Mathematics Meets Medicine: An Alignment Story

By Bernd Fischer and Jan Modersitzki

Medical imaging has come a long way since 1895, when Wilhelm Conrad Röntgen discovered x-rays. More than a century later, the phrase "going for a scan" can refer to a multitude of different medical imaging techniques that are used for diagnosis and treatment. The "old" x-rays are the backbone of computer tomography (CT) systems. Single photon emission (SPECT) and positron emission tomography (PET) rely on the properties of radioactive isotopes. Magnetic resonance imaging (MRI) exploits the well-known principles of nuclear magnetic resonance; diffusion tensor imaging (DTI) uses magnetic field-gradient pulses to track the movement of water molecules in the brain. These are just a few of the relatively new scanning techniques.

Whatever the modality, a typical scan now yields more information than ever before. The various scanners produce gigabyte-size images, often in time series consisting of hundreds or even thousands of images. Clearly, high-performance computing is increasingly important in medical imaging, and mathematics plays a crucial role in this rapidly growing discipline. Activity in this field has been especially intense in the last two decades; see, for example, [1, 5–10].

The Need for Registration

One of the many timely challenges in image processing is known variously as "registration," "fusion," "warping," or "mat-ching" (an introduction can be found in [7]). The goal is to combine the power of different imaging techniques through the automatic alignment of given images. As it turns out, registration techniques are used to advantage in many application areas, ranging from biomedical data mining to remote sensing. Though registration is a very general problem, approaches, requirements, and techniques depend strongly on the area of application.

Even when consideration is limited to medical applications, diverse scenarios arise. One class of applications is related to the fusion of data from different imaging devices—for example, the combination of multislice CT and PET in cancer imaging. PET displays areas of metabolic hyperactivity, which often indicate the presence of cancer, while the sub-millimeter resolution of multislice CT offers a clear picture of anatomical details. Combining this information yields an accurate means of locating a tumor and assessing its size and shape.

The need for registration is often caused by motion (or, more generally, transformation) of the sensed structure—in this case, the patient. Motion can occur on fine scales (breathing, heartbeat, convulsion), on medium scales (during transfer of the patient from one imaging device to another), or on large scales (drastic changes in the patient's body caused by surgery, or weight gain or loss, in scans made years apart). Uses of registration include comparison, integration, navigation, and classification.

Whether for registration (establishment of correspondences between structures in different images) or for image segmentation (division of an image into meaningful pieces), there is as yet no unified mathematical formulation of the problem. Transforming a medical registration problem into a sound mathematical model is quite challenging. We illustrate the diversity and different needs for registration with a few examples. As mentioned earlier, PET–CT fusion is useful for cancer detection. The following two examples are concerned with image-guided surgery and routine screening.

Surgery: Image-guided Liver Resection

Intraoperative imaging and image-guided surgery are increasingly important sources of registration problems. Figure 1 illustrates planning for a liver resection. Whether for tumor removal or for transplantation, liver resection gives rise to a serious optimization problem: Remove as much as possible (all tumor tissue or, in the case of transplantation, at least enough to allow the transplanted part to survive) while sparing a sufficient amount of the (patient's/donor's) liver.

This objective is inherently complicated by the complex geometry of the blood vessels that supply the liver. As shown in Figure 1 (middle), orientation during surgery is quite difficult. A time series of two-dimensional, real-time ultrasound images (Figure 1 (right)), therefore, provides valuable guidance to the surgeon.

The goal of registration is to align the preoperative plan with the actual in vivo situation, as displayed by the ultrasound images. In addition

Figure 1. *Liver surgery: Preoperative plan (left), the surgical field (middle), ultrasound-based navigation (right). Images courtesy of MeVis, Bremen (left) and HP Bruch, Clinic of Sur-gery, University of Lübeck (middle and right).*

to the low signal-to-noise ratio of ultrasound images, the registration needs to accommodate two conflicting demands: The surgeon needs highly accurate results in nearly real time.

Screening: Contrast-enhanced MR Mammography

Figure 2 shows a two-dimensional slice at a certain time in a four-dimensional time series in contrast-enhanced MR mammography. A contrast agent is injected into a vein prior to scanning to show fine structural details and highlight abnormalities. The goal is to visualize tumors and, based on the speed of the contrast agent uptake, to classify the tumors. As to be expected, this task is complicated by motion, such as the patient's heartbeat or breathing. Moreover, as uptake and diffusion of the contrast agent take place over time, the 3D images continue to change, with the changes intensifying in time. Registration is needed to compensate for the motion and to generate a clear view of the suspicious regions during the entire scanning process.
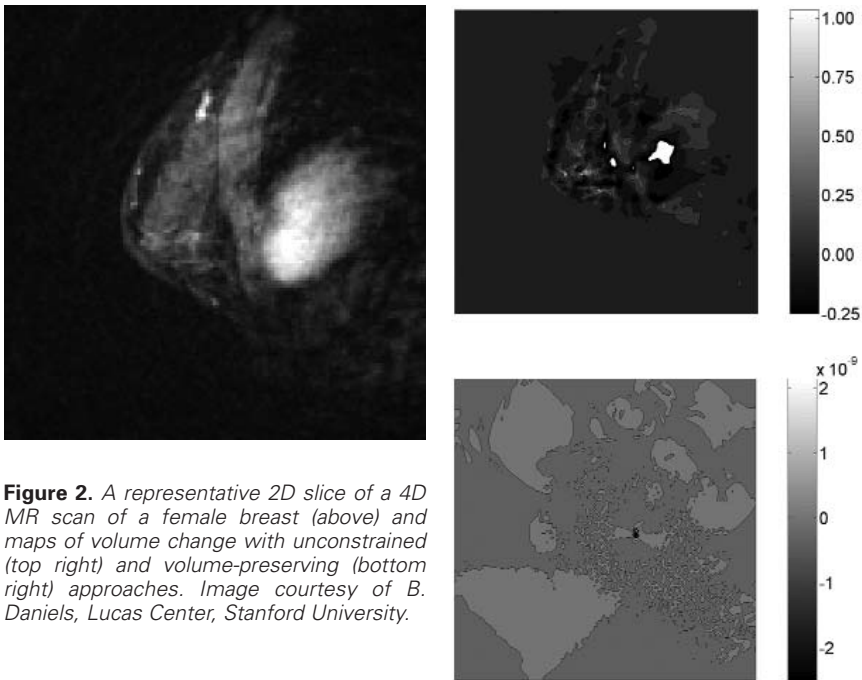

Figure 2. *A representative 2D slice of a 4D MR scan of a female breast (above) and maps of volume change with unconstrained (top right) and volume-preserving (bottom right) approaches. Image courtesy of B. Daniels, Lucas Center, Stanford University.*

Mathematical Treatment

There are many different ways to model the registration problem. We focus in this article on optimization-based methods [2], for which a generic framework is:

Minimize
$$D(T(y),R)+S(y)+\int_{\Omega}P(y)\,dx$$
subject to $y \in M$. \hfill (1)

Here, R and T denote the given "reference" and "template" images (modeled as real-valued functions over a d-dimensional domain $\Omega \subset \mathbb{R}^d$), y is the desired transformation (a d-dimensional vector field or coordinate transformation), and $T(y)$ denotes the transformed template image.

Image distance is quantified by a distance measure D (the simplest is $\|T(y)-R\|_{L_2(\Omega)}$; others are based on, for example, Kullback–Leibler distances, or Wasserstein metrics). Ideally, the distance measure mimics the eye of a physician but provides objective and reproducible results. Given that for every spatial position we supply only one piece of information (the grayscale) but ask for d quantities (the spatial transformation $y = (y_1, \dots, y_d) \in \mathbb{R}^d$), the minimization of a distance measure is obviously an ill-posed problem and regularization is inevitable (typically by rotationally invariant Sobolev seminorms). Finally, our model provides for the incorporation of constraints—soft (by a penalty P) and hard (choice of an appropriate manifold M; a trivial example is the set of affine linear transformations, $M = \{y : y = Ax + b\}$).

One of the challenging subproblems in registration is to find and model application-specific ingredients for the above optimization problem. We briefly illustrate this task for the contrast-enhanced MR mammography problem. We took an L_2 norm-based distance measure D (the underlying data is monomodal), an elastic potential as regularizer S (the female breast consists mainly of soft tissue; λ and μ describe tissue properties), and volume preservation as a hard constraint (our scheme does not allow change in the tumor volume over time). For this particular application, the ingredients of (1) read

$$D(T(y),R)=\int_{\Omega}\big(T(y(x))-R\big)^2\,dx,$$

$$S(y)=\int_{\Omega}\lambda\,|\div y|^2 +\mu\sum_{i=1}^{d}|\nabla y_i|^2\,dx,$$

$$P(y)=0$$

$$M = \{y : \det \nabla y(x) = 1,\ x \in \Omega\}.$$

We are now left with the problem of solving the resulting nonconvex optimization problem with nonlinear, differential constraints. Choosing a discretize-then-optimize approach, we perform the appropriate discretization and end up with a relatively large nonlinear optimization problem ($n = 128^3 \approx 2{,}000{,}000$ data elements, $3n$ unknowns, n constraints). Our present numerical scheme is based on an SQP approach, in which the linear systems are solved by MinRes with a Schur complement-type preconditioner and multigrid as the workhorse numerical method [3]. Multiscale, along with multilevel, strategies have been used for further regularization and efficiency [4]. Still, any suggestions for speeding up the journey would be highly appreciated.

Shared Satisfaction

Medical imaging is a fascinating research area for applied mathematicians. The challenging problems that arise not only trigger the investigation of fundamental problems in various branches of mathematics (including approximation theory, inverse problems, numerical analysis, optimization, PDEs), but at the same time provide great satisfaction for the people involved in the process. For an applied mathematician, it is quite a feeling of accomplishment to join a surgeon in planning a liver resection and then, during the procedure, to see the surgeon use images you produced. If there is a "fun factor" in science, working on medical problems that have an impact on clinical decisions certainly contributes to it. And, last but not least, medical imaging is a "sexy" topic that appeals to students.

References

[1] T.F. Chan and J.J. Shen, *Image Processing and Analysis—Variational, PDE, Wavelet, and Stochastic Methods*, SIAM, Philadelphia, 2005.

[2] B. Fischer and J. Modersitzki, *Large scale problems arising from image registration*, GAMM Mitteilungen, 27 (2004), 104–120.

[3] E. Haber and J. Modersitzki, *Numerical methods for volume preserving image registration*, Inverse Problems, 20 (2004), 1621–1638.

[4] E. Haber and J. Modersitzki, *A multilevel method for image registration*, SIAM J. Sci. Comput., 27 (2006), 1594–1607.

[5] P.C. Hansen, J.G. Nagy, and D.P. O'Leary, *Deblurring Images: Matrices, Spectra, and Filtering*, SIAM, Philadelphia, 2006.

[6] R. Kimmel, *Numerical Geometry of Images—Theory, Algorithms, and Applications*, Springer, New York, 2004.

[7] J. Modersitzki, *Numerical Methods for Image Registration*, Oxford University Press, Oxford, UK, 2004.

[8] G. Sapiro, *Geometric Partial Differential Equations and Image Processing*, Cambridge University Press, Cambridge, UK, 2001.

[9] O. Scherzer, *Mathematical Models for Registration and Applications to Medical Imaging*, Springer, New York, 2006.

[10] J. Weickert, *Anisotropic Diffusion in Image Processing*, Teubner-Verlag, Stuttgart, 1998.

Bernd Fischer is a professor and Jan Modersitzki an assistant professor (Akademischer Rat) at the Institute of Mathematics, University of Lübeck. They head the SAFIR research group on PDE-based image processing (http://www.math.uni-luebeck.de/safir/).