Shape and Radiance Estimation from the Information Divergence of Blurred Images*

Paolo Favaro and Stefano Soatto

Washington University
Department of Electrical Engineering
Electronic Signals and Systems Research Lab
One Brookings dr. 1127, St.Louis - MO 63130, USA
Tel. (314)935-7340, fax -7500, e-mail {fava,soatto}@ee.wustl.edu
Web http://ee.wustl.edu/~soatto/research

Abstract. We formulate the problem of reconstructing the shape and radiance of a scene as the minimization of the information divergence between blurred images, and propose an algorithm that is provably convergent and guarantees that the solution is admissible, in the sense of corresponding to a positive radiance and imaging kernel. The motivation for the use of information divergence comes from the work of Csiszár [5], while the fundamental elements of the proof of convergence come from work by Snyder et al. [14], extended to handle unknown imaging kernels (i.e. the shape of the scene).

1 Introduction

An imaging system, such as the eye or a video-camera, involves a map from the three-dimensional environment onto a two-dimensional surface. In order to retrieve the spatial information lost in the imaging process, one can rely on prior assumptions on the scene and use pictorial information such as shading, texture, cast shadows, edge blur etc. All pictorial cues are intrinsically ambiguous in that prior assumptions cannot be validated using the data.

As an alternative to relying on prior assumptions, one can try to retrieve spatial information by looking at different images of the same scene taken, for instance, from different viewpoints (parallax), such as in *stereo* and *motion* (note that we must still rely on prior assumptions in order to solve the correspondence problem). In addition to changing the position of the imaging device, one could change its *geometry*. For instance, one can take different photographs of the same scene with a different lens aperture or focal length. Similarly, in the eye one can change the shape of the lens by acting on the lens muscles. There is a sizeable literature on algorithms to reconstruct shape from a number of images

* This research was supported by NSF grant IIS-9876145 and ARO grant DAAD19-99-1-0139. The authors wish to thank J. C. Schotland and J. A. O'Sullivan for useful discussions and suggestions, and S. Nayar and M. Watanabe for kindly providing us with the test images used in the experimental section.

D. Vernon (Ed.): ECCV 2000, LNCS 1842, pp. 755–768, 2000.

taken with different imaging geometry (shape from defocus) or from a controlled search over geometric parameters (shape from focus) [4]. Recently the conditions under which it is possible to obtain a unique reconstruction have been derived [10].

Estimating shape from focus/defocus boils down to inverting certain integral equations, a problem known by different names in different communities: in signal processing it is "blind deconvolution" or "deblurring", in communications and information theory "source separation", in image processing "restoration", in tomography "inverse scattering". Since images depend both on the shape of the scene and on its reflectance properties – neither of which is known – estimating shape is tightly related to estimating reflectance[1]. In this paper, we consider the two problems as one and the same, and discuss the simultaneous solution of both. We choose as criterion the minimization of the information divergence (I-divergence) between blurred images, motivated by the work of Csiszár [5]. The algorithm we propose is iterative, and we give a proof of its convergence to a (local) minimum. We present results on both real and simulated images.

1.1 Statement of the Problem

Consider a piecewise smooth surface represented symbolically by σ. For instance, σ could be the parameters in a parametric class of surfaces, or it could be a smooth function such that $\sigma(x, y, z) = 0$ (note that it may not necessarily be finite-dimensional). Consider then an imaging system whose geometry can - to a certain extent - be modified by acting on some parameters $u \in \mathcal{U} \subset \mathbb{R}^k$ for some k. For instance, u could be the aperture radius of the lens and the focal length. The image at a point (x, y) in a compact subset of the plane $D \subset \mathbb{R}^2$ is obtained by integrating the energy radiated by the surface[1], which we represent as a (non necessarily continuous) positive-valued distribution R defined on σ, over a region that depends upon u. Due to the additive nature of energy transport phenomena, the image is obtained by integrating the energy distribution R against a kernel h that depends upon σ and u. We therefore write

$$I_u(x, y) = \int h_u^\sigma(x, y) dR \qquad (x, y) \in D. \tag{1}$$

Notice that, in the equation above, all quantities are constrained to be *positive*[2]: dR because it represents the radiant energy (which cannot be negative), h^σ because it specifies the region of space over which energy is integrated, and I because it measures the photon count hitting the surface of the sensor.

We are interested in estimating the shape of the surface σ and the energy distribution R to the extent possible, by measuring a number l of images obtained with different camera settings $u_1, \ldots, u_l$.

[1] Since neither the light source nor the viewer move, we do not distinguish between the radiance and reflectance of a surface.

[2] We use the term "positive" for a quantity x to indicate $x \geq 0$. When $x > 0$ we say that x is "strictly positive".

In the literature of computational vision a number of algorithms have been proposed to estimate depth from focus/defocus [1,6,7,9,11,12,13,16,18,19,20] just to mention a few. Most of the papers formulate the problem as the minimization of a least-squares norm or total variation.

The capability to reconstruct the scene's shape depends upon the energy distribution it radiates. The conditions on the radiance distribution that allow a unique reconstruction of shape have been recently characterized in [10].

1.2 Formalization of the Problem

If we collect a number of different images $I_{u_1}, \ldots, I_{u_l}$ and organize them into a vector I (and so for the kernels h), we can write

$$I(x,y) = \int h^\sigma(x,y)dR \qquad (x,y) \in D. \tag{2}$$

The right-hand side of the above equation can be interpreted as the "virtual image" of a given surface σ radiating energy with a given distribution R. We call such virtual image b:

$$b^\sigma(x,y,R) \doteq \int h^\sigma(x,y,X,Y,Z)dR(X,Y,Z) \qquad (x,y) \in D. \tag{3}$$

Note that, for images of opaque objects, the integral is restricted to their surface, and therefore can be written in the Riemannian sense [2] as

$$b^\sigma(x,y,R) \doteq \int h^\sigma(x,y,\tilde{x},\tilde{y})dR(\tilde{x},\tilde{y}) \qquad (x,y) \in D \tag{4}$$

for a suitably chosen parameterization $(\tilde{x},\tilde{y}) \in \mathbb{R}^2$. In either case, we write the integral in short-hand notation as $b^\sigma(x,y,R) \doteq \int h^\sigma(x,y)dR$. Since the image I is measured on the pixel grid, the domain D (i.e. a patch in the image) is $D = [x_1, \ldots, x_N] \times [y_1, \ldots, y_M]$, so that we have

$$I(x_i,y_j) = b^\sigma(x_i,y_j,R) \quad i = 1 \ldots N, j = 1 \ldots M. \tag{5}$$

We now want a "criterion" ϕ to measure the discrepancy between the measured image I and the virtual one, so that we can formulate our problem as the minimization of the discrepancy between the measured image and the model (or virtual) image. Common choices of criteria include the least-squares distance between I and b^σ, or the integral of the absolute value ("total variation") of their difference [3].

In order to get a "reasonable" result, the criterion ϕ should satisfy a number of requirements. Csiszár makes this notion rigorous through the axiomatic derivation of cost functions that satisfy certain consistency conditions. He concludes that, when the quantities involved are constrained to be positive[3] (such

[3] When there are no positivity constraints, Csiszár argues that the only consistent choice of discrepancy criterion is the $\mathcal{L}^2$ norm, which we have addressed in [15].

as in our case), the only consistent choice of criterion is the so-called information divergence, or *I-divergence*, which generalizes the well-known Kullbach-Leibler pseudo-metric and is defined as

$$\phi(I, b^\sigma(R)) \doteq \sum_{i,j} \left\{ I(x_i, y_j) \log \frac{I(x_i, y_j)}{b^\sigma(x_i, y_j, R)} - I(x_i, y_j) + b^\sigma(x_i, y_j, R) \right\}. \quad (6)$$

In order to emphasize the dependency of the cost function ϕ on the unknowns σ, R, we abuse the notation to write

$$\phi = \phi(\sigma, R). \quad (7)$$

Therefore, we formulate the problem of simultaneously estimating the shape of a surface and its radiance as that of finding σ and R that minimize the I-divergence:

$$\sigma, R = \arg\min_{\sigma,R} \phi(\sigma, R). \quad (8)$$

1.3 Alternating Minimization

In general, the problem in (8) is nonlinear and infinite-dimensional. Therefore, we concentrate our attention from the outset to (local) iterative schemes that approximate the optimal solution. To this end, suppose an initial estimate of R is given: R_0. Then iteratively solving the two following optimization problems

$$\begin{cases} \sigma_{k+1} \doteq \arg\min_\sigma \phi(\sigma, R_k) \\ R_{k+1} \doteq \arg\min_R \phi(\sigma_{k+1}, R) \end{cases} \quad (9)$$

leads to the minimization of ϕ, since

$$\phi_{k+1} \doteq \phi(\sigma_{k+1}, R_{k+1}) \leq \phi(\sigma_{k+1}, R_k) \leq \phi(\sigma_k, R_k) \doteq \phi_k \quad (10)$$

and the sequence ϕ_k is bounded below by zero. However, solving the two optimization problems in (9) may be an overkill. In order to have a sequence $\{\phi_k\}$ that monotonically converges it is sufficient that - at each step - we choose σ and R in such a way as to guarantee that equation (10) holds, that is

$$\begin{cases} \sigma_{k+1} \mid \phi(\sigma_{k+1}, R) \leq \phi(\sigma_k, R) & R = R_k \\ R_{k+1} \mid \phi(\sigma, R_{k+1}) \leq \phi(\sigma, R_k) & \sigma = \sigma_{k+1}. \end{cases} \quad (11)$$

2 Minimizing I-Divergence

In this section we derive an algorithm to minimize the I-divergence and prove its convergence. For simplicity, we restrict our analysis to an "equifocal imaging model", that is a model where the kernel h is translation-invariant. This corresponds to the scene being approximated, locally, by small patches of a plane parallel to the lens. This can be done to an arbitrary degree anywhere on a smooth surface away from discontinuities, which will therefore be resolvable only up to the size of the patch.

2.1 An Elementary Imaging Model

In order to obtain a simple instance of the problem, we assume that, locally for (x_i, y_j) in a patch $U \subset D$ away from discontinuities of σ, the kernel $h^\sigma(x_i, y_j, \tilde{x}, \tilde{y})$ is shift-invariant, so that we can write it as

$$h^\sigma(x_i - \tilde{x}, y_j - \tilde{y}). \tag{12}$$

Equivalently, we represent the surface as a collection of planar patches parallel to the lens. We also introduce the density corresponding to the energy distribution R and denote it with the function r defined by

$$r(x, y)dxdy \doteq dR(x, y). \tag{13}$$

Strictly speaking, r is the Radon-Nikodym derivative of R, and as such it may not be an ordinary function but, rather, a distribution of measures. We neglect such technicalities here, since they do not affect the derivation of our algorithm. The imaging process can thus be modeled (locally) as a convolution integral:

$$I(x, y) = h^\sigma * r(x, y) \qquad (x, y) \in U \subset D. \tag{14}$$

We will further assume that energy is conserved, and therefore

$$\int h^\sigma(x, y)dxdy = 1 \qquad \forall\ \sigma. \tag{15}$$

In order to further simplify the problem, we restrict our attention to geometric optics, and represent the kernel h^σ by a Gaussian with standard deviation $\sigma = \frac{d}{2}|1 - Z/Z_F|$, where Z is the depth of the scene, Z_F is the depth of the point in focus and d is the diameter of the lens (see figure 1). We choose a Gaussian kernel

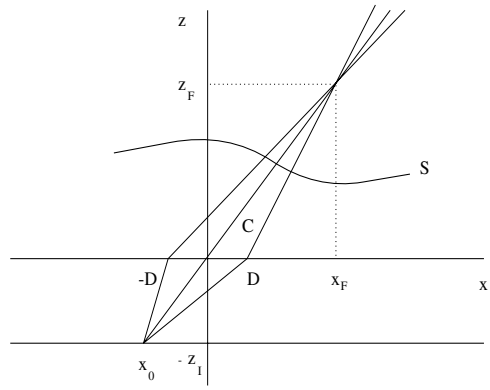

Fig. 1. *A bare-bone model of the geometry of image formation.*

not because it is a good model of the imaging process, but because it makes the

analysis and the implementation of the algorithm straightforward. The algorithm does not depend upon this choice, and indeed we are in the process of building realistic models for the kernels of commercial cameras.

2.2 Steps of the Alternating Minimization

In the imaging model just described, the "shape" of the surface is trivial and represented by a positive scalar σ that depends upon Z, the depth of the patch U. Since the first step of the minimization depends only on this parameter, we can choose any of the known descent methods (e.g. Newton-Raphson). The choice is arbitrary and does not affect the considerations that follow. Therefore, we indicate this step generically as:

$$\sigma_{k+1} = \arg\min_{\sigma_k > 0} \phi\left(\sigma_k, r\right). \tag{16}$$

The second step is obtained from the Kuhn-Tucker conditions [8] associated with the problem of minimizing ϕ for fixed σ under positivity constraints for r:

$$\sum_{i,j} \frac{h^\sigma(x_i, y_j, \tilde{x}, \tilde{y}) I(x_i, y_j)}{\int h^\sigma(x_i, y_j, \tilde{x}, \tilde{y}) r(\tilde{x}, \tilde{y}) d\tilde{x} d\tilde{y}} = \begin{cases} \sum_{i,j} h^\sigma(x_i, y_j, \tilde{x}, \tilde{y}) \ \forall\ (\tilde{x}, \tilde{y}) \mid r(\tilde{x}, \tilde{y}) > 0 \\ \leq \sum_{i,j} h^\sigma(x_i, y_j, \tilde{x}, \tilde{y}) \ \forall\ (\tilde{x}, \tilde{y}) \mid r(\tilde{x}, \tilde{y}) = 0. \end{cases} \tag{17}$$

Since such conditions cannot be solved in closed form, we look for an iterative procedure for r_k that will converge to a fixed point. Following Snyder et al. [14], we choose

$$F(\sigma, r) \doteq \frac{1}{\sum_{i,j} h^\sigma(x_i, y_j, \tilde{x}, \tilde{y})} \sum_{i,j} \frac{h^\sigma(x_i, y_j, \tilde{x}, \tilde{y}) I(x_i, y_j)}{b^\sigma(x_i, y_j, r)} \tag{18}$$

and define the following iteration:

$$r_{k+1} = r_k F(\sigma, r_k). \tag{19}$$

It is important to point out that this iteration decreases the I-divergence ϕ not only when we use the exact kernel h^σ, as it is showed in Snyder et al. [14], but also with any other kernel satisfying the positivity and smoothness constraint. This fact is proven by the following claim.

Claim 1 *Let r_0 be a non-negative real-valued function defined on $\mathbb{R}^2$, and let the sequence r_k be defined according to (19). Then $\phi(\sigma, r_{k+1}) \leq \phi(\sigma, r_k)$ $\forall k > 0$, $\forall\ \sigma > 0$. Furthermore equality holds if and only if $r_{k+1} = r_k$.*

Proof: *The proof follows Snyder et al. [14]. From the definition of ϕ in equation (6) we get*

$$\phi(\sigma, r_{k+1}) - \phi(\sigma, r_k) = -\sum_{i,j} I(x_i, y_j) \log \frac{b^\sigma(x_i, y_j, r_{k+1})}{b^\sigma(x_i, y_j, r_k)} + \sum_{i,j} b^\sigma(x_i, y_j, r_{k+1}) - b^\sigma(x_i, y_j, r_k). \tag{20}$$

From the expression of r_{k+1} in (19) we have that the second sum in the above expression is given by

$$\sum_{i,j} \int h^\sigma(x_i, y_j, x, y) r_{k+1}(x, y) dxdy - \sum_{i,j} \int h^\sigma(x_i, y_j, x, y) r_k(x, y) dxdy =$$
$$= \int H_0^\sigma(x, y) r_{k+1}(x, y) dxdy - \int H_0^\sigma(x, y) r_k(x, y) dxdy$$

where we have defined $H_0^\sigma(x, y) = \sum_{i,j} h^\sigma(x_i, y_j, x, y)$, while the ratio in the first sum is

$$\frac{b^\sigma(x_i, y_j, r_{k+1})}{b^\sigma(x_i, y_j, r_k)} = \int F(\sigma, r_k) \frac{h^\sigma(x_i, y_j, x, y) r_k(x, y)}{b^\sigma(x_i, y_j, r_k)} dxdy. \tag{21}$$

We next note that, from Jensen's inequality,

$$\log \left(\int F(\sigma, r_k) \frac{h^\sigma(x_i, y_j, x, y) r_k(x, y)}{b^\sigma(x_i, y_j, r_k)} dxdy \right) \geq \int \frac{h^\sigma(x_i, y_j, x, y) r_k(x, y)}{b^\sigma(x_i, y_j, r_k)} \log(F(\sigma, r_k)) dxdy \tag{22}$$

since the ratio $\frac{h^\sigma(x_i, y_j, x, y) r_k(x, y)}{b^\sigma(x_i, y_j, r_k)}$ can be interpreted as a probability distribution dependent on the parameters σ and r_k, and therefore the expression in (20) is

$$\phi(\sigma, r_{k+1}) - \phi(\sigma, r_k) \leq -\sum_{i,j} I(x_i, y_j) \int \log(F(\sigma, r_k)) \frac{h^\sigma(x_i, y_j, x, y) r_k(x, y)}{b^\sigma(x, y, r_k)} dxdy +$$
$$\int H_0^\sigma(x, y) r_{k+1}(x, y) dxdy - \int H_0^\sigma(x, y) r_k(x, y) dxdy.$$

The right-hand side of the above expression can be written as

$$\phi_c(H_0^\sigma(x, y) r_{k+1}(x, y), H_0^\sigma(x, y) r_k(x, y)) \tag{23}$$

where we define $\phi_c(f(x, y), g(x, y))$ as $\int f(x, y) \log \frac{f(x,y)}{g(x,y)} - f(x, y) + g(x, y) dxdy$, which can be easily verified to be a positive function for any positive f, g. Therefore, we have

$$\phi(\sigma, r_{k+1}) - \phi(\sigma, r_k) \leq 0. \tag{24}$$

Note that Jensen's inequality holds with equality if and only if $F(\sigma, r_k)$ is a constant; since the only admissible constant value is 1, then we have $r_{k+1} = r_k$, which concludes the proof.

Finally, we can conclude that the algorithm proposed generates a monotonically decreasing sequence of values of the cost function ϕ. We say that the initial conditions σ_0, r_0 are *admissible* if $\sigma_0 > 0$ and r_0 is a positive function defined on $\mathbb{R}^2$.

Corollary 1 *Let σ_0, r_0 be admissible initial conditions for the sequences σ_k and r_k defined from equations (16) and (19) respectively. Let ϕ_k be defined as $\phi(\sigma_k, r_k)$, then the sequence $\{\phi_k\}$ converges to a limit ϕ^*:*

$$\lim_{k\to\infty} \phi_k = \phi^*. \tag{25}$$

Proof: *Follows directly from equation (16), (10) and claim 1, together with the fact that the I-divergence is bounded from below by zero.*

Even if ϕ_k converges to a limit, it is not necessarily the case that σ_k and r_k do. Whether this happens or not depends on the *observability* of the model (1), which has been analyzed in [10].

3 Experiments

In this section we discuss some details that are important for implementing the algorithm just described on a digital computer, and test its performance on a set of experiments on real and simulated images.

3.1 Implementation

Since the algorithm we propose is iterative, we need to initialize it with a feasible radiance. We choose $r = I_{u_1}$, that is we choose the initial estimate of the radiance to be equal to the first image. This choice is guaranteed to be feasible since the image is positive. Since $h^\sigma(x_i, y_j, x, y)$ is discrete in the first two variables, one needs to exercise caution when performing numerical integration against kernels smaller than the unit step of the discretization (x_i, y_j). This case cannot be discounted because it occurs whenever the patch on the image that we are observing is close to be in focus. In our implementation, integrals are computed with a first order (linear) approximation as a tradeoff between speed and accuracy.

Another important detail to bear in mind is that it is necessary to choose the appropriate integration domain. The fact that we use an equifocal imaging model allows us to use the same reference frame for the image and for space, which is represented locally by a plane parallel to it. However, the image in any given patch receives contributions from a region of space possibly bigger than the patch itself. Thus we write $I(x_i, y_j) = \int h^\sigma(x_i, y_j, x, y)(r_I(x, y) + r_o(x, y))dxdy$, where r_o is the radiance outside the patch that contributes to the convolution with the kernel h^σ. In the real and synthetic experiments we always use two images, with planes focused at 529 mm and 869 mm as in the data set provided to us by Watanabe and Nayar [18]; the lens diameter is such that the maximum kernel radius is around 2.3 pixels. With these values the kernel is well approximated by a Gaussian since the radius is small compared to the image patch dimension. Therefore, we define r_o on a domain that is 3 pixels wider than the domain of r_I, which we choose to be 7×7, ending up integrating on patches of dimension 13×13.

3.2 Experiments with Synthetic Images

In this set of experiments we investigate the robustness of the algorithm to noise. Even though we have not derived the algorithm based upon a particular noise model (all the discussion is strictly deterministic), it can be shown that minimizing the I-divergence can be cast into a stochastic framework by modeling the image noise as a Poisson process (the arrivals of photons on the sensor surface).

We have generated 10 noisy image pairs and considered patches of size 7×7 pixels. Smaller patches result in greater sensitivity to noise, while larger ones challenge the equifocal approximation. We have considered additive Gaussian noise with a variance that ranges from the 1% to the 10% of the radiance magnitude, which guarantees that the positivity constraint is still satisfied with high probability. The results of the computed depths are summarized in figure 2. We iterate the algorithm 5 times at each point.

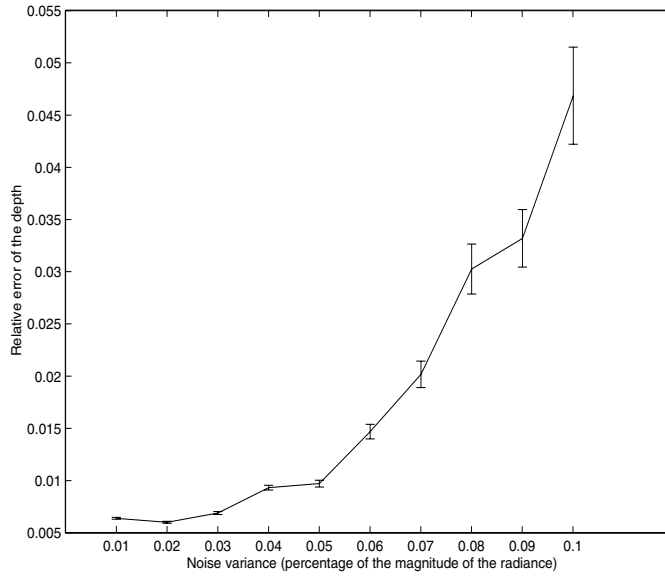

Fig. 2. *Depth error as a function of image noise, mean and std.*

As it can be seen, the algorithm is quite robust to the additive noise, even tough if the radiance is not sufficiently exciting (in the sense defined in [10]) it will not converge. All this will be seen in the experiments with real images described below.

3.3 Experiments with Real Images

We have tested the algorithm on the two images in figures 4 and 7 provided to us by M. Watanabe and S. Nayar. These images were generated by a telecentric optic (see [17] for more details) where there is no change in scale for different focus settings. A side effect is that now the real lens diameter is not constant, and therefore we need to correct our optical model according to figure 3. More precisely, we substitute the diameter d with the new diameter $D = \frac{2aZ_F}{Z_F - f}$ where a and f are indicated in figure 3. For this experiment, in order to speed up the computation, we chose to iterate the algorithm for 5 iterations and to compute depth at every other pixel along both coordinate axes. At points where the radiance is not rich enough, or where the local approximation with an equifocal plane is not valid, the algorithm fails to converge. This explains why in figure 5 some points are visibly incorrect, and in figure 8 the depth of the white background is poorly retrieved. A convergence test could be employed, although it would slow down the computation considerably.

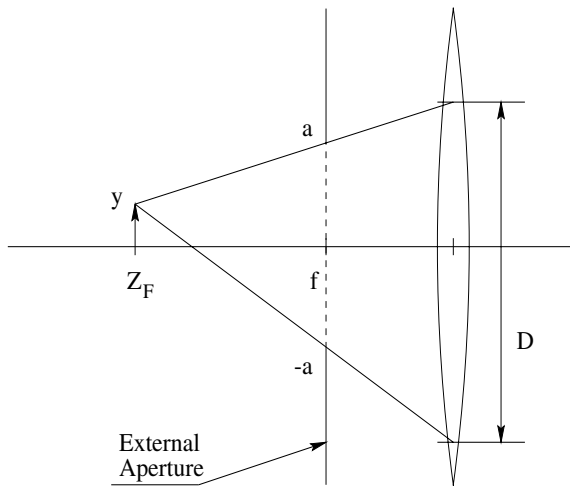

Fig. 3. *The modified diameter in the telecentric lens model.*

4 Conclusions

We have proposed a solution to the problem of reconstructing shape and radiance of a scene using I-divergence as a criterion in an infinite-dimensional optimization framework. The algorithm is iterative, and we give a proof of its convergence to a (local) minimum which, by construction, is admissible in the sense of resulting in a positive radiance and imaging kernel.

Fig. 4. *Original images: near focused (left); far focused (right). The difference between the two images is barely perceivable since the two focal planes are only 340 mm apart.*

References

1. N. Asada, H. Fujiwara, and T. Matsuyama. Edge and depth from focus. *Intl. J. of Comp. Vision*, 26(2):153–163, 1998.
2. W. Boothby. *Introduction to Differentiable Manifolds and Riemannian Geometry.* Academic Press, 1986.
3. T. Chan, P. Blomgren, P. Mulet and C. Wong. Total variation image restoration: numerical methods and extensions. *IEEE Intl. Conf. on Image Processing*, Santa Barbara, 1997.
4. S. Chaudhuri and A. Rajagopalan. *Depth from defocus: a real aperture imaging approach.*, Springer Verlag, 1999.
5. I. Csiszár. Why least-squares and maximum entropy; an axiomatic approach to inverse problems. *Annals of statistics*, 19:2033–2066, 1991.
6. T. Darell and K. Wohn. Depth from focus using a pyramid architecture. *Pattern Recognition Letters*, 11(2):787–796, 1990.
7. J. Ens and P. Lawrence. An investigation of methods for determining depth from focus. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15:97–108, 1993.
8. D. Luenberger. *Optimization by vector space methods.* Wiley, 1968.
9. J. Marshall, C. Burbeck, and D. Ariely. Occlusion edge blur: a cue to relative visual depth. *Intl. J. Opt. Soc. Am. A*, 13:681–688, 1996.
10. A. Mennucci and S. Soatto. On observing shape from defocused images. In *Proc. of the Intl. Conf. on Image Analysis and Processing*, pages 550-555, 1999.
11. S. Nayar and Y. Nakagawa. Shape from focus. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(8):824–831, 1994.
12. A. Pentland. A new sense for depth of field. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:523–531, 1987.
13. Y. Schechner and N. Kiryati. The optimal axial interval in estimating depth from defocus. In *Proc. of the Intl. Conf. of Comp. Vision*, pages 843–848, 1993.
14. D. Snyder, T. Schulz, and J. O'Sullivan. Deblurring subject to nonnegativity constraints. *IEEE Trans. on Signal Processing*, 40(5):1143–1150, 1992.
15. S. Soatto and P. Favaro. A geometric approach to blind deconvolution with application to shape from defocus. In *Proc. of the IEEE Intl. Conf. on Comp. Vision and Pattern Recognition* (in press), 2000.

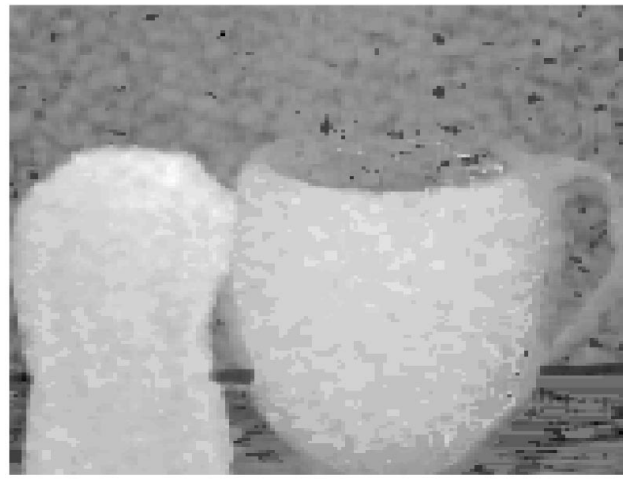

Fig. 5. *Reconstructed depth for the scene in figure 4, coded in grayscale.*

16. M. Subbarao and G. Surya. Depth from defocus: a spatial domain approach. *Intl. J. of Computer Vision*, 13:271–294, 1994.
17. M. Watanabe and S. Nayar. Telecentric optics for constant-magnification imaging. Technical Report Cucs-026-95, Columbia University, 1995.
18. M. Watanabe and S. Nayar. Rational filters for passive depth from defocus. *Intl. J. of Comp. Vision*, 27(3):203–225, 1998.
19. Y. Xiong and S. Shafer. Depth from focusing and defocusing. In *Proc. of the Intl. Conf. of Comp. Vision and Pat. Recogn.*, pages 68–73, 1993.
20. D. Ziou. Passive depth from defocus using a spatial domain approach. In *Proc. of the Intl. Conf. of Computer Vision*, pages 799–804, 1998.

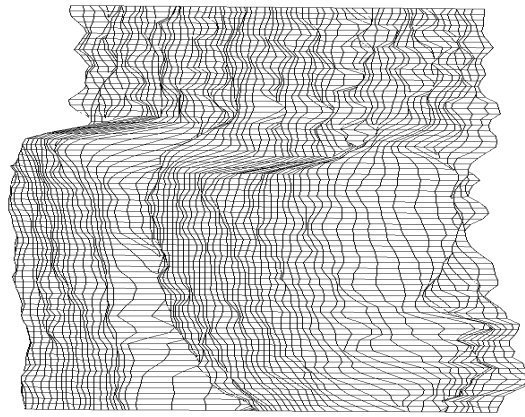

Fig. 6. *Reconstructed depth for the scene in figure 4: smoothed mesh.*

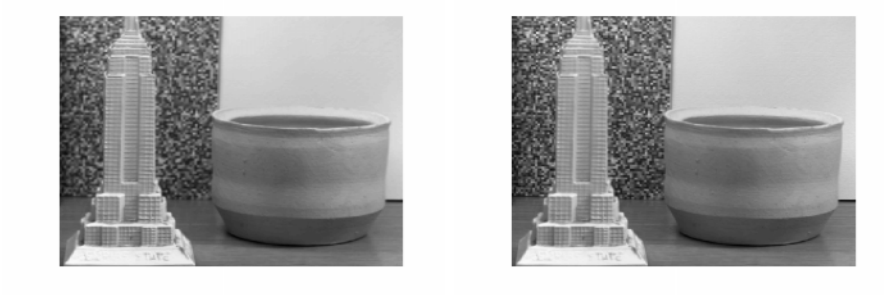

Fig. 7. *Original images: near focused (left); far focused (right). As in figure 4 the difference between the two is barely perceivable.*

Fig. 8. *Reconstructed depth for the scene in figure 7, coded in grayscale. In the uniform region of the background, the radiance is not sufficiently exciting, in the sense defined in [10]. Therefore, the algorithm cannot converge and the quality of the estimates, as it can be seen, is poor.*

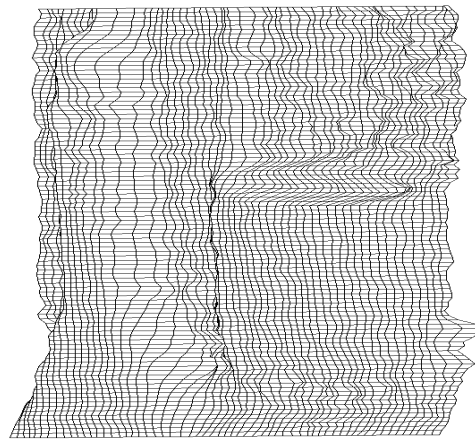

Fig. 9. *Reconstructed depth for the scene in figure 7: smoothed mesh.*