to the whole domain and $\mathbf{n}(\mathbf{x})$ is the surface normal. Here $N(\mathbf{x})$ is defined as

$$N(\mathbf{x}) = \begin{cases} \mathbf{N}_i & \text{if } \mathbf{x} \in B_\epsilon(\mathbf{x}_i) \\ \mathbf{n}(\mathbf{x}) & \text{otherwise} \end{cases}, \quad (6)$$

²The point \mathbf{x}_i is given by the intersection of the back-projected ray and the surface Γ and it is updated as the surface evolves.

where $B_\epsilon(\mathbf{x}_i)$ is a ball with radius ϵ centered around \mathbf{x}_i . The width of $B_\epsilon(\mathbf{x}_i)$ is typically chosen to enclose the nearest grid point. This definition of $N(\mathbf{x})$ will make the integrand equal to zero where there are no constraints on the normal. The functional (5) is mentioned in [10] and recently in [6]. The functional gives a gradient descent evolution equation as

$$\frac{\partial \phi}{\partial t} = \left(\nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} - \nabla \cdot N(\mathbf{x}) \right) |\nabla \phi|, \quad (7)$$

for $p = 1$, independent of the choice of $N(\mathbf{x})$. For details see [16] or the gradient descent framework in [14].

Note that for $p = 1$, measuring the normal deviation using $(1 - N(\mathbf{x}) \cdot \frac{\nabla \phi}{|\nabla \phi|})$ is equivalent to using $\frac{1}{2} |N(\mathbf{x}) - \frac{\nabla \phi}{|\nabla \phi|}|^2$, i.e. the squared difference between the desired normal and the surface normal, since $|N(\mathbf{x}) - \frac{\nabla \phi}{|\nabla \phi|}|^2 = |N(\mathbf{x})|^2 - 2N(\mathbf{x}) \cdot \frac{\nabla \phi}{|\nabla \phi|} + |\frac{\nabla \phi}{|\nabla \phi|}|^2 = 2(1 - N(\mathbf{x}) \cdot \frac{\nabla \phi}{|\nabla \phi|})$ due to the fact that the vectors are of unit length. Also note that it is easy to change the definition of $N(\mathbf{x})$ in (6) to make $N(\mathbf{x})$ differentiable without changing the results. However, this change would not make any difference in practice since, at the implementation stage, we work in a discrete setting. It is easy to show that for the motion (7), the energy $E_N(\phi, \mathcal{N}, \mathcal{X})$ decreases until it reaches equilibrium, cf. [16].

3.2 Point constraints

As for the normal constraints, we need to introduce an appropriate measure for the deviation of the surface from the point constraints. A special consideration of this measure is that it should not interfere with the normal constraints or the surface orientation. This means that it is desirable to have some form of local measure.

The surface should interpolate a given set of points \mathcal{S} . It is tempting to introduce an energy

functional similar to [19], which penalizes the deviation of the surface from the points in \mathcal{S} , as

$$E_P(\Gamma, \mathcal{S}) = \left(\int_{\Gamma} d(\mathbf{x})^p d\sigma \right)^{1/p}, \quad (8)$$

where Γ is an arbitrary surface and $d(\mathbf{x}) = \text{dist}(\mathbf{x}, \mathcal{S})$ is the distance to the set of points \mathcal{S} where the point constraints apply.

With sparse data, this has the drawback that the surface will continue to deform in the direction of the negative gradient of the distance potential $d(\mathbf{x})$ even after the points are on the surface.

This will also interfere with the local orientation of the surface, as will be shown in Section 4.

Instead, punishing the deviation of the points in \mathcal{S} from the surface Γ can be expressed using the energy functional

$$E_P(\phi, \mathcal{S}) = \frac{1}{2} \sum_{\mathbf{y} \in \mathcal{S}} \phi(\mathbf{y})^2 = \frac{1}{2} \sum_{\mathbf{y} \in \mathcal{S}} \int_{\mathbf{R}^3} \phi(\mathbf{x})^2 \delta(\mathbf{x} - \mathbf{y}) d\mathbf{x}, \quad (9)$$

which will minimize the *shortest* (squared) distance from the points to the surface if $|\nabla\phi| = 1$.

Since (9) is only evaluated on a discrete set of points in the domain and all functions are only given values at these points, the following regularization is used

$$E_P(\phi, \mathcal{S}) = \frac{1}{2} \int_{\mathbf{R}^3} \phi(\mathbf{x})^2 \delta_{\varepsilon}(\mathbf{x}, \mathcal{S}) d\mathbf{x}, \quad (10)$$

where $\delta_{\varepsilon}(\mathbf{x}, \mathcal{S}) = \sum_{\mathbf{y} \in \mathcal{S}} \delta_{\varepsilon}(\mathbf{x} - \mathbf{y})$ and $\delta_{\varepsilon}(\cdot)$ is a smoothed version of the Dirac delta function with some width ε . In practice one can use a uniform Gaussian with standard deviation ε . The evolution equation for ϕ as the gradient descent of (10) is then

$$\frac{\partial\phi}{\partial t} = -\phi(\mathbf{x})\delta_{\varepsilon}(\mathbf{x}, \mathcal{S}), \quad (11)$$

which is what was proposed in [15] in a more ad hoc fashion. This approach is also related to the point set attractors in [8] and the approach in [12]. One can show that the motion (11) indeed makes the energy $E_P(\phi, \mathcal{S})$ decrease until equilibrium, cf. [16].

3.3 Total energy

We are now ready to formulate the total energy containing the two data terms above and a regularizing term. Define the total energy E_{Tot} to be minimized as

$$E_{Tot} = E_N + E_P + \alpha E_{Smooth} , \quad (12)$$

where E_{Smooth} is a regularizing smoothness term and $\alpha > 0$ is a constant determining the amount of smoothing. The usual choice is to use surface area $E_{Smooth}(\Gamma) = \int_{\Gamma} 1 \, d\sigma$, which leads to a mean curvature motion $\frac{\partial \phi}{\partial t} = \kappa |\nabla \phi|$. Putting all this together gives

$$E_{Tot} = \int_{\Gamma} ((1 + \alpha) - N(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})) \, d\sigma + \frac{1}{2} \int_{\mathbf{R}^3} \phi(\mathbf{x})^2 \delta_{\varepsilon}(\mathbf{x}, \mathcal{S}) \, d\mathbf{x} . \quad (13)$$

The resulting PDE for the motion is then

$$\frac{\partial \phi}{\partial t} = ((1 + \alpha)\kappa - \nabla \cdot N(\mathbf{x})) |\nabla \phi| - \phi(\mathbf{x}) \delta_{\varepsilon}(\mathbf{x}, \mathcal{S}) , \quad (14)$$

as a combination of (7), (11) and the regularization term.

Note that from (14) we see that the regularizing term E_{Smooth} could have been incorporated in E_N by changing the definition of $N(\mathbf{x})$ in (6). By setting $N(\mathbf{x}) = (1 - \alpha)n(\mathbf{x})$ outside the set where the constraints apply instead of $N(\mathbf{x}) = n(\mathbf{x})$, we could have used $E_{Tot} = E_N + E_P$ instead of (12). This would amount to the same expression since then $(\kappa - \nabla \cdot N(\mathbf{x})) = (\kappa - (1 - \alpha)\nabla \cdot n(\mathbf{x})) = \alpha\kappa$ outside the points where the normal constraints apply. We choose, for the sake of clarity, to model the regularization separately from the normal alignment.

One can show that for the motion (14), if ε and α are small, the total energy E_{Tot} decreases until equilibrium under the assumption that the normal constraints do not coincide with the point set \mathcal{S} , cf. [16].

3.4 Local minima and initialization

If we want to use a gradient descent approach like the one described above, an initial surface is needed. The surface is then evolved from this initial value until an extremal of the functional is found. As is the case with many minimization problems in computer vision, our approach suffers from the fact that there might be several local minima and stationary points to the energy functional (13). This depends on the point constraints, normal constraints and the amount of regularization imposed. Minimizing (13) using the gradient descent (14) will result in finding one of these stationary points. Which stationary point is found depends on the initial values.

In practice, this leads to the problem of finding a good initial guess for the surface. Since the point constraints are only effective near the points in the set \mathcal{S} , an initial guess should interpolate these points. To do this the method described in [19] can be used. The surface is then found as a minimizer of (8) as described above. The corresponding level set evolution equation for $p = 1$ is then

$$\frac{\partial \phi}{\partial t} = (\nabla d(\mathbf{x}) \cdot \mathbf{n} + d(\mathbf{x})\kappa) |\nabla \phi| , \quad (15)$$

where, as before, κ is the mean curvature, \mathbf{n} is the normal and $d(\mathbf{x})$ is the distance potential from (8). Denote the minimizing steady state solution of (8) by ϕ_0 .

If another initial surface that still interpolates the points in \mathcal{S} is desired, weighted passive convection of the surface can be used

$$\phi = \phi_0 + c(\mathbf{x})d(\mathbf{x})|\nabla \phi_0| , \quad (16)$$

where $c(\mathbf{x}) : \mathbf{R}^3 \rightarrow \mathbf{R}$. The factor $d(\mathbf{x})$ is, as above, the distance from \mathbf{x} to \mathcal{S} and makes sure that the zero set still interpolates the points in \mathcal{S} . We have found that by varying $c(\mathbf{x})$ we can change the initial surface such that the desired local minimum is obtained.

The final algorithm can then be summarized as

1. Find initial guess for Γ , e.g. by minimizing (8).
2. Use (16), if necessary, to avoid possible undesired local minima.
3. Minimize the functional corresponding to the total energy (13) to obtain the surface.

4 Experiments

To validate the proposed approach, it has been applied to the shape from specularities problem, which was the original motivation for this work. In particular it was applied to two real examples, where the specular information was needed for a 3D reconstruction, see Figures 1 and 2. The specular reflections are selected manually and camera calibration is computed using structure from motion techniques [5]. In Figure 1 the light direction is found using an image of the camera shadow and in Figure 2 the light source is the camera flash. The computation time was 1 and 18 seconds respectively on a 2.8GHz, 2Gb RAM Pentium. Sample images and reconstructions with the part of the surface inside the convex hull of the constraint positions are shown. From these experiments it is concluded that the proposed algorithm works on this extremely sparse data. It is also seen that without both the normal and the point constraints the results would be of considerably lower quality. Lastly it is noted that the performed optimization worked well in that the constraints of the data were satisfied for synthetic data. For more details on these experiments the reader is referred to [16], where additional experiments can be found together with an error analysis on synthetic data.

To illustrate why the functional (8) is inappropriate for our purposes, experiments where this approach was compared to using (10) was performed. The results are shown in Figure 3. Using (8) for the interpolation results in a convection in the $-\nabla d$ direction of all parts of the surface where $\nabla d \cdot \nabla \phi \neq 0$. This is clearly shown in Figure 3 where the surface evolution results in the orientation of the normals to be orthogonal to ∇d .



Figure 1: (left) An image from a sequence of 38 images of a car window together with the image used to find the light direction. The two point constraints available are the two sharp corners indicated with arrows. The reflectance of the sun gave 38 normal constraints. (right) The reconstructed window, which is clearly what one would expect and substantially better than what we could achieve from the two point constraints.



Figure 2: (left) Two images from a sequence of 25 of a vacuum cleaner. The camera flash supplied 25 normal constraints alongside the 22 point constraints. (middle) The estimated 3D shape with normal and point constraints. (right) Reconstruction with only point constraints.



Figure 3: The errors introduced by using (8) instead of (10) for the interpolation of the points in \mathcal{S} . The results are after 300 iterations with the proposed point constraint (10) to the right. The underlying shape, where the constraints are sampled, is $1/8$ of a sphere.

5 Conclusions

A variational method for surface reconstruction from a set of extremely sparse point and normal constraints has been proposed. It has been applied to the shape from specularities problem and good results have been achieved from only a limited number of constraints, directly linked to the limited number of images available. It is noted that the proposed method is not limited to the shape from specularities problem, and has successfully been applied to other problems where extremely sparse point and normal constraints occur.

Acknowledgement: The authors would like to thank Niels Chr. Overgaard for useful comments on the first version of the manuscript.

References

- [1] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. In *Int. Conf. Computer Vision*, pages 591–596, Nice, France, 2003.
- [2] A. Dervieux and F. Thomasset. A finite element method for the simulation of Rayleigh–Taylor instability. In R. Rautman, editor, *Approximation Methods for Navier–Stokes Problems*, volume 771 of *Lecture Notes in Mathematics*, pages 145–158. Springer, Berlin, 1979.
- [3] O. Faugeras and R. Keriven. Variational principles, surface evolution, pdes, level set methods, and the stereo problem. *Image Processing, IEEE Transactions on*, 7(3):336–344, 1998.

- [4] M. A. Halstead, B. A. Barsky, S. A. Klein, and R. B. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline fitting of normals. In *Siggraph*, pages 335–342, 1996.
- [5] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [6] H. Jin, D. Cremers, A.J. Yezzi, and S. Soatto. Shedding light on stereoscopic segmentation. In *Proc. Conf. Computer Vision and Pattern Recognition*, Washington DC, 2004.
- [7] H. Jin, A. Yezzi, and S. Soatto. Variational multiframe stereo in the presence of specular reflections. Technical Report TR01-0017, UCLA, 2001.
- [8] K. Museth, D.E. Breen, R.T. Whitaker, and A.H. Barr. Level set surface editing operators. *ACM Transactions on Graphics*, 21(3):330–8, 2002.
- [9] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [10] S. J. Osher and R. P. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer Verlag, 2002.
- [11] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *International Journal of Computer Vision (IJCV)*, 64:31–67, 2005.
- [12] Y. Shi and W.C. Karl. Shape reconstruction from unorganized points with a data-driven level set method. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 3, pages 13–16, 2004.

- [13] J. E. Solem and F. Kahl. Surface reconstruction from the projection of points, curves and contours. In *2nd Int. Symposium on 3D Data Processing, Visualization and Transmission, Thessaloniki, Greece, 2004*.
- [14] J. E. Solem and N.C. Overgaard. A geometric formulation of gradient descent for variational problems with moving surfaces. In *The 5th International Conference on Scale Space and PDE methods in Computer Vision, Scale Space 2005, Hofgeismar, Germany*, pages 419–430. Springer, 2005.
- [15] J. E. Solem, H. Aanæs, and A. Heyden. Pde based shape from specularities. In Lewis D. Griffin and Martin Lillholm, editors, *Scale-Space Theories in Computer Vision, 4th International Conference, Scale Space 2003, Isle of Skye, UK*, volume 2695 of *Lecture Notes in Computer Science*. Springer, 2003.
- [16] J. E. Solem, H. Aanæs, and A. Heyden. A variational analysis of shape from specularities using sparse data. In *2nd Int. Symposium on 3D Data Processing, Visualization and Transmission, Thessaloniki, Greece, 2004*.
- [17] J. A. Thorpe. *Elementary Topics in Differential Geometry*. Springer-Verlag, 1985.
- [18] Anthony Yezzi and Stefano Soatto. Stereoscopic segmentation. *International Journal of Computer Vision*, 53(1):31–43, 2003.
- [19] H.K. Zhao, S. Osher, B. Merriman, and M. Kang. Implicit and non-parametric shape reconstruction from unorganized points using a variational level set method. In *Computer Vision and Image Understanding*, pages 295–319, 2000.