

A Convex Approach for Variational Super-Resolution

Markus Unger, Thomas Pock, Manuel Werlberger, Horst Bischof

Institute for Computer Graphics and Vision, Graz University of Technology, Austria

Abstract. We propose a convex variational framework to compute high resolution images from a low resolution video. The image formation process is analyzed to provide to a well designed model for warping, blurring, downsampling and regularization. We provide a comprehensive investigation of the single model components. The super-resolution problem is modeled as a minimization problem in an unified convex framework, which is solved by a fast primal dual algorithm. A comprehensive evaluation on the influence of different kinds of noise is carried out. The proposed algorithm shows excellent recovery of information for various real and synthetic datasets.

1 Introduction

The reconstruction of highly resolved images out of multiple smaller images is an important problem that occurs in surveillance, remote sensing, medical imaging, video restoration, up sampling and still image extraction. Video frames are representations of a scene that undergo arbitrary motion from frame to frame, degradation by the optical system and the digitization process. Although motion might seem to be a problem, exactly these sub-pixel shifts caused by moving objects or the camera, provide the necessary information utilized in super-resolution reconstruction. The super-resolution problem is difficult to solve because one has to deal with two kinds of image degradation: First, the camera system, that adds blur (eg. by the optical system like lenses and filters) and performs spatial integration on the sensor (often these are not square pixels). With exact knowledge of the Point Spread Function (PSF) of the camera system and sufficiently enough samples by shifted images this degradation is invertible up to the limits posed by Shannon [1]. Unfortunately the second group of degradation, namely noise, cannot be undone [2], making super-resolution an ill-posed problem. Noise occurs in different forms: Gaussian type of noise (eg. caused by sensor heating and analog processing), Outlier noise (eg. occlusions) and systematic noise (eg. quantization and compression).

The limitations of super-resolution were already studied by Kosarev [2] who stated a logarithmic dependence of quality on the signal-to-noise ratio of the input. Later, the limits on image super-resolution were analyzed by Baker and Kanade [3]. They provided experimental results using 8bit images, showing a degradation of image quality with increasing magnification. This experiments

already demonstrated that the limited dynamic range causes a dramatic loss of information. In Section 5, we will provide a deeper analysis of various types of noise. To overcome the limitations of noise some prior information might be used in the reconstruction process. Baker and Kanade [3] proposed a learning based recognition approach (called 'hallucination') to improve super-resolution on a limited class of images. More general approaches try to utilize redundant structures in images. Potter et al. [4] use a Non-local-Means algorithm for super-resolution reconstruction. A similar approach is taken by Glasner et al. in [5], where patch redundancy over various scales of a single image is used to significantly improve zooming for pictures with redundant structures.

There is also a great number of super-resolution approaches that rely on a more general prior [6]. Some of these approaches use Total Variation (TV) regularization as eg. the approach by Mitzel et al. [7]. Farsiu et al. [8], use bilateral TV (that is closely related to Non-local TV) to overcome stair-casing artifacts usually induced by TV regularization. Our approach is closely related to the variational approach of Mitzel et al. [7]. While Mitzel et al. used L^1 -norms for regularization and in the data term, we replace them in our model (see Section 3) with Huber-Norms [9]. The Huber-Norm has the advantage of smooth gray values while preserving strong edges. Another improvement over [7] is the used first-order primal dual algorithm in Section 4. Additionally we provide a comprehensive investigation of the crucial super-resolution operators used for warping, blurring and downsampling. Another energy minimization based approach is investigated by Schoenemann [10] that combines motion layer decomposition with super-resolution. He introduces an additional term that imposes regularity on the geometry of the layers.

As the input images are related by some arbitrary unknown motion, the accurate estimation of this motion is obviously very important. For small motions this can be done implicitly eg. by Non-local methods [4] or by performing joint space-time super-resolution as done by Shechtman [11]. There are also semi-implicit methods eg. using steering kernels as done by Takeda et al. [12]. For large arbitrary motion an explicit optical flow calculation is required. Fortunately dense optical flow algorithms have become very accurate and sufficiently fast [13]. We will do motion estimation by the variational optical flow proposed by Werlberger et al. [14], for which GPU-binaries are available.

Contribution: The contribution of this paper is threefold: First, we extend the variational model by Mitzel et al. [7] by the usage of the Huber-Norm and an exact choice of the crucial linear operators (Section 3). The choice and implementation of these operators as well as the discretization are described in sufficient detail that the proposed method can easily be reimplemented. Second, we provide a fast minimization procedure in Section 4. Therefore we adapt the first-order primal-dual algorithm from Pock et al. [15]. Finally, we investigate the effects of different kinds of noise on super-resolution reconstruction in Section 5. We also compare the algorithm to [7, 10], and show superior results obtained with the proposed algorithm, demonstrating its robustness for arbitrary motion and occlusions.

2 Discretization

Before describing the super-resolution model we have to make some thoughts on discretization. An image is given on a two dimensional regular Cartesian grid of the size $M \times N$:

$$\{(ih, jh) : 1 \leq i \leq M, 1 \leq j \leq N\}, \quad (1)$$

with the indices of the discrete locations given by (i, j) and the pixel size (or spacing) h . We define a finite dimensional vector space $X = \mathbb{R}^{MN}$ with a scalar product

$$\langle v, w \rangle_X^h = h^2 \sum_{i,j} v_{i,j} w_{i,j}, \quad v, w \in X. \quad (2)$$

Furthermore, we define a vector space $Y = \mathbb{R}^{MN} \times \mathbb{R}^{MN}$, with the gradient operator as a linear mapping $\nabla^h : X \rightarrow Y$ using finite differences and Neumann boundary conditions:

$$(\nabla^h v)_{i,j} = ((\delta_x^{h+} v)_{i,j}, (\delta_y^{h+} v)_{i,j})^T, \quad (3)$$

where

$$(\delta_x^{h+} v)_{i,j} = \begin{cases} \frac{v_{i+1,j} - v_{i,j}}{h} & \text{if } i < M \\ 0 & \text{if } i = M \end{cases}, \quad (\delta_y^{h+} v)_{i,j} = \begin{cases} \frac{v_{i,j+1} - v_{i,j}}{h} & \text{if } j < N \\ 0 & \text{if } j = N \end{cases}. \quad (4)$$

Given two vectors $\mathbf{p} = (p^x, p^y)^T, \mathbf{q} = (q^x, q^y)^T \in Y$ we define the scalar product as following:

$$\langle \mathbf{p}, \mathbf{q} \rangle_Y^h = h^2 \sum_{i,j} p_{i,j}^x q_{i,j}^x + p_{i,j}^y q_{i,j}^y. \quad (5)$$

Additionally we have to define a divergence operator $\text{div}^h : Y \rightarrow X$ by choosing it to be adjoint to the gradient operator in (3), and thus fulfilling the equality $\langle \nabla^h u, \mathbf{p} \rangle_Y^h = - \langle u, \text{div}^h \mathbf{p} \rangle_X^h$. Therefore, the discrete divergence operator is given as:

$$(\text{div}^h \mathbf{p})_{i,j} = (\delta_x^{h-} p^x)_{i,j} + (\delta_y^{h-} p^y)_{i,j}, \quad (6)$$

with

$$(\delta_x^{h-} p^x)_{i,j} = \begin{cases} 0 & \text{if } i = 0 \\ \frac{p_{i,j}^x - p_{i-1,j}^x}{h} & \text{if } 0 < i < M \\ -\frac{p_{i-1,j}^x}{h} & \text{if } i = M \end{cases}, \quad (\delta_y^{h-} p^y)_{i,j} = \begin{cases} 0 & \text{if } j = 0 \\ \frac{p_{i,j}^y - p_{i,j-1}^y}{h} & \text{if } 0 < j < N \\ -\frac{p_{i,j-1}^y}{h} & \text{if } j = N \end{cases} \quad (7)$$

3 The Super-resolution Model

In this section the super-resolution model is defined as a convex minimization problem, and we discuss the operators in detail. The exact minimization procedure follows in Section 4.

As input for the super-resolution algorithm, n input images $\check{f}_i \in \check{X}$ of size $\check{M} \times \check{N}$ and pixel size ξh are given. We denote the scale factor as $\xi \in \mathbb{R}^+$. The input images are warped, blurred and noisy samples of some continuous image $g : \Omega \rightarrow \mathbb{R}$. Based on the redundancy of the input images we aim to find one higher resolved super-resolution image $\hat{u} \in \hat{X}$ of size $\hat{M} \times \hat{N}$ with pixel size h . Note that we use the $\check{\cdot}$ accent to indicate that a variable belongs to the coarse (input) level with pixel size ξh and the $\hat{\cdot}$ accent to indicate everything that is related to the fine (super-resolution) level with pixel size h .

A Convex Minimization Problem: We define the super-resolution model as the following convex minimization problem:

$$\min_{\hat{u}} \left\{ \lambda \|\nabla^h \hat{u}\|_{\varepsilon_u}^h + \sum_{i=1}^n \|\mathbf{DBW}_i \hat{u} - \check{f}_i\|_{\varepsilon_d}^{\xi h} \right\}, \quad (8)$$

with regularization using the Huber-Norm [9] that is defined as:

$$\|x\|_{\varepsilon}^h = \sum_{0 \leq i, j \leq MN} |x_{i,j}|_{\varepsilon} h^2 \quad \text{and} \quad |x|_{\varepsilon} = \begin{cases} \frac{|x|^2}{2\varepsilon} & \text{if } |x| \leq \varepsilon \\ |x| - \frac{\varepsilon}{2} & \text{if } |x| > \varepsilon \end{cases}. \quad (9)$$

While Total Variation based approaches [16] favor flat gray value regions causing staircase artifacts, the Huber-Norm has the advantage of smoothing small gradients while preserving strong edges. The linear operators \mathbf{D} , \mathbf{B} and \mathbf{W} denote downsampling, blurring and warping operators and will be explained in detail in the following. The free parameter λ models the tradeoff between regularization and data term.

Warping Operator \mathbf{W} : An exact warping to align the input images is a crucial factor for super-resolution. To allow for arbitrary motion, sub-pixel accurate optical flow is required. There are already various dense approaches with very high accuracy available [13]. Throughout this paper we used the GPU-based implementation from Werlberger et al. [14]. This approach utilizes a variational minimization problem using Huber-Norm regularization and an L^1 based data term to model the optical flow constraint. For optical flow calculation we used bicubic upsamplings \hat{f}_i of the input images \check{f}_i . We denote the reference image that is used for super-resolution as \hat{f}_k with $1 \leq k \leq n$. From the image \hat{f}_k we calculate the optical flow to all input images. As a result we get n flow vector fields $\hat{w}_i : \Omega \rightarrow \mathbb{R}^2$. In the super-resolution model (8) we denoted the optical flow as a linear operator $\mathbf{W}_i : \mathbb{R}^{\hat{M}\hat{N}} \rightarrow \mathbb{R}^{\hat{M}\hat{N}}$. As \mathbf{W}_i has a size of $(\hat{M}\hat{N})^2$, direct storage and computation is not feasible. Therefore, we directly utilize the flow fields \hat{w}_i to warp the super-resolution image to the input images using bicubic interpolation. For the proposed minimization algorithm we also need the transposed operator \mathbf{W}_i^T such that $\langle \mathbf{W}_i a, b \rangle_X = \langle a, \mathbf{W}_i^T b \rangle_X$, as detailed in Section 4. This inverse warping can again be realized using the flow fields \hat{w}_i ,

but this time the input images are warped to the reference image. Therefore, in each pixel of the reference, the flow is used to find the corresponding position in the input image. The pixel values in the input image are weighted by the coefficients of the bicubic interpolation and added up in an accumulator image.

Blurring Operator \mathbf{B} : To account for blurring introduced by atmospheric effects, the optical system and the sensor with preceding filters can get very complex. More specifically one has to measure the point spread function for one specific configuration. As the aim of this paper is to present a general super-resolution algorithm one has to provide a generic model. Therefore the blurring operator \mathbf{B} is modeled by a simple Gaussian kernel. We chose the standard deviation $\sigma = \frac{1}{4}\sqrt{\xi^2 - 1}$ and 3σ for the kernel size. Again the transposed blurring operator \mathbf{B}^T is not the same as a simple Gaussian blur as which \mathbf{B} is implemented. In turn, one has to perform the operation using an accumulator to obtain correct results.

Downsampling Operator \mathbf{D} : We first define the formation of a pixel value $f_{i,j}$ at position (ih, jh) and size h^2 by calculating the mean of the continuous image $g : \Omega \rightarrow \mathbb{R}$ in the region of the pixel $f_{i,j} = \frac{1}{h^2} \int_{\Delta_{i,j}^h} g(\mathbf{x}) d\mathbf{x}$, with the pixel region $\Delta_{i,j}^h = (ih, jh) + [-\frac{h}{2}, \frac{h}{2})^2$. The input images and the super-resolution image are just samplings with different pixel size h . Since we only know the discretized input images we have to model the downsampling process in the discrete setting. This is done using a weighted area integral over the region of the low resolution pixel (see Fig. 1 for an illustration):




Fig. 1. Illustration of the downsampling process that is modeled as an area integral.

$$\check{u}_{k,l} = \frac{1}{\mu(\Delta_{k,l}^{\xi h})} \sum_{0 \leq i,j \leq \hat{M}\hat{N}} \mu(\Delta_{i,j}^h \cap \Delta_{k,l}^{\xi h}) \hat{u}_{i,j}, \quad (10)$$

with the Lebesgue measure $\mu(\Delta)$ denotes the area of the pixel region. The proposed definition of the downsampling operator \mathbf{D} has the advantage that the scale $\xi \in \mathbb{R}^+$ and is not restricted to integer values. Again, an accumulator is used to calculate \mathbf{D}^T .

4 A First Order Primal-Dual Algorithm

Recently, in [15] and [17], a first order primal-dual algorithm was proposed to solve convex-concave saddle-point problems of the type:

$$\min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y} \{ \langle K\mathbf{x}, \mathbf{y} \rangle + G(\mathbf{x}) - F^*(\mathbf{y}) \}, \quad (11)$$

with a continuous linear operator $K : X \rightarrow Y$, and $G : X \rightarrow [0, \infty)$ and $F^* : Y \rightarrow [0, \infty)$ being convex functions. Basically, the algorithm consists of taking simple gradient steps in the primal variables \mathbf{x} and dual variables \mathbf{y} . In addition a leading point $\bar{\mathbf{x}}$ is computed to ensure convergence of the algorithm [15]. The basic iterations of the algorithm are defined as

$$\begin{aligned}\mathbf{y}^{n+1} &= (1 + \sigma \partial F^*)^{-1}(\mathbf{y}^n + \sigma K \bar{\mathbf{x}}^n) \\ \mathbf{x}^{n+1} &= (1 + \tau \partial G)^{-1}(\mathbf{x}^n - \tau K^* \mathbf{y}^{n+1}) \\ \bar{\mathbf{x}}^{n+1} &= 2\mathbf{x}^{n+1} - \mathbf{x}^n,\end{aligned}\tag{12}$$

with ∂F^* and ∂G the subgradients of F^* and G . The operators $(1 + \sigma \partial F^*)^{-1}$ and $(1 + \tau \partial G)^{-1}$ denote the usual resolvent operators. The primal and dual step sizes $\tau > 0$ and $\sigma > 0$ are chosen such that $\tau \sigma L^2 \leq 1$, where $L = \|K\|$ is the norm of the linear operator K .

In order to apply the primal-dual algorithm to the proposed super-resolution model, we have to transform the minimization problem (8) into the form of (11). Using duality principles (in particular the Legendre-Fenchel transform), we arrive at the following primal-dual formulation of the super-resolution model:

$$\begin{aligned}\min_{\hat{u}} \sup_{\hat{\mathbf{p}}, \check{q}} \left\{ \langle \nabla^h \hat{u}, \hat{\mathbf{p}} \rangle_Y^h - \frac{\varepsilon_u}{2\lambda h^2} \|\hat{\mathbf{p}}\|^2 - \delta_{\{|\hat{\mathbf{p}}| \leq \lambda h^2\}} \right. \\ \left. + \sum_{i=1}^n \left(\langle \check{q}_i, \mathbf{DBW}_i \hat{u} - \check{f}_i \rangle_X^{\xi h} - \frac{\varepsilon_d}{2(\xi h)^2} \|\check{q}_i\|^2 - \delta_{\{|\check{q}_i| \leq (\xi h)^2\}} \right) \right\},\end{aligned}\tag{13}$$

with $\hat{\mathbf{p}} \in \hat{Y}$ and $\check{q} \in \check{X}$ denoting the dual variables. The indicator function δ_Σ for the set Σ given as $\delta_\Sigma(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in \Sigma, \\ \infty & \text{else.} \end{cases}$

Now, applying the basic iterations (12) to our saddle-point problem (13), we obtain the following iterative algorithm:

$$\begin{aligned}\hat{\mathbf{p}}^{n+1} &= \Pi_{B_0^{\lambda h^2}} \left(\frac{\hat{\mathbf{p}}^n + \sigma h^2 \nabla^h \bar{u}^n}{1 + \frac{\sigma \varepsilon_u}{\lambda h^2}} \right), \\ \check{q}_i^{n+1} &= \Pi_{B_0^{(\xi h)^2}} \left(\frac{\check{q}_i^n + \sigma (\xi h)^2 (\mathbf{DBW}_i \bar{u}^n - \check{f}_i)}{1 + \frac{\sigma \varepsilon_d}{(\xi h)^2}} \right), \\ \hat{u}^{n+1} &= \hat{u}^n - \tau \left(-\operatorname{div}^h \hat{\mathbf{p}}^{n+1} + (\xi h)^2 \sum_{i=1}^n (\mathbf{W}_i^T \mathbf{B}^T \mathbf{D}^T \check{q}_i^{n+1}) \right), \\ \bar{u}^{n+1} &= 2\hat{u}^{n+1} - \hat{u}^n.\end{aligned}\tag{14}$$

The projector $\Pi_{B_0^r}$ denotes the orthogonal projection to a L^2 ball with radius r . In the 2D-case this can be computed as $\Pi_{B_0^{\lambda h^2}} = \hat{\mathbf{p}} / \max\{1, \frac{|\hat{\mathbf{p}}|}{\lambda h^2}\}$. While the 1D-case $\Pi_{B_0^{(\xi h)^2}}$ results in a simple clamping to the interval $[-(\xi h)^2, (\xi h)^2]$. Based on experiments, we are using the step sizes $\tau = \frac{\xi}{\sqrt{L^2 \lambda}}$ and $\sigma = \frac{1}{\sqrt{\tau L^2}}$.

5 Experimental Results

The proposed algorithm (14) was implemented on the GPU using the CUDA framework. We will prove that the chosen model generalizes very well by using different input sources for all experiments. First we evaluated the influence of different types of noise, namely quantization noise, Gaussian noise and salt and pepper noise. In Fig. 2 the results are depicted for different scale factors. The input data was generated artificially such that the exact warping, blur and downsampling operators are known. The results are compared to the original image, and the Signal-to-noise ratio (SNR) and the Structural Similarity (SSIM) [18] were calculated. Fig. 2 demonstrates that quantization noise using an 8bit representation causes a dramatic loss of information for large scale factors, while using 16bit quantization the visual appearance for $\xi = 16$ still recovers most of the fine details. Unfortunately even weak Gaussian noise causes strong degeneration. It also shows that the proposed super-resolution model can handle a reasonable amount of outliers due to the robust Huber-Norm in the data term.



























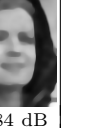







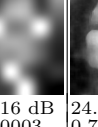

Input	Quantization				Gaussian Noise		Impulse Noise	
	Input	float	16bit	8bit	Input	$\sigma=0.01$	Input	10%
4x 								
	29.77 dB 0.94079	106.95 dB 1	47.59 dB 0.99838	39.09 dB 0.99887	29.5 dB 0.92747	34.1 dB 0.96136	16.59 dB 0.4668	33.67 dB 0.96934
16x 								
	24.92 dB 0.83238	47.95 dB 0.99834	40.21 dB 0.99166	32.58 dB 0.95394	24.84 dB 0.82358	30.65 dB 0.92639	15.62 dB 0.50452	32.01 dB 0.95136
64x 								
	21.09 dB 0.63877	39.98 dB 0.99108	35.73 dB 0.97322	30.26 dB 0.91862	21.08 dB 0.63697	27.46 dB 0.86119	14.78 dB 0.41485	24.84 dB 0.87617
256x 								
	19.88 dB 0.55709	35.32 dB 0.97172	32.58 dB 0.92639	25.98 dB 0.81105	16.52 dB 0.37576	24.78 dB 0.75237	13.16 dB 0.30003	24.73 dB 0.75077

Fig. 2. Comparison of the influence of different noise on the superresolution reconstruction for the scales 2, 4, 8 and 16 using synthetic data. The input columns always depict the input image that was chosen as reference image (using bicubic upsampling).

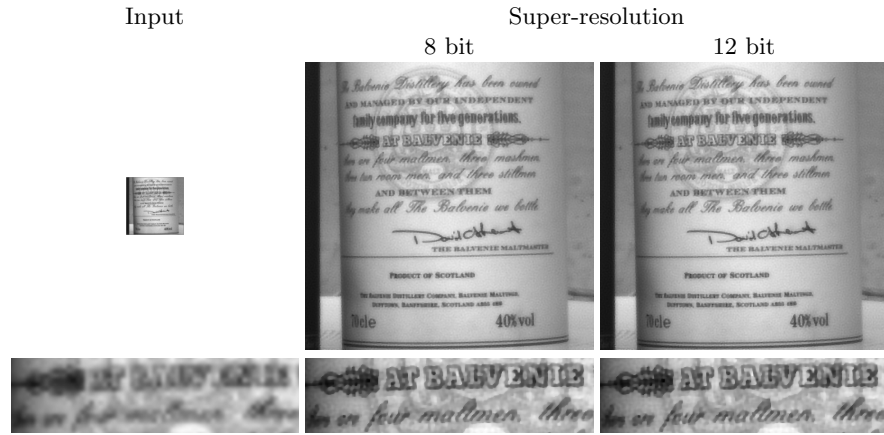


Fig. 3. Comparison of the influence of quantization using a real dataset. The bottom row depicts contrast adapted crops of the top images.

The effect of quantization noise is also investigated for real data in Fig. 3. We used a Prosilica GC1600 together with a Pentax 4.8mm 1:1.8 lens to capture a 8bit and a 12bit video. Unfortunately, in this case no notable visual gain can be seen when comparing the super-resolution results. A thorough evaluation of different cameras is left for future investigations.

In Fig. 4, a scene with strong motion is used as input. Optical Flow was calculated using [14]. To demonstrate the capability of using arbitrary scale factors, we chose $\xi = 2.67$ as scale factor. The results demonstrate that the proposed algorithm can easily handle strong occlusions due to the robust Huber-Norm. It shows that we obtain good results that are comparable to the more complex super-resolution motion layer decomposition approach of Schoenemann [10].

Finally, we compare our algorithm to the closely related work of Mitzel et al. [7]. In Fig. 5, super-resolution was done on a public car sequence¹. Note that we clearly improve the sharpness and detail recovery in the super-resolution image. Using the proposed algorithm on can easily read the cars license plate which is not the case for the other algorithm.

6 Conclusion

We presented a variational super-resolution approach that uses a robust Huber-Norm for regularization as well as in the data term. A fast primal-dual algorithm was used to solve a convex saddle-point problem. We put particular emphasis to the accurate design of the model, the used operators and discretization. The exact choice of warping, blurring and downsampling operators was discussed, and results demonstrate that we obtain state of the art results even for very

¹ <http://users.soe.ucsc.edu/~milanfar/software/sr-datasets.html>

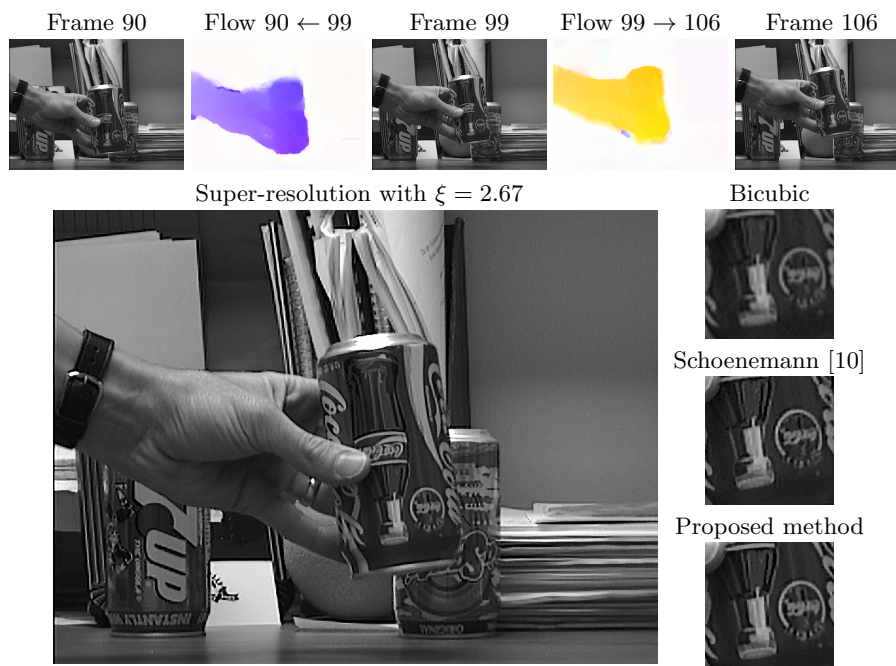


Fig. 4. Demonstration of superresolution on real data with heavy occlusions. The flow images in the first row use Middlebury color coding.

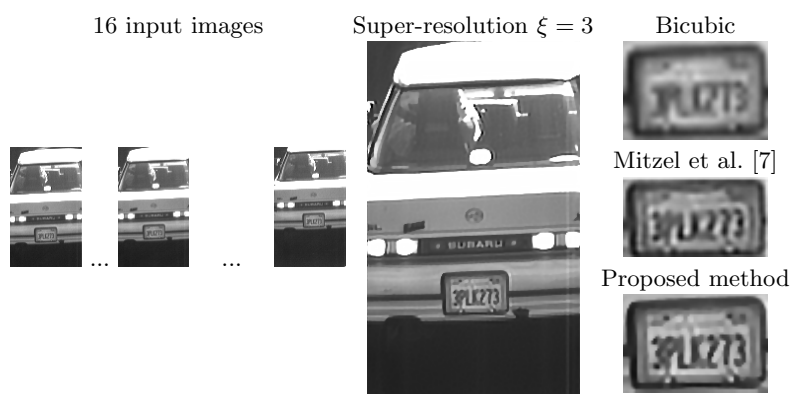


Fig. 5. Comparison to the closely related work of Mitzel et al. [7] using super-resolution to recover a license plate. (Contrast adapted for crops)

difficult scenes. The proposed approach can handle large occlusions as well as arbitrary scale factors. An extensive evaluation of the influence of noise was carried out. It showed that in theory a higher dynamic range can significantly

improve the recovered results. Though for real data, sensor noise is the limiting factor. We hope this will trigger further research on the practical implications of bit depth on superresolution. In future work we will also study more powerful regularization terms (e.g. wavelets).

Acknowledgement

This work was supported by the Austrian Research Promotion Agency within the vdQA project (no. 816003).

References

1. Shannon, C.E.: A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.* **5**(1) (2001) 3–55
2. Kosarev, E.L.: Shannon’s superresolution limit for signal recovery. *Inverse Problems* **6**(1) (1990) 55
3. Baker, S., Kanade, T.: Limits on super-resolution and how to break them. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24** (2002) 1167–1183
4. Protter, M., Elad, M., Takeda, H., Milanfar, P.: Generalizing the non-local-means to super-resolution reconstruction. In: *IEEE Transactions on Image Processing*. (2009) 36
5. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: *ICCV*. (2009)
6. van Ouwerkerk, J.: Image super-resolution survey. *Image and Vision Computing* **24**(10) (2006) 1039 – 1052
7. Mitzel, D., Pock, T., Schoenemann, T., Cremers, D.: Video super resolution using duality based TV-L1 optical flow. In: *DAGM, Jena, Germany* (2009)
8. Farsiu, S., Robinson, D., Elad, M., Milanfar, P.: Fast and robust multi-frame super-resolution. *IEEE Transactions on Image Processing* **13** (2003) 1327–1344
9. Huber, P.: *Robust Statistics*. Wiley, New York (1981)
10. Schoenemann, T.: *Combinatorial Solutions for Shape Optimization in Computer Vision*. PhD thesis, Department of Computer Science, University of Bonn, Germany (2008)
11. Shechtman, E., Caspi, Y., Irani, M.: Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27** (2005) 531–545
12. Takeda, H., Milanfar, P., Protter, M., Elad, M.: Super-resolution without explicit subpixel motion estimation. *Trans. Img. Proc.* **18**(9) (2009) 1958–1975
13. Middlebury: <http://vision.middlebury.edu/flow/> (2010) accessed 06-May-2010.
14. Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., Bischof, H.: Anisotropic Huber-L1 optical flow. In: *BMVC, London, UK* (September 2009)
15. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the mumford-shah functional. In: *ICCV*. (2009)
16. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Phys. D* **60**(1-4) (1992) 259–268
17. Esser, E., Zhang, X., Chan, T.F.: A general framework for a class of first order primal-dual algorithms for tv minimization. *UCLA CAM Report [09-67]* (August 2009)
18. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13** (2004) 600–612