Beauty with Variational Methods: An Optic Flow Approach to Hairstyle Simulation

Oliver Demetz, Joachim Weickert, Andrés Bruhn, and Martin Welk

Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Building E1.1
Saarland University, 66041 Saarbrücken, Germany
{demetz,weickert,bruhn,welk}@mia.uni-saarland.de

Abstract. Although variational models offer many advantages in image analysis, their successful application to real-world problems is documented only for some specific areas such as medical imaging. In this paper we show how well-adapted variational ideas can solve the problem of hairstyle simulation in a fully automatic way: A customer in a hairdresser's shop selects a new hairstyle from a database, and this hairstyle is automatically registered to a digital image of the customer's face. Interestingly already a carefully modified optic flow method of Horn and Schunck turns out to be ideal for this application. These modifications include an extension to colour sequences, an incorporation of warping ideas in order to allow large deformation rates, and the inclusion of shape information that is characteristic for human faces. Employing classical numerical ideas such as finite differences and SOR iterations offers sufficient performance for real-life applications. In a number of experiments we demonstrate that our variational approach is capable of solving the hairstyle simulation problem with high quality in a fully practical setting.

1 Introduction

By hairstyle simulation a customer in a hairdresser's shop can see herself wearing other hairstyles and decide whether a hairstyle suits herself or not. This process combines two images into one. The images being combined are on one hand the image of a face, and on the other hand an image of a hairstyle. The image containing the hairstyle is transparent/semi-transparent in areas where no hairs are shown.

The main problem of hairstyle simulation is not the combination itself. It is located in the fact that in general the proportions and the position of the hairstyle in its image will not fit the face. Furthermore, simple linear scalings do not solve the problem. Conventional software that is used in a hairdresser's shop requires a lot of manual interaction in order to register the desired hairstyle to a photo of the customer's face. This procedure is time-consuming and for hairdressers who lack experience with handling such a software, the results are not always optimal. Thus it would be desirable to exploit computer vision ideas in order to create a fully automised hairstyle simulation that does not require any interaction and creates a registration of high quality.

Addressing this problem is the goal of the current paper. Our work aims at computing a transformation, that, applied onto the hairstyle, makes a simple hairstyle combination possible. In order to avoid to recalculate this transformation again for each new hairstyle, all offered hairstyles have to be equally proportioned and positioned. To achieve this, they have been registered onto one reference face before. In this way a single mapping from the reference face to the customer face is sufficient to adapt any hairstyle to the customer's face. The corresponding displacement field is the transformation we are searching for. Figure 1 illustrates this procedure.

Fig. 1. Hairstyle simulation. **(a) Top left:** Reference face. **(b) Top middle:** Example of a hairstyle. Grey-chequered regions show transparency. **(c) Top right:** Hairstyle combined with the reference face. **(d) Bottom left:** Example of a customer face. **(e) Bottom middle:** Mesh representation of the resulting deformation flow field between the reference face (a) and the customer face (d). **(f) Bottom right:** Customer face combined with the new hairstyle transformed according to (e).

From a computer vision perspective, hairstyle simulation requires to find correspondences between two images. One searches for a displacement field that is used to deform one image onto the other. To this end it is highly desirable to compute a deformation field that is dense. Since it can happen that the local information is not sufficient to establish the desired correspondences, it is natural to apply variational methods that benefit from the filling-in effect of the smoothness term. Such approaches have been successfully used for computing optic flow fields in image sequences [10, 17, 8, 18], for estimating the disparity in stereo reconstruction problems [13, 12, 1], and for calculating the deformation rates for medical image registration [2, 7, 15, 19]. To our knowledge, however, there is no application of these techniques within the beauty and wellness industries. In our paper we shall see that the variational optic flow approach of Horn and Schunck [10] constitutes an ideal starting point to tackle the problem of hairstyle simulation if it is modified in order to meet some specific requirements: It has to incorporate colour information, it must be able to handle large displacements, and it must be tunable to the shape characteristics of human faces.

Our paper is organised as follows. In Section 2 we describe the basic variational optic flow model that is used for our approach. In Section 3 we discuss modifications that make the model more robust against non-optimal image acquisition. In order to cope with large deformation rates, a multiscale setting is used in which a warping strategy is applied when going from a coarser to a finer scale. This is described in Section 4. In Section 5 we sketch how we minimise the energy functional at each scale by computing the Euler-Lagrange equations, discretising them and solving the corresponding linear systems of equations in an iterative manner. Experimental results are presented in Section 6. We conclude the paper with a summary in Section 7.

2 Basic Variational Model

Our variational approach is based on an interpretation of the hairstyle simulation problem as an optic flow problem. We would like to find a displacement field that maps a reference face to a customer face. To this end we interpret these images as two subsequent frames in an image sequence. Then the displacement field between corresponding structures is nothing else than the optic flow field in this sequence. From a practical viewpoint it would be desirable to have a dense flow field within a rotationally invariant setting. This naturally suggests the use of continuous variational methods.

The oldest and simplest variational method for optic flow estimation goes back to Horn and Schunck [10]. Assume we are given some sequence $f(x, y, t)$ of greyscale images, where (x, y) denote the location and t is the time. The method of Horn and Schunck computes the optic flow field $(u, v)^\top = (u(x, y, t), v(x, y, t))^\top$ as the minimiser of the convex energy functional

$$E(u, v) \;=\; \int_\Omega \Big((f_x u + f_y v + f_t)^2 + \alpha(|\nabla u|^2 + |\nabla v|^2) \Big)\, dx\, dy \qquad (1)$$

where Ω is the rectangular image domain, subscripts denote partial derivatives, and ∇ is the spatial nabla operator. The first term of this energy functional is a data term that reflects the assumption that corresponding structures do not change their greyvalues over time (optic flow constraint), while the second term penalises spatial fluctuations of the flow field in a quadratic way. The positive regularisation parameter α determines the required amount of smoothness.

A colour image sequence may be regarded as a vector-valued function $f : \Omega \times [0, \infty) \to \mathbb{R}^3$ where the red, green and blue channels serve as components of

$f = (f_1, f_2, f_3)^\top$. If one searches for a joint optic flow field for all three channels and imposes colour constancy during the motion of image features, cf. e.g. [9], the Horn and Schunck method can be extended to

$$E(u,v) \;=\; \int_\Omega \Big(\sum_{i=1}^{3} (f_{i,x} u + f_{i,y} v + f_{i,t})^2 + \alpha (|\nabla u|^2 + |\nabla v|^2) \Big)\, dx\, dy. \qquad (2)$$

Since we work with colour images, we will make use of this adaptation.

3 Shape Weighting and Preregistration

So far we have assumed that there is a reasonable match between structures from both images throughout the image domain Ω. While this condition appears fairly reasonable for the face region, it can be expected to fail in rather common situations: Objects in the background may lead to severe mismatches as well as clothes details and jewellery worn by the customer.

This problem is solved by incorporating a-priori knowledge. We modify our model (2) such that it takes into account which parts of the image are really important to match. The face must be matched accurately while background regions are of no interest and should not influence the algorithm. To this end, we introduce a space-variant weight function $w(x,y)$ that attains high values at the face region and low values in other image areas. This leads to

$$E(u,v) \;=\; \int_\Omega \Big(w \sum_{i=1}^{3} (f_{i,x} u + f_{i,y} v + f_{i,t})^2 + \alpha (|\nabla u|^2 + |\nabla v|^2) \Big)\, dx\, dy. \qquad (3)$$

This energy functional will be the main ingredient of our hairstyle simulation approach. Interestingly it turns out that typical features of more sophisticated optic flow models (such as discontinuity-preserving regularisers and other constancy assumptions; see e.g. [18]) are not required for our present task.

The downside of the weighting approach is that, in order to apply the mask appropriately, the customer and reference faces need to be adjusted to equal position and size beforehand, because otherwise some parts of the customer face might wrongly be excluded from the matching process. That is, the adjustment already requires to match the faces. To escape this vicious circle, we make use of a remarkable observation: There are two specific details which are already reliably matched without the weighting, namely the eyes. Knowing the eye coordinates in the reference image, we can locate the eyes in the customer face via the displacement field. From these, not only the position of the customer's face but even (via the eye distance) its approximate size can be inferred. The customer image is then shifted and linearly rescaled such that the face regions approximately coincide. Running the displacement computation again, with the weighting mask, yields the final flow field.

Another source of unmatched details are illumination differences between reference and customer image. While global illumination changes turn out to be rather uncritical, the positioning of light sources constitutes a radical distinction between lighting conditions in professional studios and in "normal" environments. In a studio setting, flashlights are installed also above and next to the model, thus widely eliminating shadows. In a typical non-studio situation there is just one flashlight located above the lens of the camera, causing shadows in the background as well as, in some cases, in the face itself. Moreover, overexposed areas tend to occur near the eyebrows. To rule out the influence of these differences, we choose a pragmatic solution: We ensure that the reference image is taken under similar lighting conditions as the customer images.

Focussing at the minimisation process, it should be noted that (3) incorporates a data term that involves temporal derivatives of the image sequence. If subsequent frames are fairly different – as can be the case if they are given by the images of two different faces, and even more in the preregistration step when their location and scale may differ substantially – local derivative approximations cannot capture these large displacements anymore. A remedy to this problem is the warping approach that is described in the next section.

4 Warping

Warping approaches are frequently used techniques to handle large displacements in the context of motion estimation [3, 14, 5] or medical image registration [11, 15]. By incrementally computing the desired displacement field over multiple scales, they allow to decompose an originally difficult matching problem with large displacements into a series of simpler ones that only require the estimation of small displacements. In general, this decomposition is achieved via a coarse-to-fine strategy: Starting from a down-scaled version of the original problem, the image data and the solution are refined step by step. Thereby, previous solutions from coarser scales are used to compensate the image data for the already computed displacement field. As a consequence, only a small incremental deformation has to be estimated at each scale: the displacement field for the remaining difference problem. Once all these increments have been computed, they have to be summed up in order to form the final solution.

Let us now discuss how such an incremental coarse-to-fine strategy can actually be applied in the context of our hairstyle simulation problem. To this end, we have to define a pyramidal (hierarchical) representation of our image sequence f that contains both the customer and the reference face:

$$f^0 \rightarrow f^1 \rightarrow \ldots \rightarrow f^{k_{\text{orig}}-1} \rightarrow f^{k_{\text{orig}}} . \qquad (4)$$

Here, $k = 0$ denotes the coarsest (initial) scale of the coarse-to-fine process, while $k = k_{\text{orig}}$ stands for the finest (original) one. The ratio between two consecutive scales k and $k-1$ is given by a factor η in the interval $(0, 1)$. One should note that this hierarchical representation is also required for our weighting mask w. Moreover, since we are interested in an incremental computation of the results, we also have to decompose the deformation field u and v at each scale into

$$u^k = u^{k-1\to k} + \delta u^k \,, \tag{5}$$
$$v^k = v^{k-1\to k} + \delta v^k \,. \tag{6}$$

Here, $(u^{k-1\to k}, v^{k-1\to k})^\top$ is the already computed (and up-scaled) overall displacement field from the previous scale $k-1$, and $(\delta u^k, \delta v^k)^\top$ is the unknown deformation increment that we are looking for at the current scale k. At the coarsest scale the already known displacement field is initialised with zero, i.e. $(u^{-1\to 0}, v^{-1\to 0})^\top := (0,0)^\top$.

Due to the previous definitions we are now in the position to formulate our coarse-to-fine warping approach for the hairstyle simulation problem. It is given by the following three steps that have to be performed at each scale:

1) *Setting up the Difference Problem.* In order to derive the difference problem for the current scale, both the customer face $f(x,y,t)$ and the reference face $f(x,y,t+1)$ have to be scaled down. Moreover, the reference face has to be compensated by the already computed deformation field $(u^{k-1\to k}, v^{k-1\to k})^\top$ from the previous scale. Thus, at each scale, the following task remains to be solved: Find the deformation increment $(\delta u^k, \delta v^k)^\top$ that describes the mapping between $f^k(x,y,t)$ and $f^k(x+u^{k-1\to k}, y+v^{k-1\to k}, t+1)$.

2) *Solving the Difference Problem.* In order to compute this deformation increment, we propose the minimisation of the following energy functional:

$$\begin{aligned} E^k(\delta u^k, \delta v^k) \;=\; & \int_\Omega \Bigg(w^k \sum_{i=1}^{3} \left(f_{i,x}^k\, \delta u^k + f_{i,y}^k\, \delta v^k + f_{i,t}^k\right)^2 \\ & + \alpha \left(|\nabla (u^{k-1\to k} + \delta u^k)|^2 + |\nabla (v^{k-1\to k} + \delta v^k)|^2 \right) \Bigg) \, dx\, dy \,. \end{aligned} \tag{7}$$

As one can easily verify, this energy functional can be obtained from (3) by substituting (5)–(6). One should note that the overall displacement field $(u^{k-1\to k}, v^{k-1\to k})^\top$ from the previous scale must not occur explicitly in the data term: By using the deformation-compensated reference face for computing the temporal derivatives $f_{i,t}^k$, this displacement field is already considered implicitly in the formulation of the data term.

3) *Updating the Overall Displacement Field.* The final step at each scale is the update of the overall solution $(u^k, v^k)^\top$. To this end, the computed deformation increment has to be added to the overall displacement field from the previous scale (cf. (5)–(6)).

As soon as the update step on the finest (original) scale has been performed, the overall solution is obtained. However, due to the recursive definition of (5)–(6) it can also be computed by a simple summation of all deformation increments:

$$u = \sum_{i=0}^{k_{\mathrm{orig}}} \delta u^{i\to k_{\mathrm{orig}}} \,, \qquad v = \sum_{i=0}^{k_{\mathrm{orig}}} \delta v^{i\to k_{\mathrm{orig}}} \,. \tag{8}$$

After we have presented our warping approach for the hairstyle simulation problem in detail, let us now discuss its numerical realisation. This shall be done in the next section.

5 Numerical Realisation

In each level of our warping procedure we have to minimise an energy of type (7). At level k this convex problem has a unique minimiser that satisfies the Euler-Lagrange equations

$$J_{1,1}^k\,\delta u^k + J_{1,2}^k\,\delta v^k + J_{1,3}^k \;-\; \alpha\,\Delta\left(u^{k-1\to k} + \delta u^k\right) = 0, \tag{9}$$

$$J_{1,2}^k\,\delta u^k + J_{2,2}^k\,\delta v^k + J_{2,3}^k \;-\; \alpha\,\Delta\left(v^{k-1\to k} + \delta v^k\right) = 0 \tag{10}$$

where $\Delta = \partial_{xx} + \partial_{yy}$ is the spatial Laplace operator, and

$$J^k \;:=\; (J_{n,m}^k) \;:=\; \sum_{i=1}^{3} w_k \begin{pmatrix} f_{i,x}^k f_{i,x}^k & f_{i,x}^k f_{i,y}^k & f_{i,x}^k f_{i,t}^k \\ f_{i,y}^k f_{i,x}^k & f_{i,y}^k f_{i,y}^k & f_{i,y}^k f_{i,t}^k \\ f_{i,t}^k f_{i,x}^k & f_{i,t}^k f_{i,y}^k & f_{i,t}^k f_{i,t}^k \end{pmatrix} \tag{11}$$

denotes a generalisation of the structure tensor from [4]. These continuous equations are discretised by finite difference approximations (see e.g. [16]). We use Sobel operators for discretising the spatial derivatives within the structure tensor, and we replace the temporal derivatives by differences between both frames.

This discretisation leads to a large linear system of equations, where the system matrix is sparse, symmetric and positive definite. In this case, an iterative SOR scheme is a simple and efficient convergent method [20].

In order to calculate the temporal derivatives, we have to compensate the reference frame for the deformation already computed at the coarser level. To this end we perform backward registration and evaluate values between the grid points by bilinear interpolation. Down- and upsampling within the warping scheme is done by area-based averaging and interpolation, respectively [6].

Typical computation times for images of 512×512 on a 3.2 GHz Pentium 540 with the previously described numerical scheme are in the order of 10 seconds.

6 Experiments

In this section, we present results of our approach in order to assess its usefulness. Beyond the fundamental correctness and applicability, emphasis will be put on aspects of robustness and suitability under practical conditions, i.e. whether the model can cope with the problems that occur upon the everyday usage in a hairdresser's shop. The section is split into two parts. First, we apply the basic method with just one displacement estimation run without weighting to images taken in a studio setting. This ideal situation serves to demonstrate the fundamental applicability of our approach, and discuss important parameter settings. After that, we show how the complete method copes with images taken under realistic conditions.

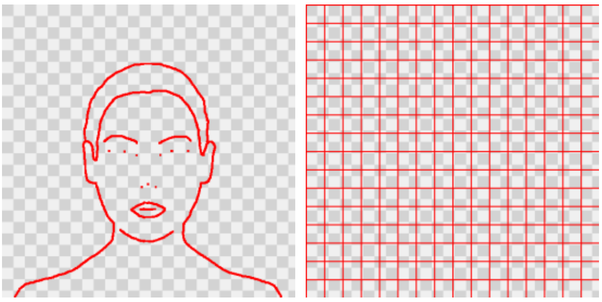

Fig. 2. Two synthetic hairstyles. Grey chequered regions show transparency. **(a) Left:** Contour representation. **(b) Right:** Mesh representation.

6.1 Test Scenario and Parameter Settings

In all tests, we use images of 512×512 pixels with 24 bit colour depth. The semi-transparent hairstyle images are given at the same resolution but with 32 bit colour depth due to the additional alpha channel. The combination of faces with hairstyles is performed by linear blending. For analytic purposes, we also use synthetic hairstyle images showing grids and contours as illustrated in Figure 2. With this image material, our scheme from Section 5 always converges to a steady state after 100 SOR iterations.

For the scaling parameter in our warping strategy, the most intuitive choice would be $\eta = \frac{1}{2}$. In practice, however, larger values such as $\eta = \frac{2}{3}$ yield better results: They slow down the refinement and thus reduce the gaps between scales.

While accuracy is a central requirement in estimating the displacement field, it has to be balanced with the conflicting goal of smoothness. The latter is equally important because the computed vector field is used to deform images. Discontinuities or steep changes would imply visible gaps or offsets in the hairstyle image, thus leading to substantially degraded results. The balance between accuracy and smoothness is controlled via the regularisation parameter α. Its influence is studied in Figure 3. If α is chosen too small, the above-mentioned discontinuities and gaps are observed. Besides, the filling-in effect is less pronounced in this case. In practice, this is observed for $\alpha < 10\,000$. Too large values for α, in contrast, lead to oversmoothed displacement fields that do not adapt well to image structures. Experimentally, one observes for $\alpha > 50\,000$ visible mismatches between face and hairstyle. A good compromise is achieved for $\alpha = 30\,000$.

Having discussed the choice of the model parameters, we demonstrate in Figure 4 that our model is capable of fitting hairstyles to different faces in good quality.

6.2 Experiments on Real-World Imagery

While the applicability of our approach is evident, we turn now to consider the robustness with respect to the perturbations that are common in realistic situations, such as illumination differences, noise, or objects in the background.

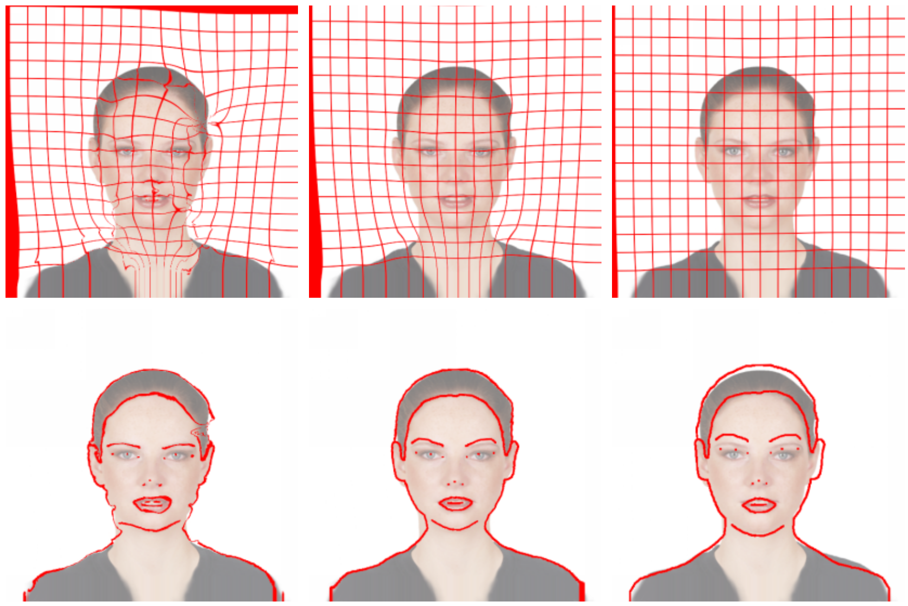

Fig. 3. Sensitivity of the solution w.r.t. the regularisation parameter α. **From left to right:** $\alpha = 1\,000$, $\alpha = 30\,000$, $\alpha = 500\,000$. **Top row:** Mesh representation. **Bottom row:** Contour representation.

Fig. 4. Sample combinations ($\alpha = 30\,000$, 100 iterations, $\eta = \frac{2}{3}$). **Top row:** Hairstyle 1. **Bottom row:** Hairstyle 2. **From Left to Right:** Customer 1, 2, 3.

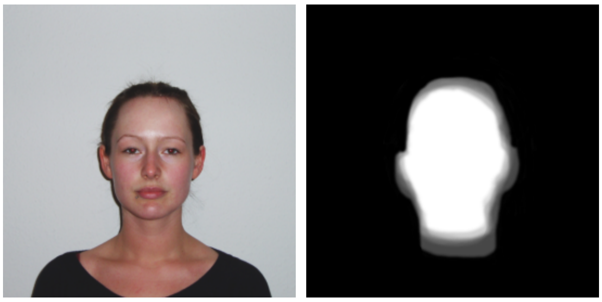

Fig. 5. **(a) Left:** Reference face used for real-world image material. **(b) Right:** Weighting mask obtained as the average of 15 different masks adopted on different example customer faces.

Let us address first the problem of different lighting conditions. As pointed out in Section 3, we aim at keeping lighting conditions similar for customer and reference images. To this end, we switch to another reference face that has been shot under lighting conditions comparable to the practical setting, i.e. using an ordinary digital camera with single flashlight in a badly lit room, see Figure 5(a).

To avoid misestimations caused by non-matched objects and structures in the customer image, we use the two-step procedure from Section 3. Our weighting mask, shown in Figure 5(b), is the average of numerous manually generated masks for individual customer faces. The results shown in Figure 6 were obtained by incorporating this preprocessing step. They demonstrate that the complete process is able to handle the kind of image material that occurs in practice.

7 Summary and Conclusions

In this paper, we have presented the complete solution of a practical visual computing problem by a variational approach, from a sound theoretical approach all the way to a fully applicable algorithm whose commercial use is imminent.

We have started out from the formulation of hairstyle simulation as a correspondence problem. The demand for a dense displacement field and the necessary balance between adaptation to image structures and smoothness of the displacement field made the variational approach with quadratic penalisation the method of choice for its solution. Within this setting, all partial problems could be addressed by specific well-founded measures. In particular, warping was employed to deal with large displacements, and a weighting function in the data term allowed to suppress the influence of spurious background and clothing details. A two-step matching procedure solved the problem of shifted and scaled images. The meaning of the few parameters is transparent, and by experimental guidance they were fixed to values that work for practical use. Careful design of the weighting function and the selection of an adequate reference face image rounded up the proceeding.

Fig. 6. Example combinations with three hairstyles. **Top:** Customer 1. **Bottom:** Customer 2. **Left to Right:** Hairstyle 1, 2, 3.

As a result, we have arrived at a method that combines results of high quality with robustness against the typical problems of the real-life setting. The algorithm processes input images taken by non-expert photographers with consumer cameras under simple lighting conditions, it does not require sophisticated user input like contours or marked features, and it achieves sufficient speed on a standard PC.

This contribution demonstrates that theoretical stringence and real-life applicability go well together. Moreover, by addressing an unprecedented field of application we want to emphasise that beyond a few well-established disciplines there are realms of visual computing applications awaiting future research efforts.

Acknowledgements. We are grateful to Jürgen Demetz (Style Concept SC KG, Beckingen) for supplying us with numerous images, and to Andrea Laick and Tina Scholl for volunteering for the real world experiments.

References

1. L. Alvarez, R. Deriche, J. Sánchez, and J. Weickert. Dense disparity map estimation respecting image derivatives: a PDE and scale-space based approach. *Journal of Visual Communication and Image Representation*, 13(1/2):3–21, 2002.
2. R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics and Image Processing*, 46(1):1–21, 1989.

3. J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. In G. Sandini, editor, *Computer Vision – ECCV '92*, volume 588 of *Lecture Notes in Computer Science*, pages 237–252. Springer, Berlin, 1992.
4. J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, August 1991.
5. T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In T. Pajdla and J. Matas, editors, *Computer Vision – ECCV 2004, Part IV*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin, 2004.
6. A. Bruhn, J. Weickert, C. Feddern, T. Kohlberger, and C. Schnörr. Variational optic flow computation in real-time. *IEEE Transactions on Image Processing*, 14(5):608–615, May 2005.
7. C. Chefd'Hotel, G. Hermosillo, and O. Faugeras. A variational approach to multimodal image matching. In *Proc. First IEEE Workshop on Variational and Level Set Methods in Computer Vision*, pages 21–28, Vancouver, Canada, July 2001. IEEE Computer Society Press.
8. I. Cohen. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, volume 1, pages 523–530, Tromsø, Norway, May 1993.
9. P. Golland and A. M. Bruckstein. Motion from color. *Computer Vision and Image Understanding*, 68(3):346–362, December 1997.
10. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
11. M. Lefébure and L. D. Cohen. Image registration, optical flow and local rigidity. *Journal of Mathematical Imaging and Vision*, 14(2):131–147, March 2001.
12. A.-R. Mansouri, A. Mitiche, and J. Konrad. Selective image diffusion: application to disparity estimation. In *Proc. 1998 IEEE International Conference on Image Processing*, volume 3, pages 284–288, Chicago, IL, October 1998.
13. R. March. Computation of stereo disparity using regularization. *Pattern Recognition Letters*, 8:181–187, October 1988.
14. E. Mémin and P. Pérez. A multigrid approach for hierarchical motion estimation. In *Proc. Sixth International Conference on Computer Vision*, pages 933–938, Bombay, India, January 1998. Narosa Publishing House.
15. J. Modersitzki. *Numerical Methods for Image Registration.* Oxford University Press, Oxford, 2004.
16. K. W. Morton and L. M. Mayers. *Numerical Solution of Partial Differential Equations.* Cambridge University Press, Cambridge, UK, 1994.
17. H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
18. N. Papenberg, A. Bruhn, T. Brox, S. Didas, and J. Weickert. Highly accurate optic flow computation with theoretically justified warping. *International Journal of Computer Vision*, 67(2):141–158, April 2006.
19. O. Scherzer, editor. *Mathematical Models for Registration and Applications to Medical Imaging.* Springer, Berlin, 2006.
20. D. M. Young. *Iterative Solution of Large Linear Systems.* Dover, New York, 2003.