Approximated Curvature Penalty in Non-rigid Registration using Pairwise MRFs

Ben Glocker[1,2⋆], Nikos Komodakis[3], Nikos Paragios[2,4], Nassir Navab[1]

[1] Computer Aided Medical Procedures (CAMP), TU München, Germany
[2] Laboratoire MAS, Ecole Centrale Paris, Chatenay-Malabry, France
[3] Computer Science Department, University of Crete, Greece
[4] Equipe GALEN, INRIA Saclay - Ile-de-France, Orsay, France
`glocker@in.tum.de`

Abstract. Labeling of discrete Markov Random Fields (MRFs) has become an attractive approach for solving the problem of non-rigid image registration. Here, regularization plays an important role in order to obtain smooth deformations for the inherent ill-posed problem. Smoothness is achieved by penalizing the derivatives of the displacement field. However, efficient optimization strategies (based on iterative graph-cuts) are only available for first-order MRFs which contain cliques of size up to two. Higher-order cliques require graph modifications and insertion of auxiliary nodes, while pairwise interactions actually allow only regularization based on the first-order derivatives. In this paper, we propose an approximated curvature penalty using second-order derivatives defined on the MRF pairwise potentials. In our experiments, we demonstrate that our approximated term has similar properties as higher-order approaches (invariance to linear transformations), while the computational efficiency of pairwise models is preserved.

1 Introduction

Non-rigid image registration is an important problem in computer vision and medical imaging. Given two images I and J, one seeks a transformation T which aligns the corresponding objects visible in the images. This is commonly solved by posing an energy minimization problem where the objective function is a sum of a matching criteria S and a regularization term R,

$$\hat{T} = \arg\min_{T} S(I, J \circ T) + \alpha R(T) \quad . \tag{1}$$

Here, α is a weighting factor controlling the influence of the regularization term. In the case of non-rigid registration, the transformation is often defined as the identity transformation plus a dense displacement field D. The new location of an image point $\mathbf{x}$ is then computed by

$$T(\mathbf{x}) = \mathbf{x} + D(\mathbf{x}) \quad . \tag{2}$$

⋆ This work is partially supported by Siemens Healthcare, Erlangen, Germany.

Regularization plays an important role due to the inherent ill-posedness of the problem [1]. A natural approach for regularization is to penalize the derivatives of the displacement field. Smoothness terms based on the first-order derivatives penalize high gradients and thus, piecewise constant deformations are favored. Such smoothness models require a proper pre-alignment by linear registration prior to the non-linear one, since penalizing the gradients is only invariant to global translation. If still some linear transformation (such as rotation or scaling) is present, penalizing the gradients might prohibit a proper non-rigid alignment. Since a perfect linear alignment is not trivial to achieve when deformations are present, one can consider to define a penalty term based on the second-order derivatives [2, 3]. Such a term penalizes high curvature in the displacement field, is invariant to linear transformations, and thus, favors deformations which are piecewise linear.

Recently, labeling of discrete Markov Random Fields has become an attractive approach for solving the problem of non-rigid image registration [4–6]. We will give a short introduction into the general framework in Section 2. Most of the methods share a similar model for the registration which is based on a pairwise MRF, i.e. is an MRF with cliques of size up to two. Then, the unary terms[5] (cliques of size one) play the role of the matching criteria, while pairwise terms are used to encode the regularization of the displacement field. In [6], the regularization is based on the norm of the displacement vector differences between neighboring control points which is an approximation of penalizing the gradients of the displacement field. The assumption is that neighboring nodes should follow a similar motion. A more robust measure is used in [4, 5] which allows more freedom on the deformation. However, this measure is still based on the gradient approximation. As remarked in [7], all these approaches penalize linear transformations such as rotation and scaling which in practice is often not desired. To this end, in [7] a regularization term based on the second-order derivatives is introduced by adding triple cliques of collinear neighboring control points to the MRF model. Each triple clique is in charge of penalizing the local curvature of the displacement field. The main problem of this approach are in fact the triple cliques, which require complex graph modifications in order to use efficient optimization techniques based on message-passing. In [7], the second-order MRF is converted to a pairwise one and then the TRW-S algorithm [8] is utilized to infer the MRF variables. Unfortunately, no running time is provided, but it is assumed to be much higher [9] than for the method proposed in [6] which uses the FastPD algorithm [10, 11] (based on iterative graph-cuts).

One may ask if it is possible to define a regularization term which has similar properties as the curvature penalty based on triple cliques while keeping the efficiency of a pairwise model.

In this paper, we investigate the use of an approximated curvature penalty term in a pairwise MRF. Our experiments demonstrate the practicability of such a regularization for non-rigid registration when the optimal transformation

[5] Please note, that in [4, 5] a decomposition of the unary terms is used but the general model is similar to [6].

contains a linear part. Compared to prior work, no higher-order cliques have to be employed for our curvature term and thus, our approach is efficient in terms of computational speed.

The remainder of the paper is organized as follows: the general framework for non-rigid registration using MRFs is described in Section 2. The proposed approximated curvature penalty is introduced in Section 3. Section 4 demonstrates the practicability of our regularization through a set of experiments, while Section 5 concludes our paper.

2 Non-rigid Registration using MRFs

Markov Random Field inference is a popular approach for parameter estimation. Given a set of parameters, one can define a graph $\mathcal{G} = (\mathcal{V}, \mathcal{C})$ consisting of a set of nodes $\mathcal{V}$ (one node per parameter) and a set of cliques $\mathcal{C}$ (where each clique is a subset of $\mathcal{V}$). Assuming that each node i takes a label l_i from a discrete set $\mathcal{L}$, the task becomes to find the optimal *labeling* $\mathbf{l}$ which minimizes

$$E_{\text{mrf}}(\mathbf{l}) = \sum_{c \in \mathcal{C}} \psi_c(\mathbf{l}_c) \quad , \tag{3}$$

which is a sum of *clique potentials* ψ_c determining the costs of certain label assignments and $\mathbf{l}_c$ is the vector of labels assigned to the parameter subset c.

The most common MRF model used in computer vision tasks (e.g. segmentation) is the first-order (pairwise) model containing at most cliques of size two. Many efficient algorithms have been proposed [12, 8, 10, 13] to solve the inference problem for this special case. For the first-order MRF the energy becomes the sum of unary and pairwise potentials

$$E_{\text{mrf}}(\mathbf{l}) = \sum_{i \in \mathcal{G}} \psi_i(l_i) + \sum_{i \in \mathcal{G}} \sum_{j \in \mathcal{N}_i} \psi_{ij}(l_i, l_j) \quad , \tag{4}$$

where $\mathcal{N}_i \subset \mathcal{G}$ defines the neighborhood system of the graph.

In case of non-rigid registration, the MRF variables correspond to locations in the image domain at which we want to estimate the motion (i.e. a displacement vector). Each discrete label is mapped to a displacement from a discretized version of the search space. For simplicity, we will denote the displacement vector associated with label l as $\mathbf{d}^l$. A simple approach would be to introduce an MRF variable for each pixel [14]. Then, the unary terms play the role of the data term or matching cost. Exemplary, we can define the costs for the sum of absolute differences (SAD) criteria based on image intensities as

$$\psi_i(l_i) = |(I(\mathbf{x}_i) - J(\mathbf{x}_i + D^{t-1}(\mathbf{x}_i) + \mathbf{d}^{l_i})| \quad , \tag{5}$$

where D^{t-1} is the dense field from the previous iteration and $\mathbf{d}^{l_i}$ the potential displacement corresponding to label l_i. The pairwise terms encode the regularization. A simple smoothness term penalizing high gradients can be defined as

$$\psi_{ij}(l_i, l_j) = \alpha \, \|(D^{t-1}(\mathbf{x}_i) + \mathbf{d}^{l_i}) - (D^{t-1}(\mathbf{x}_j) + \mathbf{d}^{l_j})\| \quad . \tag{6}$$

The optimization problem for dense registration in (1) is now completely defined as a *discrete labeling* of an MRF. The main problem for such an approach is the number of variables. One variable per pixel becomes computationally very expensive and is not feasible for large volumes in case of 3D registration. To this end, we can reduce the dimensionality of the problem by introducing a transformation model based on a sparse set of control points and an interpolation strategy. The dense displacement field in (2) is then defined as

$$D(\mathbf{x}) = \sum_{i}^{M} \eta_i(\mathbf{x})\, \mathbf{d}_i \quad , \tag{7}$$

where M is the number of control points and η_i is a weighting function (e.g. based on cubic B-splines) determining the contribution of the control point displacement $\mathbf{d}_i$ to the displacement of an image point $\mathbf{x}$. In this paper, we consider free form deformations (FFDs) [15] as the transformation model, where the control points are defined on a regular lattice and each control point has only local influence on the deformation. Let us now reformulate the matching cost (5) w.r.t. to the control points

$$\psi_i(l_i) = \sum_{\mathbf{x} \in \Omega_i} |I(\mathbf{x}) - J(\mathbf{x} + D^{t-1}(\mathbf{x}) + \mathbf{d}^{l_i})| \quad , \tag{8}$$

where Ω_i is a local image patch centered at the control point i. Intuitively, (8) can be understood as a block matching cost where the whole block Ω_i is potentially moved by $\mathbf{d}^{l_i}$. The size of the blocks is automatically defined by the distance between control points of the deformation grid. Additionally, [6] proposes a *weighted block matching* by incorporating the weighting functions η_i into the matching cost. The idea is that the influence of an image point to the matching criteria of a control point should be proportional to the contribution of that control point to the displacement of the image point. In other words, image points far away from a control point should have less influence on its cost than points in the immediate vicinity. Besides the reduction of the number of MRF variables, the block matching has additional advantages. For instance, it is straightforward to encode more sophisticated matching criteria such as correlation or mutual information which often provide more reliable matches than intensity differences. A comparison of different measures can be found in [6]. The regularization term (6) is similar as before, but now evaluated only on the control points instead of all image points

$$\psi_{ij}(l_i, l_j) = \alpha \,\|(\mathbf{d}_i^{t-1} + \mathbf{d}^{l_i}) - (\mathbf{d}_j^{t-1} + \mathbf{d}^{l_j})\| \quad , \tag{9}$$

where $\mathbf{d}_i^{t-1}$ is the displacement of control point i from the previous iteration. The final pairwise MRF energy for the non-rigid registration in (1) based on a deformation grid is then defined as

$$\begin{aligned} E_{\mathrm{mrf}}(\mathbf{l}) &= \sum_i^M \psi_i(l_i) + \sum_i^M \sum_{j\in\mathcal{N}_i} \psi_{ij} \\ &= \underbrace{\sum_i^M \sum_{\mathbf{x}\in\Omega_i} |I(\mathbf{x}) - J(\mathbf{x} + D^{t-1}(\mathbf{x}) + \mathbf{d}^{l_i})|}_{\approx S(I, J\circ T)} \\ &+ \underbrace{\sum_i^M \sum_{j\in\mathcal{N}_i} \alpha \,\|(\mathbf{d}_i^{t-1} + \mathbf{d}^{l_i}) - (\mathbf{d}_j^{t-1} + \mathbf{d}^{l_j})\|}_{\approx \alpha\, R(T)} \end{aligned} \tag{10}$$

3 Approximated Curvature Penalty

The main limitation of the registration framework based on pairwise MRFs are the constraints for regularization. The first-order cliques can only model interactions between two variables. The smoothness terms proposed so far, which all penalize high gradients on the displacement field, have the disadvantage of not being invariant to linear transformations such as rotation and scaling. Therefore, in [7] regularization is employed by introducing triple cliques which are able to encode a smoothness prior based on the discrete approximation of the second-order derivatives. The potential functions can be defined as

$$\psi_{ijk}(l_i, l_j, l_k) = c(\mathbf{d}_i^{t-1} + \mathbf{d}^{l_i}\,,\, \mathbf{d}_j^{t-1} + \mathbf{d}^{l_j}\,,\, \mathbf{d}_k^{t-1} + \mathbf{d}^{l_k}) \;\;, \tag{11}$$

$$c(\mathbf{a}, \mathbf{b}, \mathbf{c}) = \frac{1}{\delta^2} \sum_d^n (-a_d + 2b_d - c_d)^2 \;\;, \tag{12}$$

where c approximates the local curvature at location $\mathbf{b}$, b_d denotes the d-th component of the n-dimensional vector space, and δ is the control point distance. Such a smoothness term is invariant to linear transformations. The drawback of this approach is the complex handling of triple cliques. Graph modifications and insertion of auxiliary nodes are necessary in order to use efficient message-passing optimization techniques [16, 7]. The performance of message-passing algorithms in terms of computational speed is much lower than methods based on iterative graph-cuts [9, 11].

Therefore, we propose a regularization term based on second-order derivatives which works on pairwise potential functions and which we call approximated curvature penalty (ACP). Since in pairwise terms only the potential label assignment of two variables is known, we approximate the local curvature by assuming the other variables to stay fixed. In detail, for two neighboring variables i and j, we compute the approximated curvature at both locations and average them. To this end, we define different pairwise potentials depending on the axis on which the two variables are neighboring. In 2D, we have a set of potential functions for the horizontal and vertical axis

$$\begin{aligned}\psi^H_{i,j}(l_i,l_j) &= \frac{1}{2}\,\Big(c(\mathbf{d}^{t-1}_{i-1}\,,\,\mathbf{d}^{t-1}_{i}+\mathbf{d}^{l_i}\,,\,\mathbf{d}^{t-1}_{j}+\mathbf{d}^{l_j})\\ &\quad +c(\mathbf{d}^{t-1}_{i}+\mathbf{d}^{l_i}\,,\,\mathbf{d}^{t-1}_{j}+\mathbf{d}^{l_j}\,,\,\mathbf{d}^{t-1}_{j+1})\Big) \quad ,\\ \psi^V_{i,j}(l_i,l_j) &= \frac{1}{2}\,\Big(c(\mathbf{d}^{t-1}_{i-M_x}\,,\,\mathbf{d}^{t-1}_{i}+\mathbf{d}^{l_i}\,,\,\mathbf{d}^{t-1}_{j}+\mathbf{d}^{l_j})\\ &\quad +c(\mathbf{d}^{t-1}_{i}+\mathbf{d}^{l_i}\,,\,\mathbf{d}^{t-1}_{j}+\mathbf{d}^{l_j}\,,\,\mathbf{d}^{t-1}_{j+M_x})\Big) \quad ,\end{aligned} \tag{13}$$

where M_x is the number of control points on the deformation grid in horizontal direction. The definition of an additional term for the third axis in 3D registration is straightforward. In each evaluation of the ACP we determine the average of the two local curvatures by considering the displacements of four variables: i and j which are the variables with potential movement $\mathbf{d}^{l_i}$ and $\mathbf{d}^{l_j}$ and the two surrounding variables with the displacements from the previous iteration.

Considering the properties of such a smoothness prior, we claim that it allows much more flexibility on the deformation w.r.t. to linear transformations compared to other terms based on pairwise potentials. In the beginning of every registration process, the matching criteria is usually the driving force towards the correct alignment, while the regularization increases its importance on the global energy in later iterations. The incremental deformation at the end of the process is getting smaller and the ACP will favor deformations which are piecewise linear instead of piecewise constant as for the gradient penalty. The practicability of our proposed regularization is demonstrated in the following experiments.

4 Experiments

We perform several experiments which hopefully illustrate the advantages of the proposed ACP as regularization. Throughout the tests we use the FastPD algorithm for MRF inference. In the first two experiments, we generate synthetic target images by warping a source image with different linear transformations. The first one is a 60° rotation and the second one an anisotropic scaling (see Fig. 1(a) and 1(e)). For each target, four point correspondences are distributed around the image center which exactly define these linear transformations. Please note, that in these two experiments our aim is to investigate the properties of the regularization only. Therefore, we choose a geometric matching term based on perfect point correspondences. Thus, the data term is reliable and will guide the registration towards an optimal alignment of the correspondences which allows us to study the behavior of the regularization. We use a simple Euclidean distance measure as the matching criteria w.r.t the correspondences. The unary potentials are therefore defined as

$$\psi_i(l_i) = \sum_{k}^{K} \|\mathbf{p}_k - (\mathbf{q}_k + D^{t-1}(\mathbf{q}_k) + \mathbf{d}^{l_i})\| \quad , \tag{14}$$

where $\mathbf{p}_k$ and $\mathbf{q}_k$ are the corresponding points in the target and source image, respectively. Now we compare the behavior of three different regularization terms, namely the absolute vector difference (cp. (9)), the quadratic vector difference (i.e. the squared version of (9)), and our ACP defined in (13). For all registrations, we use a 7×7 deformation grid. The results for the registration with the different regularization terms are shown in Fig. 1(b-d) and (f-h). We should note, that in all cases the final mean distance for the correspondences is less than one pixel, which indicates a very good minimization of the matching costs. However, the ACP is the only method which is able to correctly regularize the deformation field towards the linear transformations.

The second part of the experiments is investigating the performance of the ACP in intensity-based registration. In fact, large rotations such as $60°$ are very unlikely to be present if a proper pre-alignment via rigid registration has been performed prior to the non-rigid one. Additionally, a block matching strategy which is mainly based on a translational search most likely fails to recover large rotations or scaling. However, in practice it is likely that a certain amount of linear transformation is still present when starting the non-rigid alignment [17]. Again, we generate two target images from a source image, both with a $25°$ rotation (cp. Fig. 1(i)). To one of the images we also add random deformation using a thin-plate spline warping [2] (cp. Fig. 1(m)). The registrations are then performed purely based on intensities using the matching criteria defined in (8). The results for the different regularization terms are shown in Fig. 1(j-l) and (n-p). Again, the ACP outperforms the gradient penalty terms in the ability of regularization towards linear transformations. In the last case of random deformation combined with rotation, the resulting transformation using ACP is very close to the ground truth. This is remarkable since only intensities are used in the matching criteria and outer control points obtain their positions solely by regularization. Additionally, when we visually inspect the warped images after registration the results for the ACP are almost perfect, while the gradient penalty terms prohibit a proper alignment due to the increasing costs for linear transformations. This is consistent with the observations in [7].

5 Conclusion

We propose a novel regularization term based on an approximated curvature penalty for pairwise MRFs. Our results demonstrate the superior performance of this approach compared to previous smoothness terms based on gradient penalties. Our regularization can successfully recover linear transformations and thus, has similar properties as a curvature penalty using triple cliques, while the computational efficiency of a pairwise MRFs is preserved. In fact, the running time using ACP increases only very little compared to the gradient penalty terms. All shown registrations are performed within a few seconds. We believe that the proposed regularization is an important extension to the MRF registration framework. Furthermore, we could show that introducing approximated terms in pairwise MRFs can lead to very promising results. Future work includes the comparison to recent advances in optimization of higher-order MRFs [18].

References

1. Tikhonov, A.: Ill-posed problems in natural sciences (1992) Coronet.
2. Bookstein, F.: Principal warps: thin-plate splines and the decomposition of deformations. In: Pattern Analysis and Machine Intelligence. (1989)
3. Modersitzki, J.: Numerical Methods for Image Registration. Oxford University Press (2004)
4. Lee, K.J., Kwon, D., Yun, I.D., Lee, S.U.: Deformable 3d volume registration using efficient mrfs model with decomposed nodes. In: British Machine Vision Conference. (2008)
5. Shekhovtsov, A., Kovtun, I., Hlavac, V.: Efficient mrf deformation model for non-rigid image matching. Comput. Vis. Image Underst. **112** (2008) 91–99
6. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through mrfs and efficient linear programming. Medical Image Analysis **12** (2008) 731–741
7. Kwon, D., Lee, K.J., Yun, I.D., Lee, S.U.: Nonrigid image registration using dynamic higher-order mrf model. In: European Conference on Computer Vision. (2008)
8. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. Pattern Analysis and Machine Intelligence **28** (2006) 1568–1583
9. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C.: A comparative study of energy minimization methods for markov random fields with smoothness-based priors. Pattern Analysis and Machine Intelligence **30** (2008) 1068–1080
10. Komodakis, N., Tziritas, G., Paragios, N.: Fast, approximately optimal solutions for single and dynamic mrfs. In: Computer Vision and Pattern Recognition. (2007)
11. Komodakis, N., Tziritas, G., Paragios, N.: Performance vs computational efficiency for optimizing single and dynamic mrfs: Setting the state of the art with primal-dual strategies. Comput. Vis. Image Underst. **112** (2008) 14–29
12. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. Pattern Analysis and Machine Intelligence **23** (2001) 1222–1239
13. Lempitsky, V., Rother, C., Blake, A.: Logcut - efficient graph cut optimization for markov random fields. In: International Conference on Computer Vision. (2007)
14. Boykov, Y., Veksler, O.: Graph Cuts in Vision and Graphics: Theories and Applications. In: Handbook of Mathematical Models in Computer Vision. Springer (2005) 79–96
15. Sederberg, T.W., Parry, S.R.: Free-form deformation of solid geometric models. In: SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques, New York, NY, USA, ACM Press (1986)
16. Potetz, B., Lee, T.S.: Efficient belief propagation for higher-order cliques using linear constraint nodes. Comput. Vis. Image Underst. **112** (2008) 39–54
17. Zikic, D., Hansen, M.S., Glocker, B., Khamene, A., Larsen, R., Navab, N.: Computing minimal deformations: Application to construction of statistical shape models. In: Computer Vision and Pattern Recognition. (2008)
18. Komodakis, N., Paragios, N.: Beyond pairwise energies: Efficient optimization for higher-order mrfs. In: Computer Vision and Pattern Recognition. (2009)

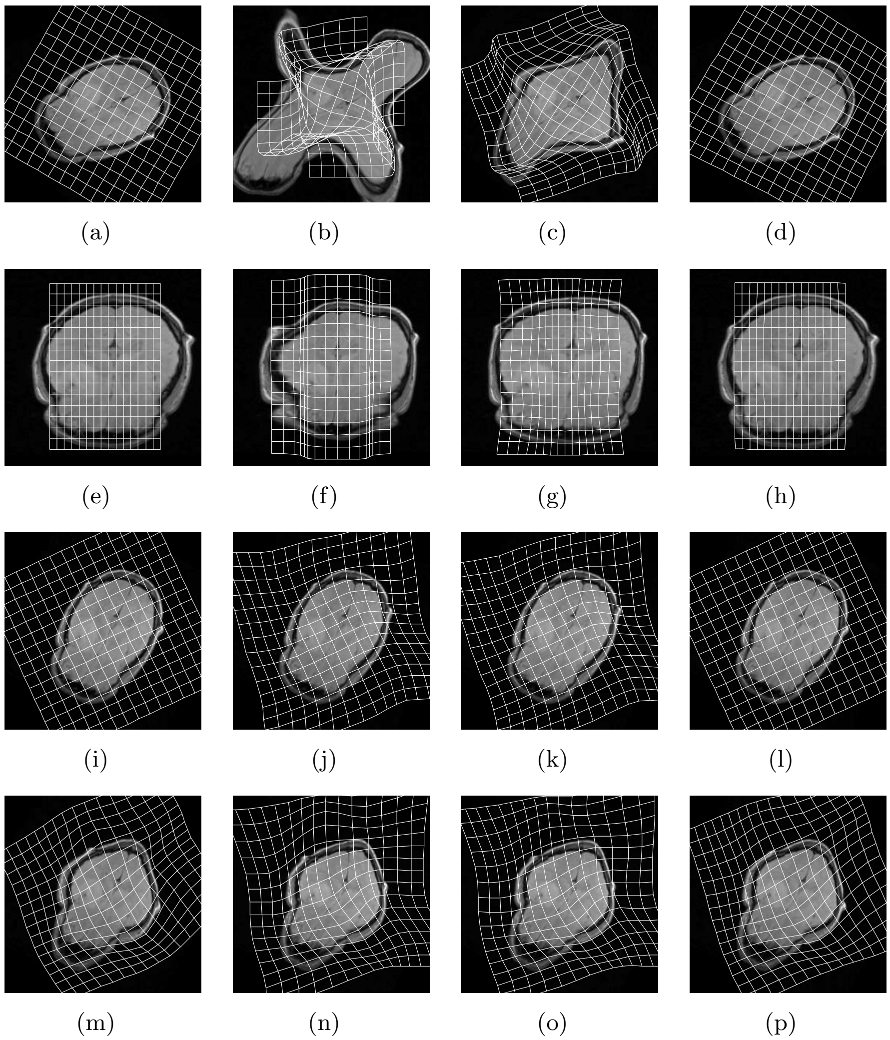

Fig. 1. First two rows: Target images on the left generated from the source image by a 60° rotation (a) and anisotropic scaling (e). The registration results are shown for the absolute difference (b,f), quadratic difference (c,g), and the approximated curvature penalty (d,h). The Euclidean distance on four point correspondences is used to define the registration data term. **Last two rows:** Target images on the left generated by a 25° rotation (i). For the last row, additional random deformation is added (m). From left to right the registration results (j-l,n-p) in the same order as shown in the upper rows. This time, no information about point correspondences is used and the registration is purely based on image intensities. Please note, that we use backward warping why the actual transformations visualized as grids appear to point in the opposite direction as the warped images.