

Automatic Prone to Supine Haustral Fold Matching in CT Colonography Using a Markov Random Field Model

Thomas Hampshire¹, Holger Roth¹, Mingxing Hu¹, Darren Boone², Greg Slabaugh³, Shonit Punwani², Steve Halligan², and David Hawkes¹

¹ Centre for Medical Image Computing, University College London, London, UK

² Department of Specialist Radiology, University College Hospital, London, UK

³ Medicsight PLC, London, UK

Abstract. CT colonography is routinely performed with the patient prone and supine to differentiate fixed colonic pathology from mobile faecal residue. We propose a novel method to automatically establish correspondence. Haustral folds are detected using a graph cut method applied to a surface curvature-based metric, where image patches are generated using endoluminal CT colonography surface rendering. The intensity difference between image pairs, along with additional neighbourhood information to enforce geometric constraints, are used with a Markov Random Field (MRF) model to estimate the fold labelling assignment. The method achieved fold matching accuracy of 83.1% and 88.5% with and without local colonic collapse. Moreover, it improves an existing surface-based registration algorithm, decreasing mean registration error from 9.7mm to 7.7mm in cases exhibiting collapse.

1 Introduction

Computed tomographic colonography (CTC) is widely considered the preferred radiological technique for detecting colorectal cancer or potentially precancerous polyps. When characterising potential polyps, the radiologist must manually match corresponding areas in the prone and supine data. However, this is a difficult, time-consuming task due to considerable deformation that occurs during repositioning [1]. Hence, a method for automatic registration of prone and supine datasets has the potential to improve radiologists' efficiency and confidence. Furthermore, accurate surface registration could improve specificity of computer-aided detection systems (CAD).

A number of methods have been proposed to find correspondence between the prone and supine positions. For example, centreline-based methods extract and align colonic centrelines by stretching and shrinking based on path geometries [2]. Anatomical landmarks can be used to help align the two datasets by first identifying a stable set of anatomical features, such as the caecum, rectum and flexures [3]. Voxel-based methods provide a further means of registration [4]. However, these methods rely to varying extents upon continuous prone-supine

colonic segmentations, free from occlusion by fluid or collapse; a scenario which occurs infrequently in daily practice, despite optimal bowel preparation [5].

Fukano et al. proposed a registration method based on haustral fold matching [6]. A second-order derivative difference filter was used to extract folds; their volume and relative positions along the centreline were used for matching. The method relied on prior automatic identification of a set of landmark locations for registration. They reported correct registration of only 65.1% of large folds and 13.3% of small folds.

Recently, methods which involve conformal mapping of the colonic surface have been proposed in order to reduce the complexity of the three-dimensional task. For example, Zeng et al. combined conformal mapping with feature matching between the prone and supine surfaces [7]. The prone and supine colonic segmentations were mapped onto five rectangle pairs. Correspondences were established using a feature matching method based upon mean curvature. The method relied on accurately determining five matching segments in the prone and supine datasets, which is difficult to achieve and may not be possible in the case of local colonic collapse. The method proposed by Roth et al. [8], aims to overcome these limitations by mapping the entire endoluminal surface to a cylinder. Dense surface correspondence was then achieved by non-rigid cylindrical B-spline registration, driven by local shape measurements. However, this method can be susceptible to mis-registration of continuous sections due to the similarities of neighbouring features.

We present a novel method for generating a set of robust landmark correspondences between the prone and supine CT data. While previous methods tried to match corresponding folds based on spatial location and size alone, e. g. [6,7], we also compare visual renderings of the colonic surface at the fold positions as well as local geometric information, without reliance on a conformal mapping. The matching problem is modelled with an MRF and the maximum a posteriori labelling solution is estimated to provide a correspondence. This method explicitly addresses the problem of colonic collapse and can provide an initialisation in order to improve methods which aim to provide a full surface correspondence.

2 Methods

2.1 Haustral Fold Segmentation

Haustral folds are elongated, ridgelike structures on the endoluminal surface which can be identified by extracting curvature measurements from a surface reconstruction of the colonic wall. The maximum and minimum values of the normal curvature at a point are called the principal curvatures, k_1 and k_2 respectively. A metric based on the principal curvatures is used to classify each vertex as fold, or non-fold: $M = k_1 - \gamma||k_2||$. This recognises that at a fold, one expects $k_1 \gg 0$ and $k_2 \approx 0$. The γ parameter penalises the metric against curvature in any direction other than in the maximum, helping to separate the folds at the teniae coli. The surface mesh is treated as a graph, with graph nodes defined




Fig. 1. External (a) and internal (b) views of segmented haustral folds with marked centres. Virtual colonoscopy views of corresponding folds in the prone (c) and supine (d)

by the mesh vertices and graph edges defined by the mesh edges. Using a virtual sink and source with the given weighting, a graph cut segmentation [9] is performed which minimises an energy function using M as a unary term and a Potts energy function smoothing term: $\delta(f_p, f_q) = (f_p \neq f_q) \cdot \zeta$. This results in a label assignment of fold or non-fold over the entire surface mesh (Fig. 1). The centre of each fold is taken as the vertex with the shortest maximum distance to any vertex lying on the border of the segmented region. An observer study was carried out using 4 data sets to evaluate the segmentation performance. Given these cases, a sensitivity of 86.5% and positive predictive value of 97.0% is achieved.

2.2 Markov Random Field Modelling

The matching of prone and supine haustral folds is formulated as a labelling problem. First, m haustral folds are detected in the supine data and these are uniquely labelled; the objective is to then assign labels to the detected prone folds, achieved by solving an MRF. The identified haustral folds in the prone data set are modelled as sites $\mathbf{S} = S_1, \dots, S_n$, each of which has an associated random variable $\mathbf{F} = F_1, \dots, F_n$ taking on a discrete label $\mathbf{f} = f_1, \dots, f_n$ taken from the set of haustral folds identified in the supine data set. A neighbourhood system $\mathbf{N} = N_i | \forall i \in \mathbf{S}$ defines the extent of local connections between sites; and a pair-wise clique defined on \mathbf{N} and \mathbf{S} , $\mathbf{C}_2 = \{i, i'\} | N_i, i \in \mathbf{S}$ allows the incorporation of a-priori knowledge of geometric dependencies between labels.

The maximum a posteriori (MAP) estimate of the optimum labelling is computed, which is equivalent to minimising the energy function:

$$\mathbf{f}^{(MAP)} = \arg \min_{\mathbf{f}} \left[\sum_{i \in \mathbf{S}} V_u(f_i) + \sum_{i \in \mathbf{S}} \sum_{j \in N_i} V_p(f_i, f_j) \right] \quad (1)$$

where $V_u(f_i)$ is the unary term, a cost function for assigning label value f_i to site S_i . The pair-wise term $V_p(f_i, f_j)$ is the cost for assigning neighbouring sites S_i and S_j to their current values.

Unary Cost Function. The aim is to calculate an $n \times m$ unary cost matrix, where $n = \|\mathbf{S}\|$ is the number of sites or prone folds, and $m = \|\mathbf{L}\|$ is the number of labels or supine folds. To calculate the cost $V_u(f_i = L_j)$ of assigning label L_j to site S_i image patches are rendered at the fold positions, visualising the internal colon wall (Fig. 1). The resulting images are then compared using a sum-of-squared-difference similarity metric. An optimisation over the external parameters of the virtual camera used to visualise the supine dataset accounts for any inaccuracies in the fold point identification. Restricting the number of degrees of freedom of camera search ensures that the camera focus remains on the correct fold. The degrees of freedom are as follows: Elevation (θ) - the fold centre and camera right vector give a position and axis about which the camera is rotated; roll (ϕ) - rotation around the camera view direction; dollying (τ) - translation along the camera right vector.

Given the three parameters θ, ϕ, τ ; the optimisation finds the local minimum in a mean sum of squared difference between the rendered images I_1 and I_2 using Powell's gradient descent method [10]. Adding a scaling parameter W allows the weighting of unary to pair-wise costs. Additionally, a constraint is added so that the matching folds must lie in a similar region. The fold centreline positions ν_s^c, ν_p^c are used to limit corresponding fold matches to a window of one fifth of the colon length. Finally, a constant unary cost α is associated with the assignment of the null label to any given node, allowing for missing labels. We define $R(I_1, I_2) = W\sqrt{MSSD(I_1, I_2|\theta, \phi, \tau)}$. The unary costs are then defined:

$$V_u(f_i = L_j) = \begin{cases} \alpha & \text{if } L_j = L_0 \\ R(I_1, I_2) - \min(R(I_1, \cdot)) & \text{if } \nu_{s \rightarrow p}^c \leq \|\nu_{full}^c\|/10 \\ \infty & \text{otherwise} \end{cases} \quad (2)$$

Pair-wise Cost Function. To improve labelling performance, geometric information about neighbouring fold positions can be used. In this work a Rotation Minimising Frame (RMF) [11] is employed to describe the relative position of each fold to its neighbours: $\boldsymbol{\nu} = [\nu^c, \nu^\theta]^T$; where ν^c is the difference in fold position along the centreline and ν^θ , the difference in angle of rotation around the centreline. This 2D parameterisation simplifies the description of the translation between corresponding pairs of folds between the prone and supine as the centreline ν^c and rotational ν^θ displacement should be similar ($\boldsymbol{\nu}_p \approx \boldsymbol{\nu}_s$). Alternatively we can state $\boldsymbol{\nu}_p = \boldsymbol{\nu}_s + \epsilon$, where ϵ represents some uncertainty, and can be modelled with a zero mean bivariate normal distribution $\epsilon \sim \mathcal{N}(0, \boldsymbol{\Sigma})$, with $\boldsymbol{\Sigma} = \text{diag}(\Sigma^c, \Sigma^\theta)$. Finally we recognise that the position of a neighbouring site becomes more uncertain as the displacement along the centreline increases. For each site pair $\{i, j\}$ we calculate $\boldsymbol{\Sigma}_{ij} = (I + (\nu_S^c \cdot \boldsymbol{\lambda}))\boldsymbol{\Sigma}_{base}$, where $\boldsymbol{\lambda} = \text{diag}(\lambda_1, \lambda_2)$ is a parameter to control the increase in positional uncertainty with centreline displacement, and $\boldsymbol{\Sigma}_{base}$ is a base covariance. With this information, a pair-wise cost for assigning neighbouring sites S_i and S_j label configurations f_i and f_j , is defined by the negative log-normal distribution: $\varphi(f_i, f_j) = (\boldsymbol{\nu}_S - \boldsymbol{\nu}_L)^T \boldsymbol{\Sigma}_{ij}^{-1} (\boldsymbol{\nu}_S - \boldsymbol{\nu}_L) / 2$. A local neighbourhood system is defined in order to enforce local geometric constraints on neighbouring fold positions.

The local neighbourhood of a site is set to be $N_i^{local} = \{i \neq j \mid \|\nu_i^c - \nu_j^c\| < t\}$, where t is a threshold distance. Pairs of sites that are separated by a local colonic collapse are removed from the neighbourhood set. A uniqueness constraint is also enforced so any two sites may not be assigned the same label. This is included in the pair-wise cost function by connecting each site with every other site in a global neighbourhood system $N_i^{global} = \{i \neq j\}$ and defining the pair-wise cost of assigning the same label to two nodes to be infinity, except in the case of a null label assignment $V_b(f_i, f_j) = \infty$ if $f_i = f_j \neq L_0$. The full pair-wise cost function is:

$$V_b(f_i, f_j) = \begin{cases} \infty & \text{if } f_i = f_j \\ \beta & \text{if } S_j \in N_i^{local} \text{ and } (f_i = L_0 \text{ or } f_j = L_0) \\ \varphi(f_i, f_j) & \text{if } S_j \in N_i^{local} \text{ and } f_i \neq L_0 \text{ and } f_j \neq L_0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

MRF Inference. The uniqueness constraint on the pair-wise costs means the problem of solving the MRF is non-submodular and restricts the possible algorithm choice for MAP inference. The Belief Propagation (BP) algorithm is suitable for this purpose. To estimate the maximum probability state configuration, the min-sum variant is used [12]. Parameters are found using a gradient ascent optimisation on training data.

3 Experimental Results

3.1 Clinical Validation

Ethical approval and informed consent was obtained to use anonymised CT colonography data. Colonic cleansing and insufflations had been performed in accordance with current recommendations [13]. A radiologist (with experience of over 500 endoscopically validated studies) used virtual colonoscopic reconstructions to identify corresponding folds in the prone and supine datasets to establish a reference standard. Any folds where a confident manual correspondence could not be established were disregarded. This resulted in a total of 1175 corresponding fold pairs over 13 datasets, 5 of which exhibited at least one local colonic collapse in one or both views (case 10 is shown in Figure 2). For a small subset of 3 randomly selected cases the reference standard was reestablished after a period of three months. Folds which have a correspondence in both sets of reference standard are used to evaluate the level of intra-observer variability, showing a 85.3% agreement.

3.2 Haustral Fold Matching

To assess the performance of the algorithm, for each case the maximum a posteriori labelling solution is compared against the reference standard described above. Table 1 shows the results for the cases with and without colonic collapse. Although the percentage of correctly labelled folds is high, at 83.1% and 88.5%




Fig. 2. External view of prone (a) and supine (b) datasets in case 10. The dotted line indicates an area of colonic collapse. (c) shows the conformally mapped colonic prone (top) and supine (bottom) surfaces at the sigmoid with shape index values. The rectangles represent unfolded cylinders. Red areas indicate a collapsed region. Black lines show detected fold correspondences of which all agree with the reference standard.

Table 1. Fold labelling performance. The number of Reference Standard (RS) points are shown. Label shows the percentage of RS points identified with a label.

	Without colonic collapse									With colonic collapse					
Case	1	2	3	4	5	6	7	8	Total	9	10	11	12	13	Total
RS Points	74	104	112	88	86	112	107	91	774	65	107	66	83	80	401
Labelled	66	97	106	84	82	92	99	88	714	62	101	63	77	51	354
Correct	49	90	98	70	74	76	91	84	632	50	78	53	74	39	294
Incorrect	17	7	8	14	8	16	8	4	82	12	23	10	3	12	60
Label(%)	89.2	93.3	94.6	95.5	95.3	82.1	92.5	96.7	92.2	95.4	94.4	95.5	92.8	63.8	88.3
Correct(%)	74.2	92.8	92.5	83.3	90.2	82.6	91.9	95.5	88.5	80.6	77.2	84.1	96.1	76.5	83.1

for cases with and without colonic collapse respectively, it is apparent that some cases show a much higher rate of accuracy than others. This is primarily due to different levels of distension causing inconsistent fold identification between data sets. Whilst the method is robust to missing data, a greater proportion of correctly labelled folds is demonstrated in cases with better fold segmentation. A second factor which contributes largely to the variance in fold pair matching in the pair-wise term, is the inaccurate estimation of the angle of rotation between folds using the centreline-based RMF estimation method, especially apparent in areas of high torsion, such as the flexures.

3.3 Initialisation of Surface Based Registration Method

There are scenarios where obtaining a one-to-one surface correspondence is required. In this case, the results of this fold matching method can be used to provide automated initialization for a surface-based registration technique [8]. Here, the fold positions are mapped onto a conformally mapped image and used to provide a linear scaling between haustral folds in the direction of the centreline, prior to the B-spline registration where full surface correspondence is established. To determine the registration error, each reference standard point is

Table 2. Surface registration initialisation with non-collapsed cases. The number of Reference Standard (RS) points are shown. Error 1 and 2 show the error of the surface-based registration without and with using points as an initialisation.

Without colonic collapse										With colonic collapse					
Case	1	2	3	4	5	6	7	8	Total	9	10	11	12	13	Total
RS Points	74	104	112	88	86	112	107	91	774	65	107	66	83	80	401
Error 1 (mm)	11.5	8.6	5.3	5.7	5.5	5.2	5.8	6.7	6.6	12.2	6.5	7.8	13.5	9.6	9.7
Error 2 (mm)	11.5	7.2	5.5	5.7	5.8	5.5	6.1	6.9	6.6	7.9	5.8	7.8	8.7	9.1	7.7
Difference (mm)	0.0	-1.4	0.2	0.0	0.3	0.3	0.3	0.2	0.0	-4.3	-0.7	0	-4.8	-0.5	-2

transformed from one dataset to the other using the registration result, and the 3D Euclidean distance between this and the corresponding reference standard point is measured. The results for cases with and without colonic collapse are shown in Table 2 using the same reference standard as in the previous experiment. It can be seen that the initialisation improves registration in cases exhibiting local colonic collapse, decreasing the mean error from $9.7mm \pm 8.7mm$ to $7.7mm \pm 7.1mm$; however in cases without a local colonic collapse, the mean error was unchanged at $6.6mm$. This shows that the fold matching technique is more robust than the surface-based registration in the case of poor insufflation (e.g. collapse); however, in well distended cases, the space to improve upon the registration of the surface-based method is limited. Using a Related Samples Wilcoxon Signed Rank Test with a significance level set at $p < 0.01$, the differences in mean errors are not significantly significant in the cases without colonic collapse ($p = 0.317$). However, in the cases with colonic collapse a statistical significance is observed ($p = 0.009$).

4 Conclusion and Future Work

Although several registration methods have been proposed, they often require manual initialisation using anatomical landmarks or are likely to fail in the presence of local colonic under-distension. Colonic collapse or segments with inadequate fluid tagging are very common in routine practice and methods must have the inherent capability to deal with this in order to be clinically useful. We propose a novel method for detecting and establishing correspondence between haustral folds in prone and supine CT colonography data sets which accurately deals with these issues. In addition, applying this method to initialise a surface-based registration technique can reduce registration error. It is clear that the unary costs are not as reliable in areas of high deformation due to the change in appearance of haustral folds. Investigation into a dynamic weighting of unary to pair-wise costs may allow for improved robustness in these situations. Future work would include the use of tenaie coli identification to provide a 'zero angle', improving the accuracy of the pair-wise costs. Although thorough clinical validation on a large dataset is ongoing, current results presented in this paper clearly demonstrate the promise of our MRF approach to match corresponding haustral folds in CTC data.

References

1. Punwani, S., Halligan, S., Tolan, D., Taylor, S.A., Hawkes, D.: Quantitative assessment of colonic movement between prone and supine patient positions during CT colonography. *British Journal of Radiology* 82(978), 475–481 (2009)
2. Wang, S., Yao, J., Liu, J., Petrick, N., Van Uitert, R.L., Periaswamy, S., Summers, R.M.: Registration of prone and supine ct colonography scans using correlation optimized warping and canonical correlation analysis. *Medical Physics* 36, 5595 (2009)
3. Näppi, J., Okamura, A., Frimmel, H., Dachman, A., Yoshida, H.: Region-based supine-prone correspondence for the reduction of false-positive cad polyp candidates in ct colonography. *Academic Radiology* 12(6), 695–707 (2005)
4. Suh, J.W., Wyatt, C.L.: Deformable registration of supine and prone colons for computed tomographic colonography. *Journal of Computer Assisted Tomography* 33(6), 902–911 (2009)
5. Taylor, S.A., Halligan, S., Goh, V., Morley, S., Bassett, P., Atkin, W., Bartram, C.I.: Optimizing Colonic Distention for Multi-Detector Row CT Colonography: Effect of Hyoscine Butylbromide and Rectal Balloon Catheter1. *Radiology* 229(1), 99 (2003)
6. Fukano, E., Oda, M., Kitasaka, T., Suenaga, Y., Takayama, T., Takabatake, H., Mori, M., Natori, H., Nawano, S., Mori, K.: Haustral fold registration in ct colonography and its application to registration of virtual stretched view of the colon. In: *Proceedings of SPIE*, vol. 7624, p. 762420 (2010)
7. Zeng, W., Marino, J., Gurijala, K.C., Gu, X., Kaufman, A.: Supine and prone colon registration using quasi-conformal mapping. *IEEE Transactions on Visualization and Computer Graphics* 16, 1348–1357 (2010)
8. Roth, H.R., McClelland, J.R., Boone, D.J., Modat, M., Cardoso, M.J., Hampshire, T.E., Hu, M., Punwani, S., Ourselin, S., Slabaugh, G.G., Halligan, S., Hawkes, D.J.: Registration of the endoluminal surfaces of the colon derived from prone and supine ct colonography. *Medical Physics* 38(6), 3077–3089 (2011)
9. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1124–1137 (2004)
10. Fletcher, R., Powell, M.J.D.: A rapidly convergent descent method for minimization. *The Computer Journal* 6(2), 163–168 (1963)
11. Klok, F.: Two moving coordinate frames for sweeping along a 3D trajectory. *Computer Aided Geometric Design* 3(3), 217–229 (1986)
12. Weiss, Y., Freeman, W.T.: On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs. *IEEE Transactions on Information Theory* 47(2), 736–744 (2002)
13. Taylor, S.A., Laghi, A., Lefere, P., Halligan, S., Stoker, J.: European society of gastrointestinal and abdominal radiology (esgar): consensus statement on ct colonography. *European Radiology* 17(2), 575–579 (2007)