

Mixture Models for Optical Flow Computation

Allan Jepson* and Michael J. Black

Department of Computer Science
University of Toronto
Toronto, Ont, Canada M5S 1A4

Abstract

The computation of optical flow relies on merging information available over an image patch to form an estimate of 2D image velocity at a point. This merging process raises a host of issues, which include the treatment of outliers in component velocity measurements and the modeling of multiple motions within a patch which arise from occlusion boundaries or transparency. We present a new approach for dealing with these issues, which is based on the use of a probabilistic mixture model to explicitly represent multiple motions within a patch. We use a simple extension of the EM-algorithm to compute a maximum likelihood estimate for the various motion parameters. Preliminary experiments indicate that this approach is computationally efficient and can provide robust estimates of the optical flow values in the presence of outliers and multiple motions.

1 Introduction

The computation of optical flow relies on merging information available over an image patch to form an estimate of 2D image velocity at a point. As the size of this neighborhood grows there is an increased likelihood that the it will span an object boundary in the scene which will result in multiple motions within the region. Multiple motions can also be the result of various kinds of transparency. In these situations, the assumption of a single motion within the region results in inaccurate estimates of the optical flow. We refer to this dilemma regarding the choice of neighborhood size as the *generalized aperture problem*. To address the problem, we relax the single-motion assumption and, instead, assume that the motion(s) within the region can be described by a probabilistic *mixture* of distributions.

We observe that, when multiple motions are present, the motion estimates within a region form distinct *clusters*. We employ a simple extension of the *EM-algorithm* [4] to isolate these clusters, estimate their

*Allan Jepson was supported by PRECARN and CIAR.



Figure 1: The **Pepsi Can** image sequence. The mixture approach is applied to the boxed region.

likelihood, and detect *outlying* measurements which do not correspond to a coherent motion.

This approach has a number of benefits. Like robust regression techniques [1], the approach allows us to robustly estimate the dominant motion within a region. But more importantly, by assuming the the motion is due to a mixture of distributions we are able to recover multiple coherent motions if they are present and reject outliers. This multiple-motion assumption is applicable at both motion boundaries and in regions containing multiple transparent motions. Additionally, information about the presence of multiple motions may prove useful for the early detection of surface boundaries from motion.

2 Mixture Models of Flow

For a given image region we attempt to model the flow in terms of a handful of smoothly varying layers. For example, $\vec{v}(\vec{x}; \vec{a})$ may represent a constant velocity field for one layer, or it could denote an affine flow where the components v_1 and v_2 are given by linear functions of the image position \vec{x} . In the first case the parameter vector \vec{a} is 2-dimensional, while it is 6-dimensional in the affine case.¹ Multiple motions within a particular

¹In practice, it is often useful to add parameters representing the uncertainty of \vec{a} .




Figure 2: Multiple motion constraint lines for the region in Figure 1 (see text).

patch are represented by selecting more than one set of parameters \vec{a} . However, note that at this stage of analysis we have not modeled *where* in the image patch each of the various models are appropriate. Thus transparent motion, with two different velocity fields realized over the whole patch, will be initially modeled in the same way as an occlusion boundary. A subsequent level of analysis is needed to determine which of these two interpretations is appropriate for a particular patch.

We seek the parameter values \vec{a}_n , $n = 1, \dots, N$ for N possibly distinct smooth fields, one for each layer. For the n^{th} layer, the probability of observing a motion constraint vector \vec{c}_k , given that the observation is at the spatial location \vec{x}_k , is modeled by the “component probability” distribution $p_n(\vec{c}_k | \vec{x}_k, \vec{a}_n)$. In addition we also have a model for outlier processes given by $p_0(\vec{c}_k)$. Finally, the probability of selecting layer n is given by the “mixture probabilities” m_n , which are treated as further parameters we need to fit. Together these pieces provide the overall probability of observing the constraint \vec{c}_k , namely

$$p(\vec{c}_k | \vec{x}_k, \vec{m}, \vec{a}_1, \dots, \vec{a}_N) = \sum_{n=0}^N m_n p_n(\vec{c}_k | \vec{x}_k, \vec{a}_n). \quad (1)$$

Here the mixture probabilities m_n , for $n = 0, 1, \dots, N$ must sum to one.

To obtain a maximum likelihood fit for the parameters m_n and \vec{a}_n , $n = 0, \dots, N$, we use a simple modification of the EM-algorithm [4]. This is a simple iterative algorithm which is guaranteed to increase the log likelihood of its fit each iteration (see [3] for details).

3 Computational Examples

To illustrate the approach, motion constraint vectors, \vec{c}_k , were computed using a phase-based approach [2]. For the 32×32 region marked in Figure 1, which is roughly centered on an occlusion boundary, a third of all the motion constraints are depicted in Figure 2a. The two “X”s in Figure 2b mark the peaks of the

fitted mixture model for this example. The method converges quickly and has clearly recovered the velocities of both sides of the occlusion boundary without difficulty.

Using (1) we can compute the *ownership probability*, namely the probability that any given constraint comes from either of the two motions. The darkness of the constraint lines in Figure 2b is proportional to the ownership probability for the first motion. A similar plot is given in Figure 2c for the second motion. Note that constraint lines that are roughly horizontal, and pass close to both peaks, have roughly equal ownership probabilities. This illustrates the competition between the various components in the mixture model for the ownership of each constraint. There are only a handful of constraint lines that are outliers which do not correspond to either motion.

The general spatial distribution of the ownership probabilities (Figure 2d) reflects the structure and location of the occlusion boundary within the patch. Darkness is proportional to the probability that the constraints at a location belong to the motion of the foreground.

For details of the method and additional experiments (including transparent motion) the reader is referred to [3].

References

- [1] M. J. Black. *Robust Incremental Optical Flow*. PhD thesis, Yale University, New Haven, CT, 1992. Research Report YALEU/DCS/RR-923.
- [2] D.J. Fleet. *Measurement of Image Velocity*. Kluwer, Boston, 1992.
- [3] A. Jepson and M. J. Black. Mixture models for optical flow computation. Technical Report RBCV-TR-93-44, University of Toronto, April 1993.
- [4] G.J. McLachlan and K.E. Basford. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker Inc., N.Y., 1988.