Robust Optical Flow Based Deformable Registration of Thoracic CT Images

Martin Urschler[1,2], Manuel Werlberger[2], Eva Scheurer[1], and Horst Bischof[2]

[1] Ludwig-Boltzmann Institute for Clinical Forensic Imaging, Graz, AT, ⋆
[2] Institute for Computer Graphics & Vision, Graz University of Technology, AT ⋆⋆

Abstract. We present an optical flow deformable registration method which is based on robust measures for data and regularization terms. We show two specific implementations of the method, where one penalizes gradients in the displacement field in an isotropic fashion and the other one regularizes by weighting the penalization according to the image gradients anisotropically. Our data term consists of the L^1-norm of the standard optical flow constraint. We show a numerical algorithm that solves the two proposed models in a primal-dual optimization setup. Our algorithm works in a multi-resolution manner and it is applied to the 20 data sets of the EMPIRE10 registration challenge. Our results show room for improvement. Our rather simple model does not penalize non-diffeomorphic transformations, which leads to bad results on one of the evaluation measures, and it seems unsuited for large deformations cases. However, our algorithm is able to perform registrations of data set sizes around 400^3 on the order of a few minutes using a dedicated CUDA based GPU implementation, which is very fast compared to other reported algorithms.

1 Introduction

Nonlinear (deformable) registration of data sets acquired at different points in time is an important research topic in medical image analysis. Image sequences of soft tissue organs like lung or liver during breathing or the beating heart often require registration algorithms to compensate for motion differences. Surveys on nonlinear registration techniques in medical imaging can be found in Maintz and Viergever [1] or Crum *et al.* [2]. In literature, one distinguishes feature- and intensity based nonlinear registration methods. Intensity based methods [3, 4] are often favored since they make use of the entire image information, however, they come at the cost of a higher computational effort.

For deformable registration problems in the context of intra-modality applications, intensity based optical flow approaches are very popular. In a variational framework setting [5, 6] one formulates the optical flow problem as an energy minimization consisting of a data and a regularization term, a technique which

⋆ martin.urschler@cfi.lbg.ac.at
⋆⋆ werlberger@icg.tugraz.at

has its roots in [7]. In this work we show an efficient deformable registration algorithm based on the optical flow constraint. We penalize the gradients of the displacement field in order to regularize the variational formulation. Our formulation is simple and straight-forward, however, compared to standard optical flow using quadratic terms, we penalize data and regularization term using the L^1 norm [8], which has the advantage to be more robust to non-Gaussian noise and to leave discontinuities in the displacement field intact. Our exact formulation is a 3D extension of the anisotropic Huber-L^1 optical flow method presented in [9], which has shown to be both accurate and extremely fast for 2D optical flow applications.

A drawback with optical flow regularization using the L^1-norm compared to a quadratic norm is its difficulty to come up with a numerical scheme to solve the underlying partial differential equation since the derivative of the L^1-norm is undefined at zero. In Section 2 we describe our optical flow formulation in detail and the numerical scheme which is based on a primal-dual algorithm. In Section 3 we describe the whole implementation framework and include the used parameter settings for the EMPIRE10 evaluation. In Section 4 we present the evaluation results of our two algorithm variants. Finally Section 5 concludes our work by discussing the results of the evaluation.

2 Algorithm Description

In this section we describe our core algorithm, the Huber-L^1 optical flow formulation and its variant that takes image gradients into consideration for regularization (anisotropic optical flow). We present the basic formulation as well as the numerical scheme that implements the minimization of the optical flow energy functional. This algorithm is embedded into a multi-resolution framework to register fixed and moving lung CT images from the EMPIRE10 evaluation. For the lung registration it also gets initialized by a very simple rigid registration that only takes translation into account. In this section we focus on the generic optical flow algorithm while the following section will provide the details of the framework that was used for the lung CT evaluation as well as its required choices of parameters.

In computer vision motion estimation via optical flow is an important topic which also has obviously found its application in medical image registration. The aim of optical flow is to compute the motion in a sequence of images, where in medical applications often only a pair of images (e.g. thorax CT scans at different states in the breathing cycle) is available. In a variational framework a possible formulation of the optical flow motion estimation is given by

$$\min_{\mathbf{u}} \int_{\Omega} \sum_{i=1}^{3} |\nabla u_i^T \mathbf{D} \nabla u_i| + \lambda \|\rho(\mathbf{u})\|_1 d\mathbf{x} \ , \tag{1}$$

where $\mathbf{u} = (u_1, u_2, u_3)^T : \Omega \to \mathbb{R}^3$ is the motion field, and

$$\rho(\mathbf{u}) = I_t + (\nabla I)^T (\mathbf{u} - \mathbf{u}^0) \tag{2}$$

is the traditional optical flow constraint (OFC). It is obtained from a linearization of the assumption that the intensities of the pixels stay constant over time. I_t is the temporal derivative of the image sequence (image pair), defined as the difference between fixed and moving image. ∇I is the spatial image gradient, and $\mathbf{u}^0$ is some given motion field. The regularization term is a total variation formulation with $\mathbf{D}$ a symmetric, positive definite diffusion tensor. The parameter λ is used to define the trade-off between data fitting and regularization.

Note that the OFC is valid only for small motion $(\mathbf{u}-\mathbf{u}^0)$. In order to account for larger motions, the entire approach has to be integrated into a coarse-to-fine framework using a multi-resolution pyramid to re-estimate $\mathbf{u}^0$. Further, we want to stress that the linearized OFC combined with the TV regularization is a convex energy formulation, so a global solution for the linearization may be obtained.

In this work we use two specific formulations of optical flow (OF), isotropic and anisotropic OF (i.e. OF_{iso}, OF_{aniso}). One can derive the isotropic formulation from (1) by setting the diffusion tensor $\mathbf{D}$ in the regularization term equal to the identity matrix

$$\min_{\mathbf{u}} \int_{\Omega} \sum_{i=1}^{3} |\nabla u_i|_\epsilon + \lambda \|\rho(\mathbf{u})\|_1 d\mathbf{x} \ , \tag{3}$$

such that the regularization is isotropic in all directions for all deformation field components. In (3) we have replaced the L^1-norm with the more robust Huber norm $|\nabla u_i|_\epsilon$, which has beneficial properties in avoiding staircasing artifacts [10, 9]. The Huber norm is defined as $|\nabla u_i|_\epsilon = \frac{|\nabla u_i|^2}{2\epsilon}$ if $|\nabla u_i| \leq \epsilon$ and $|\nabla u_i|_\epsilon = |\nabla u_i| - \frac{\epsilon}{2}$ otherwise. The advantage of the Huber norm over the widely used total variation (TV) norm for regularization is the reduction in favoring piecewise constant solutions in weakly textured areas, which leads to staircasing artifacts. As a benefit compared to quadratic regularization, still edges are preserved and not smoothed over.

As an alternative algorithm we derive the anisotropic formulation from (1) by using a diffusion tensor calculated from the image gradient

$$\min_{\mathbf{u}} \int_{\Omega} \sum_{i=1}^{3} |\nabla u_i^T \mathbf{D}_I \nabla u_i|_\epsilon + \lambda \|\rho(\mathbf{u})\|_1 d\mathbf{x} \ . \tag{4}$$

Under the assumption that gradients in the deformation field coincide with image gradients, this gives us a more accurate deformation field estimate at the cost of computing the diffusion tensor $\mathbf{D}_I$ pixel-wise as a diagonal matrix

$$\mathbf{D}_I = \begin{pmatrix} D_{I,x} & 0 & 0 \\ 0 & D_{I,y} & 0 \\ 0 & 0 & D_{I,z} \end{pmatrix} \tag{5}$$

with the components $D_{I,x} = \exp\left(-\alpha * I_x^\beta\right)$, $D_{I,y} = \exp\left(-\alpha * I_y^\beta\right)$ and $D_{I,z} = \exp\left(-\alpha * I_z^\beta\right)$. This is only an approximate diffusion tensor due to efficiency reasons, however, it is still symmetric and positive definite. Note that again we are using the robust Huber-norm in the regularization. Next we describe the numerical solver for the presented models. The definition of the diffusion tensor is the difference between our two algorithms, however, for the numerical solver this difference is only minor.

2.1 Numerical Solver & Discretization

After defining our continuous optical flow model we proceed with a discretization of the formulation and a numerical algorithm to solve the optical flow problem. We use a primal-dual convex optimization scheme to solve the linearized optical flow. This scheme is a saddle point problem where one seeks to minimize a primal and to maximize a dual variable. For our discretization, the input images I_0 (fixed image) and I_1 (moving image) with the dimensions $M \times N \times L$ are defined on the Cartesian grid $\Omega^h = \{(ih_x, jh_y, kh_z) : 1 \leq i \leq M, 1 \leq j \leq N, 1 \leq k \leq L\}$. We define the discretized optical flow $\mathbf{u}^h = (u_1^{h_x}, u_2^{h_y}, u_3^{h_z})^T \in X^h$, where the vector space $X^h = \mathbb{R}^{3MNL}$. The potentially anisotropic image spacing of medical images is taken into account using spacings h_x, h_y, h_z and the discrete pixel position $(ih_x, jh_y, kh_z) \in \Omega^h$. Discretization of the gradient ∇^h uses standard finite differences (forward differences) on the discrete lattice with Neumann boundary conditions. The primal-dual formulation requires a discretized divergence operator on the dual variable, where we use backward differences. Note that the gradient and the divergence operator are adjoint $\left\langle \nabla^h \mathbf{u}^h, \mathbf{p}^h \right\rangle \equiv - \left\langle \mathbf{u}^h, \mathrm{div^h}\, \mathbf{p}^h \right\rangle$

The discretized version of (4) on the lattice Ω^h reads

$$\min_{\mathbf{u^h}} \left\{ \left\| \nabla^h \mathbf{u}^h \right\|_\epsilon + \lambda \left\| \rho(\mathbf{u}^h) \right\|_1 \right\} \ . \tag{6}$$

Recently, it was shown in [11–13] that primal-dual approaches provide an excellent performance for solving convex-concave saddle-point problems of the form

$$\min_x \sup_y \left\{ \langle Kx, y \rangle - F^*(y) + G(x) \right\} \ , \tag{7}$$

with K a linear operator, and the convex functions F^* and G. The basic iterations of the primal-dual algorithm of [11] are defined as

$$\begin{cases} y^{n+1} &= \left(1 + \tau_d \partial F^*\right)^{-1} \left(y^n + \tau_d K \tilde{x}\right) \\ x^{n+1} &= \left(1 + \tau_p \partial G\right)^{-1} \left(x^n - \tau_p K^* y^{n+1}\right) \\ \tilde{x}^{n+1} &= 2x^{n+1} - x^n \end{cases} \tag{8}$$

Here the stepwidth for the primal and dual update is given as τ_p and τ_d, K^* denotes the adjoint operator of K and ∂ a partial derivative. In our case K resembles the gradient operator and K^* its adjoint operator, the divergence.

To gain a primal-dual saddle point problem like in (7), we apply the Legendre-Fenchel transform to (6) and obtain the optimization problem

$$\min_{\mathbf{u}^h} \sup_{\mathbf{p}^h} \left\{ \left\langle \nabla^h \mathbf{u}^h, \mathbf{p}^h \right\rangle_X - \frac{\varepsilon}{2} ||\mathbf{p}^h||_2^2 - \delta_{P^h}(\mathbf{p}^h) + \lambda \left\| \rho(\mathbf{u}^h) \right\|_1 \right\} \ . \tag{9}$$

Here we introduce the dual variable $\mathbf{p}$ defined on the convex set P which is defined as $P^h = \left\{\mathbf{p}^h \in Y^h : \left\|\mathbf{p}^h\right\|_\infty \leq 1\right\}$, where Y^h denotes the convex set $Y^h = X^h \times X^h$ and $\left\|\mathbf{p}^h\right\|_\infty$ the discrete maximum norm. Embedding the primal-dual saddle point formulation (9) into the generic formulation (7) yields $G(\mathbf{u}^h) = \lambda \left\|\rho(\mathbf{u}^h)\right\|_1$ and $F^*(\mathbf{p}^h) = \frac{\varepsilon}{2}||\mathbf{p}^h||_2^2 + \delta_{P^h}(\mathbf{p}^h)$. Based on [11], the resolvent operator for $F^*(\mathbf{p}^h)$ is given as a pointwise projection onto an L^2 ball yielding

$$\mathbf{p}^h = \left(1 + \tau_d \partial F^*\right)^{-1}\left(\tilde{\mathbf{p}}^h\right) \quad \Longleftrightarrow \quad \mathbf{p}^h_{i,j,k} = \frac{\frac{\tilde{\mathbf{p}}^h_{i,j,k}}{1+\tau_d\varepsilon}}{\max\left(1, \left|\frac{\tilde{\mathbf{p}}^h_{i,j,k}}{1+\tau_d\varepsilon}\right|\right)} \ . \tag{10}$$

With respect to $G(\mathbf{u}^h)$ the solution to the resolvent operator is given by

$$\begin{aligned}\mathbf{u}^h = \left(1 + \tau_p \partial G\right)^{-1}\left(\tilde{\mathbf{u}}^h\right) \quad \Longleftrightarrow \quad \mathbf{u}^h_{i,j,k} &= \tilde{\mathbf{u}}^h_{i,j,k} \\ &+ \begin{cases} \tau_p \lambda \nabla I_{i,j,k} & \text{if } \ \rho(\tilde{\mathbf{u}}^h_{i,j,k}) \ < -\tau_p \lambda |\nabla I|^2_{i,j,k} \\ -\tau_p \lambda \nabla I_{i,j,k} & \text{if } \ \rho(\tilde{\mathbf{u}}^h_{i,j,k}) \ > \quad \tau_p \lambda |\nabla I|^2_{i,j,k} \\ -\rho(\tilde{\mathbf{u}}^h_{i,j,k}) \frac{\nabla I_{i,j,k}}{|\nabla I|^2_{i,j,k}} & \text{if } |\rho(\tilde{\mathbf{u}}^h_{i,j,k})| \leq \quad \tau_p \lambda |\nabla I|^2_{i,j,k} \end{cases} \ .\end{aligned} \tag{11}$$

This iterative update scheme concludes our numerical implementation of the optical flow algorithms for both cases. We have to specify a maximum number of iterations in practice to work with this iterative scheme. Further, in the course of optimizing for $\mathbf{u}$ we warp the moving image from time to time to take the linearized model into account. For this purpose we perform an outer loop over a number of warps and use the current solution of $\mathbf{u}^0 = \mathbf{u}$ to warp the moving image.

3 Experimental Setup & Parameters

In the previous section we have described two variants of an optical flow algorithm and their numerical implementation. We have implemented the whole algorithm framework as well as the optical flow registration per resolution level on a CUDA-based NVidia Tesla C1060 desktop computer platform with 4GB of graphics RAM. This leads to a very efficient registration where the isotropic optical flow takes around 4 minutes on average per data set and the anisotropic variant around 6 minutes for the chosen parameter settings.

For the EMPIRE10 evaluation we utilize both algorithms and embed them into a multi-resolution framework. Our algorithm constructs a fixed and a moving image pyramid and calculates the solution of the optical flow registration on each pyramid level starting with the coarsest. Then it upsamples the displacement field $\mathbf{u}$ to propagate it to the next finer level and uses this upsampled solution as initial solution $\mathbf{u}_0$. On the coarsest resolution we perform a simple initialization of the displacement field using the difference vector of the centers of gravity of the provided segmented lung mask images. This is our only pre-registration that we perform on all of the 20 data sets of EMPIRE10. Our fixed image pyramid is constructed by first downsampling the fixed input image to $256 \times 256 \times 256$ voxels. This also defines the image size of our warped image pyramid and the pyramids for the displacement field components. We currently can not work on the full resolution as finest pyramid level, due to a 4GB memory restriction of our implementation hardware. Further, we use the moving image to construct a pyramid with the original resolution as finest pyramid level. We downsample all of our pyramids by factors of 2 and repeat this procedure until we have a total of 5 pyramid levels. Note that the moving image pyramid uses a different (finer) voxel grid than the rest of the pyramids, however due to the tricubic interpolation in the warping step, which is of course performed in physical space, not voxel space, this is not an issue. Tricubic interpolation is also used for all other warping and resampling steps in the algorithm. As a consequence of our memory restrictions and the finest fixed pyramid level, which never exceeds 256^3, we have to upsample the registration result from 256^3 to the original size of the fixed image data set. This way we will always make a certain error from the upsampling, since there is no refinement of the displacement field on the original fixed image resolution. We are currently looking into methods to parallelize the algorithm on a higher level and distribute them onto several Tesla GPUs.

In our optical flow algorithm there are a number of parameters which we either leave fixed or compute adaptively starting from some fixed value. Thus, our algorithm falls into the category of fully automatic according to the EMPIRE10 challenge rules. First, a very important parameter is the trade-off λ between regularization and optical flow constraint data term. We have chosen a value of $\lambda = 50$ in our experiments, both for isotropic and anisotropic regularization. This value gave in our experiments reasonable results, however, after investigating the EMPIRE10 evaluation results we have to say that we presumably have chosen λ too large, since the diffeomorphic behaviour of our result displacement field is not satisfactory. This indicates that we have weighted our optical flow data term too high. The specified λ is taken for the finest level of the pyramid, for coarser levels we multiply it by a factor of $1.5^{scaleLevel}$ to increase the influence of the data term.

Some more parameters that we have to set are $\epsilon = 0.01$ from the Huber norm which was set empirically and does not seem to be critical. α and β for the edge weighting are taken as $\alpha = 10, \beta = 1$, these parameters influence how large different edges are weighted in the computation of the diffusion tensor (for OF_{aniso} only). The computation of the gradients for the diffusion tensor is performed on the fixed image using central differences. The number of warps and number of iterations are chosen as 40 and 25, respectively, which means that on the finest level we perform $40 \times 25 = 1000$ iterations in total. On coarser pyramid levels we adapt these values by multiplying with $1.5^{scaleLevel}$ in order to perform more iterations since these downsampled levels are less expensive to compute. The numerical primal-dual scheme requires two timesteps $\tau_p = 2 * sqrt(\frac{1}{h_x^2} + \frac{1}{h_y^2} + \frac{1}{h_z^2})^{-1}$ and $\tau_d = \frac{\tau_p}{4*(\frac{1}{h_x^2} + \frac{1}{h_y^2} + \frac{1}{h_z^2})}$. Another implementation note is, that the update of the displacement fields is always clamped to the distance equivalent to the voxel spacing of the current pyramid level, since the optical flow model per pyramid level is not valid for more than one voxel distance. This is another reason why the course-to-fine framework is mandatory. Finally before each warp with a displacement $\mathbf{u}_0$ we perform a $3 \times 3 \times 3$ median filter on the displacement field to remove outliers.

4 Evaluation Results

The following two tables show the results of the EMPIRE10 evaluation using the algorithms OF_{iso} (see Table 2) and OF_{aniso} (see Table 1).

	Lung Boundaries		**Fissures**		**Landmarks**		**Singularities**	
Scan Pair	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**
01	0.04	18.00	15.69	28.00	17.95	29.00	13.17	34.00
02		34.00		34.00		34.00		34.00
03	0.00	5.50	0.00	12.50	0.33	6.00	0.08	31.00
04	0.00	24.00	0.00	16.50	4.86	30.00	3.00	34.00
05	0.00	13.00	0.00	16.00	0.00	5.50	0.00	27.00
06	0.00	16.00	0.00	21.00	0.33	10.00	0.01	33.00
07	0.02	16.00	7.53	28.00	6.23	25.00	4.04	33.00
08	0.00	19.00	3.56	27.00	3.22	26.00	1.17	33.00
09	0.00	23.00	0.00	6.50	0.51	3.00	0.06	32.00
10	0.00	17.00	0.00	15.00	10.33	33.00	5.50	34.00
11	0.03	15.00	2.30	27.00	2.09	23.00	2.29	33.00
12	0.00	10.00	0.00	13.50	0.00	5.00	0.00	14.50
13	0.00	11.00	0.06	5.00	0.79	6.00	0.29	33.00
14	0.06	19.00	8.46	26.00	12.78	28.00	7.47	33.00
15	0.00	8.00	0.00	7.00	0.58	2.00	0.10	31.00
16	0.01	31.00	2.06	30.00	4.40	32.00	6.68	34.00
17	0.00	24.00	0.04	12.00	0.69	7.00	0.19	32.00
18	0.06	19.00	6.87	26.00	7.06	27.00	4.61	33.00
19	0.00	14.00	0.00	12.00	0.45	4.00	0.00	31.00
20	0.01	18.00	7.93	27.00	14.09	28.00	9.00	34.00
Avg	0.01	17.72	2.87	19.50	4.56	18.17	3.03	31.67
Average Ranking Overall								21.76
Final Placement								28

Table 1. Algorithm OF_{aniso}. Results for each scan pair, per category and overall. Rankings and final placement are from a total of 34 competing algorithms.

	Lung Boundaries		Fissures		Landmarks		Singularities	
Scan Pair	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**
01	0.27	26.00	15.06	27.00	31.93	31.00	6.43	32.00
02	0.02	31.00	0.00	15.00	0.54	19.00	0.02	31.00
03	0.00	22.00	0.02	30.00	0.90	29.00	0.10	32.00
04	0.00	18.00	0.00	16.50	6.68	33.00	0.43	31.00
05	0.00	13.00	0.00	16.00	0.23	22.00	0.00	30.00
06	0.00	16.00	0.00	7.00	0.46	25.00	0.01	32.00
07	0.40	27.00	11.71	33.00	14.40	30.00	2.22	32.00
08	0.05	25.00	9.23	33.00	9.52	34.00	0.45	32.00
09	0.00	22.00	0.00	25.00	1.45	32.00	0.05	31.00
10	0.01	22.00	0.00	15.00	9.38	32.00	1.14	31.00
11	0.26	26.00	7.15	33.00	5.54	32.00	0.85	31.00
12	0.06	29.00	0.00	13.50	0.89	27.00	0.01	33.00
13	0.00	17.00	0.09	17.00	0.95	14.00	0.14	30.00
14	0.88	28.00	13.34	29.00	19.43	31.00	3.63	32.00
15	0.00	27.00	0.00	15.00	0.90	26.00	0.08	30.00
16	0.00	25.00	0.01	8.50	2.02	28.00	0.48	31.00
17	0.00	20.00	0.06	27.00	0.91	15.00	0.05	30.00
18	0.29	24.00	25.25	32.00	14.72	29.00	1.36	32.00
19	0.06	31.00	0.05	32.00	1.61	32.00	0.22	33.00
20	0.23	26.00	6.83	25.00	20.06	30.00	2.83	32.00
Avg	0.13	23.75	4.44	22.47	7.13	27.55	1.02	31.40
Average Ranking Overall								26.29
Final Placement								31

Table 2. Algorithm OF_{iso}. Results for each scan pair, per category and overall. Rankings and final placement are from a total of 34 competing algorithms.

5 Discussion & Conclusion

From the previous section one can clearly see that the overall performance of our algorithm is rather weak compared to the other methods. This is mainly due to our rather simple optical flow model, which is a straight-forward extension from a two-dimensional algorithm [9]. The most important drawback of our method is the lack of a penalty for non-diffeomorphic transformations. This can also be seen in the fourth evaluation measure where we rank at the end of the field for most of the data sets. We are currently looking into a way to include such a penalty into our convex formulation, however, this unfortunately is not straight-forward. One could explicitly model diffeomorphic transformations similar to [14] by calculating the displacement and its inverse in the optical flow model and penalizing the difference between these two transformations. However, this would at least double the time for calculating our solutions, which brings us to the main benefit of our method, the computational efficiency. We are not aware of competing algorithms which are able to perform deformable image registration as fast as our algorithm is able to, using a dedicated CUDA based GPU implementation. This is a very important feature, especially for the online registration at the workshop where 10 data sets will have to be registered in 3 hours.

Another drawback of our method can be seen from the error measures on the data sets showing large differences in the breathing cycle (scan pairs 1,4,7,8,10,14,16,18,20). In this case inhaled and exhaled lungs lead to rather different appearances, where smaller vessels vanish in the exhaled data set due to the limitations in spatial resolution. Therefore, the implicit brightness constancy assumption of optical flow is not valid anymore, and our performance drops. So, here we have to conclude that the optical flow algorithm is not useful for too large displacements in combination with disappearing vessel structures.

What we learned from the evaluation was the significantly better performance of the anisotropic model compared to the isotropic one. This implies that the image gradient information is very important for the registration and the resulting displacement fields. Performance is better although we were not able to register one data set (data set 2) due to memory problems which could not be solved in time for the offline workshop contribution.

In conclusion we want to stress that this EMPIRE10 challenge is a very important step forward in the evaluation of deformable registration algorithms. In our contribution we have observed the strengths (efficiency) and also the limitations of a simple optical flow implementation (which is a widely used model in literature) with respect to the problem of thorax CT registration. We will continue to work on improved models to perform better on the presented data sets.

References

1. Maintz, J.B.A., Viergever, M.A.: A Survey of Medical Image Registration. Medical Image Analysis **2**(1) (1998) 1–36
2. Crum, W.R., Hartkens, T., Hill, D.L.G.: Non-rigid image registration: Theory and Practice. The British Journal of Radiology - Imaging Processing Special Issue **77** (December 2004) S140–S153
3. Bajcsy, R., Kovacic, S.: Multiresolution Elastic Matching. Computer Vision, Graphics and Image Processing **46**(1) (1989) 1–21
4. Thirion, J.P.: Image matching as a diffusion Process: An analogy with Maxwell's demons. Medical Image Analysis **2**(3) (1998) 243–260
5. Aubert, G., Deriche, R., Kornprobst, P.: Computing optical flow via variational techniques. SIAM Journal on Applied Mathematics **60**(1) (2000) 156–182
6. Modersitzki, J.: Numerical Methods for Image Registration. Oxford University Press (2004)
7. Horn, B.K.P., Schunck, B.G.: Determining optical flow. Artificial Intelligence **17**(1-3) (1981) 185–203
8. Pock, T., Urschler, M., Zach, C., Beichel, R., Bischof, H.: A duality based algorithm for TV-L1-optical-flow image registration. In: Proc Int Conf Med Image Comput Comput Assist Interv (MICCAI 2007). (2007) 511–518

9. Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., Bischof, H.: Anisotropic Huber-L1 optical flow. In: Proceedings of the British Machine Vision Conference (BMVC), London, UK (September 2009)
10. Huber, P.J.: Robust Statistics. Wiley Series in Probability and Statistics. (1981)
11. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. preprint
12. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the mumford-shah functional. In: ICCV. (2009)
13. Esser, E., Zhang, X., Chan, T.: A general framework for a class of first order primal-dual algorithms for TV minimization. UCLA CAM Report 09-67 (2009)
14. Christensen, G.E., Johnson, H.J.: Consistent Image Registration. IEEE Transactions on Medical Imaging **20**(7) (2001) 568–582