

# Getting the most out of additional guidance information in deformable image registration by leveraging multi-objective optimization

Tanja Alderliesten<sup>a</sup>, Peter A. N. Bosman<sup>b</sup>, Arjan Bel<sup>a</sup>

<sup>a</sup>Dept. of Radiation Oncology, Academic Medical Center (AMC), P.O. Box 22660, 1100 DD Amsterdam, The Netherlands;

<sup>b</sup>Centrum Wiskunde & Informatica (CWI), P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

## ABSTRACT

Incorporating additional guidance information, e.g., landmark/contour correspondence, in deformable image registration is often desirable and is typically done by adding constraints or cost terms to the optimization function. Commonly, deciding between a “hard” constraint and a “soft” additional cost term as well as the weighting of cost terms in the optimization function is done on a trial-and-error basis. The aim of this study is to investigate the advantages of exploiting guidance information by taking a multi-objective optimization perspective. Hereto, next to objectives related to match quality and amount of deformation, we define a third objective related to guidance information. Multi-objective optimization eliminates the need to a-priori tune a weighting of objectives in a single optimization function or the strict requirement of fulfilling hard guidance constraints. Instead, Pareto-efficient trade-offs between all objectives are found, effectively making the introduction of guidance information straightforward, independent of its type or scale. Further, since complete Pareto fronts also contain less interesting parts (i.e., solutions with near-zero deformation effort), we study how adaptive steering mechanisms can be incorporated to automatically focus more on solutions of interest. We performed experiments on artificial and real clinical data with large differences, including disappearing structures. Results show the substantial benefit of using additional guidance information. Moreover, compared to the 2-objective case, additional computational cost is negligible. Finally, with the same computational budget, use of the adaptive steering mechanism provides superior solutions in the area of interest.

**Keywords:** Deformable image registration, guidance information, large anatomical differences, multi-objective optimization, evolutionary algorithms

## 1. INTRODUCTION

Many clinical applications could benefit from the use of deformable image registration methods. However, one solution to all deformable image registration problems does not exist. Therefore, fine-tuning methods to each clinical application is of great importance. This fine-tuning task, however, poses major challenges because currently existing methods rely on many user-defined parameter settings and it is non-trivial how to find appropriate settings for a specific clinical application. This holds especially for the weighting of different registration goals in deformable image registration (i.e., quality of image match and amount of deformation). Often the only person who is capable of performing this task is the researcher who developed and studied the method and this task is then performed by trial and error. To overcome these challenges, we started studying deformable image registration from a multi-objective optimization perspective.<sup>1-2</sup> Multi-objective optimization removes the need to set a predetermined singular combination of objectives. In multi-objective optimization, a collection of outcomes (a so-called Pareto front) that represents efficient trade-offs between the different components of the optimization function is computed, which allows for more insightful a-posteriori tuning.

Further, in our multi-objective approach we introduced a dual-dynamic transformation model (i.e., two “non-fixed” grids: one for the source- and one for the target image). This provides an elegant and powerful approach to, when combined with proper optimization techniques, tackle the most challenging deformable image registration problems, including e.g., large deformations and anatomical changes.<sup>3,4</sup>

---

Further author information: (Send correspondence to T.A. or P.A.N.B.)  
T.A.: E-mail: t.alderliesten@amc.uva.nl, Telephone: +31 (0)20 566 6886  
P.A.N.B.: E-mail: Peter.Bosman@cwi.nl, Telephone: +31 (0)20 592 4265




Figure 1. Illustration of the refinement method used in the multi-resolution scheme. As an example a 3×3 grid of points (left) is refined into a 5×5 grid of points (right). The light grey nodes and dotted edges illustrate the parts that are added.

In contrast with other work<sup>5-10</sup>, in our work the challenging problem of identifying (dis)appearing structures between image acquisitions has been considered to be part of the overall optimization process. Although this removes the need to manually fully identify (dis)appearing structures, the images themselves do not always provide sufficient information for the problem to be solved without additional guidance information. In deformable image registration, it is often desirable to incorporate additional guidance information that is known beforehand.<sup>11</sup> This is often done by adding additional constraints or cost terms to the optimization function. In common deformable image registration methods the choice between a “hard” constraint and a “soft” cost term as well as the weighting of these terms in relation to other terms in the optimization function is also typically done on the basis of trial and error.

In this paper, we study getting the most out of additional guidance information by leveraging multi-objective optimization. Further, we study the design and use of adaptive steering mechanisms to focus optimization more on areas of interest on the Pareto front.

## 2. MATERIALS AND METHODS

### 2.1 Multi-objective deformable image registration framework

Image registration can be posed as an optimization problem where different objectives (typically related to the match quality and smoothness/amount of deformation) are of interest. Commonly, registration methods optimize a single optimization function that represents a linearly weighted combination of the objectives, and sometimes also other, constraint-related, cost terms. Deformable image registration from a multi-objective optimization perspective removes the need to a-priori set a singular combination of objectives (and possibly other terms) through trial and error. The use of our framework results not in a single solution but in a Pareto front of solutions, which is a collection of solutions that are efficient trade-offs between the objectives (i.e., solutions that are all equally good (better in one objective, but worse in another)). Having a Pareto front of trade-off solutions available provides far more insight into the interplay between objectives. Moreover, a multi-objective approach is inherently more powerful than a single-objective approach because potentially not all Pareto-optimal outcomes are explored when running existing single-objective registration methods multiple times with different linear combinations of the weights for the objectives (i.e., if the Pareto front is concave).

#### 2.1.1 Objectives related to match quality and amount of deformation

One objective is related to the quality of the image match. Here we use the well-known cross-correlation similarity metric between the target image and the transformed source image.<sup>12</sup> However, since all objectives in our implementation are to be minimized, we use the negation of this value. Hence, this objective lies in the range [-1;1].

Another objective is related to the amount of the deformation (i.e., the amount of energy required to accomplish the transformation). Here we employ Hooke’s law<sup>13</sup> on the basis of the deformation vector field ( $dvf$ ). For every pixel we use the average length of the difference vector that is obtained by comparing the deformation of that pixel and each of its neighboring pixels. The final objective value then is the sum, over all pixels, of the squared average lengths. Let the set of considered pixels be denoted by  $P$ , the set of neighbors of pixel  $p$  be denoted by  $n(p)$ , and the deformation vector in the  $dvf$  associated with pixel  $p$  be denoted by  $dvf_p$ . The objective can then be defined as:

$$F_{deform} = \sum_{p \in P} \left[ \frac{1}{|n(p)|} \sum_{np \in n(p)} \|dvf_p - dvf_{np}\| \right]^2$$

### 2.1.2 Dual-dynamic transformation model

The transformation model we use, is based on triangulated grids. It is common practice to use a fixed grid for the source image and a non-fixed grid for the target image. We previously introduced a dual-dynamic transformation model that uses two non-fixed grids.<sup>3</sup> This increases transformation flexibility substantially in a manner that is especially required when tackling some of the hardest problems with large deformations and/or (dis)appearing structures. We compute the final deformation by considering corresponding pairs of triangles (i.e., in the source and target grids) and employing linear simplex interpolation. A transformation is defined by associating coordinates with each point in *two* grids. Therefore, the number of real-valued parameters to be optimized equals two times the number of grid points times the spatial dimensionality of the image (e.g., 2D or 3D).

### 2.1.3 Multi-resolution strategy

To establish a multi-resolution strategy<sup>14,15</sup>, a method is required to refine the two grids associated with a solution so as to be able to continue computation at a finer grid resolution. When using a multi-objective optimization algorithm, at the end of a single-resolution run, multiple solutions (i.e., a Pareto front) are returned. A refinement method is therefore required that refines all triangulations similarly (i.e., without completely redefining for each solution separately and differently the underlying graphs that connect the grid points). Our refinement method consists of subdividing each triangle into 4 triangles by adding points on each edge of the grid (see Fig.1).<sup>4</sup>

### 2.1.4 Optimization algorithm

Population-based methods such as evolutionary algorithms (EAs) are among the state-of-the-art in solving multi-objective optimization problems.<sup>16</sup> To perform multi-objective optimization, we use a model-based EA. More specifically, we use an Estimation-of-Distribution Algorithm (EDA) that exploits features of a problem’s structure in a principled manner via probabilistic modeling.<sup>17-20</sup> The particular EDA that we employ is iMAMaLGaM (incremental Multi-objective Adapted Maximum Likelihood Gaussian Model mixture) in which the underlying probabilistic model is a Gaussian mixture distribution.<sup>2</sup> In related work, iMAMaLGaM was shown to have excellent performance, converging to high-quality approximations of the optimal Pareto front on well-known benchmark problems.<sup>2</sup>

## 2.2 Additional guidance information

The major benefits associated with incorporating guidance information by using a third objective in a multi-objective approach are that we a-priori 1) do not have to decide if we want to incorporate the additional guidance information in a way that allows deviations (of a certain magnitude) or not (i.e., as a “hard” constraint or a “soft” cost term), and 2) do not have to define weighting factors. Moreover, taking a multi-objective approach also provides the flexibility to include *imperfect* guidance information, because solutions with little deviation from additional guidance information as well as solutions with more deviation will be computed. This significantly reduces, if not completely removes, the risk of over-fitting guidance information.

In our implementation, guidance information is stored as a set of pairs. Each pair consists of two sets, each containing one or more points associated with the source image and the target image, respectively. In this way it is possible to incorporate landmarks as well as contours. The guidance objective captures landmark/contour correspondence in terms of distances between pairs of guidance-point sets. For each pair of guidance-point sets, the points defined in the source image are transformed to the target image and vice versa. For each point, after transformation, the smallest distance is computed to the points that define the contour in the other image. The distance between two guidance-point sets is then computed in a symmetric fashion as the sum of these distances over all points in the guidance sets, normalized by dividing by the total number of points. The sum over all pairs of guidance-point sets is then to be minimized. In other words, let  $G$  be the set of pairs  $(G_s, G_t)$  of guidance-point sets, let  $T$  be the transformation function as defined by the solution and let  $d$  denote Euclidean distance. Then, the guidance objective can be formulated as follows:

$$F_{\text{guidance}} = \sum_{(G_s, G_t) \in G} \frac{1}{|G_s| + |G_t|} \left[ \sum_{g_s \in G_s} \min_{g_t \in G_t} \{d(T(g_s), g_t)\} + \sum_{g_t \in G_t} \min_{g_s \in G_s} \{d(T^{-1}(g_t), g_s)\} \right]$$

## 2.3 Adaptive zoom-in on Pareto front

Although it is theoretically interesting to see the size and shape of complete Pareto fronts, these fronts clearly also contain less interesting parts from a practical point of view (i.e., solutions with near-zero deformation). Because the algorithm allocates resources to find improvements all along the Pareto front, it may be beneficial to steer the algorithm

so as to focus more on areas of (practical) interest. To this end, we introduce an adaptive steering mechanism that shrinks the space of feasible solutions as better solutions are found over time. Because the optimization algorithm that we use includes a general-purpose constraint-handling technique based on constraint domination<sup>16</sup>, optimization then automatically proceeds to focus only on the feasible space. Specifically, for every objective separately, we define, on the basis of the best value for that objective encountered so far during the run, a range in which solutions are deemed to be feasible. All solutions for which one or more objective values do not comply with the defined ranges or thresholds are deemed to be infeasible. We used such a range for the similarity objective and for the guidance objective. Specifically, for the similarity objective, solutions that are in the range of  $[s^{best}; 0.99 \times s^{best}]$  are deemed to be feasible whereas for the guidance objective this range is taken to be  $[g^{best}; 1.5 \times g^{best}]$ . In both cases these ranges were experimentally determined to give reasonable results. The objective related to the amount of deformation was not further constrained.

## 2.4 Experiments

We consider four test cases: two deformable image registration problems with a disappearing structure and two problems with large deformations (see Fig.2). For both problem types we consider an artificial 2D example as well as a real-life example. For the real-life problem with a disappearing structure a pre-operatively and post-operatively acquired CT scan from a breast cancer patient were used. A pair of 2D slices was selected from the CT scans after rigid registration of the scans on the bony anatomy. For the real-life problem with large deformations a T2-weighted pre-operative MRI scan from a cervical cancer patient who underwent a radical hysterectomy and a 2D digital photograph of the macroscopic intersection obtained during standard pathological analysis was used. From the MRI scan a 2D sagittal slice that corresponded to the plane of macroscopic intersection was selected and the uterine cervix was segmented. Guidance information for the real-life problems was annotated by clinical experts.

A multi-resolution strategy with subsequent grid resolutions of  $5 \times 5$ ,  $9 \times 9$ , and  $17 \times 17$  was used, resulting in respectively 100, 324, 1156 parameters to be optimized for which a budget of  $1580 \cdot 10^4$  evaluations (respectively  $100 \cdot 10^4$ ,  $324 \cdot 10^4$ , and  $1156 \cdot 10^4$  evaluations in a single run on the different resolutions) was allowed.

## 3. RESULTS

Desirable behavior can be observed in Fig.2: 3D Pareto fronts are obtained that intuitively display the trade-off between the various objectives and thereby also the interplay with additional guidance information, allowing for insightful a-posteriori selection of the final desired solution. Fig.2 further shows that the use of the adaptive steering mechanism indeed contributes to finding, using the same computational budget, superior results in the area where more interesting solutions are located.

Previous results already showed that disappearing structures can be successfully handled by our framework without the use of additional guidance information. Nevertheless, because images themselves do not always sufficiently convey information about how the anatomy has changed, in case of problem 2, this primarily resulted in very local removal of the tumor area. With the possibility to include, as a third objective, an approximate indication of the area that has been removed, registration results can now be obtained that more accurately correspond to the actual changes in anatomy.

Without the use of additional guidance information it was not possible to solve problems 3 and 4 satisfactorily. In particular, for problem 3 pushing the grey bar down will not result in an improved similarity value until it has reached the target location. Therefore, proper registration cannot be achieved without the use of additional guidance information. With the use of even very little guidance information (merely a few points in the grey bar), our framework is successfully steered toward desirable solutions. We note that on problem 3, due to the minimum amount of detailed structure in the image, we used only a  $5 \times 5$  grid resolution directly instead of the multi-resolution scheme since a higher resolution grid is really unnecessary. The challenges concerning problem 4 are the strongly conflicting objectives: the best guidance match is really not the best image match based on intensity similarity. However, with the use of additional guidance information in our multi-objective framework a good solution can be intuitively selected a-posteriori by analyzing the final solutions obtained along the Pareto front. This can be done similarly for all other studied problems, without having to re-tune any algorithmic parameters whatsoever.

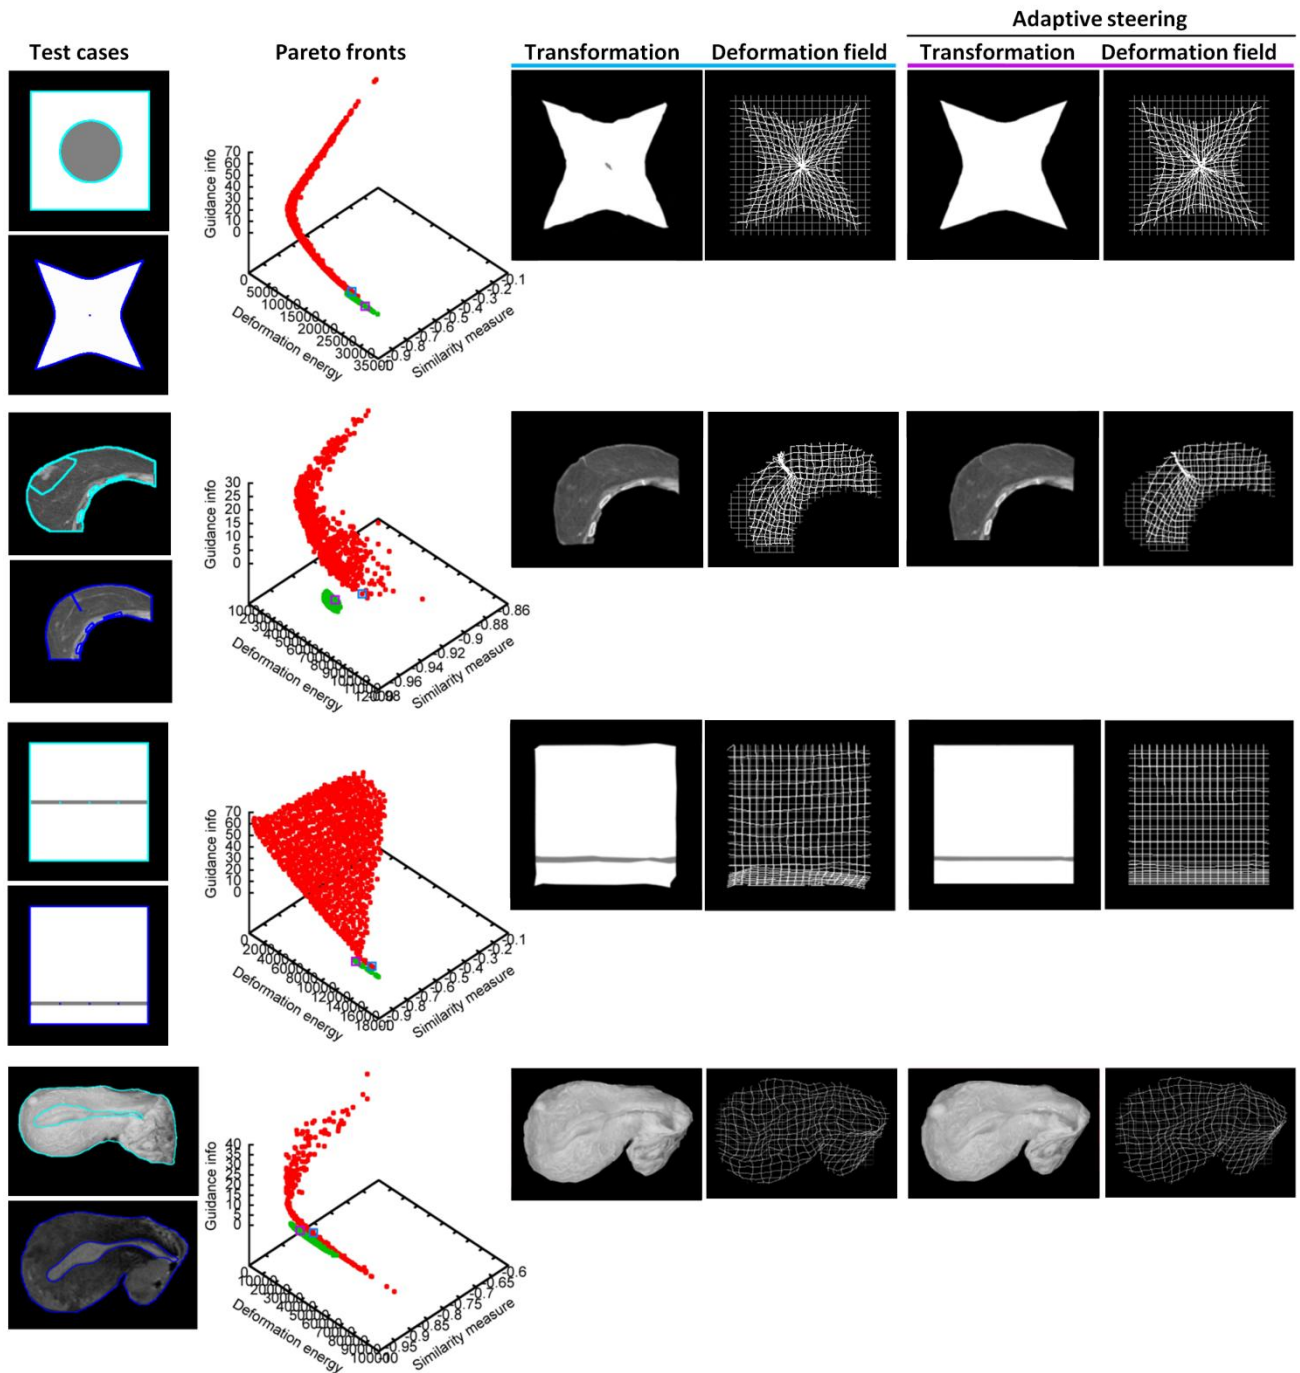


Figure 2. From top to bottom: the four test cases 1) artificial example with a disappearing structure, 2) breast CT data pre- and post-surgical removal of a tumor and some of the surrounding breast tissue (images with courtesy of The Netherlands Cancer Institute – Antoni van Leeuwenhoek Hospital), 3) artificial example with large deformation, 4) pre-surgical cervical MRI data (bottom) with corresponding macroscopic digital photo of the excised specimen (top). From left to right: source (top) and target (bottom) images with guidance information in color (contours and points), Pareto fronts (with selected solutions indicated) without (red) and with (green) adaptive steering, transformations for solutions without and with the use of adaptive steering. Further, for each transformation the deformation is illustrated via visualization of a correspondingly deformed regular rectangle grid. Note: these images do not include the actual dual-dynamic triangular grids that define the deformation field.

## 4. DISCUSSION AND CONCLUSIONS

Deformable image registration plays an important role in many clinical applications. Widespread translation of deformable image registration methods into the clinic is however still very challenging. An important underlying problem of the vast majority of existing methods is that they all rely on many user-defined parameter settings and the non-trivial task of finding appropriate settings for a specific clinical application can usually only be done by trial and error. The reason for this is that all these methods generally define the deformable image registration task as a single-objective optimization problem, i.e., one function consisting of a single linear combination of different objectives to be optimized. We therefore started investigating deformable image registration from a multi-objective optimization perspective.<sup>1-4</sup> This allows for improved understanding of interaction between different objectives, which in turn allows improved tuning of methods to specific clinical applications. Within our framework we further recently introduced the concept of using a dual-dynamic transformation model that uses two non-fixed grids and thereby increases the flexibility of transformations substantially and in a manner that is especially required when tackling very large deformations and (dis)appearing structures.<sup>3,4</sup>

In this work, we developed and studied, within our multi-objective framework, getting the most out of guidance information in deformable image registration. This is a novel contribution since this has not been studied before and because our multi-objective optimization approach eliminates the need to a-priori 1) decide if we want to incorporate the additional guidance information in a way that allows deviations or not (i.e., as a “hard” constraint or a “soft” cost term), and 2) define and tune weighting factors. Moreover, we previously illustrated that this framework is an elegant and powerful approach that shows very promising results regarding the tackling of different hard registration problems (e.g., large deformations, (dis)appearing structures).<sup>3,4</sup> In this study, we successfully illustrated that with the proposed adaptations to our framework it is capable of producing substantially improved results in a very intuitive and easy-to-use manner. Specifically, we successfully introduced the use of additional guidance information for our multi-objective deformable image registration framework by introducing a third objective. This gives unprecedented insight into the trade-off between exploiting this information and other objectives. Moreover, by treating guidance information as a separate objective, this information no longer needs to be exact for it to still be very useful in getting better registration results, without any fine-tuning of the registration algorithm itself. We furthermore showed that by adaptively steering multi-objective optimization toward a specific region of interest on the Pareto front, superior registration results could be obtained using the same computational budget.

## ACKNOWLEDGEMENTS

The authors thank M.T.F.D. Vrancken Peeters (Department of Surgical Oncology, The Netherlands Cancer Institute – Antoni van Leeuwenhoek Hospital, Amsterdam, The Netherlands), A.J.A.J. van de Schoot and P. de Boer (Department of Radiation Oncology, Academic Medical Center, Amsterdam, The Netherlands) for their contributions to this study. Financial support of this work was provided by the Dutch Cancer Society (Grant No. KWF 2012-5716).

## REFERENCES

- [1] Alderliesten, T., Sonke, J.-J. and Bosman, P. A. N., "Multi-objective optimization for deformable image registration: proof of concept," Proc. SPIE 8314, 831420, doi:10.1117/12.911268 (2012).
- [2] Bosman, P. A. N. and Alderliesten, T., "Incremental Gaussian model-building in multi-objective EDAs with an application to deformable image registration," Proc. GECCO, 241-248, doi:10.1145/2330163.2330199 (2012).
- [3] Alderliesten, T., Sonke, J.-J. and Bosman, P. A. N., "Deformable image registration by multi-objective optimization using a dual-dynamic transformation model to account for large anatomical differences," Proc. SPIE 8669, 866910, doi:10.1117/12.2006783 (2013).
- [4] Alderliesten, T., Bosman, P. A. N., Sonke, J.-J. and Bel, A., "A multi-resolution strategy for a multi-objective deformable image registration framework that accommodates large anatomical differences, " Proc. SPIE 9034, 90343G, doi:10.1117/12.2042856 (2014).
- [5] Nithianathan, S., Mirota, D., Uneri, A. and Schafer, S., "Incorporating tissue excision in deformable image registration: a modified demons algorithm for cone-beam CT-guided surgery," Proc. SPIE 7964, 796404, doi:10.1117/12.878258 (2001).

- [6] Ferrant, M., Nabavi, A., Macq, B., Black, P. M., Jolesz, F. A., Kikinis, R. and Warfield, S. K., "Serial registration of intraoperative MR images of the brain," *Med Image Anal* 6(4), 337-359 (2002).
- [7] Miga, M. I., Roberts, D. W., Kennedy, F. E., Platenik, L. A., Hartov, A., Lunn, K. E. and Paulsen, K. D., "Modeling of retraction and resection for intraoperative updating of images," *Neurosurgery* 49(1), 75-85 (2001).
- [8] Periaswamy, S. and Farid, H., "Medical image registration with partial data," *Med Image Anal* 10(3), 452-464 (2006).
- [9] Risholm, P., Samset, E. and Wells-III, W., "Validation of a nonrigid registration framework that accommodates tissue resection," *Proc. SPIE* 7623, 762319, doi:10.1117/12.844302 (2010).
- [10] Nithianathan, S., Schafer, S., Stayman, J.W., Zbijewski, W., Reh, D.D., Gallia, G.L. and Siewerdsen, J.H., "Extra-dimensional demons: a method for incorporating missing tissue in deformable image registration," *Med Phys* 39(9), 5718-5731 (2012).
- [11] Fischer, B. and Modersitzki, J., "FLIRT: A flexible image registration toolbox," *Proc. WBIR, LNCS* 2717, 261-270 (2003).
- [12] Roshni V. S. and Revathy K., "Using mutual information and cross correlation as metrics for registration of images," *JATIT* 4(6), 474-481 (2008).
- [13] Arfken, G., [Mathematical methods for physicists], Academic Press, Inc., San Diego (1985).
- [14] Bajcsy, R. and Kovačič S., "Multiresolution elastic matching," *Comput Gr Image Process* 46(1), 1-21 (1989).
- [15] Lester, H. and Arridge, S. R., "A survey of hierarchical non-linear medical image registration," *Pattern Recognition* 32, 129-149 (1999).
- [16] Deb, K., [Multi-objective optimization using evolutionary algorithms], John Wiley & Sons Inc., New York (2001).
- [17] Lozano, J. A., Larrañaga, P., Inza, I. and Bengoetxea, E., [Towards a new evolutionary computation. Advances in estimation of distribution algorithms], Springer-Verlag, Berlin (2006).
- [18] Pelikan, M., Sastry, K. and Cantú-Paz, E., [Scalable optimization via probabilistic modeling: from algorithms to applications], Springer-Verlag, Berlin (2006).
- [19] Bosman, P.A.N. and Grahl, J., "Matching inductive search bias and problem structure in continuous estimation-of-distribution algorithms," *EJOR* 185, 1246-1264 (2008).
- [20] Mühlenbein, H. and Höns, R., "The estimation of distributions and the minimum relative entropy principle," *Evolutionary Computation* 13, 1-27 (2005).