Image Registration Accuracy Estimation without Ground Truth using Bootstrap

Jan Kybic[1] and Daniel Smutek[2]

[1] Center for Machine Perception, Czech Technical University, Prague, Czech Republic
kybic@fel.cvut.cz, http://cmp.felk.cvut.cz/~kybic
[2] Faculty of Medicine I, Charles University, Prague, Czech Republic

Abstract. We consider the problem of estimating the local accuracy of image registration when no ground truth data is available. The technique is based on a statistical resampling technique called bootstrap. Only the two input images are used, no other data are needed. The general bootstrap uncertainty estimation framework described here is in principle applicable to most of the existing pixel based registration techniques. In practice, a large computing power is required. We present experimental results for a block matching method on an ultrasound image sequence for elastography with both known and unknown deformation field.

1 Introduction

Image registration algorithms [1–6] estimate a displacement field aligning corresponding features in two given input images. The common feature of most of these algorithms is that they provide a single-valued, deterministic answer. For each point from one image, the algorithm calculates the coordinates of one corresponding point from the other image. However, in practice, the estimate is always known only with limited accuracy; there is always an uncertainty associated with it.

Here we shall consider the problem of estimating this uncertainty in the case when no ground truth or a priori information is available, using only the two input images. The bootstrap method [7–9] will permit us to use the same data for estimation of both the deformation and the accuracy of this estimation.

1.1 Motivation

Knowing the accuracy of the registration result is always useful. The accuracy information permits us to judge whether and to what extent can the registration be trusted. It can be also used to identify optimal registration parameters and as a weighting factor for information fusion. It can also help us to determine the quality of the input data, so that we can discard or repeat unsuitable experiments.

In many cases the displacement field is an input to subsequent analysis. We are especially interested by the problem of elastography [10–12], specifically the

estimation of the elastic properties of tissues from ultrasound sequences [13, 14]. Knowing the uncertainty of the motion estimation will permit us to give more weight to reliably estimated points for the inverse problem solution.

1.2 Previous work

The precision of the estimated displacement has been studied experimentally using ground truth data, for computer vision [15–18] as well as for medical applications [19–21]. A 'bronze standard' [22] uses combined results of several distinct registration algorithms as a reference. Statistical properties of the estimation errors have been studied theoretically for rotations [23, 24]. Heuristic or Gaussian noise based uncertainty measures have been introduced for block matching [25] and optical flow estimation [26, 27]. Finally, for low-rank transformations (such as rigid motion), the covariance can be estimated a posteriori from a sufficiently high number of corresponding features [28].

1.3 Proposed approach

We shall describe a general procedure applicable to a large class of image registration algorithms so that they provide not only the point (crisp) estimate of the displacement field but also an estimate of the accuracy of this displacement. In contrast to the prior work we shall neither require additional test or ground truth data, nor restrict the allowed motion.

We consider images to be random processes and the input images as realizations of these random processes. If we had more acquisitions of the same scene under identical conditions, we could regard them as independent realizations of the random image processes. We could run our registration algorithm of choice on each realization and the statistical distribution of the results would give us an information about the accuracy of the registration.

Since we only have a single pair of images (one realization), we need to use a trick called bootstrap resampling (Section 1.4). More specifically, we shall randomly draw pixels from some neighborhood in the input images and treat them as if they came from independent realizations of the image processes (see Section 2.2 for more details). Then we can continue as above.

1.4 Bootstrap

Let us have N samples $\mathsf{X} = \{x_1, \ldots, x_N\}$ of a random variable $\mathcal{X}$. A bootstrap [7–9] data set (resample) $B_{\mathsf{X}}^{(\cdot)} = \{y_1, \ldots, y_N\}$ is constructed by randomly selecting N points from X, with replacement. Note that $B_{\mathsf{X}}^{(\cdot)}$ is a multiset (also known as a collection or a bag), which is unordered like a set but in which each element can have a multiplicity greater than 1. The multiset $B_{\mathsf{X}}^{(\cdot)}$ is constructed from the same elements as X, but it contains some of them more than once, some not at all. In bootstrap estimation, the resampling is repeated M times, yielding M data sets $B_{\mathsf{X}}^{(b)}$, which we shall treat as independent.

Let ϑ be some statistic of the random variable $\mathcal{X}$ (such as its mean or variance). We can calculate the estimate $\hat{\vartheta}(\mathsf{X})$ of ϑ from the samples X. Now, we might be interested in the distribution of $\hat{\vartheta}$, for example to check its accuracy or bias. To do this, we calculate $\hat{\vartheta}^{(b)} = \hat{\vartheta}(B_{\mathsf{X}}^{(b)})$ for each bootstrap set $B_{\mathsf{X}}^{(b)}$; the distribution $p_{\hat{\vartheta}}$ of $\hat{\vartheta}$ can then be approximated by an empirical bootstrap distribution $p_{\hat{\vartheta}}^{(*)}$, constructed from the samples $\hat{\vartheta}^{(1)}, \ldots, \hat{\vartheta}^{(M)}$.

2 Image registration

We shall define the registration problem more precisely and show how the bootstrap uncertainty estimation can be performed. We shall use a block matching algorithm as an example.

2.1 Problem formulation

Let us have two images $f, g : \mathbb{R}^m \to \mathbb{R}^n$. Let $\Omega \in \mathbb{Z}^m$ be a set of pixels from the image f. We assume that the two images are related by a transformation $T_\theta : \mathbb{R}^m \to \mathbb{R}^m$ described by a parameter vector $\theta \in \Theta \subseteq \mathbb{R}^d$. In other words, the point of coordinates $\mathbf{x}$ in image f and the point of coordinates $T_\theta(\mathbf{x})$ in image g correspond to the same physical point and as such are related by a statistical dependence. We assemble the pairs of corresponding pixels into a set S_θ:

$$S_\theta = \left\{ \left(f(\mathbf{x}), g\big(T_\theta(\mathbf{x})\big) \right) ;\ \mathbf{x} \in \Omega \right\}$$

and define a scalar similarity criterion J measuring the dependence between $f(\mathbf{x})$ and $g\big(T_\theta(\mathbf{x})\big)$ on Ω. This leads to the following estimate of the transformation parameters θ:

$$\hat{\theta} = \arg\min_{\theta \in \Theta} J(S_\theta) \tag{1}$$

The only restriction with respect to the usual formulation is that J operates on a multiset. Most existing pixel based image registration methods [29], [30–34] can be expressed in this form directly, whether the similarity criterion is the sum of square differences, correlation coefficient, or mutual information. The feature space can be extended (so that the elements of S are more than just pairs of pixel values) to accommodate methods that need information from a small pixel neighborhood [35], such as the image gradient [36]. The optical flow [37] and fluid registration methods [38] can also be cast into this framework.

2.2 Uncertainty estimation

Following the general bootstrap strategy, we make M bootstrap data (multi)sets $\Omega^{(b)}$ by resampling Ω. For each $\Omega^{(b)}$ we perform the minimization (1) yielding an estimate

$$\hat{\theta}^{(b)} = \arg\min_{\theta \in \Theta} J(S_\theta^{(b)}) \quad \text{with} \quad S_\theta^{(b)} = \left\{ \left(f(\mathbf{x}), g\big(T_\theta(\mathbf{x})\big) \right) ;\ \mathbf{x} \in \Omega^{(b)} \right\}$$

The probability distribution $p_{\hat{\theta}}$ of $\hat{\theta}$ can be approximated by an empirical bootstrap distribution $p^{(*)}_{\hat{\theta}}$ of the samples $\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(M)}$. We need to assume the validity of the bootstrap principle, i.e. that the distribution p_{S_θ} of S_θ is well approximated by the distribution $p^{(b)}_{S_\theta}$ of the bootstrap samples $S^{(b)}_\theta$, at least for θ in the vicinity of the optimum. The bootstrap principle is valid if the pixel values of f and $g \circ T_\theta$ are independent and identically distributed (i.i.d) in Ω. It is also valid if Ω can be partitioned into several sufficiently big classes such that in each of the classes the pixel values are i.i.d.

Since we want to make the number of bootstrap data sets M reasonably small to reduce the computational overhead, it is advisable to calculate only some simple statistics on $p^{(*)}_{\hat{\theta}}$ such as the bootstrap mean [9]

$$\mu^{(*)}_{\hat{\theta}} = \frac{1}{M} \sum_{b=1}^{M} \hat{\theta}^{(b)}$$

and covariance matrix

$$\mathbf{\Sigma}^{(*)}_{\hat{\theta}} = \frac{1}{M-1} \sum_{b=1}^{M} \big(\hat{\theta}^{(b)} - \mu^{(*)}_{\hat{\theta}}\big)\big(\hat{\theta}^{(b)} - \mu^{(*)}_{\hat{\theta}}\big)^T \tag{2}$$

or alternatively the perhaps the more relevant expected square error

$$\mathbf{C}^{(*)}_{\hat{\theta}} = \frac{1}{M} \sum_{b=1}^{M} \big(\hat{\theta}^{(b)} - \hat{\theta}\big)\big(\hat{\theta}^{(b)} - \hat{\theta}\big)^T$$

Note that the convergence of the above estimates in particular requires smoothness of $\hat{\theta}$ with respect to the input data, see [7] for details. The estimators can also be biased.

2.3 Geometrical error estimation

A mean squared geometrical error can be defined as

$$\varepsilon^2 = \operatorname*{mean}_{\mathbf{x}\in\Omega} \|T_{\hat{\theta}}(\mathbf{x}) - T_{\text{true}}(\mathbf{x})\|^2$$

where T_{true} is the true transformation. Assuming that $T_{\mu^{(*)}_{\hat{\theta}}}$ approximates T_{true} well, we can construct the bootstrap estimate for the expected value of ε^2

$$\mathrm{E}\left[\varepsilon^2\right] \approx e^2 = \frac{1}{M-1} \sum_{b=1}^{M} \frac{1}{\|\Omega\|} \sum_{\mathbf{x}\in\Omega} \|T_{\hat{\theta}^{(b)}}(\mathbf{x}) - T_{\mu^{(*)}_{\hat{\theta}}}(\mathbf{x})\|^2 \tag{3}$$

An estimate of the expected geometrical error of $\hat{\theta}$ is obtained by using $\hat{\theta}$ instead of $\mu^{(*)}_{\hat{\theta}}$. A geometrical error in a particular direction can be also easily constructed.

3 Block matching

To illustrate the ideas from the preceding section, we shall apply them on the well-known 2D block matching algorithm [18, 39, 40] in the context of motion estimation from an ultrasound sequence for elastography [13, 14]. The algorithm itself is not our main interest here; we have deliberately chosen to make it very simple, in order not to confuse the description and also to make the bootstrap calculations reasonably fast.

We divide the scalar 2D image ($n = 1$, $m = 2$) into a set of possibly overlapping blocks Ω_i. Here, for simplicity, we shall use rectangular blocks of size $w_x \times w_y$. The block centers $\mathbf{z}_i$ will be placed on a uniform Cartesian grid. For each Ω_i, we independently estimate the motion parameters θ_i by minimizing a standard sum of squared differences (SSD) criterion:

$$\hat{\theta}_i = \arg\min_{\theta\in\Theta} J(S_{i,\theta})$$

$$\text{with} \qquad S_{i,\theta} = \left\{ \left(f(\mathbf{x}), g\big(T_{i,\theta}(\mathbf{x})\big) \right) ;\ \mathbf{x} \in \Omega_i \right\}$$

$$\text{and} \qquad J(S_{i,\theta}) = \sum_{(f',g')} \left(f' - g' \right)^2 \quad \text{for} \quad (f', g') \in S_{i,\theta}$$

In each block, we shall search for a translation ($d = 2$)

$$T_\theta(\mathbf{x}) = \mathbf{x} + \theta \tag{4}$$

This way, we shall have in each block Ω_i information about the local translation.

3.1 Interpolation

The image g is interpolated using uniform cubic B-spline interpolation [41–43] which provides continuous derivatives and good approximation properties. The B-spline coefficients are calculated beforehand.

3.2 Minimization

We use a quasi-Newton type optimizer [44] — second–order information (Hessian matrix H) is used to obtain quadratic convergence. We iteratively update the estimate of H^{-1} using the BFGS strategy [44, 45]. The minimum is only sought within an a priori chosen interval of 'reasonable' values of θ.

In our application we know that the displacement at the top edge should be zero, since the ultrasound probe is in contact with the tissue there. We go through the image from top to bottom and from center to the left and to the right [46], using the value of θ found in the block immediately above or next to the current one as a starting guess.

3.3 Bootstrap accuracy estimation

For each block we evaluate the expected geometrical error e^2 (3) which in the translation case (4) simplifies to $e^2 = \operatorname{tr} \boldsymbol{\Sigma}^{(*)}_{\hat{\theta}}$ with the correlation estimator $\boldsymbol{\Sigma}^{(*)}_{\hat{\theta}}$ defined by (2).

4 Experiments

Experiments were performed on a sequence of ultrasound images acquired by a standard Philips Envisor scanner for the elastography experiments. We have used the Gammex 429 Ultrasound Biopsy Phantom[3] that mimics normal tissue and contains eleven test objects filled with low or high density gel, simulating lesions. The movement between images of the sequence is caused by varying the pressure applied on the ultrasound probe.

A key parameter for our bootstrap uncertainty estimation is the number of bootstrap resampling M. According to the literature [7], to estimate simple statistics such as correlation, $M = 50$ is often enough. Hence, we ran the following experiments with $M = 10$ and $M = 100$. Surprisingly, there was a little difference between the results. All results presented here were calculated with $M = 10$ which appears to be enough for the present application.

The algorithm was implemented in the Ocaml language[4]. Our experience suggests that a reimplementation in an optimized C code might bring another $30 \sim 50\,\%$ speed-up. The times reported are on a computer with a Pentium M 1400 MHz processor.

4.1 Synthetic displacement field

To have the ground truth information available, we have generated an artificial displacement field (Figure 1) that attempts to be simple yet similar to displacement fields encountered in real data. The right half of the images is undeformed, to create an abrupt transition. In the left half the vertical displacement increases linearly from top to bottom. The horizontal displacement increases linearly from top to bottom and from right to left; in addition there is a low amplitude harmonic component.

The block matching algorithm (Section 3) with bootstrap uncertainty estimation was run on the original grayscale images (Figure 1, top) of size 541×426 pixels. We used overlapping blocks of 29×29 pixels whose centers lie on a uniform grid with spacing of 5×5 pixels. In Figure 2 we show the recovered displacements (to be compared with Figure 1, top), the estimated and the true geometrical error.

Note that even though the amplitude of the error has not been estimated exactly, the results are in the right order of magnitude. Perhaps more importantly, the algorithm has correctly identified the areas where errors are likely — around the vertical edge where the motion is discontinuous and around the 'lesion' area where texture is missing. The registration took about 3 hours.

4.2 Window size

The bootstrap accuracy estimate permits us to optimize the parameters of the registration itself, for example the window size. Figure 3 shows the dependency

[3] `www.gamex.com`

[4] `http://caml.inria.fr`

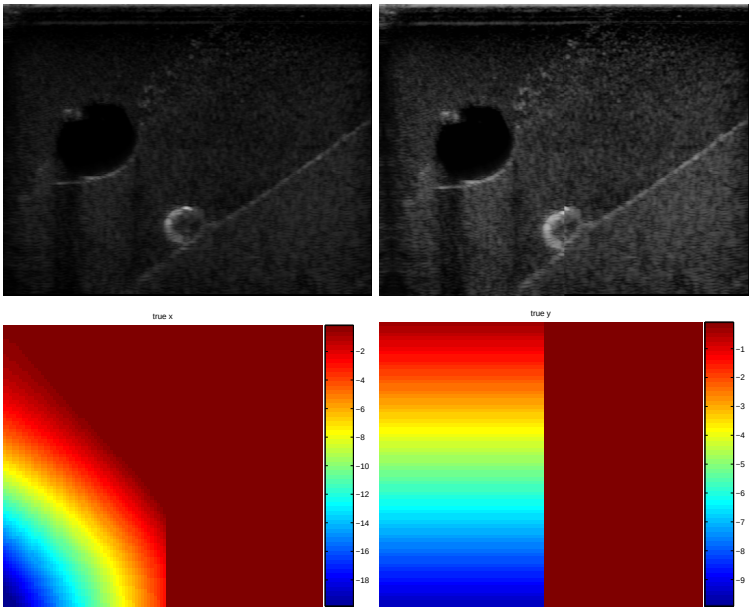

Fig. 1. The original image (top left) and its deformed version (top right). The horizontal and vertical components of the artificial deformation (bottom left and bottom right, respectively).

of the real and estimated geometric error on the window size. Even though the error estimation is not exact, it is nevertheless quite close to the reality and it permits us to correctly identify the trade-off between a window too small which provides a noisy estimate, and a window too big in which the movement is not sufficiently homogeneous and cannot be explained by a translation. One reason for the disparity is that the estimator incorrectly assumes that the true transformation is a translation. The whole experiment took 8 min.

4.3 Real images

Finally, we have taken two images from the ultrasound sequence a few frames apart and applied the registration algorithm on them (Figure 4) with the same parameters as in Section 4.1. The recovered displacements seem to be realistic and the uncertainty estimation correctly identified the problematic zone of the 'lesion' without texture. The calculation took also about 3 hours.

5 Conclusions

We have described a bootstrap-based uncertainty estimation procedure that can be applied to a wide class of pixel-based image registration algorithms. Unlike

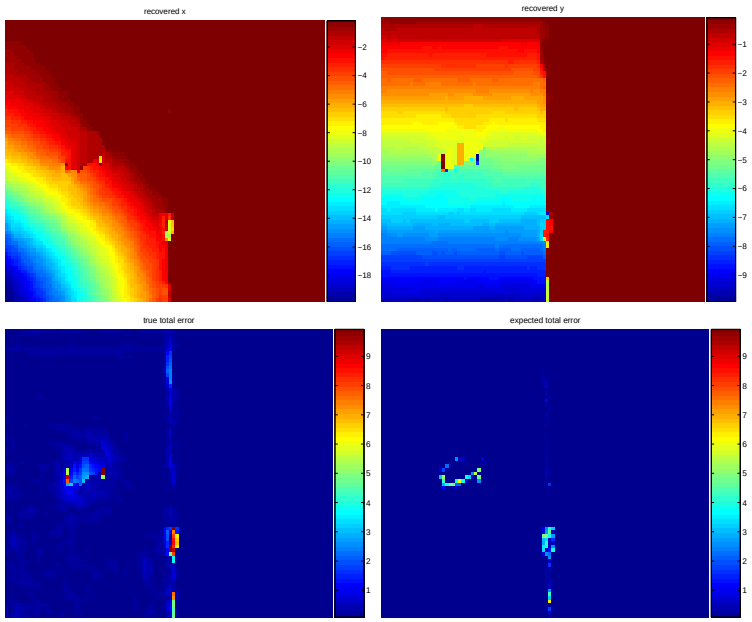

Fig. 2. The estimated horizontal and vertical displacements (top left and top right, respectively). The true geometrical error ε and its bootstrap estimate e (bottom left and bottom right, respectively). All values are in pixels.

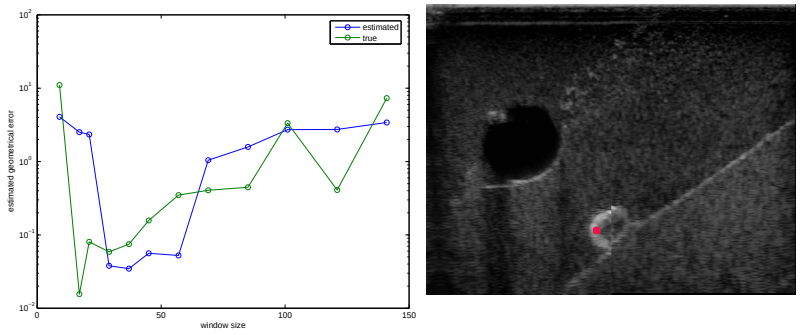

Fig. 3. The dependence of the true end expected geometric error (ε resp. e) on the window size in pixels (left). The location of the window center is marked by the red dot on the deformed image (right).

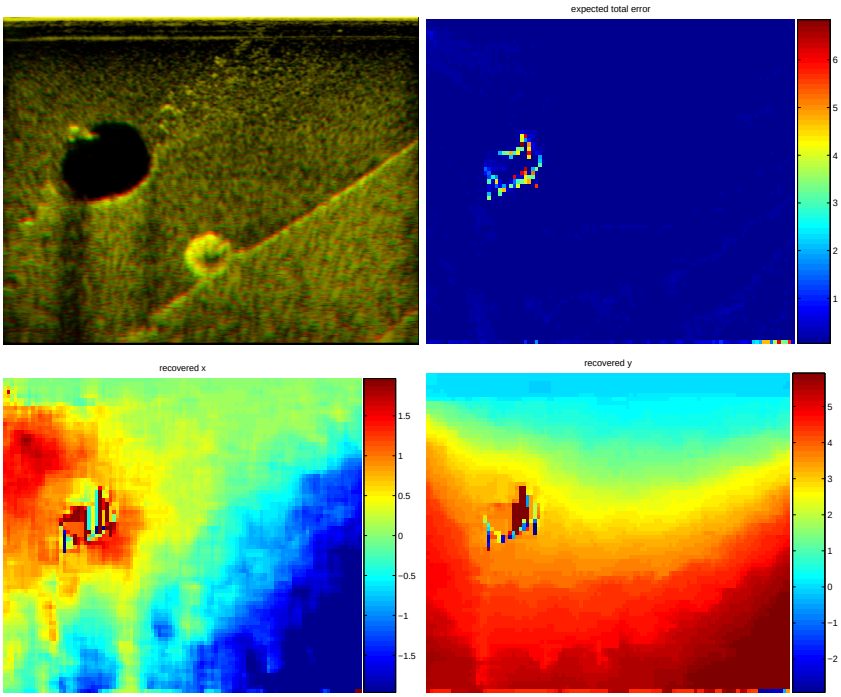

Fig. 4. Two images from the ultrasound sequence (top left, in green and red channels) and the recovered horizontal and vertical displacement (bottom left an right, respectively). The expected geometrical error e (top right).

any other registration algorithm known to us, it can estimate the registration accuracy using only the input images, yet it does not limit the class of the representable transformations, nor it assumes a specific noise or image intensity distribution. We have shown experimentally that the uncertainty estimates produced by the algorithm are reasonable.

Bootstrap algorithms can take full advantage of the abundant computing power available today. Even though for the simple registration algorithm presented here the computational time was just about acceptable, in more complex cases parallel processing might be advisable.

The algorithm presented here is only a proof of concept, there is a number of ways how to extend it, some of which were already mentioned in the introduction. Finally, note that the algorithm can only estimate the stochastic part of the error inherent to the registration itself. Therefore, even though we would be happy to see our uncertainty estimation incorporated into as many registration algorithms as possible, it is clear that there will always be a need for a careful experimental end-to-end testing of the complete registration chain.

Acknowledgements

The authors were sponsored by the Grant Agency of the Czech Academy of Sciences, grant 1ET101050403.

References

1. Brown, L.: A survey of image registration techniques. ACM Computing Surveys **24**(4) (1992) 326–376
2. Lester, H., Arridge, S.R.: A survey of hierarchical non-linear medical image registration. Pattern Recognition **32**(1) (1999) 129–149
3. Maintz, J., Viergever, M.A.: A survey of medical image registration. Medical Image Analysis **2**(1) (1998) 1–36
4. Pluim, J., Maintz, J.B.A., Viergever, M.A.: Mutual-information-based registration of medical images: A survey. IEEE Transactions on Medical Imaging **22**(8) (2003) 986–1004
5. van den Elsen, P.A., Pol, E.J.D., Viergever, M.A.: Medical image matching—A review with classification. IEEE Engineering in Medicine and Biology (1993) 26–39
6. Zitová, B., Flusser, J.: Image registration methods: a survey. Image and Vision Computing (21) (2003) 977–1000
7. Efron, B., Tibshirani, R.: An Introduction to the Bootstrap. Number 57 in Monographs on Statistics and Applied Probability. Chapman & Hall, CRC (1993)
8. Zoubir, A.M., Boashash, B.: The bootstrap and its applications in signal processing. IEEE Signal Processing Magazine (1998) 56–76
9. Duda, R.O., Hart, P.E., Stork, D.G.: Pattern classification. 2nd edn. Wiley Interscience Publication. John Wiley, New York (2001)
10. Ophir, J., Kallel, F., Varghese, T., Konofagou, E., Alam, S.K., Garra, B., Krouskop, T., Righetti, R.: Elastography: Optical and acoustic imaging of acoustic media. C. R. Acad. Sci. Paris **2**(8) (2001) 1193–1212 serie IV.
11. Ophir, J., Alam, S., Garra, B., Kallel, F., Konofagou, E., Krouskop, T., Varghese, T.: Elastography: ultrasonic estimation and imaging of the elastic properties of tissues. Proc. Instn Mech Engrs **213** (1999) 203–233
12. Washington, C.W., Miga, M.I.: Modality independent elastography (MIE): A new approach to elasticity imaging. IEEE Transactions on Medical Imaging (2004) 1117–1126
13. Kybic, J., Smutek, D.: Estimating elastic properties of tissues from standard 2d ultrasound images. In Walker, W.F., Emelianov, S.Y., eds.: Medical Imaging 2005: Ultrasonic Imaging and Signal Processing. Volume 6 of Progress in Biomedical Optics and Imaging., Bellingham, Washington, USA, SPIE (2005) 184–195
14. Kybic, J., Smutek, D.: Computational elastography from standard ultrasound image sequences by global trust region optimization. In Šonka, M., Christensen, G., eds.: Proceedings of IPMI, Lecture Notes in Computer Science. Number 3565, Heidelberg, Germany, Springer Verlag (2005) 299–310
15. Fang, J., Huang, T.S.: Some experiments on estimating the 3-D motion parameters of a rigid body from two consecutive image frames. IEEE Trans. Pattern Anal. Mach. Intell. **3**(65) (1984)
16. Snyder, M.A.: The precision of 3-D parameters in correspondence based techniques: the case of uniform translational motion in rigid environment. IEEE Trans. Pattern Anal. Mach. Intell. **5**(11) (1998) 523–528

17. Haralick, R.M., Joo, H., Lee, C.N., Zhuang, X., Vaidya, V.G., Kim, M.B.: Pose estimation from corresponding point data. IEEE Trans. Systems, Man and Cybernetics **6**(19) (1989) 1426–1446
18. Davis, C.Q., Freeman, D.M.: Statistics of subpixel registration algorithms based on spatiotemporal gradients or block matching. Optical Engineering **4**(37) (1998) 1290–1298
19. Maurer, C.J., Fitzpatrick, J., Wang, M., Galloway, R.J., Maciunas, R., Allen, G.: Registration of head volume images using implantable fiducial markers. IEEE Transactions on Medical Imaging **16**(4) (1997)
20. West, J., Fitzpatrick, J.M., Wang, M.Y., Dawant, B.M., Maurer, C.R.J., Kessler, R.M., Maciunas, R.J., Barillot, C., Lemoine, D., Collignon, A., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P.A., Napel, S., Sumanaweera, T.S., Harkness, B., Hemler, P.F., Hill, D.L.G., Hawkes, D.J., Studholme, C., Maintz, J.B.A., Viergever, M.A., Malandain, G., Pennec, X., Noz, M.E., Maguire, G.Q.J., Pollack, M., Pelizzari, C.A., Robb, R.A., Hanson, D., Woods, R.P.: Comparison and evaluation of retrospective intermodality brain image registration techniques. Journal of Computer Assisted Tomography **21**(4) (1997) 554–568
21. Jannin, P., Fitzpatrick, J., Hawkes, D., Pennec, X., Shahidi, R., Vannier, M.: Validation of medical image processing in image-guided therapy. IEEE Trans. on Medical Imaging **21**(12) (2002) 1445–1449
22. Nicolau, S., Pennec, X., Soler, L., Ayache, N.: Evaluation of a new 3D/2D registration criterion for liver radio-frequencies guided by augmented reality. In: Intl. Symp. on Surgery Sim. and Soft Tissue Model. (2003) 270–283
23. Kanatani, K.: Geometric computation for machine vision. Oxford University Press, Inc., New York, NY, USA (1993)
24. Kanatani, K.: Analysis of 3-d rotation fitting. IEEE Trans. Pattern Anal. Mach. Intell. **16**(5) (1994) 543–549
25. Anandan, P.: A computational framework and an algorithm for the measurement of visual motion. International Journal of Computer Vision **2**(3) (1989) 283–310
26. Heeger, D.J.: "optical flow using spatiotemporal filters". International Journal of Computer Vision **1**(4) (1988) 279–302
27. Simoncelli, E.P., Adelson, E.H., Heeger, D.J.: Probability distributions of optical flow. In: Proc Conf on Computer Vision and Pattern Recognition, Mauii, Hawaii, IEEE Computer Society (1991) 310–315
28. Pennec, X., Thirion, J.P.: A framework for uncertainty and validation of 3d registration methods based on points and frames. International Journal of Computer Vision **25**(3) (1997) 203–229
29. Kybic, J., Unser, M.: Fast parametric elastic image registration. IEEE Transactions on Image Processing **12**(11) (2003) 1427–1442
30. Thévenaz, P., Ruttimann, U.E., Unser, M.: A pyramid approach to subpixel registration based on intensity. IEEE Transactions on Image Processing **7**(1) (1998) 1–15
31. Thévenaz, P., Unser, M.: Optimization of mutual information for multiresolution image registration. IEEE Transactions on Image Processing **9**(12) (2000) 2083–2099
32. Szeliski, R., Coughlan, J.: Spline-based image registration. International Journal of Computer Vision **22** (1997) 199–218
33. Maes, F., Collignong, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging **16**(2) (1997) 187–198

34. Rösch, P., Weese, J., Netsch, T., Quist, M., Penney, G., Hill, D.: Robust 3D deformation field estimation by template propagation. In: Proceedings of MICCAI. (2000) 521–530
35. Rueckert, D., Clarkson, M.J., Hill, D.L.G., Hawkes, D.J.: Non-rigid registration using higher-order mutual information. In: Proceedings of SPIE Medical Imaging 2000: Image Processing. (2000) 438–447
36. Pluim, J.P.W., Maintz, J.B.A., Viergever, M.A.: Image registration by maximization of combined mutual information and gradient information. IEEE Transactions Med. Imag. **19**(8) (2000)
37. Horn, B., Schunck, B.: Determining optical flow. Artificial Inteligence **17** (1981) 185–203
38. Bro-Nielsen, M., Gramkow, C.: Fast fluid registration of medical images. In Höhne, K.H., Kikinis, R., eds.: Visualization in Biomedical Computing. Springer-Verlag (1996) 267–276
39. Morsy, A., Ramm, O.: 3D ultrasound tissue motion tracking using correlation search. IEEE Trans. Ultr. Ferro. & Freq. Cont. **20** (1998) 151–159
40. Ourselin, S., Roche, A., Prima, S., Ayache, N.: Block matching: A general framework to improve robustness of rigid registration of medical images. In: Third International Conference on Medical Robotics, Imaging And Computer Assisted Surgery (MICCAI 2000). Number 1935 in Lectures Notes in Computer Science, Springer (2000) 557–566
41. Unser, M., Aldroubi, A., Eden, M.: B-Spline signal processing: Part I—Theory. IEEE Transactions on Signal Processing **41**(2) (1993) 821–833 IEEE Signal Processing Society's 1995 best paper award.
42. Unser, M., Aldroubi, A., Eden, M.: B-Spline signal processing: Part II—Efficient design and applications. IEEE Transactions on Signal Processing **41**(2) (1993) 834–848
43. Unser, M.: Splines: A perfect fit for signal and image processing. IEEE Signal Processing Magazine **16**(6) (1999) 22–38
44. Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes in C. Second edn. Cambridge University Press (1992)
45. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale minimization. Mathematical Programming (45) (1989) 503–528
46. Brusseau, E., Fromageau, J., Rognin, N., Delacharte, P., Vray, D.: Local estimation of RF ultrasound signal compression for axial strain imaging: theoretical developments and experimental design. IEEE Engineering in Medicine and Biology Magazine **21**(4) (2002) 86–94