



HAL
open science

Data-Driven Animation of Crowds

Nicolas Courty, Thomas Corpetti

► **To cite this version:**

Nicolas Courty, Thomas Corpetti. Data-Driven Animation of Crowds. Proc. of Mirage 2007 - Computer Vision / Computer Graphics Collaboration Techniques and Applications, May 2007, Paris, France. pp.377–388. hal-00494248

HAL Id: hal-00494248

<https://hal.science/hal-00494248>

Submitted on 22 Jun 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data-Driven Animation of Crowds

Nicolas Courty¹ and Thomas Corpetti²

¹ Université de Bretagne-Sud, Laboratoire VALORIA,
56000 Vannes Cedex, France
`nicolas.courty@univ-ubs.fr`

² Université de Haute-Bretagne, Laboratoire COSTEL,
35000 Rennes Cedex, France
`thomas.corpetti@uhb.fr`

Abstract. In this paper we propose an original method to animate a crowd of virtual beings in a virtual environment. Instead of relying on models to describe the motions of people along time, we suggest to use *a priori* knowledge on the dynamic of the crowd acquired from videos of real crowd situations. In our method this information is expressed as a time-varying motion field which accounts for a continuous flow of people along time. This motion descriptor is obtained through optical flow estimation with a specific second order regularization. Obtained motion fields are then used in a classical fixed step size integration scheme that allows to animate a virtual crowd in real-time. The power of our technique is demonstrated through various examples and possible follow-ups to this work are also described.

1 Introduction

Crowds of people exhibit particular and subtle behaviors whose complexity reflects the complex nature of human beings. While computer simulation of such phenomena have made it possible to reproduce particular and singular crowd configurations, none of them have managed to reproduce, within a generic framework, the typical emergent behaviors observed within a crowd with sufficient details and at a satisfying level. In the context of animation of human-like figures, huge progress have been observed with the use of motion capture. It is now possible to use motions acquired from real performers through a variety of editing and warping operations with substantial benefits in terms of realism in the produced animation. The aim of our technique is to provide such a tool in the context of crowd animation. While other approaches try to track singular pedestrians into the flow of people, our framework is based on the hypotheses that the motions of individuals within the crowd is the expression of a continuous flow that drives the crowd motion. *This assumes that the crowd is dense enough so that pedestrians are considered as markers of an underlying flow.* In this sense, our method is more related to macroscopic simulation models (that try to define an overall structure to the crowd's motions) rather than microscopic models (that define the crowd's motions as an emergent behavior of the sum of individual displacement strategies).

Our methodology relies on an analysis/synthesis scheme which is depicted in Figure 1. First, images are extracted from a video of a real crowd. From all the pairs of successive images a vector field is computed through a motion estimation process. The concatenation of all these vector fields represent a time series which accounts for the displacement of the whole crowd along time. This ends up the analysis part. The synthesis of a new crowd animation is done by advecting particles (the pedestrians) along this time varying flow.

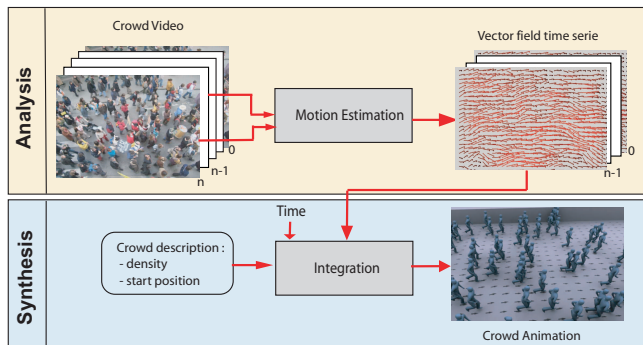


Fig. 1. Overview of the whole process

This paper is divided as follow: section 2 is a state of the art of the different existing approaches in the context of crowd simulation as well as motion estimation. Section 3 deals with the estimator used in our methodology, and section 4 presents the integration of the motion descriptor in a crowd animation controller. The last two sections present results obtained with our method along with a conclusion and perspectives for our work.

2 State of the art

The idea of using videos as an input to animation system is not new, and has already been succesfully used in the context of, for instance, facial animation [9], character animation from cartoon [8] or animal gaits [14]. Recent works showed example of re-synthesis of fluids flow from real video examples [3]. Simulating crowds from real videos falls into this challenging category of methods. First, it is interesting to understand the limitations of crowd simulation model (first part of this section). We then introduce some general issues about the motion estimation problem.

2.1 Crowd simulation

Crowd behavior and motion of virtual people have been studied and modeled in computers with different purposes: populating virtual environments, video

games, movie production or simulating emergency situations to help the design of buildings and open-spaces. The state of the art in human crowd behavioral modelling is large and can be classified in two main approaches: microscopic and macroscopic models. The models belonging to the first category are those describing the time-space behavior of individual pedestrians whereas the second category are those describing the emergent properties of the crowd.

Microscopic simulation The simplest models of microscopic simulation are based on cellular automata [6, 5]. The social force model was first introduced by Helbing [15]. It consists in expressing the motion of each pedestrian as a result of a combination of *social forces*, that repel/attract pedestrians toward each others. It has been shown that this model generates realistic phenomena as arc formations in exits or increasing evacuation time with increased desired velocities. It has been extended to account for individualities [7] or the presence of toxic gases in the environment [12]. More complex models consider each member of the crowd as autonomous pedestrians endowed with perceptive and cognitive abilities [22, 24, 23]. Those models exhibit a variety of results depending on the quality of the behavior design.

Macroscopic models Modelling a crowd composed of discrete individuals may lead to incorrect emergent global behaviors. These difficulties may be avoided by using a continuum formulation [18, 26]. Equations using the concepts of fluid mechanics have been derived in order to model such approach of human crowds. Those approaches rely on the assumption that the characteristic distance scale between individuals is much less than the characteristic distance scale of the region in which the individuals move [18]. Hence the density of the crowd has to be taken into account for those models to be pertinent. Finally several hypotheses on the behavior of each members of the crowd lead to partial derivative equations governing the flow of people.

Although crowds are made up of independent individuals with their own objectives and behaviour patterns, the behavior of crowds is widely understood to have collective characteristics which can be described in general terms. Though, macroscopic models may lack of subtleties and often rely on strong hypotheses (notably on density). Our framework propose to capture this global dynamic from real crowd video sequences. This imposes the use of motion estimation techniques.

2.2 Motion estimation

When a crowd is dense enough, the usual tracking systems like Kalman filters or stochastic filtering [13] will generate large state space that will yield a computationally too expensive problem. It is then necessary to use alternative methods to obtain the information on the dynamics of the crowd in order to characterize its behavior. This section investigates the different ways to obtain some

motion descriptors from image sequences. Many families of methods are available to measure a motion information from image sequences. One can cite for instance the parametric methods, the correlation techniques or the optical flow approaches (see [21] for a survey). These latter are known to be the most accurate to address the generic problem of estimating the apparent motion from image sequences (see for instance [27] for some presentations and [2] for comprehensive comparisons with completely different approaches). The idea of using optical flow to estimate crowd motions has recently drawn attention in the context of human activity recognition [1]. The original optical flow is based on the seminal work of Horn & Schunck [16] and is briefly described in the next paragraph.

Optical Flow The optical flow based on Horn & Schunck consists in the minimization of a global cost function \mathcal{H} composed of two terms. The first one, named “observation term”, is derived from a *brightness constancy* assumption and assumes that a given point keeps the same intensity along its trajectory. It is expressed through the well known *optical flow constraint equation* (OFCE):

$$\mathcal{H}_{obs}(E, \mathbf{v}) = \iint_{\Omega} f_1 \left[\nabla E(\mathbf{x}, t) \cdot \mathbf{v}(\mathbf{x}, t) + \frac{\partial E(\mathbf{x}, t)}{\partial t} \right] d\mathbf{x}, \quad (1)$$

where $\mathbf{v}(\mathbf{x}, t) = (u, v)^T$ is the unknown velocity field at time t and location $\mathbf{x} = (x, y)$ in the image plane Ω , $E(\mathbf{x}, t)$ is the image brightness, viewed for a while as a continuous function.

This first term relies on the assumption that the visible points conserve roughly their intensity in the course of a displacement.

$$\frac{dE}{dt} = \nabla E \cdot \mathbf{v} + \frac{\partial E}{\partial t} \approx 0. \quad (2)$$

The associated penalty function f_1 is often the L_2 norm. However, better estimates are usually obtained by choosing a “softer” penalty function [4]. Such functions, arising from robust statistics [17], limit the impact of the many locations where the brightness constancy assumption does not hold, such as on occlusion boundaries.

This single (scalar) observation term does not allow to estimate the two components u and v of the velocity. In order to solve this ill-posed problem, it is common to employ an additional smoothness constraint \mathcal{H}_{reg} . Usually, this second term enforces a spatial smoothness coherence of the flow field. It relies on a contextual assumption which enforces a spatial smoothness of the solution. This term usually reads:

$$\mathcal{H}_{reg}(\mathbf{v}) = \iint_{\Omega} f_2 [|\nabla u(\mathbf{x}, t)| + |\nabla v(\mathbf{x}, t)|], \quad (3)$$

As with the penalty function in the data term, the penalty function f_2 was taken as a quadratic in early studies, but a softer penalty is now preferred in order not

to smooth out the natural discontinuities (boundaries, ...) of the velocity field [4, 20]. Based on (1) and (3), the estimation of motion can be done by minimizing:

$$\begin{aligned} \mathcal{H}(E, \mathbf{v}) &= \mathcal{H}_{obs}(E, \mathbf{v}) + \alpha \mathcal{H}_{reg}(\mathbf{v}) \\ &= \iint_{\Omega} f_1 \left[\nabla E(\mathbf{x}, t) \cdot \mathbf{v}(\mathbf{x}, t) + \frac{\partial E(\mathbf{x}, t)}{\partial t} \right] d\mathbf{x} + \\ &\quad \alpha \iint_{\Omega} f_2 [|\nabla u(\mathbf{x}, t)| + |\nabla v(\mathbf{x}, t)|], \end{aligned} \tag{4}$$

where $\alpha > 0$ is a parameter controlling the balance between the smoothness constraint and the global adequacy to the observation assumption.

The minimization of this overall cost function enables to extract the apparent motion field between a pair of images $E(\mathbf{x}, t_1)$ and $E(\mathbf{x}, t_2)$.

Discussion It has been proved that in many image sequences and especially in fluid-like imagery, these classic assumptions are violated in a number of locations in the image plane. Even if in most of rigid-motion situations, the use of a robust penalty function enables us to recover properly the motion of pathological situations (occluding contours, ...) the usual assumptions are, unfortunately, even less appropriate in fluid imagery.

Some studies have proved that a crowd dense enough has sometimes a behavior that can be explained by some fluid mechanics laws [18]. It is then of primary interest to integrate such prior knowledge in the optical flow (in the observation term or on the regularization constraint, depending on the nature of the physical law to integrate) to obtain a technique devoted to crowd motion. In this paper, we propose to use a smoothing constraint dedicated to the capture of the significant properties of the flow from a fluid mechanics point of view. These properties are the *divergence* (linked to the dispersion of a crowd) and the *vorticity* (also named *curl*) linked to a rotation.

3 Crowd motion estimation and representation

In this section, we present the regularization used in the motion estimator to extract a reliable crowd motion information. For more details on the approach, the reader can refer to [11, 10]. Under the assumption that a dense enough crowd has a behaviour that can be modeled with some fluid mechanics laws, one can demonstrate that the usual first-order regularization functional in (3) is not adapted for fluid situations.

By using Euler-Lagrange conditions of optimality, it is indeed readily demonstrated [10] that the standard first-order regularization functional :

$$\mathcal{H}_{reg}(\mathbf{v}) = \iint_{\Omega} |\nabla u(\mathbf{x})|^2 + |\nabla v(\mathbf{x})|^2 d\mathbf{x} \tag{5}$$

is equivalent from the minimization point of view, to the so-called *div-curl* regularization [25]:

$$\mathcal{H}_{reg}(\mathbf{v}) = \iint_{\Omega} (\operatorname{div}^2 \mathbf{v}(\mathbf{x}) + \operatorname{curl}^2 \mathbf{v}(\mathbf{x})) \, d\mathbf{x}, \quad (6)$$

where $\operatorname{div} \mathbf{v} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}$ and $\operatorname{curl} \mathbf{v} = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}$ are respectively the divergence and the vorticity of the motion field $\mathbf{v} = (u, v)$.

A first-order regularization therefore penalizes the amplitude of both the divergence and the vorticity of the vector field. For a dense crowd motion estimation, this does not seem appropriate since the apparent velocity field normally exhibits compact areas with high values of vorticity and/or divergence. It seems then more appropriate to rely on a second-order div-curl regularization [25]:

$$\mathcal{H}_{reg}(\mathbf{v}) = \iint_{\Omega} (|\nabla \operatorname{div} \mathbf{v}(\mathbf{x})|^2 + |\nabla \operatorname{curl} \mathbf{v}(\mathbf{x})|^2) \, d\mathbf{x}. \quad (7)$$

This regularization tends to preserve the divergence and the vorticity of the motion field \mathbf{v} to estimate. Interested readers may refer to [11] to get precise descriptions on the optimization strategy and on associated numerical implementation issues.

The motion field \mathbf{v} is then the minimum of the following cost function (with $\bullet = (\mathbf{x}, t)$):

$$\mathbf{v}(\bullet) = \min_{\mathbf{v} \in \Omega} \iint_{\Omega} \left\{ f_1 \left[\nabla E(\bullet) \cdot \mathbf{v}(\bullet) + \frac{\partial E(\bullet)}{\partial t} \right] + \alpha \|\nabla \operatorname{div} \mathbf{v}(\bullet)\|^2 + \alpha \|\nabla \operatorname{curl} \mathbf{v}(\bullet)\|^2 \right\} \, d\mathbf{x}. \quad (8)$$

and the global crowd motion is represented as a time series of such motion fields.

4 Data-driven animation of crowds

Once the time series of motion fields has been computed, it is possible to consider this information as input data for an animation system. Let us first recall that the computed velocities correspond to a velocity *in the image space*, and our goal is to animate individualities *in the virtual world space*. Given the position of such a person in the virtual world, it is possible to get the corresponding position in the image frame along with a camera projection model. Parameters for this projection can be obtained exactly through camera calibration. We have considered as an approximation of this model a simple orthographic projection in the experiments presented in the result sections. This assumption holds whenever the camera is sufficiently far away from the crowd scene. Once this projection has been defined, animating individualities which constitute the crowd amounts to solve the classical following differential equation (with $x(t)$ the position of a person in the image frame at time t):

$$\frac{\partial x}{\partial t} = v(x(t), t) \quad (9)$$

equipped with appropriate initial condition $x(0) = x_0$ which stands for the initial positions of the individual in the flow field. In our framework we have used the classical 4-th order Runge Kutta integration scheme, which allows to compute a new position $x(t + 1)$ given a fixed timestep with an acceptable accuracy. This new position is then projected back in the virtual world frame. This process is depicted in Figure 2.




Fig. 2. Motion synthesis from flow field. The position of the crowd’s member is projected onto the flow (step 1), the integration is performed in the image frame (step 2) and then the new position is projected back in the virtual world frame (step 3).

Let us finally note that the quality of the generated animation is closely linked with the initial position of the crowd members and their density. We have used in the subsequent results curve sources that create random pedestrians along a hand-designed curve situated in the flow.

5 Results

Our approach was first tested on synthetic crowd sequences to validate the theoretical part of our work. We have also used real crowd sequences to handle real cases. Those results are presented in this section.

5.1 Synthetic example

The synthetic sequence represents a continuous flow of human beings with an obstacle (a cylinder named \mathcal{C}) in the middle of the image. It has been generated using the classical Helbing simulation model [15]. In this situation, the true motion field inside the cylinder \mathcal{C} is known (no motion, *i.e.* $\mathbf{v}(\mathbf{x} \in \mathcal{C}) = 0$). The cost function (8) being defined on the whole image plane, we need to have a particular process to deal with this specific *no-data* area. Actually, since any motion inside the area \mathcal{C} is a reliable candidate (the OFCE (1) is null everywhere), the motion estimation using relation (8) is likely to yield some incoherent results inside and outside the cylinder (due to the regularization term which spreads errors). To cope with this situation, we completely blurred the cylinder area

from an image to an other, so that the OFCE constraint is verified nowhere in \mathcal{C} . Thanks to the robust estimator f_1 used in (1), this area is not taken into account by the observation term of the estimation process. Hence, the motions fields estimated outside the cylinder are not disturbed by the ones inside \mathcal{C} . This is illustrated in Figure 3. We present an image of the sequence in Figure 3(a), the estimated motion field in Figure 3(b), a zoom of the cylinder area with and without the specific treatment proposed on this particular situation (Figure 3(c) and 3(d) respectively). Some images of the crowd animation synthesis are shown




Fig. 3. Estimation of the motion field on the synthetic example; (a): images from the original sequence; (b) the estimated motion field; (c) the motion near the cylinder estimated with a special care of this no-data area and (d) same as (c) but without a specific treatment for the cylinder. One can see that the motion near the cylinder in (d) is not totally coherent.

on Figure 4. The animation was generated thanks to a Maya plugin which defines a crowd as a set of particles and performs the synthesis described in section 4. As expected, the virtual crowd is in accordance with the underlying motion and the obstacle is correctly managed. This first example proves the ability of the proposed approach to synthesize a coherent motion from an estimated motion field. Let us now apply this technique to real data.

5.2 Real data

We present the results obtained on two real sequences. Both data have been acquired with a simple video camera with an MPEG encoder. The resulting images are hence very poor in terms of brightness: this latter is indeed sometimes constant in a squared area. It is important to note that this point is likely to disturb the motion estimation process.




Fig. 4. Some images of the synthetic crowd animation for 4 different times of the sequence.

Strike sequence The first real sequence is a video representing a strike which took place at Vannes in France. All pedestrians are walking on the same direction. Two images of the sequence can be seen on Figure 5 (a) and (b). In Figure 5 (c) and (d), we present the synthetic crowd animation obtained superimposed on the estimated motion field. One can observe that the resulting crowd animation is in accordance with the real pedestrian behaviors. Hence, on this example, our method has the advantage to synthesize correctly the observed phenomena without resorting to usual motion capture techniques. Let us now see the results on a more complicated real sequence.

Shibuya sequence The second real sequence is a video acquired in the *Shibuya crossroads* in Tokyo, Japan, which is famous for the density of people crossing the streets. Three images of the sequence can be seen on Figure 6 (a-c). This situation is complex since at least two main flows of people in opposite directions are crossing the road. It is important to observe that in this case, the underlying assumptions of our approach (a very dense crowd) are not totally respected. This example is therefore shown to evaluate the limits of our method. In Figure 6 (d-f), we present the synthetic crowd animation obtained superimposed on the estimated motion field. One can see on these figures that the two main opposite flows are correctly extracted and synthesised, despite the fact that the initial sequence was very poor in terms of quality and that our initial assumptions were not respected. The generated sequence is relatively realistic. Nevertheless, the intersection of the two groups of people is not correctly managed: some pedestrians have incoherent trajectories. This issue has two main reasons: the estimation process is locally incoherent when two people occlude each other, and there is no temporal continuity in the estimated flow. Two possibilities can be exploited to cope such a situation: the first one consists in improving the motion estimation process through a temporal smoothing of the motion field whereas the second possibility is to introduce a dynamical law in the trajectory reconstruction step. This two key points will be the scope of our further work.

6 Conclusion

In this paper, we have presented a new and original method which proposes to animate a virtual dense crowd thanks to real crowd video sequences. This is done



Fig. 5. The strike sequence. (a,b): two images of the sequence; (c,d) the corresponding animation superimposed on the estimated motion field.




Fig. 6. The Shibuya sequence. (a-c): two images of the sequence; (d-f) the corresponding animation superimposed on the estimated motion field.

using *i*) a specific motion information process applied on the input images and *ii*) an integration part to obtain the trajectories of individualities in the crowd. We applied the presented method on both synthetic and real examples. The experimental part showed the ability of the technique to synthesize reliable crowd animations but also pointed out some limitations. To improve the presented approach, some problems will be important to solve. Concerning the analysis part, the motion estimation process can be improved by introducing more specific spatio-temporal models of continuous crowd behaviours [18]. This can be done using the framework of the optical flow but it can also be relevant to explore the possibilities of the data-assimilation used for instance in meteorology [19]. This constitutes a very exciting challenge for which actually no practical and generic solution exists and stands as a very appealing alternative to tracking systems, too expensive in the context of dense crowds. An other important point would be to specify a motion estimation process for crowds that are not totally dense. Concerning the synthesis part, we aim at enriching existing simulation models by integrating this *a priori* knowledge of the behaviour of the crowd. In this context, the crowd simulation would integrate both dynamical simulation models and observed data from real sequences.

Acknowledgements The authors wish to thank Amel Achour and Gwendal Jabot for their contributions to the project, and Alvaro Cassinelli for providing us the shibuya video sequences.

References

1. E. Andrade, S. Blunsden, and R. Fisher. Modelling crowd scenes for event detection. In *Int. Conf. on Pattern Recognition, ICPR 2006*, pages 175–178, 2006.
2. J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *Int. J. Computer Vision*, 12(1):43–77, 1994.
3. K. Bhat, S. Seitz, J. Hodgins, and P. Khosla. Flow-based video synthesis and editing. *ACM Transactions on Graphics, special issue, Proceedings ACM SIGGRAPH 2004*, 23(3):360–363, 2004.
4. M. Black. Recursive non-linear estimation of discontinuous flow fields. In *Proc. Europ. Conf. Computer Vision*, pages 138–145, Stockholm, Sweden, 1994.
5. V. Blue. Cellular automata microsimulation for modeling bi-directional pedestrian walkways. *Transportation Research, Part B: Methodological*, 35(3):293–312, March 2001.
6. E. Bouvier, E. Cohen, and L. Najman. From crowd simulation to airbag deployment: particle systems, a new paradigm of simulation. *Journal of Electronic Imaging*, 6(1):94–107, January 1997.
7. A. Braun and S. Raupp Musse. Modeling individual behaviors in crowd simulation. In *Proc. of Computer Animation and Social Agents (CASA'03)*, New Jersey, USA, May 2003.
8. C. Bregler, L. Loeb, E. Chuang, and H. Deshpande. Turning to the masters: Motion capturing cartoons. *ACM Transactions on Graphics, special issue, Proceedings ACM SIGGRAPH 2002*, 21(3):320–328, 2002.

9. J. Chai, J. Xiao, and J. Hodgins. Vision-based control of 3d facial animation. In *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, pages 79–87, Grenoble, France, August 2003.
10. T. Corpetti, D. Heitz, G. Arroyo, E. Mémin, and A. Santa-Cruz. Fluid experimental flow estimation based on an optical-flow scheme. *Experiments in fluids*, 40(1):80–97, 2006.
11. T. Corpetti, E. Mémin, and P. Pérez. Dense estimation of fluid flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):365–380, March 2002.
12. N. Courty and S. Musse. Simulation of Large Crowds Including Gaseous Phenomena. In *Proc. of IEEE Computer Graphics International 2005*, pages 206–212, New York, USA, June 2005.
13. D. Crisan. Particle filters, a theoretical perspective. In N. de Freitas A. Doucet and N. Gordon, editors, *Sequential Monte-Carlo Methods in Practice*. Springer, 2001.
14. L. Favreau, L. Reveret, C. Depraz, and M.-P. Cani. Animal gaits from video. In *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, Grenoble, France, August 2004.
15. D. Helbing, I. Farkas, and T. Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(1):487–490, 2000.
16. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
17. P. Huber. *Robust Statistics*. John Wiley & Sons, 1981.
18. R. L. Hughes. The flow of human crowds. *Annual revue of Fluid. Mech.*, 20(10):169–182, 2003.
19. F.X. Le-Dimet and O. Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus*, pages 97–110, 1986.
20. E. Mémin and P. Pérez. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Trans. Image Processing*, 7(5):703–719, 1998.
21. A. Mitiche and P. Bouthemy. Computation and analysis of image motion: a synopsis of current problems and methods. *Int. J. Computer Vision*, 19(1):29–55, 1996.
22. S. Raupp Musse and D. Thalmann. Hierarchical model for real time simulation of virtual human crowds. In *IEEE Trans. on Visualization and Computer Graphics*, volume 7(2), pages 152–164. IEEE Computer Society, 2001.
23. W. Shao and D. Terzopoulos. Animating autonomous pedestrians. In *Proc. SIGGRAPH/EG Symposium on Computer Animation (SCA'05)*, pages 19–28, Los Angeles, CA, July 2005.
24. M. Sung, M. Gleicher, and S. Chenney. Scalable behaviors for crowd simulation. *Comput. Graph. Forum*, 23(3):519–528, 2004.
25. D. Suter. Motion estimation and vector splines. In *Proc. Conf. Comp. Vision Pattern Rec.*, pages 939–942, Seattle, USA, June 1994.
26. A. Treuille, S. Cooper, and Z. Popovic. Continuum crowds. *ACM Transactions on Graphics, special issue, Proceedings ACM SIGGRAPH 2006*, 25(3):1160–1168, 2006.
27. J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *Int. J. of Computer Vision*, 45(3):245–264, December 2001.