

# Hierarchical Chamfer Matching Based on Propagation of Gradient Strengths

Stina Svensson<sup>1</sup> and Ida-Maria Sintorn<sup>2</sup>

<sup>1</sup> Centre for Image Analysis, Swedish University of Agricultural Sciences,  
Lägerhyddsvägen 3, SE-75237 Uppsala, Sweden  
`stina@cb.uu.se`

<sup>2</sup> CSIRO Mathematical and Information Sciences  
Locked Bag 17, North Ryde, NSW 1670, Australia  
`Ida-Maria.Sintorn@csiro.au`

**Abstract.** A modification of the hierarchical chamfer matching algorithm (HCMA) with the effect that no binarisation of the edge information is performed is investigated. HCMA is a template matching algorithm used in many applications. A distance transform (DT) from binarised edges in the search image is used to guide the template to good positions. Local minima of a function using the distance values hit by the template correspond to potential matches. We propose to use distance weighted propagation of gradient magnitude information as a cost image instead of a DT from the edges. By this we keep as much information as possible until later in the matching process and, hence, do not risk to discard good matches in the edge detection and binarisation process.

## 1 Introduction

Chamfer matching is a template matching method based on geometric image features and can be used for both 2D and 3D images. It finds good fits between the template and edges in a search image. A generalized cost function between the edges in the search image and the template, a list of coordinate pairs corresponding to the searched pattern, is minimized. To guide the template to good positions, a distance transform (DT) is calculated from edges in the search image and the sum of the distance values hit by the superimposed template constitute the cost function. Translation, scaling, rotation, and perspective changes are for 2D images and translation, scaling, and rotation for 3D images.

In its original form, proposed by Barrow et al. in 1977 [1] as described above, chamfer matching is a fine matching algorithm which requires good start positions. In [2], Borgfors showed that using the sum of squared distances instead of the sum of distances as the cost function results in significantly fewer false minima. In [3], she embedded the chamfer matching in a resolution hierarchy. This, the hierarchical chamfer matching algorithm (HCMA), turned the original fine-matching algorithm to a fast, general and robust matching algorithm.

HCMA's usefulness is clear not the least from the number of citations (almost 200, Science Citation Index March 2006). Although introduced almost 20

years ago, it is still used in many applications today. It is described in the recent surveys of volume image registration by template matching [4] and image registration [5]. In [6], a 3D model based registration approach is taken for fusion of information from positron emission tomography (PET), magnetic resonance imaging (MRI) and magnetocardiography (MCG). For the PET-MRI registration, HCMA is used. HCMA is also used in object recognition, e.g., hand pose recognition, [7], and pedestrian detection from moving vehicles, [8].

A few suggestions of improvements to HCMA have been described. In [9] and [6], HCMA is used with an oriented DT, where not only the position of the edges in the search image, but also their orientation is taken into account. The oriented DT is used to guide the search and decreases the chance of getting stuck in false local minima. In [9], it is also suggested to use a hierarchical clustering of the template from different views to speed up the matching process. These two modifications could be incorporated in the method suggested here. Three different modifications of the DT taking into account gradient magnitude information (*saliency DTs*) are presented in [10]. The use of saliency DTs in various applications, including chamfer matching without the hierarchical structure and for 2D images only, is discussed. One of the saliency DTs presented in [10] is similar to the distance weighted propagation of gradient magnitude information suggested here. We investigate the use of distance weighted propagation of gradient magnitude information in HCMA further. We discuss how the choice of local propagation steps influences the outcome of the matching process and give some ideas on how the matching measure need to be modified to suite the cost image. Moreover, we generalise the concept to 3D images.

HCMA and our gradient magnitude based HCMA (GM-HCMA) are described for 2D images in Section 2 and Section 3, respectively. Section 3 also contains the generalisation to 3D images. The performance of GM-HCMA is illustrated on a 2D leaf-matching example in Section 4 and on a cryo-electron tomographic (cryo-ET) 3D protein identification example in Section 5. The protein identification example was inspired by [11], where the ordinary HCMA was shown to give promising results for identifying proteins in cryo-ET images. Here, we show that similar results can be achieved using GM-HCMA.

## 2 HCMA

HCMA improves the chamfer matching algorithm by embedding it in a resolution pyramid, [3], [1]. Chamfer matching is performed in the image with the lowest resolution. Relevant positions, i.e., positions giving a match measure greater than a user defined *rejection factor* (RF), are used as start positions for chamfer matching in the image at the next higher resolution level. Chamfer matching is then performed level after level until the highest resolution level, i.e., the original image, is reached. We start by describing the chamfer matching performed for each resolution level and then how the image resolution pyramid is created.

Two images are needed to perform chamfer matching: a template representing the contours of the searched pattern and a search image depicting the same type

of information, i.e., edge information, as the template. There are many possible ways to create a template. This is not the focus of this paper so we simply assume that a template is given. There are likewise several different ways to detect edges in an image. Here, we will detect edges as the gradient magnitudes derived as the sum of the responses from a set of Sobel operators, e.g., [12]. For 2D images, the operators are of size  $3 \times 3$  pixels and horizontal, vertical, and diagonal directions are taken into account. For 3D images, the operators are of size  $3 \times 3 \times 3$  voxels and directions in the  $x$ -,  $y$ -, and  $z$ -planes are considered.

An edge image is created from the gradient magnitude image by, e.g., thresholding, and a DT is calculated from the edges. In the DT, each pixel is assigned a value corresponding to its distance to the nearest edge pixel. The DT serves as a cost image which guides the template to positions giving low match measures.

When calculating a DT, the edge pixels are initially set to 0 and other pixels to  $\infty$ . Local distances from the edge pixels are propagated over the image in one forward and one backward scan. The local distances (weights) used here are 3 for an edge neighbour and 4 for a vertex neighbour, following the suggestion in [3]. Each pixel is assigned the minimum of its value and the values of its, in the current scan, already visited neighbours, each increased with their respective local distance. This method of calculating a DT is known as chamfering, which originates the name chamfer matching.

The template is represented by a set of coordinate pairs  $T = \{(X_i, Y_i) \mid X_i, Y_i \in \mathbf{Z}, i = 1, \dots, n\}$ . The search for local minimal match measures starts from a set of start positions. Each position contains information on the transformation parameters translation  $(c_X, c_Y)$ , rotation  $(\theta)$ , and scaling  $S$ . (We leave out perspective.) The template  $T$  is transformed to its current position by  $T^S = T \cdot S$ , and  $T'(X_i, Y_i) = (c_X + X_i^S \cdot \cos \theta - Y_i^S \cdot \sin \theta, c_Y + Y_i^S \cdot \sin \theta + X_i^S \cdot \cos \theta)$ . The  $(X'_i, Y'_i)$  coordinates are usually not integers and are rounded to the nearest integer value. After this,  $T'$  is superimposed on the DT image. The *match measure* used, i.e., the measure on how well  $T'$  fits edges in the underlying image, is the root mean square average (RMS) of the DT values hit by the template:

$$RMS = \frac{1}{3} \sqrt{\frac{1}{n} \sum_{i=1}^n d(X'_i, Y'_i)^2},$$

where  $d(X'_i, Y'_i)$  is the value in the DT at position  $(X'_i, Y'_i)$ , and the factor  $\frac{1}{3}$  comes from the use of  $\langle 3, 4 \rangle$  DT. RMS is preferable to the arithmetic mean as match measure in order to decrease the number of false minima, [2].

To reach a local minimum of the match measure from the current position, a steepest descent and a line search method is used [13, chapter 10.5, 11.2]. Iteratively, a step is taken in the steepest descent direction if it produces a smaller match measure. If not, the step length is halved until it does, or until the step length is so small that a local minimum can be considered reached. For each step, the steepest descent direction is calculated from all four parameters (rotation, scaling, and  $2 \times$  translation). The gradient for each parameter is approximated by centred differences, where a translation step is the side length of a pixel, a

rotation step is the angle which makes a point on the edge of the template move one pixel side in length, and a scale step is a 2% change in scale.

In HCMA, chamfer matching as described above is embedded in a resolution pyramid. The resolution pyramid is built from the binarised edges detected in the search image, which serves as the base of the pyramid ( $L_0$ ). Each block of  $2 \times 2$  pixels at level  $L_i$  is represented by a single pixel at level  $L_{i+1}$ . This means that for each higher level the size of the image is reduced by  $\frac{3}{4}$ , and the resolution is halved. The “colour” of the pixel depends on the colour of the four pixels it is assigned from. For HCMA, the highest value is chosen, i.e., an OR-pyramid is used, guaranteeing that no edges disappear between successive levels.

A DT is calculated for all levels and starting positions are then distributed at the top of the pyramid, i.e., at the lowest resolution level. The positions that after the matching at level  $L_i$  produce match measures better than RF (depending on the size of the template and the resolution level) are kept and used as starting positions at level  $L_{i-1}$ . This means that at level  $L_0$  only fine matching from relatively good starting positions is performed. If no a priori knowledge is available, starting positions for the top level of the pyramid should be evenly distributed, i.e., as a regular grid of points in the search space.

### 3 GM-HCMA

Instead of using a DT from the binarised edge image as the cost image, we suggest decreasing propagation of the gradient magnitudes themselves and using the result as the cost image. By this, we keep as much information as possible until later in the process and do not risk to loose important edges by choosing a threshold that is too high.

As the gradient magnitudes are high at edges, a good match does not correspond to a low match measure but a high one. This means that we search for local maxima in the matching process. Moreover, there is no longer a match measure corresponding to a perfect match, as 0 is in HCMA. A certain match measure can be the result of a good fit at edges with low gradient magnitudes or a rather bad fit but still hitting some very high gradient magnitudes. The same phenomenon occurs for HCMA but there only needs to be considered if very occluded objects are desirable to be found. With the new cost image, this effect will be more common and therefore additional RFs need to be considered.

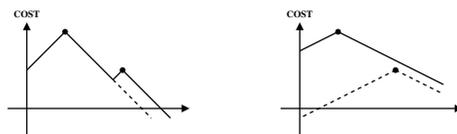
To explain the proposed GM-HCMA, we start by recalling the concept of the reverse distance transform (RDT) [14]. In RDT, decreasing local distance information is propagated over the image in one forward and one backward scan starting from a set of distance labelled seed points. Each pixel is assigned the maximum of its value and the values of its, in the current scan, already visited neighbours, each decreased by their respective local distance. As for the DT described in Section 2, good local distances are 3 for edge neighbours and 4 for vertex neighbours. Calculating a RDT can be seen as growing a ball for each seed until it has a size (radius) corresponding to the distance value of the seed.

Our aim is to spread decreasing gradient magnitude information in a shape preserving way in order to create a cost image where the values depend on the gradient magnitude as well as the relative distance, i.e., the shape of the objects in the image. We propagate the information in one forward and one backward scan in the same way as for RDT and use application dependent local “distances” (*weights*). To obtain a propagation stable under rotation, the edge and vertex neighbour weights are chosen so that the ratio is  $\frac{3}{4}$  following the  $\langle 3, 4 \rangle$  DT. The result can be thought of as smoothing the image while creating a cost image, as pixels with low gradient magnitude close to pixels with higher values will be overshadowed and not contribute to the cost image. Note that the propagation is done in full scale, which means that depending on the weights and the gradient magnitudes, the resulting cost image can have values below 0. The process to create the cost image is the same in 3D. We use face, edge, and vertex neighbour weights with the relation 3,4,5, as suggested in [15].

Calculating a resolution pyramid from the gradient magnitude image is done in the same way as in the binary case. Each pixel at a higher level is assigned the maximum value of the four pixels at the lower level it is built from. For 3D images, the resolution pyramid is created by letting each block of  $2 \times 2 \times 2$  voxels in level  $L_i$  represent a single voxel in  $L_{i+1}$ , i.e, the resolution is halved and the number of voxels is reduced by  $\frac{7}{8}$  between successive levels. For 3D images, the search space increases as translation and rotation are described by three parameters each, instead of two and one, respectively, as in the 2D case.

How the weights should be chosen depends on the application, image acquisition situation, and how much smoothing the user wants to build into the propagation. If the gradient magnitude information of the searched object is likely to be low compared to other parts of the image, high weights, resulting in a low smoothing factor, should be chosen. Using the logarithm of the gradient magnitude image, before the resolution pyramid is calculated, can also be considered, as this leads to a compression of the range of the gradient magnitudes. If the gradient magnitude information of the searched object is likely to be high compared to other parts, low weights should be chosen so that points with low gradient magnitude are overshadowed and noise is smoothed. This decreases the chances of getting stuck in false local maxima during the matching process. A 1D curve illustrating the effect of using higher and lower weights is shown in Figure 1. A higher weight (left) has the effect that the point with lower gradient magnitude prevents the propagation from the point with higher gradient magnitude. A lower weight (right) makes the point with higher gradient magnitude overshadow the point with lower gradient magnitude. If no a priori information is available, it is safer to use high weights as the influence of a point with high gradient magnitude then drops quickly, even though this means that the chance of getting stuck in false local maxima increases.

For GM-HCMA, good positions give high match measures (instead of low as for HCMA). However, it is not possible to give absolute values for how high good match measures should be as it depends on the gradient magnitude information available. Instead additional criteria can be used to help evaluate possible match



**Fig. 1.** 1D curves illustrating the effect of using higher (left) and lower (right) weights

positions. If the background is even and, hence, the gradient magnitudes of the searched object are likely to be of similar strength, the variance of the pixel values in the cost image contributing to the match measure gives a good indication on whether a position is relevant or not. In [16], a method for differentiating between positions giving low match measures when HCMA is applied to noisy images is presented. In such images, a low match measure does not necessarily mean that the position is correct and a rather high match measure does not automatically mean that the matching has failed. In the first case, it could be that the object matches parts of other objects and in the second, that the searched object is partially occluded or simply that the scene is noisy. In such cases, the variation of the cost along the template for a found position can be studied.

#### 4 Comparison of GM-HCMA and HCMA for 2D Images

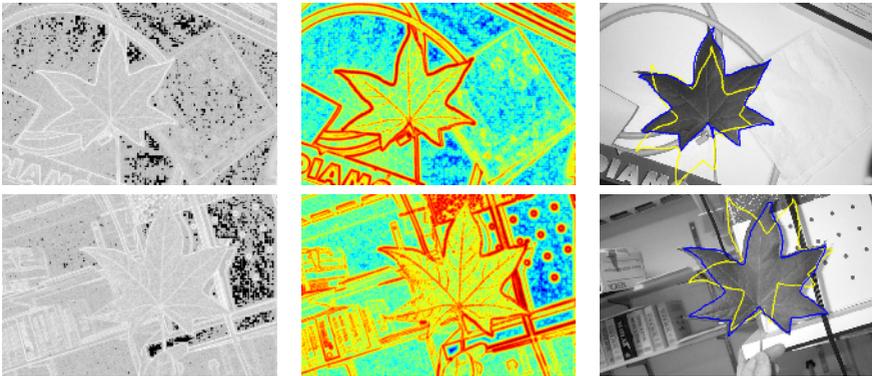
We illustrate GM-HCMA and compare it to HCMA using the two images, *leaf A* and *leaf B*, in Fig. 2. The images, both of size  $896 \times 592$  pixels, belong to a data set used for experiments by the Computational Vision Group at California Institute of Technology (<http://www.vision.caltech.edu/archive.html>). The templates used, consisting of 2494 and 1922 pixels, respectively, are also shown in Fig. 2. These have been created manually in such a way that a “perfect” position is possible for leaf B while for leaf A the template differs slightly in shape compared to the leaf in the search image. We have chosen leaf A and B to illustrate GM-HCMA as they show GM-HCMA’s advantages to HCMA as well as point out situations when GM-HCMA have difficulties (and HCMA fails).

In this example, we have used the logarithm of the gradient magnitude image to compress the range of gradient magnitudes. Assuming that we have no a priori knowledge of the gradient magnitude information for the searched object,



**Fig. 2.** Leaf A (left) and leaf B (right) and their respective templates

we have chosen weights in such a way that the influence from pixels with the highest gradient magnitude decrease to 0 over a distance corresponding to the width of the template, at the top level of the resolution pyramid. This is to make sure that edges on one side of the template do not overshadow edges on the opposite side. Hence, the weight for a step in the edge direction  $w_1$  is set to the largest value in the gradient magnitude image divided by the minimum width of the template. The weight for a step in the vertex direction  $w_2$  is  $\frac{4}{3}w_1$ . We have used five levels in the resolution pyramid and the weights  $\langle 0.32, 0.43 \rangle$  and  $\langle 0.28, 0.37 \rangle$ , respectively, on all levels. The low weights are due to the use of the logarithm of the gradient magnitude instead of the gradient magnitude itself. We have implemented GM-HCMA using MATLAB and pixel type “double” as speed is not the main issue here. Implementing GM-HCMA, using a rescaling of the logarithmic values, with 8- or 16-bit images will give approximately the same computational complexity as HCMA. For translation, rotation, and scaling step lengths and optimization process, we have followed what is described in Section 2. In Fig. 3, the logarithm of the gradient magnitude image, the cost image, and the positions giving the two highest match measures are shown for level  $L_0$ . The cost image is linearly stretched (and displayed in colour) to emphasize differences.



**Fig. 3.** Top row: leaf A. Bottom row: leaf B. Logarithm of the gradient magnitude image, the cost image, and the positions giving the two highest match measures using GM-HCMA (blue/dark for the highest and yellow/light for the second highest).

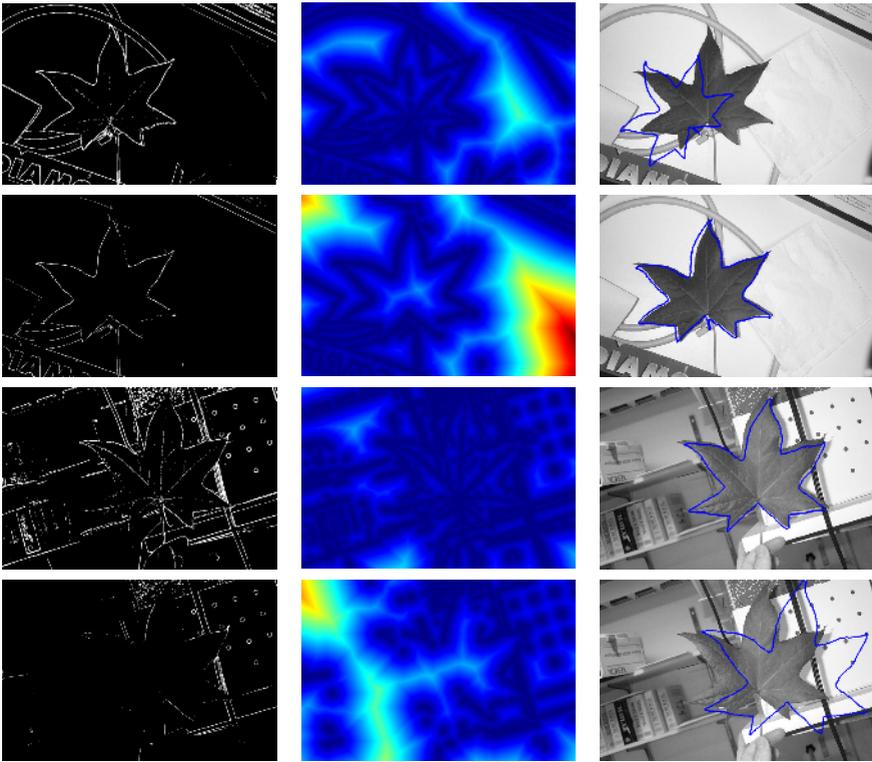
The match measures are 20.9 (leaf A) and 20.3 (leaf B) for the positions shown in blue/dark in Fig. 3. For leaf A, this position corresponds to a reliable response since the match measure for other positions are much lower (17.5 for the second position, shown in yellow/light). For leaf B, the situation is less certain as the second position, shown in yellow/light, also has match measure 20.3. Additional information is given by the variance of the pixel values in the cost image contributing to the match measure. The variance is 2.60 for the first position (and 11.0 for the second, incorrect, position) for leaf A and 7.31 and 9.46 for the first, correct, and second, incorrect, positions with match measure

20.3 for leaf B. The low variance for the first position for leaf A indicates that it is likely to be a correct match because a low variance corresponds to a good fit for *the whole* template. Higher variances (as for leaf B) can result from partial occlusion (as is the case for the first position), i.e., most of the template fits well but parts of it does not, or from a less good fit for most of the template but where parts of it hit very high gradient magnitudes. What is the cause for the higher variance can be established by studying the graph of cost variation along the template at the match position [16]. Another more automatic possibility is to use a percentile of the values hit by the template in the RMS (and/or variance) calculation, i.e., only take into account the 75% of the pixels (those with the highest costs) hit by the template. In this example, the lower variance for the first position for leaf B still give preference to that position, without performing any additional investigations. Note that the difficulties for leaf B in finding a more clear first, and correct, position is due to the fact that there are very high gradient magnitudes on one side of the searched pattern (right) and low on the other side (left). This will be even more evident when HCMA is ran for the same example, see below. We remark that a direct comparison of match measures and variances between images should not be performed unless the images are acquired identically.

For comparison, we have also ran HCMA on leaf A and leaf B. Thresholding of the gradient magnitude image was performed to generate edge images. For leaf A, the gradient magnitudes range from 0 to 2577 and for leaf B from 0 to 2603. To illustrate the effect of the binarisation step, we ran HCMA starting from two different edge images, using the thresholds 400 and 1000, for both leaf A and leaf B. For the rest, the same settings as for GM-HCMA were used. In Fig. 4, the edge images, the DTs, and the best HCMA positions are shown for resolution level  $L_0$ . The match measures for leaf A were 11.5 and 6.3, the latter corresponding to a correct position. For leaf B, the match measures were 7.6 and 26.9, the former corresponding to a correct position. This shows that to succeed, HCMA is dependent on a careful choice of threshold for the gradient magnitude image. For leaf B, we remark, to make the link to what was pointed out for the results from GM-HCMA, that the very high gradients on one side of the searched pattern (right) and low on the other side (left) makes this thresholding difficult.

## 5 GM-HCMA for Identification of Proteins

The initial goal of this work was to investigate whether it was possible to use GM-HCMA to achieve better and more stable matching results than by HCMA for a specific application. In [11], HCMA was used for shape based identification of proteins, or more specifically the Immunoglobulin G (IgG) antibody, in cryo-ET images. The analysis of cryo-ET images is of interest as the technique allows imaging of individual proteins, something which is not possible with other imaging systems. By this it is possible to study protein functional dynamics. For more information on cryo-ET experiments of the IgG antibody and the structural information possible to extract from this type of images, we refer to [17].

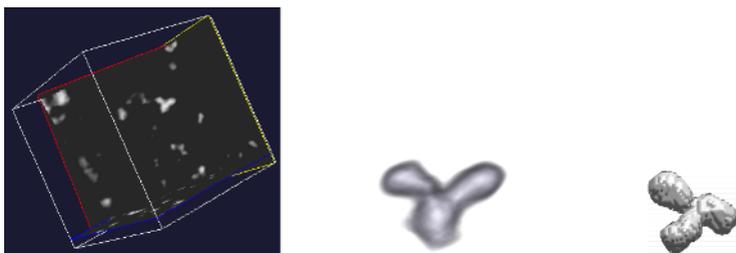


**Fig. 4.** Top rows: leaf A. Bottom rows: leaf B. Edge images created using two different thresholds (400 and 1000), the DTs, and the position giving the lowest match measure using HCMA.

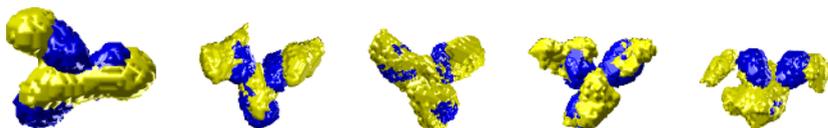
The images resulting from cryo-ET have low contrast as the micrographs used to reconstruct the tomographic images need to be acquired using a low electron dose to prevent radiation damages. The images contain not only the proteins of interest but also other structures, such as ice crystals, artifacts from the reconstruction, and other molecules, which complicate the matching. To the left in Fig. 5, a cross section of a cryo-ET image is shown. The “Y” shaped object almost centred in the cross section is an IgG antibody. This IgG antibody is shown volume rendered in Fig. 5 (middle). The cryo-ET image is of size  $150 \times 150 \times 150$  voxels with voxel size  $5.24\text{\AA}$  and a resolution of approximately 2nm.

In [11], a template created from the protein databank, [18], (PDB identification lig1) was used in the search for IgG antibodies. This template is shown in Fig. 5 (right). The experiments in [11] showed promising results. However, as the gradient magnitude information is low in cryo-ET images and, hence, the edge detection step required for HCMA is non trivial, GM-HCMA is of interest to achieve a more robust protein identification process.

We have used the same settings for the experiments as in [11], but applied GM-HCMA instead of HCMA. The size of the template was 1220 voxels, which



**Fig. 5.** Left: A cross section of a cryo-ET image with an IgG antibody shown almost in the middle. Middle: A volume rendering of the IgG antibody in the cross section. Right: The template used in GM-HCMA.



**Fig. 6.** The positions found by GM-HCMA (found object in yellow/light, template in blue/dark). All positions correspond to IgG antibodies except for the rightmost.

corresponds to the surface of an IgG antibody object of size 3600 voxels. Only two resolution levels were allowed as the objects almost disappear at lower resolution than that. Moreover, we follow the suggestion in Section 3 and use the logarithm of the gradient magnitudes. The IgG antibody consists of three roundish parts, which are approximately equal in size. We have chosen the weights in such a way that the gradient magnitudes decrease to 0 over a distance corresponding to size of one of these parts, at the top level of the pyramid. GM-HCMA was applied to four cryo-ET images, of which three are the same as the three used in [11], each containing one IgG antibody. The weights were  $\langle 2.20, 2.94, 3.67 \rangle$ ,  $\langle 2.42, 3.23, 4.04 \rangle$ ,  $\langle 2.42, 3.23, 4.03 \rangle$ , and  $\langle 2.41, 3.21, 4.02 \rangle$ , respectively.

For three of the images, the positions found by GM-HCMA, giving the highest match measure, corresponded to an IgG antibody, with match measures 2.34, 2.24, and 2.26, respectively, and variance 0.02, 0.03, and 0.03, respectively. For the fourth image, another structure was found, with match measure 2.17 and variance 0.04. The IgG antibody was listed as the second best match, with match measure 2.16 and variance 0.04. In Fig. 6, matches are shown with the template (yellow/light) translated to the positions found by GM-HCMA and superimposed on the corresponding IgG antibody (blue/dark). The rightmost subfigure shows the structure found which does not correspond to an IgG antibody. The positions are not perfect mainly due to the fact that the template does not perfectly correspond to the searched antibody.

The initial experiments shown here indicate that GM-HCMA is a robust method for finding IgG antibody candidates in cryo-ET images. The IgG antibodies are found at the position producing the highest match measure in three of the four available volumes, and at the second highest position in the fourth.

This is similar to the results in [11], with the advantage that the edge detection step can be avoided. To ensure that no IgG antibodies are missed good candidates can be presented for the molecular biologists involved in the project for visual inspection. This way of prescreening the images is of interest since the amount of data is too large to make manual search feasible. A more thorough study of GM-HCMA on cryo-ET images will be the issue for a forthcoming manuscript.

Possible application specific extensions of this work include incorporating flexibility in the template to account for the natural flexibility of the IgG antibody. However, this implies a (much) larger search space and would thereby increase the complexity of the algorithm. Another possibility is to follow the suggestion in [9] and embed hierarchical grouping of multiple templates corresponding to different molecular configurations in the matching process.

## 6 Conclusion

We have described a modification of the hierarchical chamfer matching algorithm (HCMA) [3], with the advantage that no binarisation of the gradient magnitude image is necessary. Instead, distance weighted propagation of the gradient magnitude information is used to create a cost image, resulting in a gradient magnitude based HCMA (GM-HCMA). GM-HCMA needs to be tuned for the application at hand with respect to the weights used for propagating the gradient magnitudes and the evaluation of positions found by the matching process. We have shown that GM-HCMA gives comparable or improved results to HCMA, without increasing the complexity of the algorithm or its computational cost. In Section 5, we used GM-HCMA in a real application. We could achieve the same results as in [11], but without the need of a sophisticated edge binarisation method as in [11], where HCMA was used.

## Acknowledgements

The cryo-ET images have been provided by Dr Sara Sandin, Department of Cell and Molecular Biology, Karolinska Institutet, Stockholm, Sweden (currently Division of Structural Studies, MRC Laboratory of Molecular Biology, Cambridge, United Kingdom). Stina Svensson is financially supported by Swedish Research Council (project 621-2005-5540).

## References

1. Barrow, H.G., Tenenbaum, J.M., Bolles, R.C., Wolf, H.C.: Parametric correspondence and chamfer matching: two new techniques for image matching. In: Proc. 5th International Joint Conf. on Artificial Intelligens, Cambridge, USA (1977) 659–663
2. Borgfors, G.: An improved version of the chamfer matching algorithm. In: 7th International Conference of Pattern Recognition, Montreal (1984) 1175–1177

3. Borgefors, G.: Hierarchical chamfer matching: A parametric edge matching algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(6) (1988) 849–865
4. Ding, L., Goshtasby, A., Satter, M.: Volume image registration by template matching. *Image and Vision Computing* **19** (2001) 821–832
5. Zitová, B., Flusser, J.: Image registration methods: a survey. *Image and Vision Computing* **21** (2003) 977–1000
6. Mäkelä, T., Pham, Q.C., Clarysse, P., Nenonen, J., Lötjönen, J., Sipilä, O., Hänninen, H., Lauerman, K., Kirsi, L., Knuuti, J., Katila, T., Magnin, I.E.: A 3-D model-based registration approach for the PET, MR and MCG cardiac data fusion. *Medical Image Analysis* **7** (2003) 377–389
7. Thayananthan, A., Stenger, B., Torr, P.H.S., Cipolla, R.: Shape context and chamfer matching in cluttered scenes. In: 2003 IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition. Vol. 1., IEEE Computer Society (2003) 127–133
8. Gavrilu, D.M.: Pedestrian detection from moving vehicle. In Vernon, D., ed.: *Computer Vision - ECCV 2000: 6th European Conference on Computer Vision*. Vol. 1843 of *Lecture Notes in Computer Sciences*. (2000) 37–49 Part II.
9. Olson, C.F., Huttenlocher, D.P.: Automatic target recognition by matching oriented edge pixels. *IEEE Transactions on Image Processing* **6**(1) (1997) 103–113
10. Rosin, P.R., West, G.A.W.: Saliency distance transforms. *Graphical Models and Image Processing* **57**(6) (1995) 483–521
11. Sintorn, I.M., Borgefors, G.: Shape based identification of proteins in volume images. In Kalviainen, H., Parkkinen, J., Kaarna, A., eds.: *Image Analysis: 14th Scandinavian Conference, SCIA 2005*. Vol. 3540 of *Lecture Notes in Computer Science*., Springer-Verlag (2005) 253–262
12. Gonzalez, R.C., Woods, R.E.: *Digital Image Processing*. 2 edn. Prentice-Hall, Inc. (2002)
13. Nash, S.G., Sofer, A.: *Linear and Nonlinear Programming*. Industrial Engineering Series. McGraw-Hill International Editions (1996)
14. Rosenfeld, A., Pfaltz, J.L.: Sequential operations in digital picture processing. *Journal of the Association for Computing Machinery* **13**(4) (1966) 471–494
15. Borgefors, G.: On digital distance transforms in three dimensions. *Computer Vision and Image Understanding* **64**(3) (1996) 368–376
16. Borgefors, G., Olsson, H.Å.: Localizing and identifying objects: A method for distinguishing noise, occlusion, and other disturbances. In Pietikäinen, M., Seppänen, T., eds.: *Proceedings of (2nd) Nordic Workshop on Industrial Machine Vision*, Oulu, Finland (1992) ISBN 951-42-3316-6. Also published as FOA (Swedish Defence Research Establishment) Reprint B 30284-3.4 in Sep. 1992.
17. Sandin, S., Öfverstedt, L.G., Wikström, A.C., Wränge, O., Skoglund, U.: Structure and flexibility of individual immunoglobulin G molecules in solution. *Structure* **12** (2004) 409–415
18. Berman, H.M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., Bourne, P.E.: The protein data bank. *Nucleic Acids Research* **28**(1) (2000) 235–242