

A novel model-based evolutionary algorithm for multi-objective deformable image registration with content mismatch and large deformations: benchmarking efficiency and quality

Anton Bouter ^{*a}, Tanja Alderliesten ^b, Peter A. N. Bosman ^a

^aCentrum Wiskunde & Informatica (CWI), P.O. Box 94079, 1090 GB Amsterdam, The Netherlands;

^bDept. of Radiation Oncology, Academic Medical Center (AMC), P.O. Box 22660, 1100 DD Amsterdam, The Netherlands

ABSTRACT

Taking a multi-objective optimization approach to deformable image registration has recently gained attention, because such an approach removes the requirement of manually tuning the weights of all the involved objectives. Especially for problems that require large complex deformations, this is a non-trivial task. From the resulting Pareto set of solutions one can then much more insightfully select a registration outcome that is most suitable for the problem at hand. To serve as an internal optimization engine, currently used multi-objective algorithms are competent, but rather inefficient. In this paper we largely improve upon this by introducing a multi-objective real-valued adaptation of the recently introduced Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) for discrete optimization. In this work, GOMEA is tailored specifically to the problem of deformable image registration to obtain substantially improved efficiency. This improvement is achieved by exploiting a key strength of GOMEA: iteratively improving small parts of solutions, allowing to faster exploit the impact of such updates on the objectives at hand through partial evaluations. We performed experiments on three registration problems. In particular, an artificial problem containing a disappearing structure, a pair of pre- and post-operative breast CT scans, and a pair of breast MRI scans acquired in prone and supine position were considered. Results show that compared to the previously used evolutionary algorithm, GOMEA obtains a speed-up of up to a factor of ~1600 on the tested registration problems while achieving registration outcomes of similar quality.

Keywords: Deformable image registration, multi-objective optimization, evolutionary algorithms, partial evaluations, content mismatch, large anatomical differences

1. INTRODUCTION

Many state-of-the-art algorithms, such as the `elastix` toolbox¹, designed for the Deformable Image Registration (DIR) problem are very efficient and capable of producing good results for certain applications, but only produce a single registration outcome. The reason for this is that although DIR encompasses multiple objectives, e.g., similarity and deformation smoothness objectives, these objectives are often condensed into a single optimization function through a weighted sum of the objectives. Setting these weights appropriately can be very problem specific. Moreover, manually fine-tuning them is difficult and time-consuming. Using a multi-objective approach removes the requirement of setting weights manually via trial and error a priori, because such an approach results in a set of non-dominated solutions, i.e., a set where no single solution is better in every objective than any other solution in this set. Each solution in this so-called Pareto set represents a different trade-off between the objectives of interest. The solution that is deemed the most appropriate for the problem at hand can manually be selected from this Pareto set a posteriori. This set of solutions is highly intuitive to navigate, contrary to tuning weights a priori as is current practice. Previously², a multi-objective Evolutionary Algorithm (EA) was used within a recently proposed multi-objective framework for DIR because EAs are among the state-of-the-art for multi-objective optimization³. Although recent research on DIR has shown that multi-objective EAs are capable of obtaining excellent results comparable to algorithms that are currently used in practice, such multi-objective EAs have been used without many problem-specific enhancements, making them very slow, potentially even prohibiting practical use.

* E-mail: Anton.Bouter@cwi.nl; Telephone: +31 (0)20 592 4014

In the aforementioned recent framework for multi-objective DIR, a dual-dynamic grid transformation model was proposed in order to deal with large (dis)appearing structures, which is an issue that is addressed by few studies⁴⁻⁶. Using this model does however double the number of variables to optimize. This large number of variables impacts the efficiency of optimization, especially for the multi-objective EA that has been used in the aforementioned recent multi-objective framework for DIR: iMAMaLGaM⁷. The inefficiency is largely caused by the fact that, to ensure the most robust performance on a large variety of problems without making any assumptions on these problems, dependencies are assumed to exist between all pairs of variables. It was previously indeed demonstrated that disregarding all dependencies in DIR leads to inferior results⁸. However, taking into account *all* dependencies is only required for very particular, often artificial, problems. In DIR, however, grid points are strongly dependent on grid points that are adjacent, because these points directly influence the mapping of pixels from the source image to the target image, but are only weakly dependent on grid points that are remote.

To improve on the results of iMAMaLGaM, we exploit the implicit dependency structure that is present in DIR. To do so, we introduce a real-valued adaptation of the recently introduced Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA)⁹ that uses a prescribed dependency model to incrementally improve parts of solutions. Excellent results have recently been obtained with GOMEA on a wide range of discrete optimization problems⁹. Real-Valued GOMEA (RV-GOMEA) adopts mechanisms used in iMAMaLGaM in order to be able to solve optimization problems that have real-valued variables, such as DIR. In this paper, we describe how RV-GOMEA was tailored problem-specifically for DIR.

2. MATERIALS AND METHODS

2.1 Dual-dynamic grid transformation model

In a frequently used transformation model¹⁰, the source and the target image are overlaid with a regular grid with the same topology. With a regular triangulated grid model, a deformed image is calculated by mapping the contents of each source triangle to the shape and location of the corresponding target triangle. Deformation can therefore be achieved by modifying the coordinates of target grid points. Instead of overlaying the source image with a fixed grid and the target image with a moving grid, the dual-dynamic grid transformation model² overlays both the source and the target image with a moving triangulated grid of the same topology. Increasing the size of a triangle while decreasing the size of the corresponding triangle in the opposing grid naturally supports large deformations, including (dis)appearing structures. The dual-dynamic transformation model does however increase the model complexity, because the number of variables to optimize, i.e., for both grids: the number of grid points times the spatial dimensionality of the image, is doubled.

2.2 Multi-objective DIR

Using a multi-objective approach for DIR allows the utilization of multiple objectives of interest without setting their weights a priori. In this approach, one solution is said to dominate a different solution when it is better in at least one objective of interest, and not worse in any other objective. The final result of a multi-objective optimization approach consists of a so-called Pareto set, i.e., a set where no solution dominates any other solution. Such a Pareto set of solutions defines a so-called Pareto front of optimal trade-offs between the objectives of interest. A multi-objective framework for DIR using iMAMaLGaM was previously successfully introduced². iMAMaLGaM is a multi-objective EA that maintains a population of promising solutions and keeps track of non-dominated solutions in the so-called elitist archive. During each generation, a selection of the best solutions in the population is performed, based on which a probabilistic model¹¹⁻¹² is estimated. Specifically, a normal mixture probability distribution is estimated with maximum-likelihood. New solutions are sampled from this probability distribution, driving the optimization procedure. The final result of iMAMaLGaM is a so-called approximation set, which is the combined set of non-dominated solutions in the population and the elitist archive. This approximation set describes the best possible approximation of the optimal Pareto front based on the solutions found by the EA.

2.3 RV-GOMEA

A key difference between RV-GOMEA and iMAMaLGaM is that RV-GOMEA has a much higher selection pressure as a result of incrementally improving parts of existing so-called parent solutions. The parts to improve are explicitly defined in a so-called dependency, or linkage, model. Instead of estimating one normal mixture probability distribution,

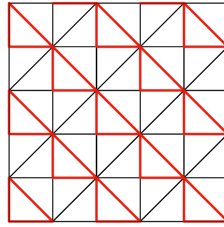


Figure 1. The selection of triangles and, at the borders, line segments, included in the dependency model, i.e., which are resampled during each generation of RV-GOMEA, highlighted in red in a triangulated 6×6 grid.

as is done in iMAMaLGaM, RV-GOMEA estimates such a distribution for each subset of variables in the dependency model separately. Solutions are then improved by sampling new values for these subsets of variables in the linkage model independently. Only if these newly sampled values lead to an improved solution, the newly sampled values are accepted, otherwise the parent solution is returned to its previous state. Because solutions are now only partially altered, it is not required to re-evaluate the entire solution to assess the contribution of a partial alteration. Such partial evaluations are an important reason for the practical efficiency of RV-GOMEA. Finally, to avoid manual tuning of the population size parameter, RV-GOMEA uses a population-sizing-free scheme¹³, which interleaves generations of independent instances of the EA with different population sizes, which are initially small and grow over time.

To apply RV-GOMEA to DIR we have studied the use of various dependency models. A good model should represent the minimum number of dependencies required to be able to solve the problem efficiently. We use a dependency model in which each element describes all coordinates of a single triangle. This means that solutions are incrementally modified by sampling new coordinates for the single triangles (either in the source or the target triangulation) that are included in the dependency model. To avoid including certain grid points an excessive number of times in the dependency model, we only include a subset of triangles. We have chosen to select the subset of triangles according to the highlighted triangles in Figure 1, because this selection has little overlap between triangles and it is easily generalizable to higher-dimensional grids. A small number of triangles at the borders of the grid are cropped to just line segments, but this does not harm the optimization process. Initial experiments have shown that this model leads to good results. Moreover, for DIR, efficient partial evaluations require that objective values can be updated by only re-evaluating contributions made by individual triangles (δ) that were modified. For this purpose, each objective is defined as a sum over the set of triangles Δ_s in the source grid and the triangles Δ_t in the target grid. Whenever any point of any triangle is moved, all objectives for this triangle must be re-evaluated. This leads to a large number of triangles having to be re-evaluated after the coordinates of a single triangle are resampled, as illustrated in Figure 2. In these illustrations, we show an arbitrary triangulated grid and a triangle highlighted in red. If the highlighted triangle were to be resampled, the objective functions of all triangles shaded in blue would be affected and would have to be re-evaluated to find the objective values of the entire grid. The two illustrations in Figure 2 show that the re-evaluation of adjacent triangles produces the most significant effect on small-dimensional grids, because a larger fraction of triangles has to be recomputed after the coordinates of one triangle are resampled. RV-GOMEA uses a slightly different dual-dynamic grid transformation model than iMAMaLGaM. The triangulation model for the recent most publication of iMAMaLGaM for DIR did not have grid points constrained to the borders². Conversely, for RV-GOMEA the outer points of the grid are constrained to be on the borders of the image so as to ensure that the entire image is covered by the transformation grid. Covering the entire non-empty area of the image is required for a valid transformation, but maintaining this condition would require more problem-specific enhancements and would increase the complexity of each partial evaluation.

2.4 Objective functions

Three objective functions, namely the similarity objective, the deformation magnitude objective, and the guidance error objective, are used for the optimization process. The guidance error objective is optional, and is only used when guidance information is supplied. All three objectives are now defined as a sum over the set of triangles Δ_s in the source grid and the triangles Δ_t in the target grid, because this allows the total objective values to be efficiently updated when the contribution of any single triangle to these objectives is changed.

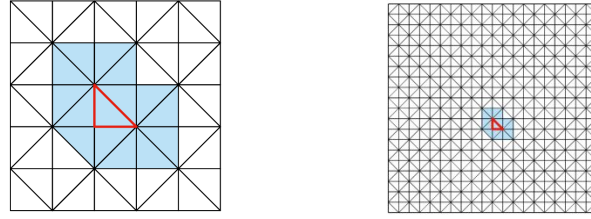


Figure 2. Triangulated 6×6 and 21×21 grids, indicating in blue for which triangles the objective values would change if the coordinates of the triangle highlighted in red would be resampled.

Similarity objective

In this paper, the similarity of the transformed source image and the target image is evaluated by computing the mean squared difference in pixel intensities, which is to be minimized. We note that, although not considered here, other notions of (dis)similarity often used in DIR can be evaluated partially also. For a single triangle δ , we iterate over the pixels inside this triangle, denoted $px(\delta)$ and compute the squared difference between a pixel's intensity in the one image and the bilinear-interpolated intensity at this pixel's corresponding position in the other image. The correspondence of positions in opposing images is straightforwardly governed by the correspondence of triangles in the triangulations defined over the opposing images. To obtain values that are independent of the image resolution, normalization is applied by dividing by the total number of pixels. Defining the (bi-linearly interpolated) intensity of a point p as $I(p)$ and its corresponding position in the opposing grid as p^c , we have:

$$F_{dissimilarity} = \frac{1}{|\Delta_s| + |\Delta_t|} \left[\sum_{\delta_s \in \Delta_s} \left[\sum_{p_s \in px(\delta_s)} (I_s(p_s) - I_t(p_s^c))^2 \right] + \sum_{\delta_t \in \Delta_t} \left[\sum_{p_t \in px(\delta_t)} (I_t(p_t) - I_s(p_t^c))^2 \right] \right]$$

Deformation magnitude objective

We employ Hooke's law¹⁴ based on the mean squared difference of edge lengths between edges e in the one grid and their corresponding edges e^c in the opposing grid. The final objective value of one triangle is then the squared sum of edge-length differences, normalized by dividing by three times the total number of triangles, i.e.,:

$$F_{deformation} = \frac{1}{3(|\Delta_s| + |\Delta_t|)} \left[\sum_{\delta_s \in \Delta_s} \left[\sum_{e_s \in edges(\delta_s)} (\|e_s\| - \|e_s^c\|)^2 \right] + \sum_{\delta_t \in \Delta_t} \left[\sum_{e_t \in edges(\delta_t)} (\|e_t\| - \|e_t^c\|)^2 \right] \right]$$

Guidance error objective

This objective is only used when guidance information is supplied, which is a set of tuples of contours or landmarks $G = \{(G_s, G_t)_1, \dots, (G_s, G_t)_k\}$ that predefine corresponding points or lines in the source and target images. The guidance objective aims to minimize the distance between these pairs of contours/landmarks. For all pixels on a contour within a certain triangle δ , denoted $G(\delta)$, the minimal distance to a point on the opposing contour is calculated. The objective value is the total sum of these minimal distances, normalized by the total number of pixels on the contour, again computed symmetrically, and summed over all pairs of contours, i.e.,:

$$F_{guidance} = \sum_{\delta_s \in \Delta_s} \left[\sum_{(G_s, G_t) \in G} \frac{1}{|G_s| + |G_t|} \sum_{p_s \in G_s(\delta_s)} \min\{d(p_s^c, p_t)\} \right] + \sum_{\delta_t \in \Delta_t} \left[\sum_{(G_s, G_t) \in G} \frac{1}{|G_s| + |G_t|} \sum_{p_t \in G_t(\delta_t)} \min\{d(p_t^c, p_s)\} \right]$$

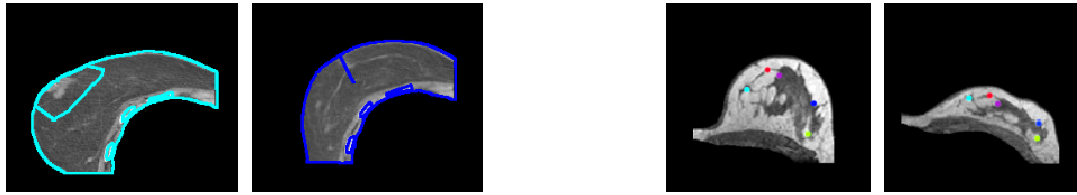


Figure 3. Left: the guidance information used for the pre- and post-operative benchmark problem. Right: the annotated landmark locations used for (calculation of the registration accuracy for) the prone-supine benchmark problem. Corresponding locations are annotated using the same color.

Feasibility constraints

If any two edges within a triangulation intersect, certain pixels are located within the perimeter of more than one triangle, meaning that the triangulation is infeasible. A technique known as constraint domination¹⁵, by which feasible solutions are always preferred over infeasible solutions, is used to deal with infeasible solutions. For this technique a constraint-violation value is computed that describes the degree of infeasibility of a solution, where a constraint-violation value of 0 defines a feasible solution. To ensure efficiency by exploiting partial evaluations in RV-GOMEA, constraint-violation values also must be computable per triangle. Therefore, we check whether a grid point p is inside the polygon bounded by the edges between the neighbors of p . If p lies outside this polygon, the triangulated grid has intersecting edges and is therefore infeasible. The total constraint-violation value is the number of grid points that violate this constraint. After generating a subset of new points, each of the newly generated points and each of their neighbors have to be re-checked to update the constraint-violation value. Note that this is a linear-time calculation as opposed to the quadratic-time method previously used in iMAMaLGaM to check all pairs of edges for intersections.

2.5 Experiments

In the first experiment, we run RV-GOMEA and iMAMaLGaM on two problems with a content mismatch. In particular, we consider an artificial benchmark problem and a pair of 2D slices taken from pre- and post-operatively acquired breast CT scans from a breast cancer patient. The selection of the pair of slices from the breast CT scans was done after rigid registration on the bony anatomy. For both problems, the guidance error objective was used. Guidance information for the real-life problem was annotated by a clinical expert, and is shown in Figure 3. Previous work demonstrated that both problems can be adequately solved by iMAMaLGaM in a multi-resolution scheme², but for the sake of comparison, in this experiment we considered the use of RV-GOMEA and iMAMaLGaM using only single grid-resolution schemes with dimensions 6×6 , 11×11 , and 21×21 . Each run was terminated after $\ell \cdot 10^4$ evaluations where ℓ is the number of variables, respectively being 144, 484, and 1764. For this purpose, each of RV-GOMEA's partial evaluations was counted as a single evaluation in iMAMaLGaM.

In a second experiment, we used a multi-resolution scheme² in RV-GOMEA with resolutions 6×6 , 11×11 , and 21×21 on two 2D slices of a breast MRI scan acquired from a healthy volunteer in prone and supine position, not including guidance information. The selection of this pair of slices was done after rigid registration on the bony anatomy. A time limit equal to the number of variables in seconds is used for each resolution, totaling close to 40 minutes.

We compare the speed and quality of results of RV-GOMEA with that of iMAMaLGaM². For a notion of Pareto front quality, we use the well-known hypervolume metric¹⁶ that calculates the total volume of the objective space that is dominated by a Pareto front. In our experiments, we compare the hypervolumes of the approximation sets of the two algorithms. Because the volume calculated by the hypervolume metric would otherwise be infinitely large, the space that is considered by the hypervolume metric is bounded by the so-called Nadir point, for which we use 1.1 times the worst encountered value in each objective. Due to this, hypervolumes of different benchmark problems or different grid resolutions cannot directly be compared. By considering the time it takes two algorithms to reach a certain hypervolume, we can define a notion of speed-up. We evaluate the quality of registration results on the pair of prone-supine MRI scans with a set of five landmarks annotated by a clinical expert, shown in Figure 3, allowing us to compute the target registration error (TRE), which is the mean distance between the deformed landmark locations and the predefined target landmark locations. The purpose of this metric is similar to that of the guidance objective, but it is not used as an objective during optimization.

All breast CT and MRI scans are provided through the courtesy of The Netherlands Cancer Institute – Antoni van Leeuwenhoek Hospital.

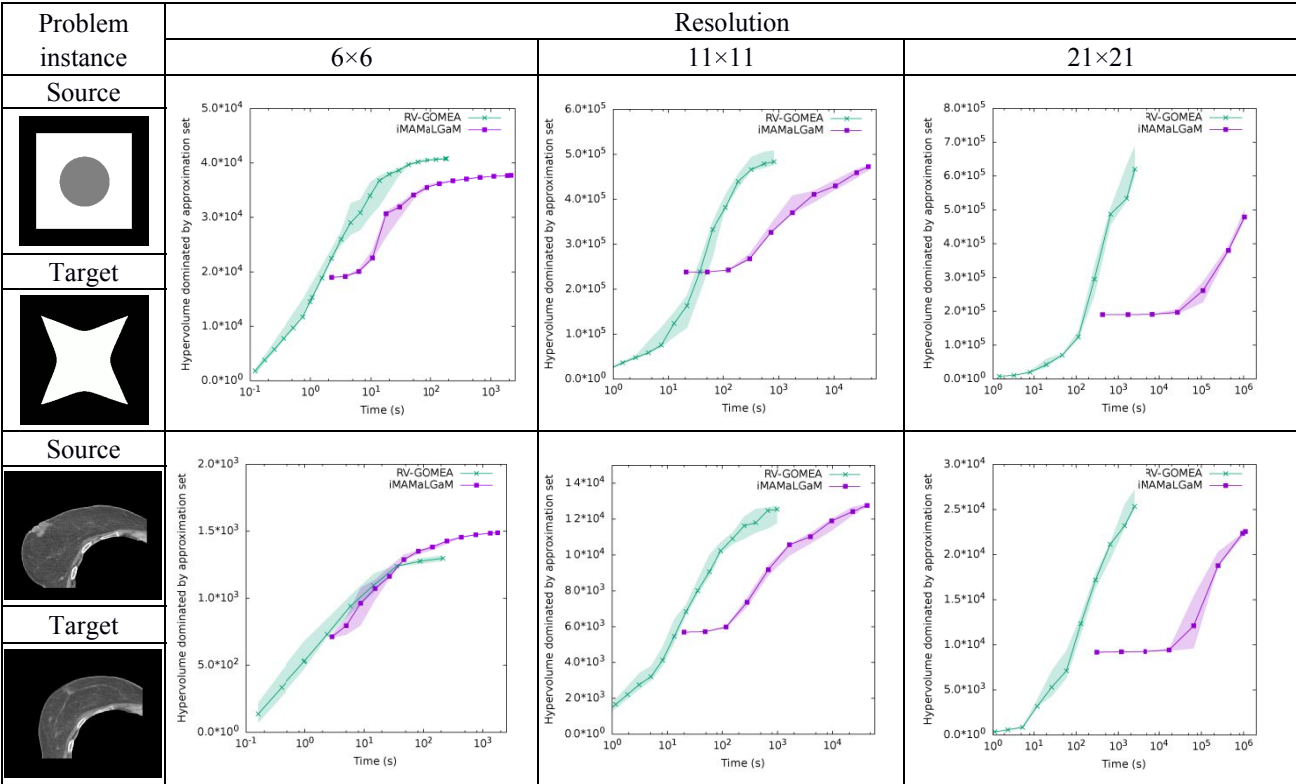


Figure 4. Hypervolumes of single-resolution runs with 6×6, 11×11, and 21×21 grids, on an artificial problem instance with a disappearing structure (top row) and a set of pre- and post-operative breast CT scans (bottom row).

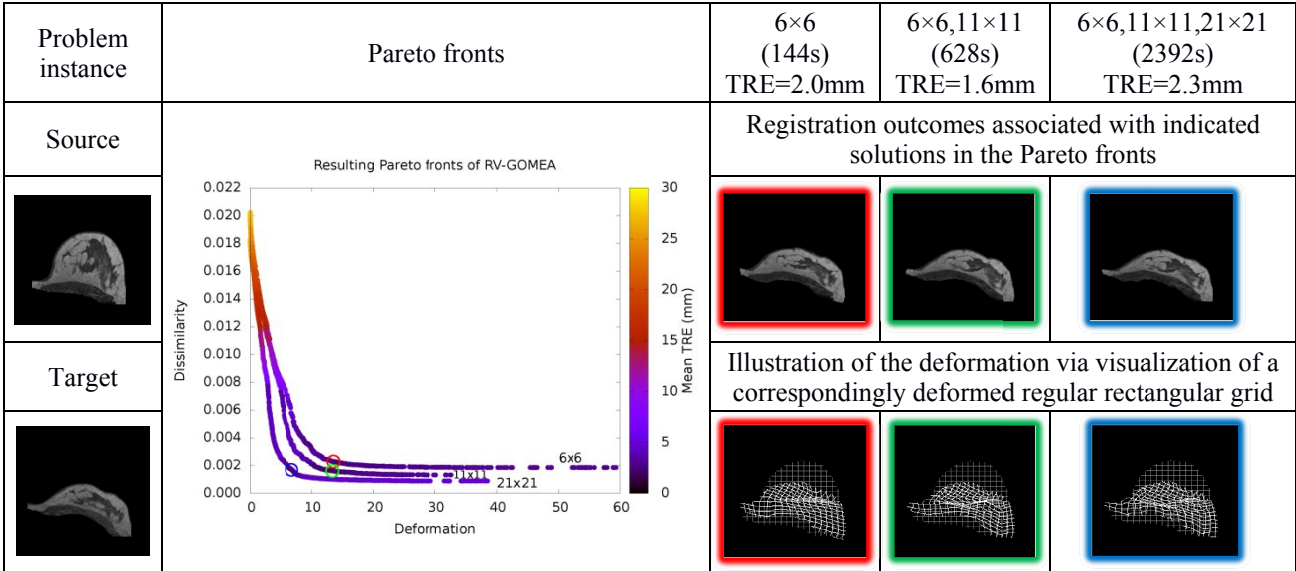


Figure 5. Results on the prone-supine breast MRI test case. The Pareto fronts are displayed for the single, double, and triple resolution runs where every run started from the 6×6 grid resolution. In each Pareto front, the solution with the minimal mean TRE is encircled and displayed on the right along with its deformation applied to a uniform grid, its mean TRE, and the run-time. Note: the actual dual-dynamic triangular grids that define the deformation field are not shown.

3. RESULTS AND DISCUSSION

For the first experiment, the hypervolumes interpolated over 4 runs are displayed in Figure 4 where the shaded areas demark the range between the best and worst run. The median run is displayed as the plotted line. Speed-up is then calculated by comparing the time required to reach the minimum hypervolume obtained by either iMAMaLGaM or RV-GOMEA upon termination. For the artificial problem, speed-up factors of 114, 102, and 1641 are obtained on resolutions of 6×6 , 11×11 , and 21×21 , respectively. For the pre- and post-operative problem, speed-up factors of 0.25, 31, and 940 are obtained on the respective resolutions. The apparent inferior performance of RV-GOMEA for the 6×6 grid resolution on the pre- and post-operative problem is caused by the fact that in RV-GOMEA we use a slightly different transformation model where the outer points of the grid (i.e., 20 out of 36 points for a 6×6 regular grid resolution) are constrained to be on the borders of the image. For higher grid resolutions, the impact of this difference is negligible. Also note that iMAMaLGaM achieves a higher hypervolume in its first generation, which is caused by the relatively small initial population size of RV-GOMEA due to its population-sizing-free scheme.

Results for the second experiment are shown in Figure 5. This figure shows the Pareto fronts after each stage of the multi-resolution scheme for RV-GOMEA. Each solution is color-coded with regard to its mean TRE and the solution with the minimal mean TRE is encircled. After running the 6×6 grid resolution for only 144 seconds, a solution was already found with a mean TRE of 2.0 mm. For higher resolutions the mean TRE does not get much better, but the deformation does become smoother. The position of the encircled solution makes it clear that minimizing dissimilarity does not always lead to higher-quality solutions. Instead, highly deformed, over-fitted solutions with unnatural deformations are obtained. Moreover, because the mean TRE is very irregular along the Pareto front, searching by trial and error for a linear combination of weights to use in a single-objective approach that linearly weights the objectives is very likely to get stuck in a local optimum of weights with a high risk of obtaining an unsatisfactory registration outcome. This highlights the strength of a multi-objective approach.

Our results could be improved even further by enabling a previously introduced adaptive steering approach². During optimization, this technique purges certain solutions from the Pareto set in order to direct the algorithm towards the most interesting part of the Pareto front, i.e., the part where each objective value of a solution is close to the best known values of this objective function. RV-GOMEA could further be combined with the recently introduced smart grid-initialization¹⁷ technique that was observed to obtain an additional speed-up of a factor between 10 and 100.

4. CONCLUSIONS

The advantage of a multi-objective approach is that it results in a set of solutions representing different trade-offs between the objectives of interest, from which a desirable registration outcome can be selected much more insightfully compared to manually tuning weights of a traditional, weighted-objectives, single-objective approach.

Multi-objective DIR with a dual-dynamic transformation model to account for large anatomical differences has high potential, but using off-the-shelf state-of-the-art multi-objective EAs it is prohibitively slow for real-world practice. In this paper, we introduced a new multi-objective EA for DIR. We showed that by using a problem-specific tailored multi-objective implementation of the recently introduced EA known as RV-GOMEA, a large speed-up can be obtained. We observed a speed-up of up to a factor of ~ 1600 compared to the previously used algorithm, scaling to even higher factors as the dimensionality of the transformation grid increases.

Because of its substantially improved scalability, RV-GOMEA opens the door to applying multi-objective EAs to accurately solve DIR problems in 3D with dual-dynamic transformation grids to support large anatomical variations and content mismatches.

ACKNOWLEDGEMENTS

The authors thank the following people from The Netherlands Cancer Institute – Antoni van Leeuwenhoek Hospital for their contributions to this study: K. Pirpinia (Department of Radiation Oncology), C.E. Loo (Department of Radiology), and M.T.F.D. Vrancken Peeters (Department of Surgical Oncology). Further, the authors thank C. Witteveen (Algorithmics Group, Department of Software Technology, Faculty of Engineering, Mathematics and Computer Science, Delft University of Technology).

REFERENCES

- [1] Klein, S., et al. "Elastix: a toolbox for intensity-based medical image registration." *IEEE Trans. Med. Imaging* 29(1), 196-205 (2010).
- [2] Alderliesten, T., Bosman, P.A.N. and Bel, A., "Getting the most out of additional guidance information in deformable image registration by leveraging multi-objective optimization" *Proc. SPIE* 9413, 94131R, doi:10.1117/12.2081438 (2015).
- [3] Deb, K. [Multi-objective optimization using evolutionary algorithms], Vol. 16. John Wiley & Sons, (2001).
- [4] Periaswamy, S. and Farid, H., "Medical image registration with partial data," *Med. Image Anal.* 10(3), 452-464 (2006).
- [5] Risholm, P., Samset, E. and Wells-III, W., "Validation of a nonrigid registration framework that accommodates tissue resection," *Proc. SPIE* 7623, 762319 (2010).
- [6] Nithiananthan, 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).
- [7] Rodrigues, S.F., Bauer, P. and Bosman, P.A.N., "A novel population-based multi-objective CMA-ES and the impact of different constraint handling techniques" *Proc. GECCO*, 991-998, doi: 10.1145/2576768.2598329 (2014).
- [8] 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).
- [9] Thierens, D. and Bosman, P.A.N., "Optimal mixing evolutionary algorithms." *Proc. GECCO*, 617-624, doi: 10.1145/2001576.2001661 (2011).
- [10] Crum, W.R., Hartkens, T. and Hill, D.L., "Non-rigid image registration: theory and practice." *Br. J. Radiol.* 77, 140-153 (2004).
- [11] 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).
- [12] Pelikan, M., Sastry, K., and Cantú-Paz, E., [Scalable optimization via probabilistic modeling: from algorithms to applications], Springer-Verlag, Berlin (2006).
- [13] Pereira, J.C. and Lobo, F.G., "A Java Implementation of Parameter-less Evolutionary Algorithms." *arXiv preprint, arXiv:1506.08694* (2015).
- [14] Arfken, G., [Mathematical methods for physicists], Academic Press, Inc., San Diego (1985).
- [15] Deb, K., Pratap, A. and Mevarivan, T., "Constrained test problems for multi-objective evolutionary optimization" *Proc. EMO*, 284-298, doi: 10.1007/3-540-44719-9_20 (2001).
- [16] Zitzler, E. and Thiele, L. "Multiobjective optimization using evolutionary algorithms-A comparative case study" *Proc. PPSN*, 292-301, doi: 10.1007/BFb0056872 (1998).
- [17] Bosman, P.A.N. and Alderliesten, T., "Smart grid initialization reduces the computational complexity of multi-objective image registration based on a dual-dynamic transformation model to account for large anatomical differences" *Proc. SPIE* 9784, 978447, doi:10.1117/12.2217011 (2016).