Learning Discriminative Chamfer Regularization

Pradeep Yarlagadda *
pradeep.yarlaga@iwr.uni-heidelberg.de
Angela Eigenstetter *
aeigenst@iwr.uni-heidelberg.de
Björn Ommer
ommer@uni-heidelberg.de

Interdisciplinary Center for Scientific Computing (IWR)
University of Heidelberg
Germany

Chamfer matching is an effective and widely used technique for detecting objects or parts thereof by their shape. However, a serious limitation is its susceptibility to background clutter. The primary reason for this is that the presence of individual model points in a query image is measured independently. A match with the object model is then represented by the sum of all the individual model point distance transformations. Consequently, i) all object pixels are treated as being independent and equally relevant, and ii) the model contour (the foreground) is prone to accidental matches with background clutter.

As demonstrated by Attneave [1], and various experiments on illusionary contours, object boundary pixels are not all equally important due to their statistical interdependence. Moreover, in dense background clutter the points on the model have a high likelihood to find good spurious matches [1, 3]. However, any arbitrary model would match to such a cluttered region, which consequently gives rise to matches with high accidentalness. Chamfer matching only matches the template contour and thus fails to discount the matching score by the accidentalness, i.e., the likelihood that this is a spurious match.

We take account of the fact that boundary pixels are not all equally important by applying a discriminative approach to chamfer distance computation, thereby increasing its robustness. Let $T = \{\mathbf{t}_i\}$ and $Q = \{\mathbf{q}_j\}$ be the sets of template and query edge map respectively. Let $\phi(\mathbf{t}_i)$ denote the edge orientation of the edge point $\mathbf{t}_i$. For a given location $\mathbf{x}$ of the template in the query image, directional chamfer matching [2] finds the best $\mathbf{q}_j \in Q$ for each $\mathbf{t}_i \in T$, thus resulting in a matching cost $p_i^{(T,Q)}(\mathbf{x})$.

$$p_i^{(T,Q)}(\mathbf{x}) = \min_{\mathbf{q}_j \in Q} |(\mathbf{t}_i + \mathbf{x}) - \mathbf{q}_j| + \lambda |\phi(\mathbf{t}_i + \mathbf{x}) - \phi(\mathbf{q}_j)| \quad (1)$$

Adjacent template pixels are statistically dependent and, thus, we do average (1) over the direct neighbors of pixel i. The resulting $\bar{p}_i$ are then used to learn the importance of contour pixels.

While learning the weights for individual pixels improves the robustness of template matching, chamfer matching is still prone to accidental responses in spurious background clutter. To estimate the accidentalness of a match, a small dictionary of simple background contours T_{bg} is utilized. Rather than placing background contours at a fixed single location, i.e., at the center of the model contour as in [3], background elements are trained to focus at locations where, relative to the foreground, typically accidental matches occur.

Let $d_{DCM}^{(T,Q)}(\mathbf{x})$ denote the directional chamfer distance between Q and T with a relative displacement $\mathbf{x}$. To measure where clutter typically interferes with the model contour we compute $d_{DCM}^{(T_{bg},T)}$ between each background contour T_{bg} and the object template T. We consider placements of the background contour with better (lower) chamfer matching score to be more important since they occur on or close to the model contour. In order to weight these matching locations higher we create a mask $M^{(T_{bg},T)}(\mathbf{x})$

$$M^{(T_{bg},T)}(\mathbf{x}) = 1 - d_{DCM}^{(T_{bg},T)}(\mathbf{x}) \quad (2)$$

To describe the background matching costs for a hypothesis in a robust way we build weighted histograms over chamfer matching scores $d_{DCM}^{(T_{bg},Q)}$ obtained from matching a background contour T_{bg} with the query image Q. Let $B(\bar{\mathbf{x}})$ be the bounding box region with center $\bar{\mathbf{x}}$ for a specific placement of the foreground template T in the query image Q. For each foreground hypothesis we build weighted histograms $h^{(T_{bg},Q)}$ over the directional chamfer matching scores $d_{DCM}^{(T_{bg},Q)}$ in the corresponding bounding box region. The weights introduced in (2) are used to weight the histogram votes. Therefore chamfer matching scores $d_{DCM}^{(T_{bg},Q)}$ are weighted according to their position relative to the foreground template. Each histogram consists of K bins where $\mathcal{M}_k$ is the range of the kth bin and $k = 1, \dots, K$. A histogram bin $h_k^{(T_{bg},Q)}$ is defined as

$$h_k^{(T_{bg},Q)} = \sum_{\substack{\mathbf{x} \in B(\bar{\mathbf{x}}) \\ d_{DCM}^{(T_{bg},Q)}(\mathbf{x}) \in \mathcal{M}_k}} M^{(T_{bg},T)}(\mathbf{x}), \quad (3)$$

for each background contour T_{bg} on a certain position of the foreground template T in the query image Q.

For each object hypothesis we build a feature vector $f_i = [\bar{p}_1 \dots \bar{p}_L \, h_1 \dots h_G]$ consisting of the average pixel cost $\bar{p}_i$ and the corresponding background histograms h_i, where L is the number of template edge pixels and G is the number of background contours.

Finally, a max-margin classifier is employed to learn the co-placement of all background contours and the foreground template. This classifier yields a regularized distance function d_{RDCM}

$$d_{RDCM}^{(T,Q)}(\mathbf{x}) = 1 - \Big(\sum_i \alpha_i \mathcal{K}(f_j, S_i) + b \Big). \quad (4)$$

$\mathcal{K}$ denotes the kernel used in the SVM. b denotes the offset. S_i, α_i denotes the support vectors and their respective coefficients.

Our approach is easily integrated into an off-the-shelf directional chamfer matching approach and it shows significant improvements over state-of-the-art chamfer matching on standard benchmark datasets. The qualitative and quantitative results are detailed in the paper.

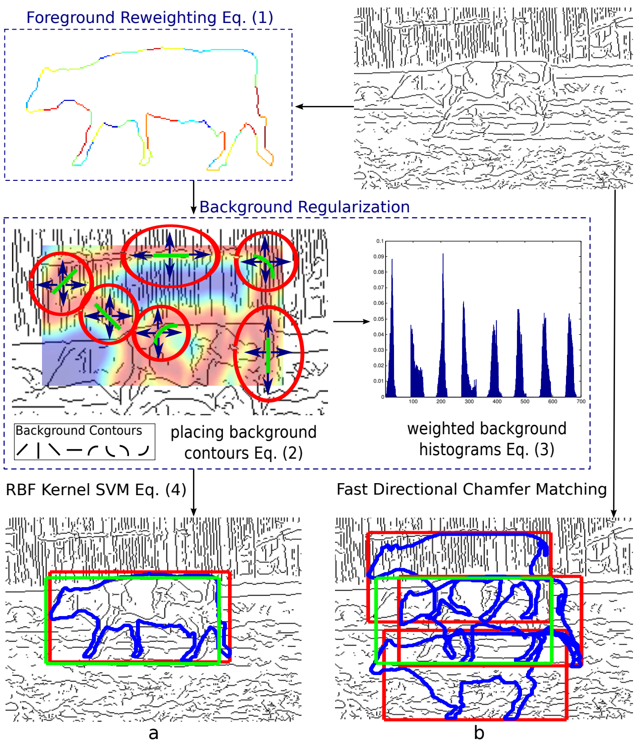

Figure 1: Comparison of a) reguarized chamfer matching with b) directional chamfer matching.

[1] F. Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183–193, 1954.

[2] M. Liu, O. Tuzel, A.Veeraraghavan, and R. Chellappa. Fast directional chamfer matching. *CVPR*, 2010.

[3] T. Ma, X. Yang, and L.Latecki. Boosting chamfer matching by learning chamfer distance normalization. *ECCV*, 2010.

* Both authors contributed equally to this work.