Uncertainty-Based Feature Learning for Skin Lesion Matching Using a High Order MRF Optimization Framework

Hengameh Mirzaalian[1], Tim K. Lee[1,2,3], and Ghassan Hamarneh[1]

[1] Medical Image Analysis Lab, Simon Fraser University
[2] Cancer Control Research, BC Cancer Agency
[3] Department of Dermatology and Skin Science, University of British Columbia
{hma36,hamarneh}@sfu.ca, tlee@bccrc.ca

Abstract. We formulate the pigmented-skin-lesion (PSL) matching problem as a relaxed labeling of an association graph. In this graph labeling problem, each node represents a mapping between a PSL from one image to a PSL in the second image and the optimal labels are those optimizing a high order Markov Random Field energy (MRF). The energy is made up of unary, binary, and ternary energy terms capturing the likelihood of matching between the points, edges, and cliques of two graphs representing the spatial distribution of the two PSL sets. Following an exploration of various MRF energy terms, we propose a novel entropy energy term encouraging solutions with low uncertainty. By interpreting the relaxed labeling as a measure of confidence, we further leverage the high confidence matching to sequentially constrain the learnt objective function defined on the association graph. We evaluate our method on a large set of synthetic data as well as 56 pairs of real dermatological images. Our proposed method compares favorably with the state-of-the-art.

1 Introduction

The presence of a large number of pigmented skin lesions (PSL) is a strong predictor of malignant melanoma [7]. Since detecting newly appearing, disappearing, and changing PSL is important for early detection of the disease, many dermatologists advocate total-body photography for high-risk patients (Figure 1(a)). However, manual inspection and matching of PSLs is a subjective, tedious, and error prone task. A computer vision system for tracking the corresponding PSLs greatly improves the matching process, thereby easing the workload on dermatologists while also improving matching accuracy and removing operator variability [7]. There exists limited works on automating the matching between lesions. Huang and Bergstresser developed a PSL matching algorithm based on a Voronoi decomposition of the image space [3]. Yet, their method does not deal with the presence of the newly appearing or disappearing PSLs. Prednia and White performed affine registration between the two sets of PSLs [8]. However, their method does not take into account the elastic deformation of the human back. Roning and Riech defined a set of geometric properties as a similarity metric to find the corresponding PSLs. Their method requires manually determining two initial matches

N. Ayache et al. (Eds.): MICCAI 2012, Part II, LNCS 7511, pp. 98–105, 2012.

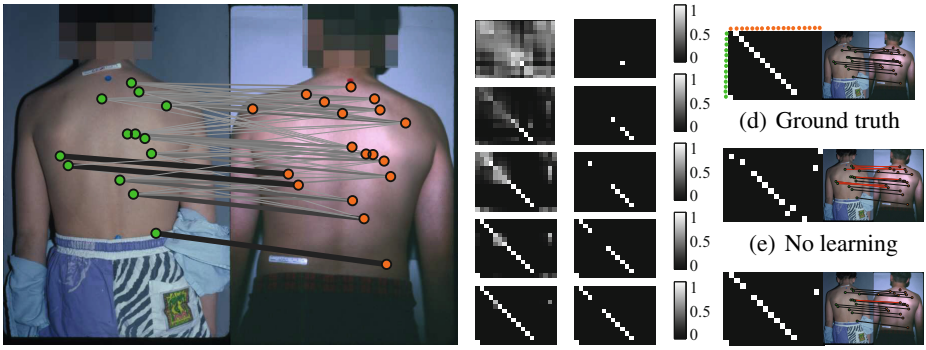

Fig. 1. (a) Example back images of the same subject at two different times. The green and red dots are overlaid at the PSL's coordinates. The thickness and the color of the edges encode the matching confidence between the connected points; the thicker and darker the line, the higher the confidence (please refer to Section 2.3 for more details). The five rows in (b) and (c) represent the output of five iterations of the learning step. (b) The probabilistic solutions. (c) The selected high confidence matchings. The ground truth matching is shown in (d). (e) and (f) show the estimated matching without and with the learning step. Wrong matches are shown in red on the back images. It can be noticed that the unsupervised learning step improves the matching accuracy (i.e. less red lines).

[9]. The authors in [7] computed the matching probabilities of the edges of two graphs representing the spatial distribution of the two PSL sets. They then extracted pointwise probabilities utilizing the marginalization matrix of the computed pairwise matchings. However, they did not make use of high-order term to the PSL matching. Recently, there have been several works on high order graph matching, combining both appearance similarity and geometric compatibility [1,10,11,12].

Compared with the previous works on PSL matching, we present the first application of high-order term to PSL matching. Our approach is most closely related to the work of Zeng et. al [12], who formulate a non-rigid surface registration problem as a high order graph matching problem and extract the matchings by solving a corresponding pseudo-boolean function. Their matching cost function depends on the feature appearance and geometric compatibility of the pair-wise and triplet-wise correspondences (Section 2.1). To solve their non-convex optimization problem, they make use of the dual-decomposition (DD) approach, similar to the work by Torresani et. al [10]. Our method differs from those in [10,12] in several ways. First, we relax the labels to continuous variables. By interpreting the relaxed labeling as a measure of confidence, we sequentially leverage the high confidence matchings via a self-learning approach to learn the features of the association graph (Section 2.3). We further propose to add a novel entropy energy term encouraging solutions with low uncertainty. We evaluate our method on a large set of synthetic data (hundreds of pairs) as well as 56 pairs of real dermatological images. The experimental results confirm the usefulness of adopting the entropy term and the unsupervised learning procedure (Section 3).

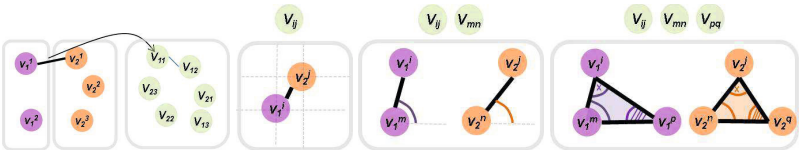

(a) $\mathcal{G}_1$, $\mathcal{G}_2$, and $\mathbb{G}$ (b) (c) (d)

Fig. 2. (a) Association graph $\mathbb{G}$ for graph matching between two graphs, $\mathcal{G}_1$ and $\mathcal{G}_2$. Each node in $\mathbb{G}$ represents a connection between a point in $\mathcal{G}_1$ and a point in $\mathcal{G}_2$. The matching problem between $\mathcal{G}_1$ and $\mathcal{G}_2$ is formulated as a labeling problem for $\mathbb{G}$. (b-d) illustrate examples of the unary, binary, and ternary terms used in the MRF-based labeling cost function in (1) (please refer to Section 2.1 for further details).

2 Method

Let us denote the PSLs coordinates of the l^{th} image by a graph $\mathcal{G}_l(\mathcal{V}_l, \mathcal{E}_l, \mathcal{C}_l)$, $l \in \{1, 2\}$, consisting of a set of nodes $\mathcal{V}_l$ ($|\mathcal{V}_l| = N_l$), edges $\mathcal{E}_l \subset \mathcal{V}_l \times \mathcal{V}_l$, and cliques $\mathcal{C}_l \subset \mathcal{V}_l \times \mathcal{V}_l \times \mathcal{V}_l$. We define a set of intra and inter-edges between the graphs to encode features related to the nodes connected by the edges and the cliques. An intra edge $\mathcal{E}_{lm,ln} \in \mathcal{G}_l$ connects the m^{th} vertex $\mathcal{V}_{lm}$ to the n^{th} vertex $\mathcal{V}_{ln}$, where $n \neq m$. An inter-edge $\mathcal{E}_{1m,2n}$ connects $\mathcal{V}_{1m}$ to $\mathcal{V}_{2n}$. Our aim is to find a mapping $\Pi(\mathcal{V}_1) \rightarrow \mathcal{V}_2$.

The matching problem (i.e. finding the mapping Π) can be formulated as a graph labeling problem. To this end, given $\mathcal{G}_l|_{l=1,2}$, we first construct their association graph $\mathbb{G}(\mathbb{V}, \mathbb{E}, \mathbb{C})$, in which each vertex in $\mathbb{V}$ corresponds to an inter-edge, e.g. $\mathbb{V}_{1m,2n} = \mathbb{V}_{mn} \leftrightarrow \mathcal{E}_{1m,2n}$ ($|\mathbb{V}| = N_1 N_2$) (Figure 2(a)). The matching problem can then be solved by binary labeling, x, of $\mathbb{G}$ [10]. A correspondence $\Pi(\mathcal{V}_{1m}) \rightarrow \mathcal{V}_{2n}$ is active iff $x(\mathbb{V}_{mn}) = 1$ and 0 otherwise. The details describing the objective function for binary labeling is provided in Section 2.1.

Compared with [10], we solve the matching problem as a relaxed (fuzzy) labeling, i.e. $x \in [0, 1]$. We interpret the fuzzy labels as a measure of confidence. The high confidence matchings are then extracted for unsupervised learning of the features of the association graph (Section 2.3).

Let us denote the label by $\ell \in \{0, 1\}$. Then, $x^\ell(\mathbb{V})$ represents our confidence in $\mathbb{V}$ having the label ℓ. Since we have the following equality $x^0(\mathbb{V}) = 1 - x^1(\mathbb{V})$ in our framework, we denote x^1 by x for simplicity.

2.1 MRF-Based Binary Labeling

MRF-optimization seeks the labeling x_p for each vertex $\mathbb{V}_p$ of graph $\mathbb{G}(\mathbb{V}, \mathbb{E}, \mathbb{C})$ by optimizing an energy function of the form:

$$E(x) = w_u \sum_{p \in \mathbb{V}} \phi_x(x_p) + w_b \sum_{(p,q) \in \mathbb{E}} \phi_{xx}(x_p, x_q) + w_t \sum_{(p,q,z) \in \mathbb{C}} \phi_{xxx}(x_p, x_q, x_z) \quad (1)$$

where ϕ_x is the unary term which measures the likelihood of labeling a vertex with a specific label disregarding the labels of any of the neighbours; and ϕ_{xx} and ϕ_{xxx} are

regularization terms penalizing different label configurations of neighboring vertices. w's are the weights of the different terms.

We define our unary term as a weighted sum of the two energy terms:

$$\phi_x(x_{ij}) = w_u^{(1)}\phi_x^{(1)}(x_{ij}) + w_u^{(2)}\phi_x^{(2)}(x_{ij}). \tag{2}$$

$\phi_x^{(1)}(x_{ij})$ measures the dissimilarity between the appearance descriptors of $\mathcal{V}_{1i}$ and $\mathcal{V}_{2j}$, denoted by $\mathcal{F}_{\mathcal{V}_{1i}}$ and $\mathcal{F}_{\mathcal{V}_{2j}}$:

$$\phi_x^{(1)}(x_{ij}) = x_{ij}d_u(\mathcal{V}_{1i}, \mathcal{V}_{2j}) + (1 - x_{ij})(1 - d_u(\mathcal{V}_{1i}, \mathcal{V}_{2j})), \tag{3}$$

$$d_u(A, B) = \sum_{r=1}^{R} |\mathcal{F}_A(\alpha_r) - \mathcal{F}_B(\beta_r)|, \quad \alpha = [\alpha_1\alpha_2...\alpha_R], \beta = [\beta_1\beta_2...\beta_R]$$

α, β are the indices of $\mathcal{F}_A$ and $\mathcal{F}_B$, which are compared to each other in computing d_u in (3), and are given by:

$$\mathcal{F}_{\mathcal{V}_{lm}} = \mathbf{x}(\mathcal{V}_{lm})\mathbf{1}_{1\times N_l+1} - [\mathbf{0}, \mathbf{x}(\mathcal{V}_{l1}), \mathbf{x}(\mathcal{V}_{l2}), ..., \mathbf{x}(\mathcal{V}_{lN_l})], \;\; l \in \{1, 2\}. \tag{4}$$

where $\mathbf{x}_{2\times 1}$ is the normalized coordinate of the PSLs resulting from applying the skin back-template proposed in [7]. In the first iteration, α and β in (3) are initialized with 1. Therefore, d_u measures the Euclidean distance between $\mathcal{V}_{1i}$ and $\mathcal{V}_{2j}$, i.e. $d_u\Big(\mathcal{F}_{\mathcal{V}_{1i}}(\alpha = 1), \mathcal{F}_{\mathcal{V}_{2j}}(\beta = 1)\Big) = |\mathbf{x}(\mathcal{V}_{1i}) - \mathbf{x}(\mathcal{V}_{2j})|$, and later on, as explained in Section 2.3, α and β in (3) will be updated in a sequential learning step to include more entries of $\mathcal{F}$ in computing d_u.

$\phi_x^{(2)}(x_{ij})$ in (2) is our new *entropy* term, which is used to encourage the cost function towards solutions with low entropy or low uncertainty:

$$\phi_x^{(2)}(x_{ij}) = -\Big(x_{ij}\log_2 x_{ij} + (1 - x_{ij})\log_2(1 - x_{ij})\Big) \approx x_{ij}(1 - x_{ij}). \tag{5}$$

Equation (5) shows a quadratic approximation term achieved using a second order Taylor expansion. We treat x as a probability when calculating Shannon's entropy although we didn't present our method in a formal probabilistic framework. Nevertheless, the intuition of having higher uncertainty as x nears 0.5 and lower uncertainty as x gets close to 1 or 0 still holds.

To measure compatibility between pairwise correspondences, we use:

$$\phi_{xx}(x_{ij}, x_{mn}) = x_{ij}x_{mn}d_b(\overrightarrow{\mathcal{V}_{1i}\mathcal{V}_{1m}}, \overrightarrow{\mathcal{V}_{2j}\mathcal{V}_{2n}}) + (1 - x_{ij}x_{mn})(1 - d_b(\overrightarrow{\mathcal{V}_{1i}\mathcal{V}_{1m}}, \overrightarrow{\mathcal{V}_{2j}\mathcal{V}_{2n}}))$$

$$\overrightarrow{\mathcal{V}_{li}\mathcal{V}_{lm}} = \mathbf{x}(\mathcal{V}_{li}) - \mathbf{x}(\mathcal{V}_{lm}), \; d_b(\overrightarrow{A}, \overrightarrow{B}) = \omega_b^1|1 - \frac{\overrightarrow{A}.\overrightarrow{B}}{|\overrightarrow{A}||\overrightarrow{B}|}| + \omega_b^2||\overrightarrow{A}| - |\overrightarrow{B}||. \tag{6}$$

d_b evaluates the length and direction agreement between the line segments $\overrightarrow{A}$ and $\overrightarrow{B}$. w_b^1, w_b^2 weight the direction and length terms.

To measure the compatibility in corresponding triplets, e.g. triangles $\mathcal{T}_1 = \mathcal{V}_{1i}\widehat{\mathcal{V}_{1m}\mathcal{V}_{1p}}$ and $\mathcal{T}_2 = \mathcal{V}_{2j}\widehat{\mathcal{V}_{2n}\mathcal{V}_{2q}}$, we use:

Algorithm 1.1. Our proposed MRF-based point matching algorithm.

Input: Two point sets $\mathcal{V}_1$ and $\mathcal{V}_2$; the spatial coordinates of the points: $\mathbf{x}(\mathcal{V}_1)$ and $\mathbf{x}(\mathcal{V}_2)$.
Output: A mapping between the vertices: $\Pi(\mathcal{V}_1) \rightarrow \mathcal{V}_2$.
Initialization: Construct $\mathbb{G}(\mathbb{V}, \mathbb{E})$ (Section 2); construct $\mathcal{F}_{\mathcal{V}_{lm}}$ (4); $\tau = 0.9$; $\alpha = 1$; $\beta = 1$; $w_u^{(1)} = 0.04$; $w_u^{(2)} = 0.1$; $w_b^{(1)} = w_b^{(2)} = 0.04$; $w_t^{(1)} = w_t^{(2)} = 0.02$.
Compute d_b (6), d_t (7), and $d_u\Big(\mathcal{F}_{\mathcal{V}_{1i}}(\alpha = 1), \mathcal{F}_{\mathcal{V}_{2j}}(\beta = 1)\Big) = |\mathbf{x}(\mathcal{V}_{1i}) - \mathbf{x}(\mathcal{V}_{2j})|$.
Define the objective function $E(x) = \mathbf{func}(\phi_x^{(1)}, \phi_x^{(2)}, \phi_{xx}, \phi_{xxx})$ (Section 2.1).
Optimize $E(x)$. ▷ e.g. apply SP [11] or TIP [1] to maximize $x = \max_x E(x)$.
$(\mathcal{A}, \mathcal{B}) = \{(i,j)|x_{ij} > \tau\}$. ▷ $\mathcal{A}$ and $\mathcal{B}$ are the indices of the high confidence nodes in $\mathbb{G}$.
if $\mathcal{A} + 1 = \alpha$ & $\mathcal{B} + 1 = \beta$ ▷ The high confidence nodes do not change any more.
    $\Pi \leftarrow$ Hard matching obtained by discretizing $X = [x_{ij}]$.
else
    $\alpha \leftarrow \mathcal{A} + 1$, $\beta \leftarrow \mathcal{B} + 1$.
    Compute $d_u\Big(\mathcal{F}_{\mathcal{V}_{1i}}(\alpha), \mathcal{F}_{\mathcal{V}_{2j}}(\beta)\Big) = \sum_{r=1}^{|\mathcal{A}|+1} |\mathcal{F}_{\mathcal{V}_{1i}}(\alpha_r) - \mathcal{F}_{\mathcal{V}_{2j}}(\beta_r)|$.
    Go to step 5.

$$\phi_{xxx}(x_{ij}, x_{mn}, x_{pq}) = x_{ij}x_{mn}x_{pq}d_t(\mathcal{T}_1, \mathcal{T}_2) + (1 - x_{ij}, x_{mn}, x_{pq})(1 - d_t(\mathcal{T}_1, \mathcal{T}_2)$$

$$d_t(\mathcal{T}_1, \mathcal{T}_2) = w_t^1|area(\mathcal{T}_1) - area(\mathcal{T}_2)| + \sum_{i=1}^{3} w_t^2|\measuredangle\mathcal{T}_1^i - \measuredangle\mathcal{T}_2^2| \quad (7)$$

d_t measures the difference between the area and the angles of the triangles. The weights w_t^1 and w_t^2 encode the trade off between preserving areas vs. angles.

2.2 Solving for the PSL Matching via MRF Optimization

Since we bootstrap our PSL matching from the high confidence matches in Section 2.3, we restrict our work to the relaxed version of the problem, while having the entropy term discouraging high uncertainty. We explore: (i) tensor power iteration (TPI) [1], and (ii) successive projection (SP)[1] [11] optimization methods (Section 3). Both TPI and SP provide a soft solution considering global constraints $\sum_i x_{ij} \leq 1$ and $\sum_j x_{ij} \leq 1$ to ensure partial matching and to avoid multiple matchings. Note that in Section 3 the results are provided using TPI.

2.3 Self-Learning

As shown in [6], learning the parameters that control the graph matching is important for improving the matching accuracy. The authors in [6] learn the weights w in (1) using gradient descent-based approach. We instead learn an improved objective function by encoding into the unary term new geometric information from the current high confidence matching. In the learning step of our method, we update α and

[1] The SP algorithm is applied to the marginalization matrix computed based on the probability of matching the edges and the cliques [11].

Table 1. Comparison between the different methods in terms of the optimization domain, energy terms, and the self-learning (SL) characteristics

Method	Optimizer Soft vs. Hard	$\phi_x^{(1)}$	$\phi_x^{(2)}$	ϕ_{xx}	ϕ_{xxx}	SL	Objective Function	
CVPR09 [7]	SP (Soft)	✓	×	✓	×	×	MRF1	MRF1 = func$(\phi_x^{(1)}, \phi_{xx})$
CVPR08 [11]	SP (Soft)	✓	×	✓	✓	×	MRF2	MRF2 = func$(\phi_x^{(1)}, \phi_{xx}, \phi_{xxx})$
ECCV08 [10]	DD (Hard)	✓	×	✓	×	×	MRF1	MRF1EN = func$(\phi_x^{(1)}, \phi_x^{(2)}, \phi_{xx})$
CVPR10 [12]	DD (Hard)	✓	×	✓	✓	×	MRF2	MRF2EN = func$(\phi_x^{(1)}, \phi_x^{(2)}, \phi_{xx}, \phi_{xxx})$
PAMI11 [1]	TPI (Soft)	✓	×	✓	✓	×	MRF2	
Proposed	TPI (Soft)	✓	✓	✓	✓	✓	MRF2EN	

β in (3), which indicate the indices of $\mathcal{F}$ that should be considered in measuring d_u. As shown in Algorithm 1.1, given the current high confidence matching $x(\mathbb{V}_{\mathcal{AB}})$, i.e. $(\mathcal{A}, \mathcal{B}) = \{(i,j)|x_{ij} > \tau\}$, where τ is a confidence-threshold and $|\mathcal{A}| = |\mathcal{B}| = R$ and R is the total number of the high confidence points, α and β in (3) are updated: $\alpha = \mathcal{A} + 1$ and $\beta = \mathcal{B} + 1$. Therefore,

$$d_u\Big(\mathcal{F}_{\mathcal{V}_{1i}}(\alpha), \mathcal{F}_{\mathcal{V}_{2j}}(\beta)\Big) = \sum_{r=1}^{R+1} |\mathcal{F}_{\mathcal{V}_{1i}}(\alpha_r) - \mathcal{F}_{\mathcal{V}_{2j}}(\beta_r)| \tag{8}$$

The α_r-th entry of $\mathcal{F}_{\mathcal{V}_{1k}}$ represents the distance between the vertex $\mathcal{V}_{1k}$ and $\mathcal{V}_{1\alpha}$. In fact, we are effectively diffusing the binary term to the unary term, since this entry in $\mathcal{F}$ is related to the length agreement between the edges. Figure 1 shows examples of the selected high confidence mappings at different iterations.

3 Results

Given a ground truth matching Π^*, and an estimated mapping Π obtained by discretizing the estimated fuzzy solution $X = [x_{ij}]$ (e.g. applying simple thresholding or the Hungarian algorithm [1,4], where $\Pi_{ij} = 1$ is interpreted as a mapping $\Pi(\mathcal{V}_{1i}) = \mathcal{V}_{2j}$), we use the following error measurement to evaluate the quality of the estimated mapping: $\Delta = \sum |\Pi - \Pi^*|/(N_1 N_2)$. We evaluate our method on synthetic data as well as 56 pairs of real images [2]. Note that we identify the PSLs' coordinates on our real data manually and the number of PSLs in our dataset is varied between 3 and 60.

Our synthetic data follows a setup similar to [7]. A cloud of n_c points are generated. The corresponding points in the second set are constructed by perturbing the n_c points. Then, different number of outliers n_o^1 and n_o^2 (representing disappearing and newly appearing PSLs) are added to the two sets.

In Table 1, we analyse our method and five state of the art point matching algorithms in terms of different characteristics. In summary, CVPR09 [7], CVPR08 [11], ECCV08 [10], CVPR10 [12], PAMI11 [1], and our method, can be implemented by setting the objective function in the form of MRF1, MRF2, or MRF2EN mentioned in Table 1, and applying different optimization approaches. For example, we can arrive to PAMI11 [1] by setting the objective function to MRF2 and using the TPI optimizer. To study

the effectiveness of the entropy term (5), we compare the matching errors resulting from using a given function, with and without the entropy term; i.e. compare MRF1 vs. MRF1EN and MRF2 vs. MRF2EN. The results in Figure 3 indicate that adding the entropy term can lead to lower error. The effect of applying different iterations of the self-learning procedure is shown in Figure 4. The results confirm the usefulness of our unsupervised learning from high confidence matches. Note that the errors are gradually decreasing by increasing the number of the iterations.

A comparison between the point matching methods: CVPR09 [7], CVPR08 [11], PAMI11 [1], and our method on the real data is shown in Figure 5. Note that all the methods are fed with the normalized coordinates of the PSLs resulting from applying the skin back-template proposed in [7]. It can be seen that the lowest error is resulting from MRF2EN+SL, i.e. the results of augmenting MRF2EN with the learning procedure.

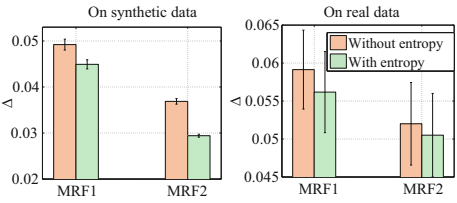

Fig. 3. Usefulness of the entropy term in (5): the y-axis represents the matching error resulting from applying different objective functions, without (orange bars) and with (green bars) the entropy term

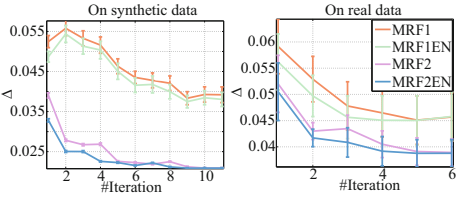

Fig. 4. Learning procedure: the y-axis represents the matching error. Different colors correspond to different objective functions. The results indicate that iteratively applying the unsupervised learning procedure (Section 2.3) leads to lower error.

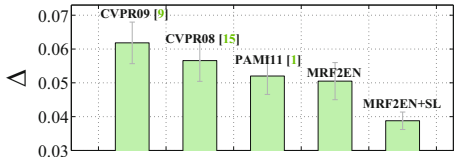

Fig. 5. Comparison between the point matching methods CVPR09 [7], CVPR08 [11], PAMI11 [1], and our method on the real data. Note that MRF2EN+SL has the lowest error.

4 Conclusion

We formulate the PSL matching problem in dermoscopic images as the relaxed labeling of the corresponding association graph in a high order MRF optimization framework. We add a novel entropy term to the objective function encouraging the cost function towards solutions with low uncertainty. We also propose to learn the objective function in a sequential framework by leveraging the high confidence matching of the fuzzy solutions. Although we evaluate the usefulness of the entropy term and the learning procedure on a specific application, the same idea can be used to extend other existing point matching algorithms.

This work can be extended in a number of ways. As mentioned in Section 2.3, for example, the learning step can be generalized for the binary and ternary terms of the matching objective function.

References

1. Duchenne, O., Bach, F., Kweon, I., Ponce, J.: A tensor-based algorithm for high-order graph matching. IEEE TPAMI 1(99), 1–13 (2011) 99, 102, 103, 104
2. Gallagher, R., Rivers, J., Lee, T., Bajdik, C., McLean, D.I., Coldman, A.: Broad-Spectrum Sunscreen Use and the Development of New Nevi in White Children: A Randomized Controlled Trial. JAMA 283(22), 2955–2960 (2000) 103
3. Huang, H., Bergstresser, P.: A new hybrid technique for dermatological image registration. IEEE BIBE, 1163–1167 (2007) 98
4. Kuhn, H.W.: The Hungarian method for the assignment problem. Naval Research Logistic Quarterly 2, 83–97 (1955) 103
5. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: IEEE ICCV, vol. 2, pp. 1482–1489 (2005)
6. Leordeanu, M., Zanfir, A., Sminchisescu, C.: Semi-supervised Learning and Optimization for Hypergraph Matching. In: IEEE ICCV (2011) 102
7. Mirzaalian, H., Hamarneh, G., Lee, T.: Graph-based approach to skin mole matching incorporating template-normalized coordinates. In: IEEE CVPR, pp. 2152–2159 (2009) 98, 99, 101, 103, 104
8. Perednia, D.A., White, R.G.: Automatic registration of multiple skin lesions by use of point pattern matching. CMIG 16(3), 205–216 (1992) 98
9. Roning, J., Riech, M.: Registration of nevi in successive skin images for early detection of melanoma. In: IEEE ICPR, vol. 1, pp. 352–357 (1998) 99
10. Torresani, L., Kolmogorov, V., Rother, C.: Feature Correspondence Via Graph Matching: Models and Global Optimization. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part II. LNCS, vol. 5303, pp. 596–609. Springer, Heidelberg (2008) 99, 100, 103
11. Zass, R., Shashua, A.: Probabilistic graph and hypergraph matching. In: IEEE CVPR, pp. 1–8 (2008) 99, 102, 103, 104
12. Zeng, Y., Wang, C., Yang, W., Gu, X., Samaras, D., Paragios, N.: Dense non-rigid surface registration using high-order graph matching. In: IEEE CVPR, pp. 382–389 (2010) 99, 103