A Lagrangian Particle Dynamics Approach for Crowd Flow Segmentation and Stability Analysis

Saad Ali
Computer Vision Lab
University of Central Florida
sali@cs.ucf.edu

Mubarak Shah
Computer Vision Lab
University of Central Florida
shah@cs.ucf.edu

Abstract

This paper proposes a framework in which Lagrangian Particle Dynamics is used for the segmentation of high density crowd flows and detection of flow instabilities. For this purpose, a flow field generated by a moving crowd is treated as an aperiodic dynamical system. A grid of particles is overlaid on the flow field, and is advected using a numerical integration scheme. The evolution of particles through the flow is tracked using a Flow Map, whose spatial gradients are subsequently used to setup a Cauchy Green Deformation tensor for quantifying the amount by which the neighboring particles have diverged over the length of the integration. The maximum eigenvalue of the tensor is used to construct a Finite Time Lyapunov Exponent (FTLE) field, which reveals the Lagrangian Coherent Structures (LCS) present in the underlying flow. The LCS divide flow into regions of qualitatively different dynamics and are used to locate boundaries of the flow segments in a normalized cuts framework. Any change in the number of flow segments over time is regarded as an instability, which is detected by establishing correspondences between flow segments over time. The experiments are conducted on a challenging set of videos taken from Google Video and a National Geographic documentary.

1. Introduction

Management of large gatherings of people at events such as religious festivals, parades, concerts, football matches, etc., pose significant challenges to public safety management officials. Often these gatherings involve movement of crowds in confined spaces such as city streets, overhead bridges, or narrow passageways. Figure 1 depicts some example scenarios. It is quite obvious that an incident free management of such huge gatherings is a daunting task, simply due to shear number of people involved in these events. One way to reduce the incidence of any catastrophic event in situations involving large crowds is through better coordination and remodelling of expected bottleneck areas.

Figure 1. Example scenarios involving thousands of people. (a) A scene from New York City marathon. (b) A large crowd participating in a political rally in Los Angeles. (c) Pilgrims circling around Kabba in Mecca.

However, numerous occurrences of stampedes in the recent past have shown that better coordination between public safety organizations and remodelling alone cannot solve the problem of managing large crowds. For instance, 270 pilgrims were killed at Jamarat bridge in Mecca in May 1994 and another 251 were killed in February 2004. This led to redesigning of the approach towards the bridge and exit points, but unfortunately in January 2006, 345 lives were again lost at the same bridge due to a stampede.

Over the last few years, computer vision based algorithms have been integrated into wide area surveillance systems for in areas such as city streets, subway stations, malls, etc. However, one common weakness among these systems is their inability to handle crowded scenes. As soon as the density of objects in the scene increases, a degradation in their performance in terms of object detection, tracking and event detection, is observed. Limited research effort, if any, has been spent in building computer vision systems that can model high density scenes and provide useful information to public safety officials. One reason for the lack of effort in this direction is the complexity of the problem.

This research effort is a first step towards building an automated monitoring system capable of modelling high density crowd scenes. For this paper, the term *modelling* includes tasks of segmenting dominant crowd flows present in the scene, and detection of any abnormalities that may arise in these flows. Our proposed approach for solving this modelling problem starts by treating moving crowds as an aperiodic dynamical system, which is manifested by a time dependent flow field. Such a flow field in a general scene will consist of regions with qualitatively different dynam-

1-4244-1180-7/07/$25.00 ©2007 IEEE

ics, reflecting the motion patterns emerging from the spatio-temporal interactions of the participants between each other, as well as with the physical world. These emerging motion patterns which have dynamical and physical interpretations in a given scene are referred to as 'flow segments'.

Now, in order to locate these flow segments, we propose a flow segmentation framework which makes use of recent advances in the areas of nonlinear dynamical systems [2], fluid dynamics [4][6], and turbulence theory [1][11]. The basis of the idea is to use Lagrangian Particle Dynamics to uncover the spatial organization of the flow field by examining clouds of particles as they mix and get transported under the action of the flow field generated by the crowd motion. The motivation behind the idea is the observation that the trajectories generated from the advection of particles through a flow reveal representative characteristics of flow such as locations of the barriers, mixing properties, sources, sinks, etc. Since the flow fields we are dealing with are generated by moving crowds, their characteristics will have a direct relationship with the physical properties of the scene and the crowd such as physical barriers in the scene, heading direction, number of crowd segments, locations at which segments merge or split, etc.

The key theoretical notion that we are going to use is the existence of *Coherent Structures* (CS) [6] in fluid flows which can be discovered by fluid particle advection. Roughly speaking, CS are separatrices that influence the kinematics of a particle cloud over finite time intervals, and divide the flow into dynamically distinct regions where all particles within the same region have a similar fate, i.e., coherent behavior. The notion of coherent structure is extendable to flow fields generated by crowds where they map to the boundaries of different crowd segments. Intuitively, 'CS' are to flow data what 'edges' are to image data. Note that when CS are studied in terms of quantities derived from trajectories, they are often named as Lagrangian Coherent Structures (LCS). The LCS are located through a Lyapunov Exponent approach, resulting in a FTLE field over the flow domain. Lyapunov exponents measure the exponential rate of convergence or divergence between two particles. It has been shown by Haller [6] that CS are local maximizing curves of this field appearing as ridges. These ridges can be treated as edges separating flow segments of different dynamics from each other. At this point, we formulate the problem of locating physically and dynamically meaningful segments as segmentation of a scalar field in a normalized cut framework.

The second goal of this paper is the detection of any changes in the dynamics of flow from its learned pattern. Our formulation for flow segmentation allows us to achieve this task by simply detecting presence of new flow segments from one time instant to the next. Recall that the difference in the dynamics of any part of the flow gives rise to CS at the location where the change in dynamics is happening. Therefore, any part of the flow which shows change in its previously known dynamics will give rise to new CS, which will be segmented out using the FTLE field. By detecting these new segments, we can pinpoint which part of the flow is deviating from its normal flow behavior.

2. Related Work

Recently, the computer vision community has started taking interest in addressing different research problems related to the scenarios involving large crowds of people. The focus so far has been on the tasks of crowd detection, and detection and tracking of individuals in the crowd. Representations based on xt slices of spatio-temporal video volume [7], shape and color models of individuals [9], boundary contours [3], and trajectories of interest points [5] have been used for this purpose. A dew approaches [13][8] which explicitly count the number of people in the crowd have also been proposed. The main difference of our work with respect to this body of literature is that we are attempting to model the dynamics of the crowd as one global entity, instead of focusing on the individuals making up the crowd. Recently, Chan et. al. [16] proposed to segment videos of crowded environments using a dynamic texture based representation. However, our prime focus is on using the optical flow information for this task without making any assumptions about the nature of the underlying generative process. Another related work is by Sand et. al., [17] which proposes a particle based framework for motion computation. As opposed to their application of particles for motion estimation, we use particles for high level motion interpretation in the form of physically and dynamically meaningful segments. In next section we describe the mathematical notations and key concepts using the nomenclature used in [10].

3. Definitions and Notations

Let a compact set $D \subset \mathbb{R}^2$ be the domain of the crowd flow under study. Given a time dependent velocity field $\mathbf{v}(\mathbf{x},t)$ defined on D and satisfying C^0 and C^2 continuity in time and space, respectively, a trajectory $\mathbf{x}(t : t_0, \mathbf{x}_0)$ starting at point $\mathbf{x}_0$ at time t_0 is a solution of

$$\dot{\mathbf{x}}(t; t_0, \mathbf{x}_0) = \mathbf{v}(\mathbf{x}(t; t_0, \mathbf{x}_0), t), \quad \mathbf{x}(t_0; t_0, \mathbf{x}_0) = \mathbf{x}_0. \tag{1}$$

The C^0 assumption is required to keep the flow field smooth. A trajectory $\mathbf{x}(t : t_0, \mathbf{x}_0)$ of a particle depends on the initial position $\mathbf{x}_0$ and the initial time t_0. From continuity constraints of $\mathbf{v}(\mathbf{x},t)$, it follows that $\mathbf{x}(t : t_0, \mathbf{x}_0)$ will be C^1 in time and C^2 in space. The solution of a dynamical system described in Equation 1 can be viewed as a map, which takes points from their position $\mathbf{x}_0$ at time t_0 to their position at time t and is referred to as a *flow map*. It is denoted by $\phi_{t_0}^t$ and satisfies

$$\phi_{t_0}^t : D \to D : \mathbf{x}_0 \mapsto \phi_{t_0}^t(\mathbf{x}_0) = \mathbf{x}(t; t_0, \mathbf{x}_0). \tag{2}$$

Additionally, flow map $\phi_{t_0}^t$ satisfies following properties

$$\phi_{t_0}^{t_0}(\mathbf{x}) = \mathbf{x}, \quad (3)$$

$$\phi_{t_0}^{t+s}(\mathbf{x}) = \phi_s^{t+s}(\phi_{t_0}^s(\mathbf{x})) = \phi_t^{t+s}(\phi_{t_0}^t(\mathbf{x})). \quad (4)$$

3.1. Finite Time Lyapunov Exponent Field

The Lyapunov exponent is an asymptotic quantity which measures the extent to which infinitely close particles separate in an infinite amount of time. In fluid dynamics literature, the finite-time Lyapunov exponent is often used to quantify the mixing and dispersion of particles instead of the infinite time Lyapunov exponent. Finite time exponents depend on the initial positions of the trajectories and the length of the integration of the trajectories, and they reveal local properties of the flow, i.e., the properties which depend on the finite time integration of particle trajectories. When finite-time Lyapunov exponent analysis is performed over a grid of particles, it generates a finite time Lyapunov Exponent (FTLE) field. In our formulation, a FTLE field $\sigma_T(\mathbf{x}_0, t_0)$ can be computed using the flow map $\phi_{t_0}^{t_0+T}$. The amount of stretching around the trajectory of the particle can be understood by considering the evolution of the perturbed point $\mathbf{y} = \mathbf{x} + \delta\mathbf{x}(0)$, where $\delta\mathbf{x}(0)$ is infinitesimal and arbitrary oriented. After a time interval this perturbation becomes [10][1]:

$$\begin{aligned}\delta\mathbf{x}(T) &= \phi_{t_0}^{t_0+T}(\mathbf{y}) - \phi_{t_0}^{t_0+T}(\mathbf{x}) \\ &= \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}\delta\mathbf{x}(0) + O(\|\delta\mathbf{x}(0)\|^2).\end{aligned}$$

We linearize the perturbation by dropping the higher order terms and calculate the magnitude of deformation along the particle path by using following relation

$$\Delta = \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}^* \cdot \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}, \quad (5)$$

where superscript * refers to transpose operator. Δ is known as a finite time version of Cauchy-Green deformation tensor while quantity $\frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}$ is the spatial gradient tensor of the flow map. Note that maximum stretching between infinitesimally close particles occurs when $\delta\mathbf{x}(0)$ is aligned with the eigenvector of maximum eigenvalue of Δ. Then, from [10] we know that the largest finite time Lyapunov exponent with a finite integration time T corresponding to point $\mathbf{x} \in D$ at time t_0 is

$$\sigma_{t_0}^T = \frac{1}{T} ln\sqrt{\lambda_{max}(\Delta)}. \quad (6)$$

3.2. Lagrangian Coherent Structures

As mentioned earlier, the LCS reveal underlying flow structures that are typically not evident from the raw velocity field. In the case of an aperiodic velocity field, LCS can be located directly from the FTLE field, where they appear as ridges. The relationship between ridges in the FTLE field and LCS can be explained in the following way. If a flow is experiencing qualitatively different dynamics in two regions of the flow, then we expect a coherent motion of particles within each region and the eigenvalues of Δ to be close to 1, an indication that the fate of nearby particles is similar inside the sub-regions. At the boundary of two regions of qualitatively different dynamics, particles will move in an incoherent fashion. This will create much higher eigenvalues in the direction normal to the boundary, resulting in prominent ridges in the FTLE field. In this work, we are not interested in developing any precise mathematical definition of LCS, since information about their location in the given flow is enough for us to proceed with our flow segmentation process.

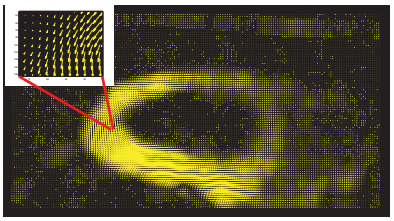

Figure 2. Mean flow field corresponding to the sequence displayed in Fig. 1(c). Inset shows a zoomed in version of the field.

4. Implementation

In this section we describe the implementation details of our flow segmentation framework.

4.1. Flow Field Calculation

Given a video sequence, the first task is to estimate the flow field. We employ a scheme consisting of block based correlation in the fourier domain, where displacement between blocks of consecutive frames is calculated by locating peaks in the correlation surfaces. This process is repeated for all of the possible blocks in the given frame. Local outliers are replaced by using an adaptive local median filtering. A typical size of the block used is 16x16. Optical flow fields calculated over n frames are averaged to obtained one mean field M. Finally, T mean fields are stacked together to obtain one block of mean fields. We use the symbol B_i^{i+T} to represent a stack of flow fields $M_i, M_{i+1}, .., M_{i+T}$, for frames i to frame $n \times (i+T)$. Fig. 2 shows a snapshot of the mean flow computed for the Mecca sequence (Fig. 1(c)).

4.2. Particle Advection

The task is to carry out particle advection under the influence of the stacked flow field B_i^{i+T}, corresponding to the time interval t to $i+T$. To perform this step, a grid of particles is launched over the first mean field M_i in B_i^{i+T}. The Lagrangian trajectory $[x(t+T;t,x_0,y_0), y(t+T;t,x_0,y_0)]$ corresponding to a particle at grid location (x_0, y_0) is computed by numerically solving the ordinary differential equations, $\frac{dx}{dt} = u(x,y,t)$, $\frac{dy}{dt} = v(x,y,t)$, subject to the initial

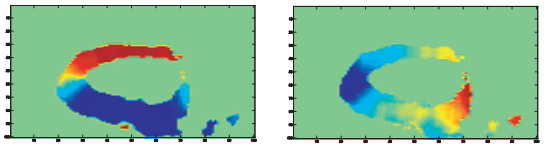

Figure 3. Intermediate spatial gradient maps during the advection process. (a) Spatial gradient of flow map of x-coordinate of particle grid i.e. $\frac{d\phi_x}{dx}$ (b) Spatial gradient of flow map of x coordinate of particle grid with respect to y-axis i.e. $\frac{d\phi_x}{dy}$.

conditions $[x(0), y(0)] = (x_0, y_0)$. Here, $t + T$ is the later time up till which we want to compute the trajectory. We use the fourth order Runge-Kutta-Fehlberg algorithm along with cubic interpolation [12] of the velocity field to solve this system. The above process is repeated for each block B_i^{i+T} of the video. Note that we do not reseed the particles if they go out of the image bounds.

4.3. Flow Maps and FTLE Field

A pair of flow maps, namely ϕ_x and ϕ_y, is maintained for the grid of particles at each time instant. The first map, ϕ_x, keeps track of how the x coordinate of the particle is changing, and similarly, ϕ_y keeps track of the y coordinate. At the start, these maps are populated with the initial positions of the particles, which are the pixel locations at which the particle is placed. When the particles are advected under the influence of B_i^{i+T} using the method described in Section 4.2, the positions of the particles are updated in these maps until the end of the integration time length T.

The FTLE field is computed from these maps, first by taking their spatial gradients as, $\frac{d\phi_x}{dx}$, $\frac{d\phi_x}{dy}$, $\frac{d\phi_y}{dx}$ and $\frac{d\phi_y}{dy}$. This step can be easily accomplished by using a finite differencing approach for taking derivatives. Fig. 3 shows an example of $\frac{d\phi_x}{dx}$ and $\frac{d\phi_x}{dy}$ computed for the Mecca sequence shown in Fig. 1(c). Finally, the Cauchy-Green deformation tensor is computed by plugging spatial gradients of the flow maps in Equation 6. Fig. 4 shows the FTLE field computed for two different sequences where the ridges are clearly visible. These ridges will be used to separate out the flow segments with different dynamics.

4.4. FTLE Field Segmentation

The FTLE field is a scalar field capturing the underlying flow dynamics and geometry. For segmenting this scalar field we employ the normalized cuts algorithm, first proposed by Shi et. al [14]. Our segmentation procedure is composed of two main steps. The first step involves an over-segmentation of the given FTLE field. In the second step, we merge the segments whose boundary particles have similar behavior in Lyapunov sense. The reason for using this scheme is the observation that graph based segmentation algorithms require the user to specify how many segments to return. This condition is too restrictive for a system that

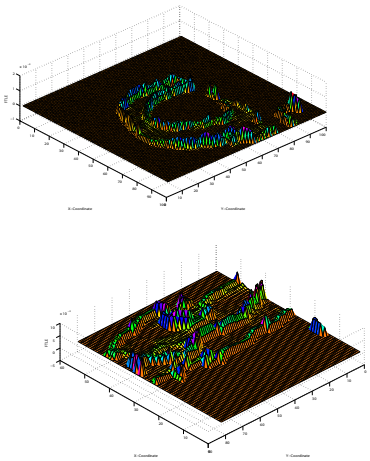

Figure 4. (a) FTLE field for the Mecca sequence. Ridges are prominent at the edges of the circle along which the pilgrims are moving, emphasizing the fact that region inside the ridges has a coherent motion. (b) FTLE field for the sequences shown in Fig. 6(b). Ridges are prominent at the edges of the flow as well as between the flows of different dynamics. The sequence has two groups of people moving side by side.

is envisaged to be deployed in crowded situation with arbitrary settings. By over-segmenting the FTLE field, we are making the algorithm responsible for finding the exact number of flow segments. We used the standard procedure where the weight $w(i, j)$ between nodes is calculated from the FTLE field.

Let us assume that the over-segmented field S has k number of segments. Let s_1 and s_2 be the two neighboring segments and P_1 and P_2 be the particles belonging to s_1 and s_2, which are at the shared boundary. Let the size of P_1 and P_2 be n and m respectively. Fig. 5 shows an example of these two sets of particles for a pair of segments. In order to decide whether to merge these two segments, the Lyapunov exponent χ is computed between each pair of particles in P_1 and P_2 using the following equation first proposed by [15],

$$\chi(t) = \frac{1}{t} ln \frac{d(t)}{d(0)}, \tag{7}$$

where $d(t)$ is the distance between two particles at time t, initially separated by $d(0)$. The measure of Lyapunov divergence between the two sets of particles is now computed by

$$lypdiv(s_1, s_2) = \frac{1}{nm} \sum_{i=1}^{m} \sum_{j=1}^{n} \frac{1}{t} ln \frac{d_{i,j}(t)}{d_{i,j}(0)}, \tag{8}$$

where nm is the normalizing factor and $d_{i,j}(t)$ is computed by advecting the particles i and j and measuring the Euclidian distance between them. The segments are merged if $lypdiv(s_1, s_2)$ is less then some threshold. Note that before starting the merging process, we can remove those segments where the magnitude of the flow is zero. We call such segments 'vacuum segments'. Examples of over-segmentation and subsequent merging are shown in Fig. 6.

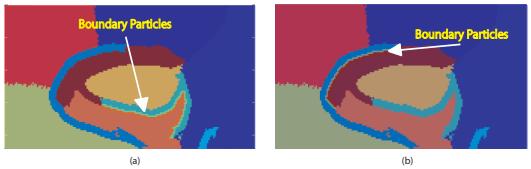

Figure 5. Boundary particles are highlighted which are used for measuring Lyapunov divergence between the segments.

5. Flow Instability

Given flow segments, we define the problem of locating flow instability as detecting the change in the number of flow segments. Recall that in our flow segmentation framework, boundaries of the flow segments having different dynamics are reflected as LCS in the FTLE field. Due to this formulation any changes in the dynamic behavior of the flow will cause new LCS to appear in the FTLE field at exactly those locations where the change happens. These new LCS will eventually give rise to new flow segments which were not there before. By detecting these new flow segments, we will be in a position to identify the locations in the scene where the flow is changing its behavior.

For detecting new flow segments, we establish correspondence between the flow segments which are generated from two consecutive blocks of video. Let us denote these as blocks A and B. The shape of a flow segment is represented by a Gaussian probability distribution of the spatial coordinates of pixels belonging to the flow segment. The mean of the Gaussian distribution is initialized with the mean of the spatial coordinates, while the variance is initialized with the variance of the coordinates of the boundary pixels. A voting scheme is employed for establishing correspondence between flow segments of block A and B. Each pixel within a flow segment of block B votes for one of the flow segments in block A. A flow segment from block B corresponds to a flow segment in block A if the majority of the pixels from a segment in B have voted for that flow segment in A. A flow segment in B whose correspondence cannot be established is 'flagged' as an instable flow region.

6. Experiments and Discussion

The approach is tested on videos taken from the stock footage web sites (Getty-Images, Photo-Search) and Video Google, representing high density crowd and high density traffic scenes. The second set of videos are extracted from the National Geographic documentary, titled 'Inside Mecca'. For each video, a flow field is computed using the algorithm described in Section 4.1. The values used for parameters n and T are kept between the range of 5 to 10. For advection, the resolution of the particle grid is kept the same as the number of pixels on which the flow field was calculated. The FTLE field is computed from the spatial gradient tensor of the flow maps using Equation 6. For segmentation the number of input segments to the normalized cuts algorithm is varied from 12 to 20 segments. The output of the normalized cuts algorithm is then merged using the the Lyapunov divergence measure described in Equation 8. This process is repeated for each block of frames to obtain the flow segmentation at consecutive time intervals.

Fig. 6 shows some of the qualitative results of flow segmentation. The first sequence shown in Fig. 6(a) is extracted from the documentary 'Inside Mecca'. It shows thousands of people circling around the Kabba in an anticlockwise direction. Here the notion of a 'physically and dynamically meaningful flow segment' implies that the entire group of people circling around Kabba are part of the same flow segment, since they are performing exactly the same task. The flow field for this sequence (Fig. 2) offers a unique challenge. All of the flow vectors within the circle have different directions, which means a simple clustering of these vectors will not assign all of the vectors to the same cluster, when in fact, they all belong to one cluster. However, we are able to handle this situation by quantifying the dynamics of the crowd flow using Lagrangian particles. The LCS previously shown in Fig. 4(a) show that the dynamics of the crowd are preserved by emphasizing the boundaries of the coherent flow regions.

The second sequence (Fig. 6(f)) shows a high density traffic scene. Vehicles are moving in two opposite directions on the main section of the highway, while a third group of vehicles is merging onto the main highway from the ramp. The challenge is to find the right segmentation of the flow generated by the traffic on the ramp. Based on our formulation, all of the particles which have the same fate (or destination) are part of the same flow, which means the flow generated by the traffic on the ramp should be part of the flow generated by the lanes on the right hand side of the highway. The result in Fig. 6(f) shows that our method is able to handle the aforementioned situation and has accurately

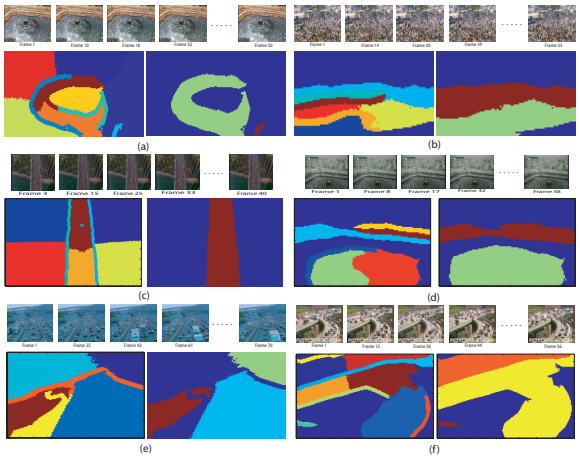

Figure 6. The flow segmentation results on some of the sequences. For each result, the image on the left shows the over-segmented field, while the image on the right shows the final segmentation mask obtained after the merging. Top row in each image shows the sequence of frames itself.

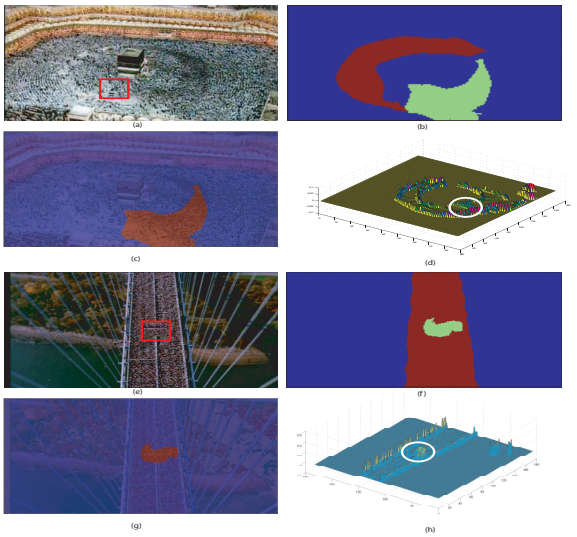

Figure 7. Results of the instability detection algorithm. (a,e) Bounding box is showing the synthetic instability. (b,f) Flow segmentation. (c,g) Detected instable region. (d,h) LCS around the detected instability.

segmented out the flow segments which are physically and dynamically meaningful in this scene.

The second set of experiments is performed for the detection of flow instability for which we inserted synthetic instabilities into the original video sequences. For each experiment, the original video sequence is used to compute the flow segments corresponding to the normal flow of the crowd and the spatial pdf is initialized for each segment as described in section 5. After the learning stage, the flow segmentation is performed on the block of frames that contain the synthetic instability. Synthetic instability itself is created by randomly placing a bounding box within the flow, and flipping or rotating it to change the behavior of the flow at that location. The correspondence of the flow segments generated from the synthetic sequence is established with the learned set of segments, using the procedure described in Section 5. Fig. 7 show the results of these experiments. In the case of the first sequence (Fig. 7(a-d)), the instability has caused a barrier in the flow, which resulted in the break up of the original segment into two parts, as shown in Fig. 7(b). The segment for which the correspondence can not be established is flagged as a potential instable flow region, as shown in Fig. 7(c). The emergence of new LCS in the FTLE field, as shown in Fig. 7(d) (circled in white), validates the observation that any change in the dynamics of the flow will result in new LCS that can be used to locate the instabilities. The second sequence shown in Fig. 7(e-h) captures a bird's eye view of a New York City marathon, where again our algorithm was able to locate and flag the potential problem region.

7. Conclusion

We have proposed a mathematically exact framework based on Lagrangian particle dynamics for crowd flow segmentation and flow instability detection. The framework is tested on a challenging set of videos taken from different web sources. Future directions include improving the flow instability detection by incorporating properties of individual particles into the spatial representation of flow segments.

Acknowledgement: This research was funded in whole by the U.S. Government VACE program.

References

[1] G. Haller et.al, "Lagrangian Coherent Structures and Mixing in Two Dimensional Turbulence", Physica D, 147, 2000.

[2] G. Haller, "Finding Finite-Time Invariant Manifolds in Two Dimensional Velocity Data", Chaos, 10(1), 2000.

[3] P. Tu et. al., "Crowd Segmentation through Emergent Labeling", In ECCV Workshop SMVP, 2004.

[4] A. C. Poje et. al, "The Geometry and Statistics of Mixing in Aperiodic Flows", Physics of Fluids, 11, 1999.

[5] G. Brostow et. al., "Unsupervised Bayesian Detetcion of Independent Motion in Crowds", IEEE CVPR, 2006.

[6] G. Haller, "Distinguished Material Surfaces and Coherent Strcutures in Three Dimensional Fluid Flows", Physics of Fluids, 2000.

[7] P. Reisman, "Crowd Detection in Video Sequences", IEEE Intelligent Vehicles Symposium, 2004.

[8] V. Rabaud et. al., "Counting Crowded Moving Objects", IEEE CVPR, 2006.

[9] T. Zhao et. al., "Bayesian Human Segmentation in Crowded Situations", IEEE CVPR03, 2003.

[10] S. C. Shadden et. al, "Definition and Properties of Lagrangian Coherent Structures from Finite Time Lyapunov Exponents in Two Dimensional Aperiodic Flows", Physica D, 212, 2005.

[11] G. Lapeyre et. al, "Characterization of Finite Time Lyapunov Exponents and Vectors in Two Dimensional Turbeulence", Chaos, 2002.

[12] F. Lekien et. al, "Tricubic Interpolation in Three Dimensions", Journal of Numerical Methods and Engineering, 63(3), 2005.

[13] D. Yang et. al., "Counting People in Crowds with a Real-Time Network of Simple Image Sensors", ICCV, 2003.

[14] J. Shi et. al, "Normalized Cuts and Image Segmentation", IEEE PAMI, 2000.

[15] P. Bryant et. al., "Lyapunov Exponents from Observed Time Series", Physics Review Letters, 65(13), 1990.

[16] A. B. Chan et. al., "Mixtures of Dynamic Textures", IEEE ICCV, 2005.

[17] P. Sand et. al., "Particle Video: Long-Range Motion Estimation using Point Trajectories", IEEE CVPR, 2006.