

Towards Robust Identification and Tracking of Nevi in Sparse Photographic Time Series

Jakob Vogel^a, Alexandru Duliu^a, Yuji Oyamada^{a,c}, José Gardiazabal^{a,d},
Tobias Lasser^{a,b}, Mahzad Ziai^e, Rüdiger Hein^e, Nassir Navab^a.

^a Computer Aided Medical Procedures (CAMP), Technische Universität München, Germany;

^b Institute of Biomathematics and Biometry, Helmholtz Zentrum München, Germany;

^c School of Fundamental Science and Engineering, Waseda University, Japan;

^d Department of Nuclear Medicine, ^e Department of Dermatology and Allergology,
Klinikum Rechts der Isar, Technische Universität München, Germany;

ABSTRACT

In dermatology, photographic imagery is acquired in large volumes in order to monitor the progress of diseases, especially melanocytic skin cancers. For this purpose, overview (macro) images are taken of the region of interest and used as a reference map to re-localize highly magnified images of individual lesions. The latter are then used for diagnosis. These pictures are acquired at irregular intervals under only partially constrained circumstances, where patient positions as well as camera positions are not reliable. In the presence of a large number of nevi, correct identification of the same nevus in a series of such images is thus a time consuming task with ample chances for error. This paper introduces a method for largely automatic and simultaneous identification of nevi in different images, thus allowing the tracking of a single nevus over time, as well as pattern evaluation. The method uses a rotation-invariant feature descriptor that uses the local neighborhood of a nevus to describe it. The texture, size and shape of the nevus are not used to describe it, as these can change over time, especially in the case of a malignancy. We then use the Random Walks framework to compute the correspondences based on the probabilities derived from comparing the feature vectors. Evaluation is performed on synthetic and patient data at the university clinic.

Keywords: Biomedical Imaging, Dermatology, Feature Descriptor, Robust Matching, Random Walks

1. INTRODUCTION

Skin cancer, basal cell carcinoma, squamous cell carcinoma and melanoma are the most common types of cancer diagnosed in the United States. Melanoma represent a small subset of those, but it is the most deadly cutaneous neoplasm and the incidence is increasing by 4.1% per year, faster than any other malignancy.^{1,2}

There are several methodologies to differentiate a benign nevus from a melanoma, including the ABCDE test (Asymmetry, Border, Color, Diameter, Evolution).² The ABCD conditions can be evaluated directly (even automatically as in³), but to check the evolution of a nevus, the physician needs to compare it with a previously acquired image of the same nevus.

It is thus very important to re-locate the nevi in order to do a meaningful comparison. This scheme of relocate and compare, where the lesion itself might have changed, but the surroundings provide guidance, is conceptually similar to the endoscopic in-situ optical biopsies.^{4,5} While the acquisition methods and images are vastly different, the macro (localization) to micro (evaluation) approach are present in both problems.

During dermatological examination of a patient, macro images (see fig. 1b) are acquired at each visit and used as a “map” for relocation, while microscopic images (e.g. from a dermoscope) are used for the actual diagnosis. For the macro images, illumination and camera position are only approximately reproducible. Furthermore, the patient may be positioned differently, thus deforming the skin in a smooth, non-linear way.

Further author information: E-mail: duliu@in.tum.de




Figure 1: (a) Example imaging scenario. (b), (c) Sample of a real data set: left shoulder, back, full of nevi, with red box marking a region of interest (ROI). (d) False color overlay: green channel and red channel are taken from different photos of ROI. Arrows denote distortions.

The error observable between images nominally showing the same region of interest will consequently be a mixture of perspective and non-linear distortions (see fig. 1d). However, the general topology of the nevi will not change despite the deformations altering local distances, and the dermatologist will use the neighborhood of a certain nevus to identify it on different images, among other features.

When reproducing this process on a computer, a robust description of the neighborhood of a certain nevus is, therefore, vital. For realizing robust identification, we need to solve two research issues: how to represent such visual features (feature descriptors) and how to match/identify the represented features (feature matching). Computer vision approaches generally use visual features such as corners, edges, key-points, and texture.⁶ Almost all of the recently developed feature descriptors use the texture of the target objects, specifically image gradients.⁷ Considering our target, nevi on human skin, this is difficult as skin is quite homogeneous. Potential visual features on a nevus are its size, color, and shape. However, these features are sensitive to environmental changes, e.g. color is sensitive to lighting change, shape and size are sensitive to perspective changes, and the nevus itself might change over time. In particular, it is not advisable to rely on a purely distance-based description as this would fail due to perspective changes. For less-textured objects, one potential idea is to use the distribution of key-points.^{8,9} Their descriptors compute a feature vector of a key-point by using its neighboring key-points. However, distortions of the neighborhood e.g. due to skin deformations or perspective changes are problematic.

In the remainder of this article we introduce a rotation-invariant per-nevus descriptor, which is robust to perspective changes, non-linear deformations and changes in the shape and color of the nevus. This allows to largely automate the matching process, only requiring the dermatologist's help when the identification is not sufficiently secure. The proposed approach is easily extensible to additionally incorporate other features as well as prior knowledge.

2. METHODS

We assume to have several photographic images of a patient, all showing about the same region of interest, and a segmentation of the nevi in each image. Furthermore, we have the centroid $\mathbf{c}_i^t \in \mathbb{R}^2$ and an approximate radius $r_i^t \in \mathbb{R}$ for every “nevus” region i .

Using such centroids and radii only, the method consists of a local neighborhood descriptor for nevi, and a probabilistic matching method to find correspondences within images taken at different times. An overview of the whole pipeline is shown in fig. 2.

2.1 Local Neighborhood Descriptor

The proposed feature descriptor characterizes a nevus i at time t using its local neighborhood (we omit most references to t in this section for clarity) as a histogram $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,b})^T$ of b buckets. The histogram, like a local “radar screen”, gives approximate headings and rough distance estimates to the visible neighbors. Each bin of the histogram $x_{i,k}$ corresponds to a discretized direction $k \in [0, 2\pi)$ and accumulates contributions of nevi




Figure 2: Given images at t_0 and t_1 , feature vectors $\mathbf{x}$ are computed for every nevus in both images. These are used to construct a bipartite graph connecting all unknown to all reference nevi. Then, Random Walks computes a matrix P containing identity probabilities. Finally, unique and reliable pairs of corresponding nevi are selected.

located along the respective direction. The contribution of a particular neighbor j to bucket k within the horizon of nevus i is based on two criteria, a *distance-based weight* and the *relative angular overlap*.

Given the distances l_{ij} between nevus i and j (blue line in fig. 3b), the weighting function $w_{ij} = (1/l_{ij})^{0.5}$ assigns large impact to very close neighbors only, while distance differences of farther neighbors get levelled. l_{ij} can optionally be normalized such that all mean distances to the respective m closest neighbors of all nevi are equal.

Algorithm 1 Pseudocode for radar screen computation

```

1: for all  $(i, j) \ 1 \leq i, j \leq N \setminus i = j$  do
2:    $l_{ij} = \|\mathbf{c}_j - \mathbf{c}_i\|_2$ 
3:    $w_{ij} = (1/l_{ij})^{0.5}$  or  $(1/\text{normalize}(l_{ij}))^{0.5}$ 
4:   for all  $k$  do
5:     compute relative angular overlap  $o_{ijk}$ .
6:      $x_{ik} += w_{ij} * o_{ijk}$ 
7:   end for
8: end for

```

The second criterion, angular overlap, is computed from two circle sectors: Every nevus j defines a sector within the horizon of nevus i (gray fan in fig. 3b), spanned by the centroid $\mathbf{c}_i$ and the two extremal points of nevus j on the axis orthogonal to the difference vector. More formally, the latter two are given by $\mathbf{c}_j \pm r_j \mathbf{n}_{ij}$ where $\mathbf{n}_{ij}$ is the normal to the difference vector $\mathbf{d}_{ij} = \mathbf{c}_i - \mathbf{c}_j$. Then, the absolute angular overlap $\hat{o}_{ijk}$ relative to the bucket's sector (red area in Fig. 3b) is computed, and the relative angular overlap o_{ijk} is then the ratio of $\hat{o}_{ijk}$ over bucket width.

The product of w_{ij} and o_{ijk} is accumulated into the bucket $x_{ik} = \sum_j w_{ij} \cdot o_{ijk}$ over all neighbors j (Fig. 3a).

2.2 Comparing Descriptors

Considering the relatively sparse structure of nevi, we assume two to be identical if their “Radar” Histograms (feature vectors) show about the same features. In order to compute the probability estimate $P[i \equiv j]$ of a




Figure 3: Computing the feature vector of nevus i from its m nearest neighbors.

match between nevi i and j , we first rotate the “Radar” Histogram of j until we have maximal coincidence (i.e. minimizing $\|x_i - x_j\|_2$) and then compute P via a Gaussian weight function of $\|x_i - x_j\|_2$ as:

$$P[i \equiv j] = \exp \left(-\beta \min_k \|x_i - \text{shift}(\mathbf{x}_j, k)\|_2 \right). \quad (1)$$

β is the scaling factor controlling the width of the Gaussian bell curve and $\text{shift}(\mathbf{x}_j, k)$ is a helper function that rotates the histogram $\mathbf{x}_j$ by k places, i.e. $x_l = x'_{l+k} \forall l \leq b - k$ and $x_l = x'_{l+k-b}$ else.

Furthermore, in order not to be too sensitive to aliasing effects encountered while computing the overlaps, we smooth the descriptors before computing (1) by convolving the descriptor with a discrete approximation of a Gaussian bell. In particular, the filtering routine considers the circular structure of the descriptor and “wraps around”.

2.3 Simultaneous Identification using Random Walks

The major problem in a purely putative approach is the fact, that a single nevus at time t_0 can be assigned to multiple nevi at time t_1 , and vice versa. It is thus necessary to evaluate the correspondences simultaneously, and we propose to use an adaption of the Random Walks framework.¹⁰

Assume that all n_0 nevi at time t_0 are considered to be fix, and that we need to find correspondences at time t_1 out of a set of n_1 nevi. We can then create a weighted bipartite graph by connecting every nevus at time t_0 to every nevus at time t_1 (see fig. 2), and assign the probability (1) as respective weight to every edge. We label every nevus at time t_0 with a unique label, and compute the label probability distributions for all nevi at time t_1 using Random Walks.

As a result, we obtain a matrix $P \in \mathbb{R}^{n_0 \times n_1}$ where entry P_{i_0, j_1} contains the probability that nevus j_1 at time t_1 corresponds to nevus i_0 at time t_0 . In particular, every column p_{j_1} gives a probability distribution for every possible label that could be assigned to nevus j_1 .

2.4 Match Extraction

Once the probability matrix P is obtained, the proposed identification algorithm assigns a label to each nevus. In contrast to the original application of Random Walks, image segmentation, we can not just assign the most likely label because a single nevus at time t_1 can be assigned to multiple nevi at time t_0 , and vice versa. Therefore, the proposed algorithm makes sure to assign every label only once and omits unreliable matches.

We initialize two index vectors, $\eta_0 = (1, \dots, n_0)$ for the rows and $\eta_1 = (1, \dots, n_1)$ for the columns. Then, we apply the following iterative greedy strategy:

1. Find the maximal probability $p_{\max}$ in P at indices $i_{\max}, j_{\max}$.
2. Extract the column $\mathbf{p}_{j_{\max}}$ and extract the second-largest probability p_{second} from that distribution. Compute the trust value $\tau = p_{\max}/p_{\text{second}}$, and abort if this value is too close to 1, where both are almost equally likely. In an interactive system, we can now ask the human expert for help.
3. If our selection is sufficiently secure, store that nevus $\eta_0(i_{\max})$ at time t_0 and $\eta_1(j_{\max})$ and time t_1 are corresponding.
4. Delete the entire row $i_{\max}$ and column $j_{\max}$ from matrix P , and delete the respective components from the row and column index vectors η_0 and η_1 .
5. Fix the probability distributions by individually normalizing every column sum of P to 1.




Figure 4: Segmented nevi at t_0 and t_1 are drawn on an image as red and green blobs respectively. Correct matches are marked by cyan lines, while incorrect matches are marked by yellow lines.

3. EXPERIMENTS AND RESULTS

For the evaluation of our method, we studied synthetic data as well as real patient data. The methods have been implemented using Matlab, and were executed on a state-of-the-art portable computer. The matching pipeline runs within seconds for all data.

The main aim is to identify correct correspondences. Given a sufficient number of correspondences, the dermatologist can relocate the nevi in question. We, therefore, count the number of *true positives* (TP) (true correspondences) and *false positives* (FP) (wrong correspondences) and compute the *positive predictive value* (*precision*) ($TP/(TP + FP)$). The results are stated as means $\pm$ standard deviation over the 50 data sets each.

3.1 Synthetic Data

Synthetic data was created by randomly placing nevi in 50 images, and subsequently deforming the dataset perspective, mapping it onto a curved surface or distorting it non-linearly using thin-plate splines with randomly displaced control points. The deformation parameters were chosen to realistically reproduce the situation observable in clinical practice in dermatology.

Deformation type	Perspective	Curved	Non-linear
PPV (using l_{ij})	$94.51 \pm 16.86\%$	$99.18 \pm 2.08\%$	$96.35 \pm 8.85\%$
PPV (using normalized l_{ij})	$96.29 \pm 10.54\%$	$99.36 \pm 1.94\%$	$95.92 \pm 8.30\%$

3.2 Patient Data

Real data was acquired at our university dermatology clinic. Viewing angle and image resolution are comparable to commercial digital dermoscopes. We segment the nevi following the concepts of Eigen decomposition-based methods for vessel enhancement¹¹ by evaluating the Hessian's eigenvalues at every pixel and several scales. The value is a measure for how much the image resembles a dark blob at the given pixel.

We conducted an experiment using clinical data with three patients, one patient with breathing deformations and two cases with camera viewpoint changes. There are six image pairs for each view change case, and four image pairs for the breathing case. Fig. 4 shows cropped sub-images from the results. For quantitative evaluation, we also show the average PPV for each patient in the following table.

Deformation	Breathing	View change 1	View change 2
PPV	$94.85 \pm 4.99\%$	$84.65 \pm 3.84\%$	$75.82 \pm 16.41\%$

Investigating the incorrect and missing matches more closely, there are two main causes: neighbors only visible in one of the images (i.e. impossible to match) and either of the images have strong perspective distortion. In the worst case (Fig. 4c), the images share only a small part and the view angle change is more than 45 degrees along the elevation axis. In the remaining cases, image segmentation errors have led to additional blobs wrongly identified as nevi and thus matching errors.

4. DISCUSSION AND CONCLUSION

We have presented a local descriptor for nevi based on the local neighborhood, and a matching algorithm. The evaluations on synthetic and clinical patient data both are very promising. The method has the potential to automatically identify large numbers of nevi with limited human interaction, thus enabling dermatologists to quickly navigate to certain nevi. Due to the probabilistic approach, the proposed method's architecture has several advantages in prospect to a future clinical program.

First, it is very simple and natural to extend the framework beyond pure location information by including domain specific heuristics such as nevus size or color. Second, we can provide a trust estimate for most decisions. In particular, we can automatically decide when the help of a human expert is required. In such a case, based on the probability distribution, we can even provide a small set of equally likely alternatives, thus keeping the effort low. Third, expert knowledge can also be included very elegantly in non-error cases. For instance, information about the surgical removal of a certain nevus can be included by removing the respective edges from the graph.

Even though the system has been presented as matching problem between two images, it is simple to extend the method to several images by adding additional layers of nodes to the graph.

For a real clinical solution, the most important aspect (which has been largely omitted in this article) is a reliable segmentation routine. Multi-spectral imaging has been used with huge success for comparable problems, and an eventual solution may use such concepts.¹² It may also be interesting to interleave the segmentation and matching processes. For instance, a matching may be used to compute a deformation field that can be used to transform the location of an unmatched nevus and start a guided segmentation there.

REFERENCES

- [1] Wolff, T., Tai, E., and Miller, T., "Screening for Skin Cancer: An Update of the Evidence for the U.S. Preventive Services Task Force," *Ann Intern Med* **150**(3), 194–198 (2009).
- [2] Rager, E. L., Bridgeford, E. P., and Ollila, D. W., "Cutaneous melanoma: update on prevention, screening, diagnosis, and treatment," *Am Fam Physician* **2005**;72:269-76 (2005).
- [3] Wighton, P., Sadeghi, M., Lee, T. K., and Atkins, M. S., "A Fully Automatic Random Walker Segmentation for Skin Lesions in a Supervised Setting," in [*Springer LNCS 5762 (MICCAI)*], 1108–1115 (2009).
- [4] Allain, B., Hu, M., Lovat, L. B., Cook, R. J., Vercauteren, T., Ourselin, S., and Hawkes, D. J., "A System for Biopsy Site Re-targeting with Uncertainty in Gastroenterology and Oropharyngeal Examinations," in [*Springer LNCS 6362 (MICCAI)*], 514–521 (2010).
- [5] Atasoy, S., Mateus, D., Meining, A., Yang, G.-Z., and Navab, N., "Targeted Optical Biopsies for Surveillance Endoscopies," in [*Springer LNCS 6893 (MICCAI)*], 83–90 (2011).
- [6] Szeliski, R., [*Computer Vision Algorithms and Applications*], ch. Feature detection and matching, 183–209, Springer (2010).
- [7] Lowe, D. G., "Distinctive image features from scale-invariant keypoints," *Int J Comput Vision* (2) (2004).
- [8] Nakai, T., Kise, K., and Iwamura, M., "Camera-based document image retrieval as voting for partial signatures of projective invariants," in [*Intl. Conf. on Document Analysis and Recognition (ICDAR)*], 379–383 (2005).
- [9] Uchiyama, H. and Saito, H., "Random dot markers," in [*IEEE Virtual Reality Conference*], 35–38 (2011).
- [10] Grady, L. J., "Random Walks for Image Segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.* (11) (2006).
- [11] Frangi, A. F., Niessen, W. J., Vincken, K. L., and Viergever, M. A., "Multiscale Vessel Enhancement Filtering," in [*Springer LNCS 1496 (MICCAI)*], 130–137 (1998).
- [12] Diebele, I., "Melanoma-nevus differentiation by multispectral imaging," *Journal of Biomedical Optics* (2011).