Published in final edited form as:

Comput Vis Image Underst. 2011 April ; 115(4): 531–540. doi:10.1016/j.cviu.2011.01.001.

Optimal-Flow Minimum-Cost Correspondence Assignment in Particle Flow Tracking

Alexandre Matov¹, Marcus M. Edvall², Ge Yang^{1,3}, and Gaudenz Danuser^{1,*}

¹ Department of Cell Biology, The Scripps Research Institute, La Jolla, CA 92037

² Tomlab Optimization Inc., Pullman, WA 99163

³ Lane Center for Computational Biology & Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

A diversity of tracking problems exists in which cohorts of densely packed particles move in an organized fashion, however the stability of individual particles within the cohort is low. Moreover, the flows of cohorts can regionally overlap. Together, these conditions yield a complex tracking scenario that can not be addressed by optical flow techniques that assume piecewise coherent flows, or by multiparticle tracking techniques that suffer from the local ambiguity in particle assignment. Here, we propose a graph-based assignment of particles in three consecutive frames to recover from image sequences the instantaneous organized motion of groups of particles, i.e. flows. The algorithm makes no a priori assumptions on the fraction of particles participating in organized movement, as this number continuously alters with the evolution of the flow fields in time. Graph-based assignment methods generally maximize the number of acceptable particles assignments between consecutive frames and only then minimize the association cost. In dense and unstable particle flow fields this approach produces many false positives. The here proposed approach avoids this via solution of a multi-objective optimization problem in which the number of assignments is maximized while their total association cost is minimized at the same time. The method is validated on standard benchmark data for particle tracking. In addition, we demonstrate its application to live cell microscopy where several large molecular populations with different behaviors are tracked.

Index Terms

Particle tracking; vector field; multi-directional flows; graph algorithms; multi-objective optimization

1 Introduction

Computer vision algorithms for tracking particle flow have applications in fields as diverse as machine and stereo vision [1,2], perceptual grouping [3], video surveillance [4,5], meteorology [6], and astronomy. Recent reviews of particle tracking methods can be found

* Corresponding Author's current address: Harvard Medical School, 240 Longwood Ave, Boston, MA 02115, Gaudenz_Danuser@hms.harvard.edu, tel: +1 617 432 7941.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

in [7]. Experimental fluid dynamics is another field with extensive tracking applications [8]. Here, particle-based analyses of flow fields are often complemented with texture-based approaches that do not rely on image segmentation [9], yet assume piecewise coherent flow conditions [10]. Less known by the computer vision community, particle tracking is also a very prominent field in biophysics, where the goal is to follow the motion of single molecules and molecular complexes [11-13]. It is in this context that we have begun to develop particle tracking methods.

Particle-based tracking requires the formulation of a motion model that defines an explicit or implicit cost of particle correspondences between frames and an assignment process. Multiple-hypothesis tracking (MHT) computes the most probable particle path by achieving globally optimal assignments both in space and time [14,15]. However, this approach leads to a combinatorial explosion in cases where many particles move in close proximity. Greedy optimization schemes have been proposed as alternatives. They reduce the computational cost by compromising temporal globality [16-19]. In live cell microscopy particles generally originate from fluorescently labeled molecules incorporated in sub-cellular structures. They have low stability and move in very dense, complex motion patterns, rendering the application of MHT and most of its approximations impossible. Thus, we sought to develop a method that captures the instantaneous state of particle flows from a minimal number of consecutive frames. We require particles to be present for a minimum of three time points and assume that between three time points they move at constant speed and preserve directionality. This permits definition of a local motion model similar to those in [16,20]. Additionally, very often sub-populations of particles undergo organized movements, while others perform a random walk. The relative size of the populations and the number of sub-populations with distinct directions of organized motion are not known a priori. To extract these populations, we complement the local motion model with a motion model that accounts for regional coherence among particle movements.

Like other particle tracking methods [16,21-24], we employ graph-based approaches to assign corresponding particles between frames. In general, these methods maximize the number of assignments and only then minimize the assignment cost. However, when particles are dense and unstable, maximizing the number of assignments tends to generate a large number of false positives. The risk of over-linking is further increased in data with a substantial portion of particles present for only one frame, and where multiple particle flow fields interpenetrate, both leading to high particle clutter. Under such conditions it is advantageous to identify a subset of particles with high confidence correspondences while leaving out uncertain links. On the other hand, to capture the spatial details of the particle motion field the number of assignments should still be maximized. Balancing the rejection of false positive assignments without generating a high number of false negative assignments is one of the key challenges in real-world tracking problems. The main contribution of this paper consists of proposing a multi-objective optimization method that determines the trade-off between the two competing requirements to establish robust particle correspondences over short time windows and thus extracts on-the-fly information on the instantaneous organization of particle flows in complex scenes. We compare the performance of this approach to related work on standard benchmark examples for particle tracking. We then introduce measures for the complexity of a tracking problem and demonstrate the performance of the proposed optimal-flow minimum-cost particle assignment on significantly more difficult problems on a particle image velocimetry example with ground truth and on fluorescence microscopy sequences capturing multi-directional, unstable flows in the foreground of a significant background population of particles with unorganized movements.

2 Tracking Problem Statement

Consider a set of N_f particles $\mathbf{P}_f = \{p_1, \dots, p_{N_f}\}$ detected in frame f . Our goal is to define a particle flow field in this frame with instantaneously smooth motion that is represented by a set of tracks connecting each one of a subset of particles in $\mathbf{P}_f$ to one and only one particle in the set $\mathbf{P}_{f-1}$ of N_{f-1} particles detected in frame $f-1$ and to one and only one particle in the set $\mathbf{P}_{f+1}$ of N_{f+1} particles detected in frame $f+1$. To identify these tracks we assemble a graph $G(\mathbf{V}, \mathbf{E})$ with a source $s \in \mathbf{V}$, a sink $r \in \mathbf{V}$, vertices $\mathbf{V}_P \in \mathbf{V}$ comprising the union of particles detected in the three consecutive frames, i.e. $\mathbf{V}_P = \mathbf{P}_{f-1} \cup \mathbf{P}_f \cup \mathbf{P}_{f+1} = \{p_1, \dots, p_N\}$ with $N = N_{f-1} + N_f + N_{f+1}$, and vertices $\mathbf{V}_T \in \mathbf{V}$ comprising M candidate triplets $\mathbf{V}_T = \{t_1, \dots, t_M\}$ (Fig. 1a,b). Each candidate triplet t_i consists of one particle in $\mathbf{P}_{f-1}$, one particle in $\mathbf{P}_f$, and one particle in $\mathbf{P}_{f+1}$, thus defining a candidate particle track over three consecutive frames. The construction of candidate particle tracks requires that particle displacements between frames are less than a prior estimate of the maximal particle velocity times the time interval between frames (Fig. 1a; see also Section 3). The graph contains three types of edges. First, edges $(s, t_i) \in \mathbf{E}$ between source and any one of the triplets. They have a capacity $a(s, t_i) = 3$ and a cost $c(s, t_i)$. The cost is derived from the motion smoothness of the three-frame track associated with triplet t_i . Second, edges $(t_i, p_h) \in \mathbf{E}$ connect triplets to the three contributing particles. Third, edges $(p_h, r) \in \mathbf{E}$ connect particles to the sink. Flow conservation throughout the graph

requires that
$$a(s, t_i) = \sum_{\forall p_h, a(t_i, p_h) \neq 0} a(t_i, p_h) = 3$$
 and $a(t_i, p_h) = a(p_h, r) = 1$. The cost values associated with the second and the third type of edges is null, i.e. $c(t_i, p_h) = c(p_h, r) = 0$.

The first question is how many triplets $\hat{M} \leq M$ with mutually exclusive particles exist in the set $\{t_1, \dots, t_M\}$. Given the graph $G(\mathbf{V}, \mathbf{E})$, answering this is equivalent to solving the maximum source-to-sink flow [25], which yields one or multiple subgraphs $\{G'_1, \dots, G'_K\}$. Each subgraph contains a subset of edges and vertices (i.e triplets and particles) carrying a maximum flow $3\hat{M}$. The second question is then which of the subgraphs $\hat{G}$, with $\hat{G} \in \{G'_1, \dots, G'_K\}$, has overall minimal cost. This is equivalent to solving the minimum cost flow problem [25]. Under conditions of dense and unstable particle fields enforcing maximum flow tends to select many false positive triplets at the expense of a smaller set of true positive triplets. Whereas false positives could be eliminated posterior to the track identification by criteria of motion smoothness, this data cleaning would not allow the recovery of true positives that were missed in fulfilling the condition that triplet selection must be comprised of mutually exclusive particles.

To increase the robustness of particle track selection, we propose to relax the maximum-flow condition to find a flow $3M^*$ with $M^* < \hat{M}$ triplets. We refer to this procedure as an *optimal-flow minimum-cost* algorithm. Specifically, the objective of the *optimal-flow minimum-cost* particle assignment is to identify a subgraph G^* with a minimal overall cost carrying a reduced (less than the maximum) flow so that low confidence triplets are eliminated. Evidently, this problem has a trivial solution $M^* = 0$. Thus, we introduce as the second objective the maximization of the source-to-sink flow carried by G^* :

$$G^* = \arg \min_{G'' \in G} \left(\sum_{G''(\mathbf{V}, \mathbf{E}_i)} c(s, t_i) \right) \&\& \arg \max_{G'' \in G} \left(\sum_{G''(\mathbf{V}, \mathbf{E}_i)} a(s, t_i) \right)$$

where

$$\sum_{G^*(\mathbf{V}, \mathbf{E}_i)} a(s, t_i) = 3M^* \quad (1)$$

Of note, by formulating the multi-objective problem Eq. 1 the sequential procedure of the maximum-flow minimum-cost computation is replaced by a competition of selecting only high-confidence assignments against generating as many assignments as possible. The triplets in G^* are then defining the particle tracks that represent the instantaneous particle flow in the three-frame sequence.

3 Construction of Particle Triplets

To build the triplets $\{t_1, \dots, t_M\}$ we apply the following three steps:

1. Build a set of candidate assignments between particles in frame f and frame $f-1$ for which $\|\mathbf{x}(f) - \mathbf{x}(f-1)\| \leq r_0$ and a set of candidate assignments between particles in frame f and frame $f+1$ for which $\|\mathbf{x}(f+1) - \mathbf{x}(f)\| \leq r_0$. The search radius r_0 defines the largest particle displacement allowed between consecutive frames.
2. Construct candidate triplets by pairing all candidate assignments from the first and from the second set that share the same particle in frame f . Unpaired candidate assignments are ignored in the graph construction.
3. For each candidate triplet, compute a cost function that penalizes changes of track direction and speed.

3.1 Selection of the Search Radius

We introduce a search radius r_0 to limit the search space for triplet construction. Values for r_0 can be derived from prior knowledge of the maximum speed of particles. If this information is unavailable, we define the radius so that the number of candidate triplets is approximately one order of magnitude greater than the average number of particles per frame. Accordingly, a large majority of particles participates in more than one triplet. Increasing the search radius beyond this point has virtually no effect on the cost distribution of the triplets selected by a maximum-flow minimum-cost computation (Fig. 1c). At this point it is likely that all true three-frame tracks are included in the set of candidate triplets and the triplet selection is sufficiently constrained by the condition of mutual exclusivity. Any larger search radii generate additional candidate triplets that merely increase computational cost of the solution without contributing true positives.

3.2 Cost Computation

We apply a local and regional motion model to define the cost of a triplet. The cost function iteratively weighs the influences of the two models based on the posterior distribution of model variables that is generated by the selected candidate tracks (see below). The iterative updating of model weights is terminated when the population of selected triplets varies less than 5% between iterations. For most tracking problems with a few hundred particles 3 – 4 iterations are sufficient for convergence because stable statistics can be drawn from the large number of selected triplets.

3.2.1 Local Motion Model—The angle α_i between the first and the second segment of each triplet and the difference between their respective lengths δ_i define the variables of the local motion model, analogous to the motion model proposed by Sethi and Jain [20] (Fig. 1d). In large particle fields both variables have a zero mean and are normally distributed. The latter requirement can be further relaxed to unskewed distributions. In our implementation deviations from these conditions are monitored during the iterations and tracking solutions are rejected in case these assumptions are not fulfilled. To initiate the tracking, values of the distribution variances, σ_α^2 and σ_δ^2 , are defined by an initial assumption

of the particle behavior. Subsequently, the values are updated based on the distributions of the selected triplets, see below.

3.2.2 Model of Regional Flow Organization—After a first selection of triplets by optimal-flow minimum-cost computation (see Section 4) the preferential direction of particle flow is determined in each location of the image by anisotropic vector field filtering of the resulting particle displacements between frames $f-1$ and $f+1$ [26]. If the data contains interpenetrating flow fields, the filtering step is preceded by directional clustering and classification of particles into subfields (see Section 4.3). The filtered particle flow fields are employed in subsequent iterations to define an additional term in the cost function that penalizes regional flow incoherence. For each triplet the alignment of the corresponding three-frame track with the filtered particle flow field is determined by the angle β_i between local flow direction and the vector from the start to the end of the track (Fig. 1d). By construction of the filtered particle flow, β_i is a random variable with zero mean and a symmetric distribution.

3.2.3 Cost Defined by Local and Regional Motion Models—We then assume that the three variables characterizing the local and regional motion models are mutually independent random variables. Thus, the cost of a triplet is expressed as the Mahalanobis distance between $[\alpha_i, \delta_i, \beta_i]$ and the zero mean average of each of the distributions:

$$c_i(s, t_i) = \frac{1}{\sigma_\alpha^2} \alpha_i^2 + \frac{1}{\sigma_\delta^2} \delta_i^2 + \frac{1}{\sigma_\beta^2} \beta_i^2. \quad (2)$$

Importantly, the balancing of the weights between the three components of the cost function is entirely data-driven and requires no user-adjustment between different data sets.

4 Optimal-Flow Minimum-Cost Computation

To select from the candidate triplets those fulfilling the optimal-flow minimum-cost condition in Eq. 1 we perform six steps:

1. Calculate the maximum source-to-sink flow that can be carried by the graph defined in Fig. 1b.
2. Compute the maximum-flow minimum-cost solution.
3. Construct the Pareto optimal front (POF) [27] to the bi-objective problem Eq. 1 with equidistant spacing on both objective axes.
4. Define the optimal-flow from the location on the POF closest to the utopia point and recompute the minimum-cost solution enforcing optimal-flow.
5. Cluster the track vectors associated with the selected triplets into multiple flow fields and/or filter the tracks to calculate particle flow fields.
6. Recompute the cost of the triplets in the graph, taking into account the updated particle flow field for computing the cost contribution of the regional motion model and the updated distributions of $[\alpha, \delta, \beta]$ for computing the relative weights between the components in Eq. 2.

Steps 2 – 6 are repeated iteratively until the change in selected triplets is $<5\%$. At each iteration, the maximum-flow minimum-cost solution is recomputed using the updated costs for the triplets and a new POF is constructed. Of note, the set of candidate triplets depends

only on the search radius and thus is not affected by cost updates. As discussed in the following sections steps 2 – 4 can be accomplished at near real-time speed on a standard PC, even for several thousand triplets, using linear programming and an approximation of the POF. Generally, three iterations suffice to obtain a solution for which over 95% of the selected triplets are identical to the selection of the previous iteration.

4.1 Maximum-Flow Minimum-Cost Computation to Initialize Optimal-Flow Minimum-Cost Computation

To ensure that the maximum-flow minimum-cost computation in the graph of Fig. 1b does not generate residual capacity in any of the vertices linking the source to the selected triplet, we reformulate the underlying minimum-cost flow problem as a binary programming problem [25]:

$$\begin{aligned} \hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \{0,1\}} (\mathbf{c}^T \mathbf{x}) \\ \text{subject to:} \\ \mathbf{A}\mathbf{x} = \mathbf{b} \end{aligned} \quad (3)$$

The definition of the sparse adjacency matrix $\mathbf{A}$ and of the supply-demand vector $\mathbf{b}$ are given in Table 1. The first M elements of the binary solution vector $\hat{\mathbf{x}}$ indicate the minimum cost selection of M' triplets for any given source-to-sink flow $3M' \leq 3\hat{M}$ through the graph, with $\hat{x}_i = 1$ if the triplet is selected and $\hat{x}_i = 0$ if it is excluded from the solution. Importantly, by the definition of $\mathbf{A}$ each selected triplet receives three flow units from the source and passes one unit on to each of its three associated particles. Flow conservation is maintained for vertices representing individual particles in the three consecutive frames. To supply the graph with sufficient flow for the selection of M' triplets, we set $\mathbf{b}_1 = 3M'$. To enforce flow conservation in all vertices, the graph must be drained completely at the sink, hence the last element $\mathbf{b}_{M+N+2} = -3M'$, while all other entries of $\mathbf{b}$ are zero. Flow conservation in all vertices implies that triplets are never split up.

To determine $\hat{M}$, i.e. the maximum number of mutually exclusive triplets in the graph, Eq. 3 is first solved with equal cost for all edges and an inequality constraint $\mathbf{b}_l \leq \mathbf{A}\mathbf{x} \leq \mathbf{b}_u$ with $\mathbf{b}_{l,1} = 0, \mathbf{b}_{u,1} = 3M$ and $\mathbf{b}_{l,M+N+2} = -3M, \mathbf{b}_{u,M+N+2} = 0$ defining the upper bound for supply and demand in source and sink. Solving Eq. 3 in a second step with $\mathbf{b}_1 = -\mathbf{b}_{M+N+2} = 3\hat{M}$ yields a maximum-flow minimum-cost selection of triplets.

4.2 Multi-Objective Optimization Based on Pareto Optimality between Flow and Cost Reduction

By maximizing the flow prior to identifying the minimum-cost configuration, solutions with a large number of triplets are favored over solutions with fewer triplets, although the additional triplets may be very costly (Fig. 2a). Slight relaxation of the maximum flow condition allows the solution to converge to a selection of triplets with dramatically decreased overall cost (Fig. 2b). Figure 2c shows the cost distribution for triplets in a real-world data example (see Figure 5) for various levels of flow reductions. Starting with the cost distribution produced by the maximum-flow minimum-cost solution, a sub-supply-demand $\mathbf{b}_1 = -\mathbf{b}_m = 3M'$, $M' < \hat{M}$ in Eq. 3 substantially reduces the number of high-cost triplets. Critically, the removal of high-cost triplets rescues low-cost triplets, which under maximum-flow conditions are sacrificed for the benefit of selecting the largest possible number of triplets. Thus, the number of triplets in the lowest-cost bin is higher for all sub-maximum flow conditions. The relationship between flow reduction and cost reduction is non-linear

but convex. In the beginning, a slight decrease in flow results in a significant drop of the cost. Beyond a certain optimal flow, the improvement in the total cost is insignificant vis-à-vis the loss of triplets. Removal of too many triplets can even result in elimination of low-cost triplets (Fig. 2c).

The search for the optimal flow defines a bi-objective optimization problem with two competing goals:

$$x^* = \arg \min_{x \in \{0,1\}} (c^T x) \quad \&\& \quad \arg \max_{x \in \{0,1\}} \left(\sum_{i=1}^M x_i \right). \quad (4)$$

The second goal is equivalent to minimizing the number of triplets $R = \hat{M} - M^*$ removed from the maximum-flow solution. Thus, Eq. 4 can be reformulated as a convex single objective problem by applying the weighted sum method [28],

$$\begin{aligned} x^* &= \arg \min_{x \in \{0,1\}} (\alpha c^T x + 3R(x)) \\ \text{subject to:} & \\ Ax &= b \\ b_1 &= -b_{M+N+2} = 3(\hat{M} - R) \end{aligned} \quad (5)$$

where α defines the weight between cost reduction and loss of triplets.

We set $\alpha^{-1} = E(c^T \hat{x} / \hat{M})$ and estimate $E(c_i) \approx \text{median}(c_i)$, $\forall i$ with $\hat{x}_i = 1$, where $\hat{x}$ denotes the solution of the initial maximum-flow minimum-cost problem; thus, $E(\hat{c}_i) = E(\alpha c_i) = 1$. Underlying the choice $E(c_i) \approx \text{median}(c_i)$ is the assumption that the majority of the triplets selected by maximum-flow minimum-cost computation have costs similar to the triplets to be selected by optimal-flow minimum-cost computation. Under this transformation, removal of a triplet with average cost triplet results in equal steps on the axis of flow reduction (Fig. 2d horizontal) and on the axis of total cost reduction (Fig. 2d vertical). Removal of a high-cost triplet yields a larger cost than flow reduction, whereas removal of a low-cost triplet yields less cost than flow reduction. Thus,

$$\begin{aligned} \tilde{x}(R) &= \arg \min_{x \in \{0,1\}} (\tilde{c}^T x) \\ \text{subject to:} & \\ Ax &= b \\ b_1 &= -b_{M+N+2} = 3(\hat{M} - R), \forall R \in \{0, \dots, \hat{M}\} \end{aligned} \quad (6)$$

describes the convex POF for the two competing objectives in Eq. 4 (Fig 2d, dashed line). The optimal solution requires that cost and flow reduction are equal, hence

$$\left| \frac{\partial \tilde{x}(R)}{\partial R} \right| = 1. \quad (7)$$

The solution to Eq. 7 is geometrically defined as the location on the POF closest to the origin, i.e. the Utopia point (Fig. 2d, diamond). For large problems with thousands of particles, computation of the exact POF according to Eq. 6 would be computationally expensive. Thus, we opted to construct an approximation that avoids re-calculating the minimum-cost flow problem for every cost reduction. To achieve this we first rank all triplets in the order of ascending normalized cost. Then, the optimal flow is determined as (Fig. 2d, square)

$$M^* = \arg \min_{0 < M' \leq \widehat{M}} \left(\sum_{i=1}^M \widetilde{c}_i + 3(\widehat{M} - M') \right),$$

with $\widetilde{c}_i \leq \widetilde{c}_{i+1} \forall i, i=1 \dots \widehat{M}.$

(8)

The final selection of triplets under optimal-flow minimum-cost conditions is performed by solving Eq. 3 subject to a supply-demand vector with $\mathbf{b}_1 = -\mathbf{b}_{M+N+2} = 3M^*$.

The approximation requires minimum-cost flow computation only twice. Hence, the optimal-flow minimum-cost triplet assignment is accelerated by two to three orders of magnitude. The distance between the exact and the approximate POF depends on the application. By construction, the exact front is below the approximate front. From this follows, the point on the approximate POF closest to the utopia point is always associated with a larger flow reduction than the exact solution. This means that the approximation never introduces additional false positives over the exact solution, the main objective underlying the design of our algorithm.

4.3 Filtering and Mapping of Particle Flow Fields

For graphical representations of flow fields and to enforce regional flow coherence in the computation of the cost of each triplet (see 3.2.2.) we spatially filter the tail-to-head vectors of the selected optimal-flow minimum-cost triplets. To account for the possibility of interpenetrating flow fields we first identify clusters of vectors with distinct directionality. Initially, the number of clusters and their angular orientation is determined using an expectation-maximization algorithm [29]. For reasons of speed we omit this step for image sequences with a-priori known number of flow. Subsequently, a k-means algorithm, modified to account for the cyclic distribution of vector directions, is used to assign triplets to a particular cluster. Each of the flow vector clusters is smoothed individually by fast anisotropic Gaussian filtering [26]. If the density of vectors per cluster and time-point is insufficient for spatial filtering, but the particle flows are quasi-stationary at the time scale of the movie, optimal-flow minimum-cost triplets from multiple consecutive three-frame tracking solutions are integrated.

5 Results

We first benchmarked the optimal-flow minimum-cost particle tracking on standard particle tracking data sets used in the computer vision literature and then analyzed particle image velocimetry (PIV) benchmarking data sets with known ground truth and high-resolution light microscopy time-lapse sequences of dynamic molecular assemblies. Besides the much higher particle density, these sequences set a significantly harder test case for tracking due to the low SNR of the raw movies, the instability of the particles, and the interpenetration of multiple flow fields.

Table 2 indicates parameters to score the level of difficulty of the presented tracking problems. In all examples we set the search radius two times the mean particle displacement

between consecutive frames. According to these parameters the first three examples could be tracked without a motion model. Rapid movement with symmetric motion can be tracked using the local motion model (example 4), while dense (example 5) and interpenetrating (example 6) flows require the use of both local and regional motion models. A particularly critical parameter for scoring the difficulty of a tracking problem is how many particles participate in more than one candidate assignment, i.e. in our case triplets (Table 2, last column). The results in Fig. 3 – 5 indicate the participation distributions for each of the problems.

5.1 Validation on Common Benchmark Data Sets for Particle Tracking

In this section we present four tracking examples (data kindly provided by Khurram Shafique) with associated difficulty scores listed in Table 2. The reader is encouraged to compare the results presented in Fig. 3 with those presented in [17]. The ‘rotating ball’ sequence exhibits numerous particle occlusions, leading to false positive triplets. Also, our tracker produces some errors caused by image border effects. For the ‘fish’ and ‘bird’ sequences our algorithm has produced no more than one false positive per ten frames (see Videos 1 and 2). Overall, the tracking results represent well the multidirectional flow of the particle ensemble as well the particular motion of individual particles. In all scenes a small fraction of false positives ($\ll 1\%$) is picked up by the algorithm, associated with random triplets of very low cost, i.e. straight trajectories of constant speed. Although it is not the goal of our algorithm to reconstruct complete trajectories, the high quality of triplets permitted a seamless amendment of triplets extracted from consecutive runs of the algorithm, first on frames 1,2,3, then on 2,3,4, etc. Tracks in the ‘rotating dish’ data [16] follow a ratio of particle mean displacement to nearest neighbor distance roughly three times higher than the first three sequences (Table 2, second column). While for the sequences Fig. 3a – c particles participate in 1-6 candidate triplets, 87% of the particles on the ‘rotating dish’ participate in multiple triplets, some in up to 20 triplets (histogram insets for each panel of Fig. 3). Nevertheless, optimum-flow minimum-cost tracking resolved 99.5% true positives and 0.5% false negatives, respectively (Fig.3d). The false negatives are an inherent consequence of optimal-flow computation in noise-free sequences, where some low cost tracks may be excluded for the benefit of robust rejection of high cost tracks. Indeed, the result does not contain a single false positive.

5.2 Particle Image Velocimetry Data

To further investigate the breakpoint of the algorithm relative to a ground truth we analyzed simulated particle image velocimetry data using the simulation software [30]. We generated images with more than 1300 particles with a mean particle displacement between frames of 4.5 pixels and a maximum particle displacement of 9 pixels. With this, the ratio between average particle displacement and average nearest neighbor distance between particles within a frame was >1 (Table 2), i.e. particles generally move beyond the position of their nearest neighbor particles. This renders the tracking problem significantly more challenging than any of the problems in the previous section. Moreover, to test the performance of the algorithms in rejecting true negatives, we perturbed 50% of the simulated particle positions within the range $[-9, 9]$ pixels to generate tracks that do not belong to the flow field. We set the search radius to 9 pixels. On average, particles participate in 24 candidate triplets, with some particles in as many as 120 (see histogram Figure 4b). To appreciate the difficulty of this tracking problem, the reader is encouraged to watch Video 3. Figure 4a presents the results from optimal-flow minimum-cost tracking. True positives are shown in blue, true negatives in black. Magenta and red colors indicate false negatives and false positives, respectively. The self-adaptive weights of the cost function converged to a solution within 3 iterations.

The specificity, i.e. the ratio true negatives to true negatives plus false positives, is 96%. The sensitivity, i.e. the ratio true positive to true positives plus false negatives, is 75%. While this performance seems relatively low at first sight, it is inherently connected to the design of the optimal-flow minimum-cost algorithm. By construction, false positives are avoided at the expense of generating some false negatives. However, in large and dense particle fields the loss of ~20% of particle trajectories is uncritical for the reconstruction of particle flows, if few of the accepted trajectories have errors. Overall, the high specificity demonstrates a robust performance of the tracker in resolving strongly perturbed, dense particle flows.

5.3 Application in Live Cell Imaging of Sub-Cellular Structural Dynamics

One area of tracking applications where high specificity is preferred over high sensitivity is in light microscopic imaging of sub-cellular structures. Most often, these structures generate diffraction-limited particles of very high density, i.e. the distances between particles are comparable to the particle size proper. A live cell imaging modality that generates such data is Fluorescent Speckle Microscopy (FSM) [31]. Fluorescent speckles are diffraction-limited images of fluorophore clusters that are randomly incorporated into macromolecular assemblies. Figure 5a shows an FSM image of the microtubule polymer of a metaphase meiotic spindle [32], a molecular machine that forms during cell division for the segregation the replicated DNA into two daughter cells. The polymers in this structure undergo bi-polar flux with areas of dense, anti-parallel sliding (see Video 4 for side-by-side comparison of the raw data (left) and the time evolution of the two anti-parallel particle flows (right)). Robust tracking and mapping of this flux is critical for cell biological analyses of the spindle architecture and the mechanisms that mediate DNA segregation. FSM imaging of the meiotic spindle yields time-lapse sequences with hundreds to thousands (>1000 in this example) of point features moving past one another at distances comparable to the size of the point feature.

5.3.1 Scale Invariant Speckle Detection—To identify speckles in each frame of the time-lapse sequence we used a three-step detection approach [33] based on scale-space theory [34]. In the first step we applied a band-pass filter to select the image frequencies associated with the speckles while attenuating frequencies associated with image noise or with longer range variations of the diffuse background. The filter was implemented as a difference of Gaussians (DoG). As speckles represent a diffraction-limited image feature the higher frequency cut-off was defined by a Gaussian $G(\sigma_1)$ with $\sigma_1 = 0.21\lambda/NA \cdot P_{xy}$ derived from the microscope point spread function [35]. Here, λ denotes the emission wavelength of the fluorescent label, NA the numerical aperture of the objective lens, and P_{xy} the pixel size in object space. The lower frequency cut-off was defined empirically by a Gaussian $G(\sigma_2)$ with $\sigma_2 = 1.1 \cdot \sigma_1$. In the second step we applied unimodal thresholding [36] to the DoG image intensity histogram to remove pixels that are unlikely to belong to speckles. After thresholding candidate speckle images were collected by connected component labeling. To further eliminate from this set of image objects false positives associated with protein aggregates or out-of-focus structures, which usually have a brightness and size different from speckles, we compared in a third step the unfiltered image intensities $I_i(x,y)$ of individual objects against the mean speckle image $I_{\text{speckle}}(x,y)$ estimated by spatial alignment and averaging of all extracted objects. Objects that significantly deviated from this mean image were detected again by unimodal thresholding of the histogram of the sum of the squared intensity differences between the pixels of individual objects and the pixels of the mean speckle image. As a result of the feature detection procedure, a list of speckle was obtained for each frame including the position of the object centroid (example shown on the lower half of Fig. 5a).

5.3.2 Tracking of Anti-Parallel Speckle Flows—Intrinsic to the mechanism of speckle formation in the spindle these particles are unstable with an average presence in 3 to 4 frames [37]. Three-frame tracks produced by optimal-flow minimum-cost tracking and subsequently separated by motion direction are displayed in Figure 5b. Fig. 5c shows the filtered vector field indicating the average local particle movement over three frames. Of note, the displacement of these particles exceeds the distance between particles and in the mid-zone opposing flows densely interlace. In addition, the density of particles fluxing northward (red) increases towards the north pole of the spindle and vice versa for southward fluxing particles (yellow). Together with the low particle stability, such spatial inhomogeneity adds significant difficulty to the tracking problem.

Figs 5d-f show the distribution of α , δ , and β from which the weights of the local and regional motion models were derived. Interestingly, the distribution of α is significantly broader than that of β , suggesting that the directional changes between consecutive frames are higher for individual tracks than the overall deviation of a track from the general direction of the particle flow. This characteristic can be appreciated in the supplementary video raw data and is related to fluctuation of individual microtubules in an overall stable polymer scaffold. While this behavior was unknown before applying the tracking, and thus could not be enforced by user inputs, the adaptive adjustment of the cost function robustly uncovered this behavior. The weights converged in 3 iterations.

Figs 5h and 5i depict a raw image and the optimal-flow minimum-cost tracking of a tri-polar particle flow field in a sub-region of a meiotic spindle (Video 5 shows the raw images of the spindle (left), and the cropped and enlarged area of the three overlapping particle flow fields (right)). This data is included as an example to illustrate the adjustment of the local and regional motion models to an a priori unknown number of overlapping flow directions. Again, because of the strict rejection of false positives, the iterative classification of particle tracks locked in with three clusters although there is no region in this small field of view with fewer than three motion directions. This permitted the reconstruction of a complex flow configuration, allowing further biophysical analysis of this peculiar polymer dynamics.

6 Conclusion

We have developed an algorithm for the resolution of multiple intersecting and/or interlacing particle flow fields of high density and with low stability of individual particles. The approach captures the instantaneous state of particle flows by consideration of three time points at the time and a self-calibrating cost function, permitting robust tracking in a wide range of applications. The robustness of the tracking relies on i) the use of a single user-defined control parameter, the search radius; ii) fast convergence of an iterative solution that adjusts all other control parameters on-the-fly; and iii) particle appearances and disappearances are handled reliably during particle assignments into tracks. The distinguishing concept of the approach is to explicitly balance the selection of false positive and false negative assignments; in contrast to most existing particle-based tracking methods which implicitly maximize the number of assignments. Generally, this introduces numerous false positive assignments, or it requires the introduction of a cost threshold beyond which assignments are rejected. By performing multi-objective optimization local and regional motion models can be imposed adaptively to ever-changing particle behaviors without over-linking particles.

Practical applications to illustrate the performance of optimal-flow minimum-cost particle tracking were drawn from microscope image sequences capturing the dynamics of molecular machines in cells. These data pose substantial challenges in terms of robustness requirements, complexity of interdigitating flow fields, and spatiotemporal variation of

particle behavior within the flow field. Tracking in unstructured crowded scenes has recently gained attention in computer vision for the surveillance of human or vehicle motion in vulnerable areas such as airports, train stations or highways [38]. These problems have very similar characteristics to the cell biological applications discussed in this paper. Although it is outside the scope of this report to formally compare optimal-flow minimum-cost particle tracking in surveillance, it is worth pointing out that the presented methods will bring several strengths to these types of applications. First, in contrast to the algorithm proposed by Ali *et al.* [39] which is limited to one flow direction per location in a scene, our method can resolve several intermingled flows and it is robust also in presence of a large fraction of stationary or randomly-moving particles. Second, our algorithm requires as few as three frames to capture the changes in the configuration of multidirectional flow fields. If the motion models are known, for instance straight equidistant particle steps between frames, the weights of the cost function can be fixed. This will increase the speed of the solution several-fold compared to our currently iterative cost function adjustment, allowing real-time tracking of particle flow fields. Thus, it could be used to predict on the fly anomalous behaviors or congestions associated with a security alert at an airport terminal or with the arrival of a new train in a highly frequented station. On the other hand, our algorithm relies on an explicit detection of objects for tracking. This may limit its use in scenes with low granularity, a problem that has been successfully addressed by Rodriguez et al. [40] using optical flow methods. Future work will focus on testing the here proposed method on this problem and to combine it with texture segmentation that relax the requirement for discrete particle detection.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work is supported by the grant U01 GM 67230 to G.D. The authors would like to thank Khurram Shafique for benchmarking data.

References

1. Barron J, Fleet DJ, Beauchemin S. Performance of Optical Flow Techniques. *Int'l J Comp Vision*. 1994; 12:43–77.
2. Micheli ED, Torre V, Uras S. The accuracy of the computation of optical flow and of the recovery of motion parameters. *IEEE Trans Pattern Anal Mach Intell*. 1993; 15:434–447.
3. Medioni G, Cohen I, Bremond F, Hongeng S, Nevatia R. Event Detection and Analysis from Video Streams. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2001; 23:873–889.
4. Hu W, Tan T, Wang L, Maybank S. A Survey on Visual Surveillance of Object Motion and Behaviors. *IEEE Transactions on Systems, Man, and Cybernetics*. 2004; 34:334–352.
5. Morris BT, Trivedi MM. A Survey of Vision-Based Trajectory Learning and Analysis for Surveillance. *IEEE Transactions on Curcuits and Systems for Video Technology*. 2008; 18:114–1127.
6. Wu QX. A correlation-relaxation labeling framework for computing optical flow - template matching from a new perspective. *IEEE Trans Pattern Anal Mach Intel*. 1995; 17:843–853.
7. Yilmaz A, Javed O, Shah M. Object tracking: A survey. *ACM Computing Surveys*. 2006; 38
8. Chetverikov D. Applying feature tracking to Particle Image Velocimetry. *International Journal of Pattern Recognition and Artificial Intelligence*. Jun.2003 17:487–504.
9. Keane RD, Adrian RJ. Theory of Cross-Correlation Analysis of Piv Images. *Appl Sci Res*. Jul.1992 49:191–215.

10. Okamoto K. Particle cluster tracking algorithm in particle image velocimetry. *JSME Int'l J Series B*. 1998; 41:151–154.
11. Kalaidzidis Y. Multiple objects tracking in fluorescence microscopy. *Journal of Mathematical Biology*. Jan.2009 58:57–80. [PubMed: 18478231]
12. Meijering E, Dzyubachyk O, Smal I, van Cappellen WA. Tracking in cell and developmental biology. *Seminars in Cell & Developmental Biology*. Oct.2009 20:894–902. [PubMed: 19660567]
13. Meijering E, Smal I, Danuser G. Tracking in molecular bioimaging. *IEEE Signal Proc Mag*. May. 2006 23:46–53.
14. Reid DB. An algorithm for tracking multiple targets. *IEEE T Automat Contr*. 1979; 24:843–854.
15. Blackman SS. Multiple Hypothesis Tracking for Multiple Target Tracking. *IEEE A&E Systems Magazine*. 2004; 19:5–18.
16. Veenman CJ, Reinders MJT, Backer E. Resolving Motion Correspondence for Densely Moving Points. *IEEE T Pattern Anal*. 2001; 23:54–72.
17. Shafique K, Shah M. A Noniterative Greedy Algorithm for Multiframe Point Correspondence. *IEEE T Pattern Anal*. 2005; 27:51–65.
18. Chetverikov D, Verestoy J. Feature Point Tracking for Incomplete Trajectories. *Computing*. 1999; 62:321–338.
19. Genovesio A, Liedl T, Emiliani V, Parak WJ, Coppey-Moisand M, Olivo-Marin JC. Multiple particle tracking in 3-D+t microscopy: Method and application to the tracking of endocytosed quantum dots. *IEEE T Image Process*. May.2006 15:1062–1070.
20. Sethi JK, Jain R. Finding Trajectories of Feature Points in a Monocular image sequence. *IEEE Trans Pattern Anal Mach Intell*. 1987; 9:56–73.
21. Blackman, SS.; Popoli, R. Design and Analysis of Modern Tracking Systems. Norwood, MA: Artech House; 1999.
22. Jonker R, Volgenant A. A Shortest Augmenting Path Algorithm for Dense and Sparse Linear Assignment Problems. *Computing*. 1987; 38:325–340.
23. Vallotton P, Ponti A, Waterman-Storer CM, Salmon ED, Danuser G. Recovery, visualization, and analysis of actin and tubulin polymer flow in live cells: A Fluorescence Speckle Microscopy Study. *Biophys J*. 2003; 85:1289–1306. [PubMed: 12885672]
24. Jaqaman K, Loeferle D, Mettlen M, Kuwata H, Grinstein S, Schmid SL, Danuser G. Robust single particle tracking in live-cell time-lapse sequences. *Nat Methods*. 2008; 5:695–702. [PubMed: 18641657]
25. M, TM.; Ahuja, RK.; Orlin, JB. Network Flows: Theory, Algorithms and Optimization. New Jersey: Prentice-Hall, Inc.; 1993.
26. Geusebroek JM, Smeulders AWM, van de Weijer J. Fast anisotropic gauss filtering. *IEEE Trans Image Processing*. 2003; 12:938–943.
27. Statnikov, RB.; Matusov, JB. Multicriteria Optimization and Engineering. New York: Chapman & Hall; 1995.
28. Rao, SS. Engineering Optimization: Theory and Practice. New York: Wiley & Sons; 1996.
29. Figueiredo MAT, Jain AK. Unsupervised learning of finite mixture models. *Ieee Transactions on Pattern Analysis and Machine Intelligence*. 2002; 24:381–396.
30. Okamoto K, Nishio S, Saga T, Kobayashi T. Standard Images for Particle-Image Velocimetry. *Meas Sci Technol*. 2000; 11:685–691.
31. Danuser G, Waterman-Storer CM. Quantitative Fluorescent Speckle Microscopy of Cytoskeleton Dynamics. *Ann Rev of Biophys Biomol Struct*. 2006; 35:361–87. [PubMed: 16689641]
32. Waterman-Storer CM, Desai A, Bulinski JC, Salmon ED. Fluorescent speckle microscopy, a method to visualize the dynamics of protein assemblies in living cells. *Curr Biol*. 1998; 8:1227–1230. [PubMed: 9811609]
33. Matov A, Applegate K, Kumar P, Thoma CR, Krek W, Danuser G, Wittmann T. Analysis of Microtubule Dynamic Instability Using a Plus End Growth Marker. *Nature Methods*. 2010; 7:761–768. [PubMed: 20729842]
34. Lindeberg T. Feature detection with automatic scale selection. *Int'l J Comp Vision*. Nov.1998 30:79–116.

35. Thomann D, Rines DR, Sorger PK, Danuser G. Automatic fluorescent tag detection in 3D with super-resolution: application to the analysis of chromosome movement. *J Microsc - Oxford*. 2002; 208:49–64.
36. Rosin PL. Unimodal thresholding. *Pattern Recognition*. 2001; 34:2083–2096.
37. Yang G, Houghtaling BR, Gaetz J, Liu JZ, Danuser G, Kapoor TM. Architectural Dynamics of the Meiotic Spindle Revealed by Single-Fluorophore Imaging. *Nat Cell Biol*. 2007; 9:1233–1242. [PubMed: 17934454]
38. Hu WM, Tan TN, Wang L, Maybank S. A survey on visual surveillance of object motion and behaviors. *Ieee Transactions on Systems Man and Cybernetics Part C-Applications and Reviews*. Aug.2004 34:334–352.
39. Ali, S.; Shah, M. Floor Fields for Tracking in High Density Crowd Scenes. The 10th European Conference on Computer Vision (ECCV); Marseille, France. 2008. p. 1-14.
40. Rodriguez, M.; Ali, S.; Kanade, T. Tracking in Unstructured Crowded Scenes. IEEE International Conference on Computer Vision (ICCV); Kyoto, Japan. 2009.

Biographies

Alexandre Matov received his Ph.D. in computational cell biology from the Scripps Research Institute in La Jolla, California in 2009. He has completed postgraduate programs in communication systems at EPF Lausanne and in biophysics at ETH Zurich in Switzerland, and holds a M.S. degree in telecommunications from the Technical University of Varna, Bulgaria. During his dissertation and postdoctoral work, he developed computer vision algorithms for the detection of fluorescent features and particle tracking in video microscopy image sequences. His biomedical research focus has been on quantitative measurements and modeling of cytoskeleton dynamics in living cells. He is a member of IEEE and ASCB.

Marcus Edvall received his M.S. degree in chemical engineering with engineering physics from Chalmers University of Technology (Goteborg, Sweden) in conjunction with Imperial College London in 2000. From 2000 to 2003, he worked as a consultant for the pulp and paper industry with model predictive control systems. Since 2003, he has been the business and development manager of Tomlab Optimization Inc. His primary interest is to accelerate research activities for software development in the area of global black-box optimization and mixed-integer programming.

Ge Yang received his B.S degree in mechanical engineering and B.S. degree in electrical engineering from Tsinghua University (Beijing, China) in 1991. He received his M.S. and Ph.D. degrees in mechanical engineering from the University of Minnesota Twin Cities in 2002 and 2004, respectively. He pursued his postdoctoral training in computational cell biology at the Scripps Research Institute (La Jolla, CA). Since 2009, he is an assistant professor of computational biology and biomedical engineering at Carnegie Mellon University. His research interests are in computational biology, fluorescence imaging, and bioimage informatics.

Gaudenz Danuser received his doctoral degree in computer vision from ETH Zurich, Switzerland. Since his postdoctoral training, he has been interested in the application of image processing and computer vision to the discovery of molecular mechanisms of normal and diseased cell biological functions. He is a professor of cell biology at Harvard Medical School. He has been a member of the editorial boards of IEEE Transactions on Image Processing and the Biophysical Journal between 2003 and 2008. Currently, he serves as a regular member of the NIH study section on Microscopic Imaging.

Fig. 1.

Definition of a graph for the assignment of particles over three consecutive frames. (a) Definition of candidate matches of particles based on circular windowing of particle displacements; (b) Derivation of a graph from candidate matches in a). Numbers on graph edges indicate the capacity of the conduit and the cost to push flow through it; (c) Cost distribution of triplets in the data set shown in Fig. 5 for different search radius settings. The larger the search radius, the more triplets are included in the graph. However, cost distributions after maximum-flow minimum-cost optimization are almost identical (with the exception of very low costs, see text), indicating that the topological constraints in the source-to-sink flow are sufficiently selective to render the tracking essentially independent of the number of false candidate matches in the initial network; (d) Definition of geometric elements of the local and regional motion model.

Fig. 2.

Multi-Objective Optimization. (a) Example where maximizing network flow, i.e. the number of triplets, generates incorrect tracks. (b) Relaxation of the maximum-flow condition generates triplets with much less overall cost. (c) Cost distribution of triplets for various levels of flow reduction (using the data in Fig. 5). Inset: Selection of triplets under maximum-flow condition (dashed) and under optimum-flow condition (solid). One of the triplets identified by maximum-flow solution is preserved in the optimum-flow condition whereas two false positive triplets of the maximum-flow solution are rearranged to form a single triplet moving anti-parallel to the first one. This example is extracted from the vector field shown in Fig. 5. (d). Bi-objective optimization between flow reduction and total cost minimization (see text for details on the exact and approximate Pareto optimal front).

Fig. 3.

Benchmarking on standard particle tracking test data sets used in the computer vision literature. a) Rotating ball; b) Flock of birds; c) School of fish; d) Rotating dish. In all cases three-frame tracks are appended over time-points and overlaid to the first frame of the test image sequence. See videos 1 and 2 for animated presentations of tracking results in b) and c). Panels arranged in order of increasing tracking difficulty (see Table 2). The histograms show how many particles participate in 1 to 20 candidate triplets. For the first three examples only a few particles participate in the formation of more than one triplet.

Fig. 4.

(a) Benchmarking on particle-image velocimetry data set, generated by <http://piv.vsj.or.jp/piv/java/img2-e.html>. The example contains 1300 particles per frame. 50% of all particle coordinates were perturbed. Blue arrows - true positive; black arrows - true negative; magenta arrows - false negative; red arrows - false positive. Inset shows the simulated particles in the tracked region. (b) Histogram showing the number of triplets each particle participates in.

Fig. 5.

Tracking bi- and tri-directional microtubule polymer flux in metaphase mitotic spindles of *Xenopus leavis* egg extracts imaged by fluorescent speckle microscopy (see videos 4 and 5 for raw data): (a) Speckle detection. (b) Three-frame tracks selected by optimum-flow minimum-cost computation. Clustered tracks and interpolated flow fields are overlaid in (c). (d) – (f) Distributions of α , δ , and β which are used to derive the weights of the cost function Eq. 2. (g) Histogram of speckle participation in candidate triplets for the 2-polar spindle. (h) Image of a defective meiotic spindle with three poles. (i) Directionally separated three-frame tracks in boxed region where three polymer flux fields intercept. (j) Histogram of speckle participation in candidate triplets for the 3-polar spindle.

Table 1
Construction of the adjacency matrix A and the demand vector b .

Table 2

Parameters to score the difficulty of a tracking problem. Mean NN denotes the mean nearest neighbor distance between particles in a frame.

	Participating Particles per Frame	Mean Displacement / Mean NN	Search Radius / Mean NN	# Triplets Each Particle is in [mean / max]	# Candidate Triplets	% Particles Participating in >1 Triplets
1) Rotating Ball	176	0.2	0.5	1.1 / 3	178	3%
2) Flock of Birds	38	0.3	0.5	1.1 / 2	27	4%
3) School of Fish	131	0.3	0.6	1.5 / 6	151	44%
4) Rotating Dish	80	0.8	1.3	7.1 / 20	555	87%
5) Particle Image Velocimetry	1304	1.1	2.2	24.1 / 120	25128	98%
6) 3-Polar Mitotic Spindle	299	0.4	0.9	2.3 / 12	567	63%
7) 2-Polar Mitotic Spindle	880	0.7	1.3	2.6 / 18	1688	69%