

A Nonlinear Projection Scheme for Fast Rigid Registration

V. Daum¹, D. Hahn¹, and J. Hornegger¹

Abstract—Intensity based rigid registration algorithms commonly employed for medical image fusion are based on the iterative optimization of a pixel-by-pixel distance measure defined on the images. As medical images grow larger in size due to advanced scanner technology, evaluating such similarity measures is no longer computationally efficient. In order to overcome the inherent limitations of the standard approach we propose a new, nonlinear projection scheme that enables a very fast evaluation of the distance between two images. Current state-of-the-art projection schemes decompose a six dimensional search space into three dimensional subspaces. The proposed approach, however, yields a complete decomposition into 1D subspaces. The optimization on these subspaces is highly efficient and does not require a reprojection. This scheme is therefore suitable for 2D and 3D registrations, and it is able to cope with subvolume matching problems. Furthermore, the use of modern graphics hardware allows for a highly efficient implementation. Experiments show that computation times can be reduced to less than 10 seconds with the proposed approach for 256^3 sized volumes.

Index Terms—rigid registration, fusion, projection, optimization

I. INTRODUCTION & MOTIVATION

RIGID registration algorithms are required in many clinical scenarios, e.g. tumor therapy monitoring, digital subtraction angiography or as a pre-registration for a subsequently applied non-rigid registration algorithm. A registration algorithm calculates a spatial transform from one image's coordinate system to another, hence allowing to align corresponding image content. In the case of rigid registration this transform consists of six degrees of freedom (DOF), three translational and three rotational parameters. Rigid registration approaches are currently heavily used in medical practice. All share the need for being as fast, precise and robust as possible. In tumor therapy monitoring, for instance, an automatic registration algorithm has to deal with a mono-modal 3D optimization problem that is defined by an objective function. This function - in the most general case - is based solely on the similarity between the image intensities. There exists a huge variety of intensity similarity measures for different kinds of applications. Some are only suitable for mono-modal cases, while others can also deal with different types of multi-modal problems. They all have in common that a straightforward implementation leads to a pixel-by-pixel distance measure evaluation which is prohibitively expensive for large images. Therefore, standard approaches for solving

nonlinear optimization problems (e.g. Newton, Gradient Descent, Levenberg-Marquardt) that are applied directly to the objective function tend to be rather computationally expensive when it comes to real clinical data. Although multi-resolution approaches help to alleviate this problem to a certain amount, the highest accuracy can only be achieved by performing some additional optimization on the best resolution of the data. For a clinical workflow this may take too much time, even on the latest workstations used nowadays.

In order to make a fast optimization possible, we present a novel projection based parameter decomposition scheme that speeds up the registration process significantly. Compared to state-of-the-art projection based algorithms the proposed approach is able to deal with subvolume matching problems and also yields a complete decomposition of the six DOF problem into disjoint 1D optimizations. While solving the 1D optimization problems, the images do not have to be reprojected because the optimized parameter is related linearly to the projection space. Using orthogonal projections, this holds only true for translational parameters. Therefore, the projection scheme is extended by introducing a circular and cylindrical projection geometry for 2D and 3D images, respectively. Additionally, the proposed approach is suited for a fast implementation on standard graphics hardware.

II. RELATED WORK

Widely established or freely available, intensity based, rigid registration algorithms require the computation of an image similarity measure defined on all pixels of the overlap region between a reference and a template image [1], [2], [3]. Although the partial derivatives of this measure, with respect to the transform parameters, can often be analytically formulated, its evaluation on a pixel-by-pixel basis is very time-consuming. In addition, it is well known from theory that the conditioning of optimization techniques in parameter spaces decreases with the number of dimensions. In rigid image registration, the six DOF parameter optimization problems (rotation and translation in 3D) are usually nonlinear and contain local optima. Therefore, multi-level approaches are commonly used in order to address these problems and to speed up the registration. Nonetheless, to achieve the highest accuracy in terms of spatial alignment of corresponding image content it is necessary to carry out at least some final iterations of the optimization on the best resolution of the images. As this usually involves some sort of interpolation, the evaluation of the similarity measure gradients becomes very time consuming.

In 2000 and 2002, Hornegger and Niemann [4], [5] proposed a new method to reduce the complexity of the opti-

¹V. Daum, D. Hahn and J. Hornegger are with the Friedrich-Alexander-University of Erlangen-Nuremberg (FAU), Institute of Pattern Recognition, Martensstr. 3, 91058 Erlangen, Germany

mization by introducing projections (marginalizations). The problem described in their article is to find the position and orientation of 3D objects in space for an object recognition application. The authors describe how appropriate projections onto lower dimensional sub-spaces of the objective function lead to a reduced number of DOF. This enables noteworthy faster optimizations on the disjoint sub-spaces compared to the full-fledged problem. The solution can then be found iteratively by repeating the reduced DOF optimizations in the lower dimensional spaces. This principle was applied to 3D image registration by Khamene et.al. [6]. They iteratively project the 3D volumes onto 2D images along the coordinate axes and hence eliminate a subset of the parameters for each subproblem. Although the presented results are quite impressive, their approach still couples three DOF (for the 2D registrations). Their algorithm therefore cannot reduce the complexity for 2D image registration problems. The authors also did not address how the projections can be performed in case of subvolume matching situations, which imposes limits on the general applicability of the orthogonal projection approach.

III. METHODS

In the following we will shortly introduce the standard approach and related drawbacks. The subsequent subsections will illustrate the basic ideas of using projections for the evaluation of image similarities and how they can help to accelerate the registration. Finally we extend the projection based registration approach to support subvolume problems by optimally fitting the projection geometry into the overlap domain.

A. Standard Approach

As previously mentioned, the rigid, image intensity based registration corrects for six DOF misalignments between a reference image R and a template image T . For the sake of simplicity, only the sum of squared differences (SSD) is considered as distance measure between the two images:

$$\hat{\Phi} = \underset{\Phi}{\operatorname{argmin}} \int_{\Omega} \mathcal{D}[R, T, \Phi](\mathbf{x}) \, d\mathbf{x} \quad (1)$$

$$\Phi(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{t} \quad (2)$$

$$\mathcal{D}[R, T, \Phi](\mathbf{x}) = \frac{1}{2}(R(\mathbf{x}) - T(\Phi(\mathbf{x})))^2 \quad (3)$$

Here, Φ denotes the six DOF (2) consisting of a rotation matrix $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ and a translation vector $\mathbf{t} \in \mathbb{R}^3$. It maps spatial positions $\mathbf{x} \in \mathbb{R}^3$ from the template T to the reference image R . The SSD similarity measure \mathcal{D} is defined only within the overlap domain Ω that is dependent on the current transform Φ .

The parameters of this transform are usually estimated by a nonlinear optimization approach. Finding an extremum of (1) is equivalent to finding the root of its derivative, which consists of a sum of partial derivatives of (3).

$$\begin{aligned} \nabla_{\mathbf{a}} \mathcal{D}[R, T, \Phi](\mathbf{x}) &= \\ &= (T(\Phi_{\mathbf{a}}(\mathbf{x})) - R(\mathbf{x})) \nabla_{\mathbf{a}} T(\Phi_{\mathbf{a}}(\mathbf{x})) \\ &= (T(\Phi_{\mathbf{a}}(\mathbf{x})) - R(\mathbf{x})) (\nabla_{\mathbf{a}} \Phi_{\mathbf{a}}(\mathbf{x}))^T \nabla_{\mathbf{x}} T(\Phi_{\mathbf{a}}(\mathbf{x})) \\ &\stackrel{!}{=} 0 \end{aligned} \quad (4)$$

The partial derivatives given by (5) with respect to the parameters, which are contained in the six dimensional parameter vector \mathbf{a} , can be combined into a vector that determines the direction to the ascent of the distance measure. This information is used during the optimization in order to find the direction towards the desired optimum.

The drawback of this standard approach becomes obvious when looking at the integral contained in (1): this means that the derivative has to be evaluated on the entire overlap domain Ω . In practice, an implementation has to iterate over this domain and perform an interpolation and a computation of (5) for each voxel.

B. Axis Aligned Projections

The general idea of the aforementioned parameter space reduction algorithms is to eliminate DOF by generating a projection that depends linearly on the parameters. This principle is illustrated in the Fig. 1 taken from [4]. The figure shows a

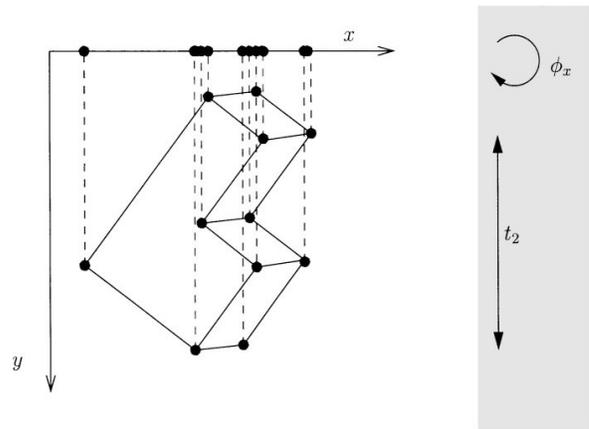


Fig. 1. Illustration of linear projection of features onto the x -axis (from [4]). The projection remains the same if the object is translated along the y -axis or rotated around the x -axis. If the object is translated parallel to the x -axis, the projection is translated accordingly on the x -axis. It is not necessary to reproject, but the projection can be translated directly instead.

projection of 3D object features onto the x -axis. A translation of the 3D features along the projection direction, as well as a rotation around the x -axis, does not affect the projection. In the presented article, the image voxels may be regarded as such features. In general, shifting an image along a direction that is orthogonal to the projection axis translates its projection by the same amount and vice versa. An optimization within these reduced parameter spaces does not require reprojections in the above cases. Any other rotation of the image, however, would make a reprojection necessary. As far as only translations are considered, an axis aligned projection down to 1D images is sufficient for the entire optimization. This is illustrated in the left image of Fig. 2.

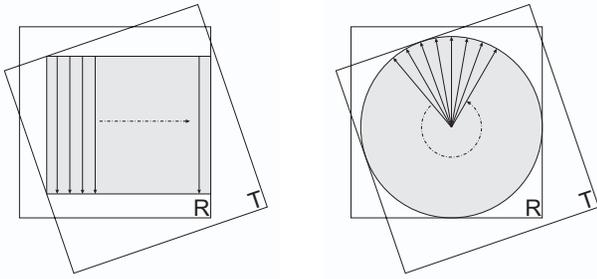


Fig. 2. In 2D the projection scheme is achieved by orthogonal projections along the coordinate axes for the translational parameters (left). A nonlinear projection onto a circle for the rotational parameter is shown on the right. The projection geometry (square or circle) is fit into the overlap domain of the reference image R and template image T .

C. Circular Projection

For the complete projection of the 2D or 3D registration problem domain to 1D disjoint parameter subsets it is necessary to have a projection geometry for the rotational parameters that exhibits the same properties as the axial projection scheme in III-B for the translational parameters. A nonlinear projection along the rays from the center of a circle onto its boundary is therefore proposed for the rotational parameter optimization in 2D problems (see the right illustration of Fig. 2). Hence, rotations of the image around the center of the circle result in a shift of the projection on the circle boundary without the need for a reprojection. If the same circle geometry is chosen for projecting both the reference and template image, the parameter for a rotation around the circle's center can therefore be optimized in the 1D projections without the need to look at the original data. Only one circle projection for each image is required for several optimization steps in the 1D space.

Combining both the axis aligned and the circular projections for the case of 2D registration problems yields an iterative optimization scheme that is highly efficient with respect to the number of projections performed. One projection has approximately the same computational complexity as one evaluation of the standard distance measure derivative on the entire overlap domain (5).

D. Cylinder Projection

In 3D, one would assume that the analog projection geometry for a circle in 2D is a sphere. It turns out, however, that the parameters on a spherical surface are nonlinearly related and feature problems at the poles. Modifying one parameter on the surface of the sphere therefore results in a nonlinear warping on the entire projection image domain. Therefore, a suitable projection geometry in 3D is a cylinder, which again leads to a linear relationship between the parameters. A rotation around the axis of the cylinder, similar to the rotation around the center of the circle in the 2D case, only shifts the projection on the surface of the cylinder. A translation along the axis of the cylinder, again similar to the translation along an axis orthogonal to the projection direction in 2D, results in a shift on the surface of the cylinder along its axis. In both

cases no reprojections are necessary. If the unrolled cylinder surface is viewed as a 2D image, a rotation around the cylinder axis in 3D can be regarded as a translation of the projection image along its x -axis with periodical boundary conditions. This applies analogously to a translation of the image along the cylinder axis in 3D, which leads to a translation of the projection image along its y -axis. An example is given in Fig. 3.

This unrolled surface of the cylinder can then be further linearly projected along its axes, which results in a fully decoupled sequence of 1D optimizations. One cylinder geometry allows the optimization of one rotational and one translational parameter. In order to cover the entire six DOF parameter space, three cylinders (aligned with the coordinate axes) are necessary. Since the cylinder projections are not independent of each other, it is inevitable to perform the optimization iteratively by projecting on each of the cylinders in an alternating order.

Since performing projections along rays is very similar to volume rendering techniques, it is very suitable for evaluating it with graphics hardware. The entire projection approach can be implemented using trilinear filtering with specially adapted shading programs for the evaluation of the projection rays. Additionally, it is possible to incorporate transfer functions or segmentation results (e.g. see [7]).

E. Fitting of the Projection Geometry

Medical image registration often involves subvolume problems, where only a part of the content of the first image is shown in the other. If the projection does not address this, the projections of the larger image contain additional information that is missing in the others. An example for this problem is given in Fig. 4. The projection area is acquired by a maximum fit of the projection geometry into the overlap domain. The resulting optimization problem for the fit of the projection geometry is solved by sequential quadratic programming [8]. By maximizing the projection area, we also maximize the amount of information used during the registration. Therefore, we retain the ability to deal with subvolume registration problems while transferring the maximal amount of information into the projection images.

IV. RESULTS

The proposed algorithm has been implemented for 2D and 3D on the CPU. In addition, results have been acquired with a graphics processing unit (GPU) accelerated version. The following comparison of the algorithm with an ITK based rigid registration [9] on the CPU was performed on an Intel Pentium M 2.26 GHz machine with 2 GByte RAM and an NVidia GeForce 7800 Go. The 2D CPU registration of two 256^2 medical images takes approximately 10 seconds using a common ITK rigid registration approach and no more than 0.3 seconds with the proposed approach (no GPU acceleration). In 3D, an ITK registration of two 256^3 volumes lasts more than 10 minutes. Our approach on the CPU takes approximately 3 minutes. The GPU optimized implementation takes about 40 milliseconds to perform one cylinder projection on a 256^3

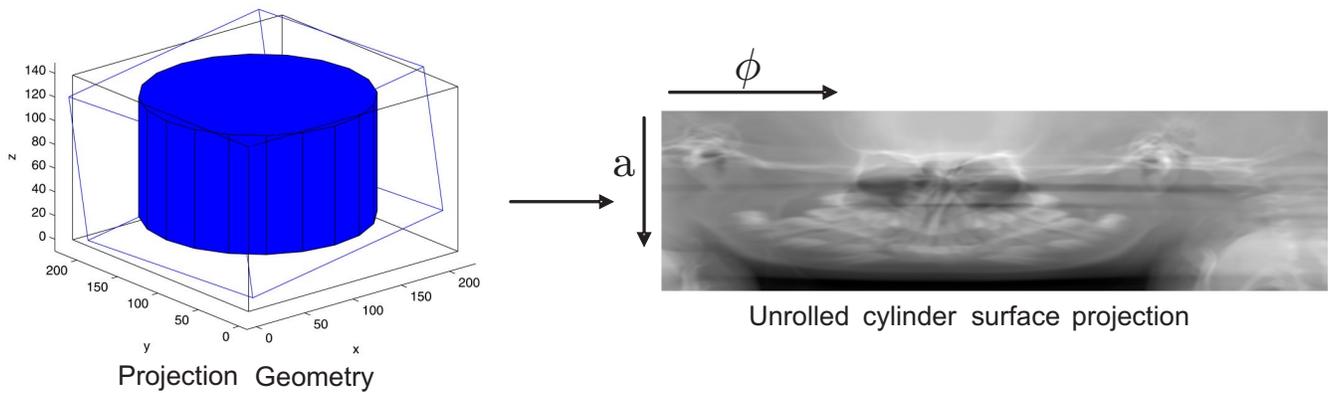


Fig. 3. The illustration on the left shows the cylindrical projection geometry fitted into the overlap domain. The right hand side depicts the unrolled surface of the cylinder projection from a 3D CT head dataset. Its x -axis corresponds to a rotation parameter around the cylinder axis denoted by ϕ , whereas the y -axis corresponds to a translation parameter along the cylinder axis given by a .

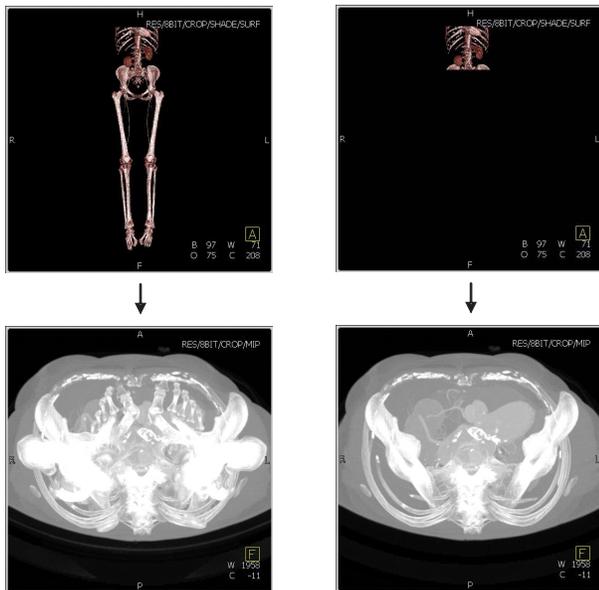


Fig. 4. If the region of overlap is not incorporated into the projection scheme, the resulting projections are biased. The top row of this figure shows the two images that are projected along their longitudinal axes. The bottom row depicts the according projections. The structures that are only visible in the left image (i.e. legs, hip, etc) contribute to the projection image and lead to additional information that imposes a problem on a subsequent registration that utilizes these images.

dataset, which results in a total runtime for the registration of less than 10 seconds. A typical result is shown in Fig. IV. Additionally, we believe there is still room for improvement in our GPU implementation. The current bottleneck is the processing of the 2D unrolled cylinder surfaces. These images are currently transferred from GPU to CPU memory to perform the optimization, which is no longer needed if the entire approach is executed on the GPU.

V. CONCLUSIONS

Projection based image intensity registration methods yield large performance gains compared to commonly used algo-

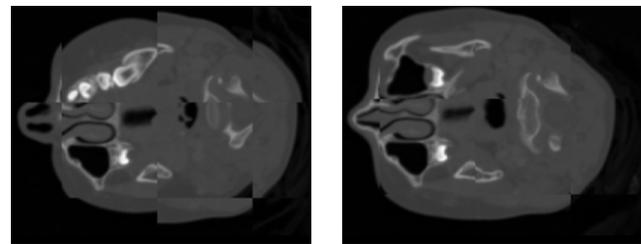


Fig. 5. A slice from the 3D checkerboard images resulting from a registration of a tumor pre- with post-therapy 3D head CT dataset. The left image shows the images before the registration and on the right after registration.

gorithms. We have shown that the proposed nonlinear projection scheme decouples the rotational as well as the translational parameters from their respective projections, which is the major drawback of prior algorithms. The ability to perform subvolume registrations is retained through a nonlinear fitting of the projection geometries. The entire algorithm can be ported to modern graphics hardware, which highly accelerates its performance.

ACKNOWLEDGMENT

The authors acknowledge the financial support of the International Max Planck Research School for Optics and Imaging (IMPRS-OI), the Deutsche Forschungsgemeinschaft (DFG), grant number SFB, TP C10 and HipGraphics for providing the volume rendering software InSpace. We also gratefully acknowledge the financial support of this work by the Erlangen Graduate School in Advanced Optical Technologies (SAOT).

REFERENCES

- [1] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration," *Physics in Medicine and Biology*, vol. 46, no. 3, pp. R1–R45, 2001.
- [2] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, no. 1, pp. 1–36, 1998.
- [3] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta, "Automated image registration: I. general methods and intrasubject, intramodality validation," *J. Comput. Assist. Tomogr.*, vol. 22, no. 1, pp. 139–152, 1998.

- [4] J. Hornegger and H. Niemann, "Probabilistic modeling and recognition of 3-D objects," *International Journal of Computer Vision*, vol. 39, no. 3, pp. 229–251, 2000.
- [5] J. Hornegger, V. Welker, and H. Niemann, "Localization and classification based on projections," *Pattern Recognition*, vol. 6, no. 35, pp. 1225–1235, 2002.
- [6] A. Khamene, R. Chisu, W. Wein, N. Navab, and F. Sauer, "A novel projection based approach for medical image registration." in *WBIR*, ser. Lecture Notes in Computer Science, J. P. W. Pluim, B. Likar, and F. A. Gerritsen, Eds., vol. 4057. Springer, 2006, pp. 247–256.
- [7] D. A. Hahn, J. Hornegger, W. Bautz, T. Kuwert, and W. Roemer, "Unbiased rigid registration using transfer functions," in *Proc. of SPIE on Medical Imaging*, vol. 5747, 2005, pp. 151–162.
- [8] A. Wächter and L. T. Biegler, "On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [9] L. Ibáñez, W. Schroeder, L. Ng, and J. Cates, *The ITK Software Guide*, 2nd ed. Kitware Inc., 2005. [Online]. Available: <http://www.itk.org>