

Image-based 3D Modeling via Cheeger Sets

Eno Töppe^{1,2}, Martin R. Oswald¹, Daniel Cremers¹, Carsten Rother²

¹ Technische Universität München, Germany

² Microsoft Research, Cambridge, UK

Abstract. We propose a novel variational formulation for generating 3D models of objects from a single view. Based on a few user scribbles in an image, the algorithm automatically extracts the object silhouette and subsequently determines a 3D volume by minimizing the weighted surface area for a fixed user-specified volume. The respective energy can be efficiently minimized by means of convex relaxation techniques, leading to visually pleasing smooth surfaces within a matter of seconds. In contrast to existing techniques for single-view reconstruction, the proposed method is based on an implicit surface representation and a transparent optimality criterion, assuring high-quality 3D models of arbitrary topology with a minimum of user input.

1 Introduction

1.1 Single-View Reconstruction

Generating models of the three-dimensional world from sets of images is at the heart of Computer Vision. An interesting limiting case is the problem of single view reconstruction – a highly ill-posed problem where stereo and multiview concepts like point correspondence and photo-consistency cannot be applied. Nevertheless, it is an important problem: In many applications we may only have a single image of the scene, and yet we may want to interactively extract solid 3D models of respective objects for virtual and augmented reality applications, or we may want to simply render the same scene from a novel vantage point or with different illumination based on estimates of the geometric structure.

Human observers have an excellent ability to generate plausible 3D models of objects around them – even from a single image. To this end, they partially rely on prior knowledge about the geometric structures and primitives in their world. Yet, they also generate plausible models of objects they have never seen before. It is beyond the scope of this work to contemplate on the multitude of criteria the human visual system may be employing for solving the single view reconstruction problem. Instead, we will demonstrate that for a large variety of real-world images very simple extremality assumptions give rise to convincing 3D models. The key idea is to compute a silhouette-consistent weighted minimal surface for a user-specified volume. In this sense, the proposed formulation is closely related to the concept of *Cheeger sets* – sets which minimize the ratio of area over volume [1].




Fig. 1. The proposed method generates convincing 3D models from a single image computed by fixed volume weighted minimal surfaces. Colored lines in the input image mark user input, which locally alters the surface smoothness. Red marks low, yellow marks high smoothness (see section 4.4 for details).

1.2 Related Work

Existing work on single view reconstruction and on interactive 3D modeling can be grouped into two classes based on the choice of mathematical surface representation, namely explicit surface representation and implicit surface representations. Some of the pioneering works on single view reconstruction are those of Criminisi and coworkers [2, 3] on generating three-dimensional models of architecture from single images by exploiting the perspective structure of parallel lines and other aspects of man-made environments. Different aspects of the reconstruction have been emphasized among related works. Horry et al. [4] aim for pleasant 3D visual effects that do not result in high quality meshes. Hoiem et al. [5] are similar in this respect but they try to fully automate the reconstruction process. Also related to the field are easy-to-use tools like *Teddy* [6] and *Fiber-Mesh* [7] that have pioneered sketch based modeling but are not image-based. Note that there are also approaches that seek to reconstruct height fields [8] and are therefore not suited for getting closed 3D surfaces.

All of the above works are using explicit surface representation – while surface manipulation is often straight forward and a variety of cues are easily integrated leading to respective forces or constraints on the surface, there are two major limitations: Firstly numerical solutions are generally not independent of the choice of parameterization. And secondly, parametric representations are not easily extended to objects of varying topology. While Prasad et al. [9] were able to extend their approach to surfaces with one or two holes, the generalization to objects of arbitrary topology is by no means straight forward. Similarly, topology-changing interaction in the FiberMesh system requires a complex remeshing of the modeled object leading to computationally challenging numerical optimization schemes.

A first effort in single view reconstruction using an implicit representation was recently proposed by Oswald et al. [10]. There the authors combined a minimal surface constraint with a data term that favored the object thickness to be proportional to the distance to the boundary.

Despite a number of convincing results, the latter work suffers from several drawbacks: Firstly, imposing a thickness proportional to the distance from the curve is very strong and not always a correct assumption. Secondly, the modeling required a large number of not necessarily intuitive parameters controlling the data term. Our work is different from [10] in that we impose exact volume consistency and do not require additional tuning parameters.

All cited works on single view reconstruction have in common that they revert to inflation heuristics in order to avoid surface collapsing. These techniques boil down to fixing absolute depth values, which undesirably restrict the solution space. A precursor to volume constraints are the volume inflation terms pioneered for deformable models by Cohen and Cohen [11]. However, no constant volume constraints were considered and no implicit representations were used.

1.3 Contribution

In this paper, we revisit the problem of single view reconstruction. We will show that one can compute silhouette-consistent weighted minimal surfaces for a user-prescribed volume using convex relaxation techniques. To this end, we revert to an implicit representation of the surface given by the indicator function of its interior (sometimes referred to as voxel-occupancy). In this representation, the weighted minimal surface problem is a convex functional and relaxation of the binary function leads to an overall convex problem. In addition, we will show that the volume constraint amounts to a convex constraint which is easily integrated in the reconstruction process. We show that the relaxed indicator function can be binarized so that we obtain a surface which firstly has exactly the user-specified volume and secondly is within a computable energetic bound of the optimal combinatorial solution.

The convex optimization is solved by a recently proposed provably convergent primal-dual algorithm enabling *interactive* reconstruction within seconds. We show on a variety of real-world images that the simple extremality condition of a *fixed-volume minimal surface* gives rise to convincing 3D models for a large variety of real-world images, comparing favorably to alternative approaches.

To the best of our knowledge this is the first work on convex shape optimization with guaranteed volume preservation.

2 Variational Formulation

Assume we are given the silhouette of an object in an image as returned by an interactive segmentation tool ¹. The goal is then to obtain a smooth 3D model of the object which is consistent with the silhouette. How should we select the correct 3D model among the infinitely many that match the silhouette? Clearly,

¹ For brevity and since it is not part of our contribution we will not detail the graph cut based interactive segmentation algorithm we use. Instead we refer to representative work in the field [12].

we need to impose additional information, at the same time we want to keep this information at a minimum since user interaction is always tedious and slow. In the following, we will show that merely specifying the object’s volume and computing a minimal surface of given volume is sufficient to give rise to a family of plausible 3D models.

We propose a solution that comes in two flavors: one is formulated with a soft the other with a hard volume constraint. We then go into detail on the fast optimization of the resulting energy which finally leads to an interactive user interface for single view reconstruction.

2.1 Implicit Weighted Variational Surfaces

We are given an image plane Ω which contains the input image and lies in \mathbb{R}^3 . As part of the image we also have an object silhouette $\Sigma \subset \Omega$. Now, we are seeking to compute reconstructions as minimal weighted surfaces $S \subset \mathbb{R}^3$ that have a certain target volume V_t and are compliant with the object silhouette Σ :

$$\min \int_S g(s) ds \quad (1)$$

$$\text{subject to} \quad \pi(S) = \Sigma \quad (2)$$

$$\text{Vol}(S) = V_t \quad (3)$$

where $\pi : \mathbb{R}^3 \rightarrow \Omega$ is the orthographic projection onto the image plane Ω , $g : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ is a smoothness weighting function, $\text{Vol}(S)$ denotes the volume enclosed by the surface S and $s \in S$ is a surface element. In the following we will gradually derive an implicit representation for the above problem.

We begin by replacing the surface S with its implicit binary indicator function $u \in BV(\mathbb{R}^3; \{0, 1\})$, where BV denotes the functions of bounded variation [13]. The desired minimal weighted surface area is then given by minimizing the total variation over a suitable set U of feasible functions u :

$$\min_{u \in U} \int g(x) |\nabla u(x)| d^3x \quad (4)$$

where ∇u denotes the derivative in the distributional sense. Eq. (4) favors smooth solutions. However, smoothness is locally affected by the function $g(x) : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ which will be used later for modeling.

How does the set U of feasible functions look like? For simplicity, we assume the silhouette to be enclosed by the surface. Then all surface functions that are consistent with the silhouette Σ must be in the set

$$U_\Sigma = \left\{ u \in BV(\mathbb{R}^3; \{0, 1\}) \mid u(x) = \begin{cases} 0, & \pi(x) \notin \Sigma \\ 1, & x \in \Sigma \end{cases} \right\} \quad (5)$$

Still, solving (4) with respect to the set U_Σ of silhouette consistent functions will result in the silhouette itself. In the following section we will show a way to avoid this trivial solution.

2.2 Volume Constraint

In order to inflate the solution of (4) we propose to use a constraint on the size of the volume enclosed by the minimal surface. We formulate this both as a soft- and as a hard constraint and discuss the two approaches in the following.

Hard Constraint. By further constraining the feasible set U_Σ one can force the reconstructed surface to have a specific target volume V_t . We regard the problem

$$\min_{u \in U_\Sigma \cap U_V} E(u) \quad \text{where} \quad E(u) = \int g(x) |\nabla u(x)| d^3x \quad (6)$$

$$\text{and} \quad U_V = \left\{ u \in BV(\mathbb{R}^3; \{0, 1\}) \mid \int u(x) d^3x = V_t \right\} \quad (7)$$

where U_V denominates all reconstructions with bounded variation that have the specific volume V_t .

Soft Constraint. For the sake of completeness we also consider the soft formulation of the volume constraint. One can add a ballooning term to (4):

$$E_V(u) = \lambda \left(\int u(x) d^3x - V_t \right)^2 \quad (8)$$

The integral quadratically punishes the deviation of the surface volume from a certain target volume V_t . In contrast to the constant volume constraint above, this formulation comes with an extra parameter λ which is why in the following we will focus on (6) instead.

Different approaches to finding V_t can be considered. In the implementation the optimization domain is naturally bounded. We choose V_t to be a fraction of the volume of this domain. In a fast interactive framework the user can then adapt the target volume with the help of instant visual feedback. Most importantly, as opposed to a data term driven model volume constraints do not dictate where inflation takes place.

2.3 Fast Minimization

In order to convexify the problem in (6) we make use of a relaxation technique [14]. To this end we relax the binary range of functions u in (5) and (7) to the interval $[0, 1]$. In other words we replace U_V and U_Σ with their respective convex hulls U_V^r and U_Σ^r . The corresponding optimization problem is then convex:

Proposition 1. *The relaxed set $U^r := U_\Sigma^r \cap U_V^r$ is convex.*

Proof. The constraint in the definition of U_V is clearly linear in u and therefore U_V^r is convex. The same argument holds for U_Σ . Being an intersection of two convex sets U^r is convex as well.

One standard way of finding the globally optimal solution to this problem is gradient descent, which is known to converge very slowly. Since optimization speed is an integral part of an interactive reconstruction framework, we employ a recently proposed significantly faster and provably convergent primal-dual algorithm published in [15]. The scheme is based on the weak formulation of the total variation:

$$\min_{u \in U^r} \int g(x) |\nabla u| d^3x = \min_{u \in U^r} \sup_{|\xi(x)|_2 \leq g(x)} \left\{ \int -u \operatorname{div} \xi \, d^3x \right\} \quad (9)$$

Optimization is done by alternating a gradient descent with respect to the function u and a gradient ascent for the dual variable $\xi \in \mathcal{C}_c^1(\mathbb{R}^3; \mathbb{R}^3)$ interlaced with an over-relaxation step on the primal variable:

$$\begin{cases} \xi^{k+1} = \Pi_{|\xi(x)|_2 \leq g(x)}(\xi^k + \tau \cdot \nabla \bar{u}^k) \\ u^{k+1} = \Pi_{U^r}(u^k + \sigma \cdot \operatorname{div} \xi^{k+1}) \\ \bar{u}^{k+1} = 2u^{k+1} - u^k \end{cases} \quad (10)$$

where Π_A denotes the projection onto the set A . Projection of ξ is done by simple clipping while that of the primal variable u will be detailed in the next paragraph. The scheme (10) is numerically attractive since it avoids division by the potentially zero-valued gradient-norm which appears in the Euler-Lagrange equation of the TV-norm. Moreover, it is parallelizable and we therefore implemented it on the GPU. On a volume of 63x47x60 voxels the computation takes only 0.47 seconds.

Projection Scheme. The projection Π_{U^r} in (10) needs to ensure three constraints on u : Silhouette consistency, constant volume and $u \in [0, 1]$. In order to maintain silhouette consistency (5) of the solution we restrict updates to those voxels which project onto the silhouette interior excluding the silhouette itself.

Still we need to enforce the other two constraints. An iterative algorithm which computes the Euclidean projection of a point onto the intersection of arbitrary convex sets is the one of Boyle and Dykstra [16]. It is fast for a low number of convex constraints and converges provably to the projection point. In our case step i of this algorithm reduces to two separate projections for volume and range

$$\begin{cases} u_V^i = u_R^{i-1} - v_V^{i-1} + \frac{V_d}{N} \\ v_V^i = u_V^i - (u_R^{i-1} - v_V^{i-1}) \end{cases} \quad (11)$$

$$\begin{cases} u_R^i = \Pi_{[0,1]}(u_V^i - v_R^{i-1}) \\ v_R^i = u_R^i - (u_V^i - v_R^{i-1}) \end{cases} \quad (12)$$

where we initialize u_R with the current u^k in (10) and v_R, v_V with zero. $\Pi_{[0,1]}(u)$ simply clips the value of u to the unit interval and V_d is the difference between the target volume V_t and the current volume of the values $u_R^{i-1} - v_V^{i-1}$. N is the number of voxels in the discrete implementation.




Fig. 2. The two cases considered in the analysis of the material concentration. On the *left* hand side we assume a hemi-spherical condensation of the material. On the *right* hand side the material is distributed evenly over the volume.

2.4 Optimality Bounds

Having computed a global optimal solution u_{opt} of (9), the question remains how we obtain a binary solution and how the two solutions relate to one another energetically. Unfortunately no thresholding theorem holds, which would imply the binary optimality of the thresholded relaxed optimum for arbitrary thresholds.

Nevertheless we can construct a binary solution u_{bin} as follows:

Proposition 2. *The relaxed solution can be projected to the set of binary functions in such a way that the resulting binary function preserves the user-specified volume V_t .*

Proof. It suffices to order the voxels x by decreasing values $u(x)$. Subsequently, one sets the value of the first V_t voxels to 1 and the value of the remaining voxels to 0.

Concerning an optimality bound the following holds:

Proposition 3. *Let u_{opt}^r be the global optimal solution of the relaxed energy and u_{opt} the global optimal solution of the binary problem. Then*

$$E(u_{bin}) - E(u_{opt}) \leq E(u_{bin}) - E(u_{opt}^r) . \quad (13)$$

3 Theoretical Analysis of Material Concentration

As we have seen above, the proposed convex relaxation technique does not guarantee global optimality of the binary solution. The thresholding theorem [14] – applicable in the unconstrained problem – no longer applies to the volume-constrained problem. While the relaxation naturally gives rise to a posteriori optimality bounds, one may take a closer look at the given problem and ask why the relaxed volume labeling u should favor the emergence of solid objects rather than distribute the prescribed volume equally over all voxels.

In the following, we will prove analytically that the proposed functional has an energetic preference for material concentration. For simplicity, we will consider the case that the object silhouette in the image is a disk. And we will

compare the two extreme cases of all volume being concentrated in a ball (a known solution of the Cheeger problem) compared to the case that the same volume is distributed equally over the feasible space (namely a cylinder) – see Figure 2. Note that in the following proof it suffices to consider the volume only on one side of the silhouette.

Proposition 4. *Let u_{sphere} denote the binary solution which is 1 inside the sphere and 0 outside – Fig. 2, left side – and let u_{cyl} denote the solution which is uniformly distributed (i.e. constant) over the entire cylinder – Fig. 2, right side. Then we have*

$$E(u_{sphere}) < E(u_{cyl}), \quad (14)$$

independent of the height of the cylinder.

Proof. Let R denote the radius of the disk. Then the energy of u_{sphere} is simply given by the area of the half-sphere:

$$E(u_{sphere}) = \int |\nabla u_{sphere}| d^2x = 2\pi R^2. \quad (15)$$

If instead of concentrated to the half-sphere, the same volume, i.e. $V = \frac{2\pi}{3}R^3$, is distributed uniformly over the cylinder of height $h \in (0, \infty)$, we have

$$u_{cyl}(x) = \frac{V}{\pi R^2 h} = \frac{2\pi R^3}{3\pi R^2 h} = \frac{2}{3} \frac{R}{h}. \quad (16)$$

inside the entire cylinder, and $u_{cyl}(x) = 0$ outside the cylinder. The respective surface energy of u_{cyl} is given by the area of the cylinder weighted by the respective jump size:

$$E(u_{cyl}) = \int |\nabla u_{cyl}| d^2x = \left(1 - \frac{2R}{3h}\right) \pi R^2 + \frac{2R}{3h} (\pi R^2 + 2\pi R h) = \frac{7}{3} \pi R^2 > E(u_{sphere}). \quad (17)$$

4 Experimental Results

Having detailed the idea of variational implicit weighted surfaces and their fast computation, in this section we will study their properties and applicability within an interactive reconstruction environment. We will compare our approach to methods which resort to heuristic inflation techniques and finally show that appealing and realistic 3D models can be generated with minimal user input.

4.1 Cheeger Sets and Single View Reconstruction

Solutions to (6) are Cheeger sets, i.e. minimal surfaces for a fixed volume. In the simplest case of a circle-shaped silhouette one therefore expects to get a ball. Fig. 4 demonstrates that in fact round silhouette boundaries (in the unweighted case) result in round shapes.




Fig. 3. By simply increasing the target volume with the help of a slider, the reconstruction is intuitively inflated. Due to a highly parallelized implementation the result can be computed almost instantly. In this example the initial rendering of the volume with $175 \times 135 \times 80$ voxels took 3.9 seconds. Starting from there each subsequent volume adaptation took only about 1 second.




Fig. 4. The proposed Cheeger set approach favors minimal surfaces for a user-specified volume. Therefore the reconstruction algorithm is ideally suited to compute smooth, round reconstructions.

4.2 Fixed Volume vs. Shape Prior

Many approaches to volume reconstruction incorporate a shape prior in order to avoid surface collapsing. A common heuristic is to use a distance transform of the silhouette boundary for depth value estimation. We show that the fixed-volume approach solves several problems of such a heuristic.

Fig. 5 shows that it is hard to obtain ball-like surfaces with a silhouette distance transform as a shape prior. Another issue is the strong bias a shape prior inflicts on the reconstruction resulting in cone-like shapes (see Fig. 6) and inhibiting the flexibility of the model. The uniform fixed-volume approach fills both gaps while exhibiting the favorable properties of the distance transform (as seen in Fig. 8). With the results in Fig. 5 and 6 we directly compare our method to [10] and [9], in which the reconstruction volume is inflated artificially.

4.3 Varying the Volume

Apart from the weighting function of the TV-norm (see next section), the only parameter we have to determine for our reconstruction is the target volume V_t .




Fig. 5. Using a silhouette distance transform as shape prior the relation between data term (*second from left*) and reconstruction (*third from left*) is not easy to assess for a user. With only one parameter our method delivers more intuitive and natural results.




Fig. 6. In contrast to the approach in [10] (*center*), the proposed method (*right*) does not favor a specific shape and generates more pleasing 3D models. Although in the center reconstruction the dominating shape prior can be mitigated by a higher smoothness, this ultimately leads to the vanishing of thin structures like the handle.

The effect on the appearance of the surface can be witnessed in Fig. 3. One can see that changing the target volume has an intuitive effect on the resulting shape. This is important for a user driven reconstruction.

4.4 Weighted Minimal Surface Reconstruction

So far all presented reconstructions came along without further user input. The weight $g(x)$ of the TV-norm in (9) can be used to locally control the smoothness of the reconstruction: with a low $g(x)$, the smoothness condition on the surface is locally relaxed, allowing for creases and sharp edges to form. Conversely setting $g(x)$ to a high value locally enforces surface smoothness. For controlling the weighting function we employ a user scribble interface. The parameter associated to each scribble marks the local smoothness within the respective scribble area and is propagated through the volume along projection direction. What we show in Fig. 7 is that with this tool not only round, but other very characteristic shapes can be modeled with minimal user interaction.

The air plane in Fig. 1 represents an example, where a parametric shape prior would fail to offer the necessary flexibility required for modeling protrusions. Since our fixed-volume approach does not impose points of inflation, user input can influence the reconstruction result in well-defined ways: Marking the wings as highly non-smooth (i.e. low $g(x)$) effectively allows them to form.




Fig. 7. The proposed approach allows to generate 3D models with sharp edges, marked by the user as locations of low smoothness (see section 4.4). Along the red user strokes (*second from left*) the local smoothness weighting is decreased.




Fig. 8. Volume inflation dominates where the silhouette area is large (bird) whereas thin structures (twigs) are inflated less.

Note that apart from Fig. 1, 7 the adaptation of the target volume was the only user input for all experiments.

5 Conclusion

We presented a novel framework for single view reconstruction which allows to compute 3D models from a single image in form of Cheeger sets, i.e. minimal surfaces for a fixed user-specified volume. The framework allows for appealing and realistic reconstructions of curved surfaces with minimal user input. The combinatorial problem of finding a silhouette-consistent surface with minimal area for a user defined volume is solved by reverting to an implicit surface representation and convex relaxation. The resulting convex energy is optimized globally using an efficient provably convergent primal-dual scheme. Parallel GPU implementation allows for computation times of a few seconds, allowing the user to interactively increase or decrease the volume. We proved that the computed surfaces are within a bound of the optimum and that they exactly fulfill the target volume. On a variety of challenging real world images, we showed that the pro-

posed method compares favorably over existing implicit approaches, that volume variations lead to families of realistic reconstructions and that additional user scribbles allow to locally reduce smoothness so as to easily create protrusions.

Acknowledgments

We thank Andrew Fitzgibbon and Mukta Prasad for fruitful discussions on single view reconstruction.

References

1. Cheeger, J.: A lower bound for the smallest eigenvalue of the laplacian. In: *Problems in analysis*. Princeton Univ. Press, Princeton, N.J. (1970)
2. Liebowitz, D., Criminisi, A., Zisserman, A.: Creating architectural models from images. In: *Proc. EuroGraphics*. Volume 18. (1999) 39–50
3. Criminisi, A., Reid, I., Zisserman, A.: Single view metrology. *Int. J. Comput. Vision* **40** (2000) 123–148
4. Horry, Y., Anjyo, K.I., Arai, K.: Tour into the picture: using a spidery mesh interface to make animation from a single image. In: *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, New York, NY, USA, ACM Press/Addison-Wesley Publishing Co. (1997) 225–232
5. Hoiem, D., Efros, A.A., Hebert, M.: Automatic photo pop-up. *ACM Trans. Graph.* **24** (2005) 577–584
6. Igarashi, T., Matsuoka, S., Tanaka, H.: Teddy: a sketching interface for 3d freeform design. In: *SIGGRAPH '99*, New York, NY, USA, ACM Press/Addison-Wesley Publishing Co. (1999) 409–416
7. Nealen, A., Igarashi, T., Sorkine, O., Alexa, M.: Fibermesh: designing freeform surfaces with 3d curves. *ACM Trans. Graph.* **26** (2007) 41
8. Zhang, L., Dugas-Phocion, G., Samson, J.S., Seitz, S.M.: Single view modeling of free-form scenes. In: *Proc. of CVPR*. (2001) 990–997
9. Prasad, M., Zisserman, A., Fitzgibbon, A.W.: Single view reconstruction of curved surfaces. In: *CVPR*. (2006) 1345–1354
10. Oswald, M.R., Toeppe, E., Kolev, K., Cremers, D.: Non-parametric single view reconstruction of curved objects using convex optimization. In: *Pattern Recognition (Proc. DAGM)*, Jena, Germany (2009)
11. Cohen, L.D., Cohen, I.: Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. on Patt. Anal. and Mach. Intell.* **15** (1993) 1131–1147
12. Rother, C., Kolmogorov, V., Blake, A.: GrabCut: interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* **23** (2004) 309–314
13. Ambrosio, L., Fusco, N., Pallara, D.: *Functions of bounded variation and free discontinuity problems*. Oxford Mathematical Monographs. The Clarendon Press Oxford University Press, New York (2000)
14. Chan, T., Esedoglu, S., Nikolova, M.: Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics* **66** (2006) 1632–1648
15. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the piecewise smooth mumford-shah functional. In: *IEEE Int. Conf. on Computer Vision*, Kyoto, Japan (2009)
16. Boyle, J.P., Dykstra, R.L.: An method for finding projections onto the intersection of convex sets in Hilbert spaces. **37** (1986) 28–47