

Image Retrieval in Forensics: Application to Tattoo Image Database

Jung-Eun Lee, Wei Tong, Rong Jin, and Anil K. Jain

Michigan State University, East Lansing, MI 48824
{leejun11, tongwei, rongjin, jain}@cse.msu.edu

Abstract The continuing growth of and increasing dependence on forensic image databases require fast and reliable image matching and retrieval techniques. We present a content-based image retrieval (CBIR) system for a particular forensic image database, namely a large collection of tattoo images. The system employs a local point descriptor to represent images, and, given a query tattoo image, it retrieves near-duplicate images from a large-scale database. Despite the high retrieval accuracy of the system, the performance heavily relies on the quality of query images. If query images are of low quality, features extracted from the query are noisy and not sufficiently discriminative, resulting in poor retrieval performance. In this paper, we improve the robustness of the system, especially for low quality query images, which, consequently, improves the overall retrieval performance. We introduce effective weighting schemes for matching local keypoints as well as utilize metadata to further improve the retrieval performance. Experimental results on a database of 100,000 images show that our system has excellent retrieval performance with a top-20 retrieval accuracy of 90.5%.

Keywords: Near-duplicate image retrieval, forensic databases, biometrics, tattoo images

1. Introduction

Whether in passports, credit cards, laptops, or mobile phones, automated methods of identifying citizens through their anatomical features or behavioral traits have become a common feature of modern life. *Biometric recognition*, or simply *biometrics*, refers to the automatic recognition of individuals based on their anatomical and/or behavioral characteristics [1]. One of the most well known biometric traits is fingerprints. The success of automatic fingerprint systems in law enforcement and forensics around the world has prompted the use of biometrics in various civil identification systems. For example, in 2007 alone, US-VISIT (U.S. Department of Homeland Security Immigration and Border Management System) [2] collected fingerprint and face images of over 46 million visitors to the United States.

While tremendous progress has been made in biometrics and forensics, there are many situations where the primary biometric traits (i.e. fingerprint, face, and iris) alone are not able to identify an individual with sufficiently high accuracy. This is especially true when the image quality is poor (e.g., blurred or off-central pose in a surveillance camera) or a print of only a portion of the finger is available, as in the case of latent fingerprints lifted at crime scenes. In the case of face recognition, the matching performance severely degrades under pose, lighting and expression variations, occlusion, and aging. In such cases, it is critical to acquire supplementary information to assist in the identification procedure. Based on this rationale, the Federal Bureau of Investigation (FBI) is developing the Next Generation Identification (NGI) system for identifying criminals [3]. In addition to utilizing additional biometric modalities, such as palmprint and iris, to augment evidence provided by fingerprints, the NGI system will also include *soft biometric traits* (e.g. scars, marks, and tattoos, collectively referred to as SMT).

Soft biometric traits are characteristics that provide some identifying information about an individual, but lack the distinctiveness and permanence to sufficiently differentiate between two individuals [1]. Since soft biometric traits help narrow down the identity of a suspect or a victim in forensics investigations, many law enforcement agencies collect and maintain such information in their databases. It is thus not surprising that the FBI collection standard includes prominent scars, marks, and tattoos if they are present on a subject's body. In spite of the value of soft biometrics in forensics, putting them to practical use has been difficult. Unlike primary biometric traits, there is a

very large variability in pattern types in many of the soft biometric traits. While a primary biometric trait has its own unique physical representation (e.g. ridge patterns and minutiae in fingerprints; eyes, nose, and lip in faces; texture in irises), in contrast, tattoo images often consist of objects with varying shapes, color, and texture (Figure 1), making it challenging to effectively represent them. This is the main reason why relatively little effort has been made for automatic matching and retrieval of tattoo images.

Among the various soft biometric traits, tattoos have been considered one of the most important pieces of evidence. Tattoos provide more discriminative information for identifying a person than the traditional demographic indicators such as age, height, race, and gender [4]. In addition, since many individuals acquire tattoos in order to be identified as distinct from others, to display their personality, or to exhibit a membership in a group (see Figures 1(c)-1(e)), the analysis of tattoos often leads to better understanding of an individual's background and membership in various organizations. In this paper, we present an automatic image retrieval system for a large tattoo image database. Although the current system is focused on tattoo images, the design of the system can be easily adapted to other forensic image databases, such as shoeprints and gang graffiti images.



Figure 1. Tattoos for identification: (a) a tattoo on a suspect of several crimes, (b) tattoos of a victim of the 2004 Asian Tsunami, and (c)-(d) gang membership tattoos of the *Mexikanemi Mafia* gang, a well-known gang in Texas. Note the large intra-class variability in the same gang's membership tattoos (c)-(e).

2. Tattoo Image Retrieval

Tattoos engraved on the human body have been successfully used to assist human identification in forensics. This is not only because of the increasing prevalence of tattoos¹, but also due to their impact on other methods of human identification such as visual, pathological, or trauma-based identification. Tattoo pigments are embedded in the skin to such a depth that even severe skin burns often do not destroy a tattoo; tattoos were used to identify victims of the 9/11 terrorist attacks and the 2004 Asian tsunami [4] (Figure 1(b)). Criminal identification is another important application because tattoos often contain hidden meaning related to a suspect's criminal history, such as gang membership, previous convictions, years spent in jail etc. (see Figures 1 and 2).

Law enforcement agencies routinely photograph and catalog tattoo patterns for the purpose of identifying victims and suspects (who often use aliases). The ANSI/NIST-ITL1-2011 standard [5] defines eight major classes (i.e. human, animal, plant, flag, object, abstract, symbol, and other) and a total of 70 subclasses (e.g. male face, cat, narcotics, American flag, fire, figure, national symbols, and wording) for categorizing tattoos. A search of a typical tattoo image database currently involves matching the class label of a query tattoo with the labels for the tattoos in the database. The current practice of matching tattoos based on the manually assigned ANSI/NIST class labels has the following limitations:

¹ A study published in the Journal of the American Academy of Dermatology in 2006 reported that about 36% of Americans in the age group 18 to 29 have at least one tattoo [6].

- class label does not capture the semantic information in tattoo images,
- there are millions of tattoo images maintained by law enforcement agencies,
- tattoos often contain multiple objects and cannot be classified appropriately into the ANSI/NIST classes,
- tattoo images have large intra-class variability, and
- ANSI/NIST classes are not complete for describing new tattoo designs.



Figure 2. Tattoo images from the Michigan State Police database.

In order to overcome the limitations of the current practice of keyword-based tattoo matching, we have developed an automatic tattoo matching and retrieval system, called Tattoo-ID [7,8,9]. This system has been licensed to *MorphoTrak*, which plans to release a commercial version of Tattoo-ID [10]. To the best of our knowledge, Tattoo-ID is the first prototype of an operational system for tattoo image matching and retrieval. While Acton and Rossi [11] also proposed a tattoo matching and retrieval system based on global features (i.e. color and shape), their system was evaluated on high quality web-downloaded images where query images were synthetically generated from the gallery images. We have already shown [7] that global features used in [11] are not adequate to match tattoo images in operational databases.

3. The Tattoo-ID System

Tattoo-ID is based on content-based image retrieval (CBIR) [12], where the goal is to find the images from a database that are nearly duplicates of the query image. Although general-purpose CBIR systems have only limited retrieval performance due to the well known problem of semantic gap [12], CBIR systems have been shown to be quite effective for near-duplicate image retrieval [12], which fits in well with the objective of tattoo image retrieval. Tattoo-ID extracts keypoints from images using Scale Invariant Feature Transform (SIFT) [13], and uses matching algorithm [8,9] to measure the visual similarity between two images; the database images with the largest similarities to the query are retrieved. We choose SIFT because it yields the best performance for tattoo matching and retrieval compared to both the global image features (e.g. color, shape, and texture), and the other local descriptors (e.g. SURF, GLOH, and Harris Laplace [14]). More information about Tattoo-ID can be found in [7,8,9].

To objectively evaluate the performance of Tattoo-ID, we constructed a database of 64,000 tattoo images provided by the Michigan State Police (see Figure 2). The tattoo images were cropped to extract the foreground and suppress the background. To construct the query set, we manually identified 1,000 images in the database that have near duplicates. These duplicates are introduced in the database due to multiple arrests of the same person at different times or multiple photographs of the same tattoo taken at a booking time (see Figures 3 and 8). One of the duplicates is used as a query to retrieve the other duplicate(s) in the database. To examine the robustness of our system, we further augmented the 64,000 tattoo images with 36,000 randomly selected images from the ESP game database [15]. The retrieval performance of Tattoo-ID is evaluated by the Cumulative Matching Characteristics (CMC): for a given rank position N , its CMC score is computed as the percentage of queries whose matched images are found in the top- N retrieved images. Our previous work [9] has shown that Tattoo-ID is able to correctly retrieve the duplicate

tattoos in the top 20 images (i.e. $N=20$) for 85.6% of queries and the average retrieval time per query is ~191 seconds on an Intel Core 2, 2.66 GHz, 3 GB RAM processor (see Figure 3). In addition, an unsupervised ensemble ranking approach is proposed in [9] to manage the scalability problem; the approach achieves similar retrieval accuracy, (i.e. 85.9% rank-20 accuracy), at a significantly reduced retrieval time (i.e. 14.7 seconds/query).



Figure 3. Tattoo-ID retrieval examples. Each row shows a query tattoo (with the number of keypoints), top-7 retrieved images, and the associated matching score (number of matching keypoints). Note that three duplicates were retrieved from the database for query 1, and two duplicates retrieved for query 2.

3.1. Ugly Tattoos

While the overall retrieval accuracy of Tattoo-ID is quite good, the performance drops off significantly if query images are of low quality (Figure 4). For example, when images have low contrast, uneven illumination, or small tattoo size, only a small number of keypoints are extracted from the images, making it difficult to perform the matching. If tattoo images are covered by heavy body hair, the majority of keypoints are extracted from body hair, not from the tattoos. These noisy keypoints lead to a number of false matches and, consequentially, low retrieval accuracy. We refer to the images with limited retrieval performance as *ugly tattoo images*, following the nomenclature introduced for poor quality latent fingerprint images in the NIST-SD27 database.

To systematically evaluate the performance of Tattoo-ID for ugly tattoos, a subset of 252 ugly tattoo images was extracted from the 1,000 query images as follows:

1. query tattoo for which the correct duplicate cannot be retrieved in the top 20 ranks, or
2. query tattoo for which the matching score of the first retrieved image is small (<10) and the top-10 retrieved images have similar matching scores (the standard deviation of the top-10 matching scores is less than 0.1).

Figure 5 compares the retrieval performances of Tattoo-ID against 748 typical quality and 252 ugly quality queries. Compared to the typical quality queries (i.e. 97.7% rank-20 accuracy), the 252 ugly quality queries show significantly lower retrieval performance (i.e. 49.6% rank-20 accuracy). In this paper, we aim at improving the robustness of the system, especially for the low quality images, and, consequentially, improving the overall retrieval performance.



Figure 4. Examples of ugly quality tattoos and the number of extracted keypoints: (a)tattoo with low contrast, (b)tattoo with uneven illumination, (c)small tattoo size, (d)tattoos faded and covered with hair, and (e)tattoo covered by substantial body hair.

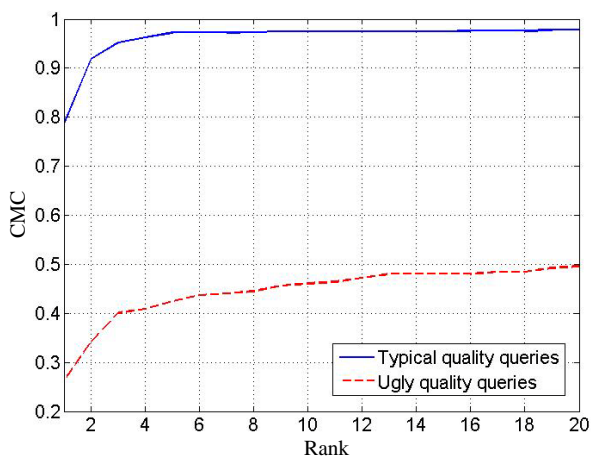


Figure 5. Retrieval performances for typical and ugly quality queries.

4. Enhancements to Tattoo-ID

We have improved the system performance by (i) developing more robust similarity measures, and (ii) utilizing the metadata associated with tattoo images. We discuss these enhancements in detail in this section.

4.1. Robust Similarity Measures

Due to the low image contrast and/or vagueness of faded tattoos, there are a number of spurious keypoints extracted that lead to many false matches. To address this challenge, we developed two strategies to improve the robustness of the similarity measure, i.e. *symmetric matching* and *weighted keypoint matching*.

Symmetric matching. To measure the similarity between a query image I_q and a database image I , denoted by $S(I_q, I)$, we compute the number of keypoints from I_q that match with the keypoints from I [13]. A keypoint K_q^i from I_q is considered to be matched to a keypoint from I , if the ratio of the shortest and the second shortest distance from K_q^i to the keypoints from I , is smaller than a predefined threshold γ ($\gamma = 0.49$). This similarity measure is *asymmetric*, i.e. $S(I_q, I) \neq S(I, I_q)$. One shortcoming of the asymmetric similarity measure is that it may produce many false matches, particularly if there is a keypoint in the database image I whose descriptor is very similar to that of several keypoints in I_q . We address this limitation by developing a symmetric similarity measure for a pair of

images I_q and I as follows: (i) compute the asymmetric match scores between I_q and I , and, between I and I_q , resulting in two sets of matched keypoint pairs, denoted by $M(I_q|I)$ and $M(I|I_q)$, (ii) compute the symmetric similarity measure, denoted by $S_S(I_q, I)$, as the number of matched keypoint pairs that appear in both sets, i.e., $S_S(I_q, I) = |M(I_q|I) \cap M(I|I_q)|$. Note that $S_S(I_q, I) = S_S(I, I_q)$. The symmetrization step allows us to remove some of the false matches.

Weighted keypoint matching. This approach tries to reduce the effect of false matches by introducing two sets of weights to the keypoints in a query image. It is based on the following two intuitions. First, if a keypoint K_l in a gallery image I is matched to multiple keypoints from a query image, we consider these multiple keypoints in the query image to be indistinctive and assign them low weights in the similarity measure. We refer to this weight as *local distinctiveness*. Second, if a keypoint K_q^i finds its matches from many different gallery images, we consider it to be indistinctive and assign it a low weight. We refer this weight as *global distinctiveness*. More specifically, suppose a query image I_q has l keypoints, $K_q = \{K_q^1, K_q^2, \dots, K_q^l\}$, and there are N_G images in the gallery G . Let $m^i(I)$ be the number of keypoints in K_q that are mapped to the same keypoint in a gallery image I as K_q^i , and n^i be the number of images in the gallery G where K_q^i finds its matched keypoints. Given $m^i(I)$ and n^i , the similarity between a query image I_q and a database image I , denoted by $S_W(I_q, I)$, is computed as follows:

$$S_W(I_q, I) = \sum_i^l x_i \left(\frac{1}{m^i(I)} \cdot \log \frac{N_G}{n^i} \right) \quad \text{where } x_i = \begin{cases} 1, & \text{if } K_q^i \text{ is matched} \\ 0, & \text{otherwise} \end{cases}$$

Figure 6 compares the retrieval performance of the asymmetric similarity, the symmetric similarity, and weighted keypoint matching on the database of 100,000 images with 1,000 query images that were described in Section 3. We observe that both the symmetric matching and weighted keypoint matching improve the retrieval performance. The average rank-20 accuracy is improved from 85.6% to 86.3% by the symmetric matching and to 88% by the weighted keypoint matching (Figure 6(b)). More noticeable improvements are observed for the ugly query images (Figure 6(a)), where the average rank-20 accuracy is improved from 49.6% to 51.8% by the symmetric matching and to 57% by the weighted keypoint matching. Finally, compared to the symmetric matching, the weighted keypoint matching is more effective. According to the student-t test (at the level of 5%), all the improvements are statistically significant. Overall, our result indicates that a soft weighting approach is more robust to false matches than a hard threshold approach such as the symmetric matching.

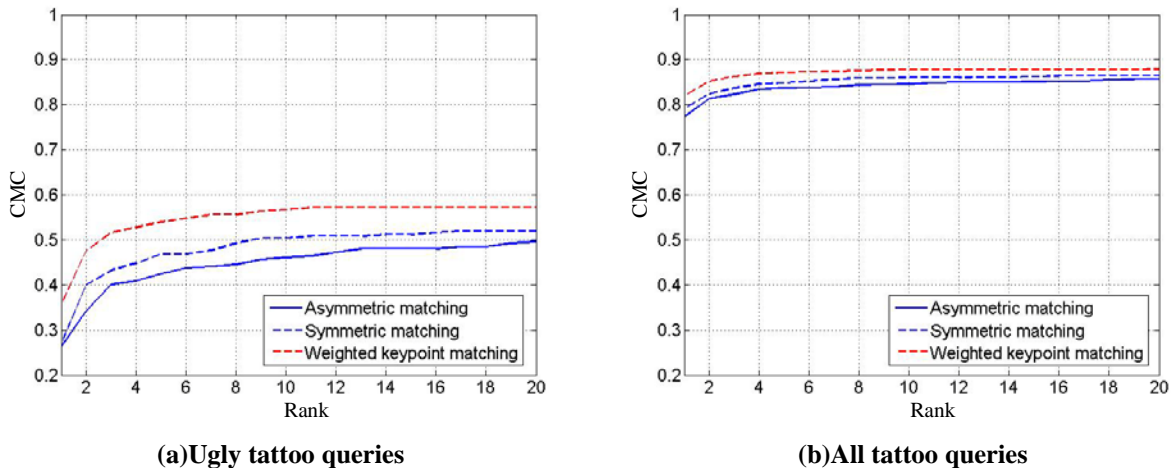


Figure 6. Retrieval performances for (a) 252 ugly quality queries and (b) all the 1,000 tattoo queries with the robust similarity measures.

4.2. Metadata Utilization

In order to further improve the retrieval performance, we evaluate the utility of metadata for tattoo image retrieval. We created a collection of tattoo images with manually assigned metadata. Due to substantial manual labor needed to label the images, we randomly selected 21,000 tattoo images from the 64,000 tattoo images in our database, including the 1,000 queries and their near-duplicate images, for manual annotation. The labeling was done by 12 subjects who were Michigan State University students. On average, each subject was asked to annotate about 3,500 images in two ways: using up to four ANSI/NIST major classes and his/her own keyword(s). The average number of classes assigned per a tattoo image is two and that of free keywords is ~ 3.5 . Each image is annotated by two subjects, and the final result is formed by merging the annotations from the two subjects. By performing spell check and word stemming, the final number of unique free keywords is 2,019. Recall that the number of ANSI/NIST major classes is eight. We use this collection of manually annotated tattoo images to examine the effect of metadata.

To utilize the ANSI/NIST-based metadata (eight major classes), we implemented a two-stage matching scheme: (i) select a subset of database tattoos that shared at least one class label with the query tattoo, (ii) perform keypoint-based image matching only for the selected subset. The retrieval results for 252 ugly quality tattoo queries and all the 1,000 tattoo queries are shown in Figure 7. We observe that in both the cases, the introduction of ANSI/NIST class labels leads to a significant drop in the retrieval performance. This is because each ANSI/NIST class covers a wide range of tattoo types. Consequently, “similar” tattoo images may be assigned to different classes, making it difficult to match tattoo images based on their class assignments (see Figure 8). This limitation of the ANSI/NIST major classes leads us to explore the free keyword annotation for improving tattoo image retrieval performance.

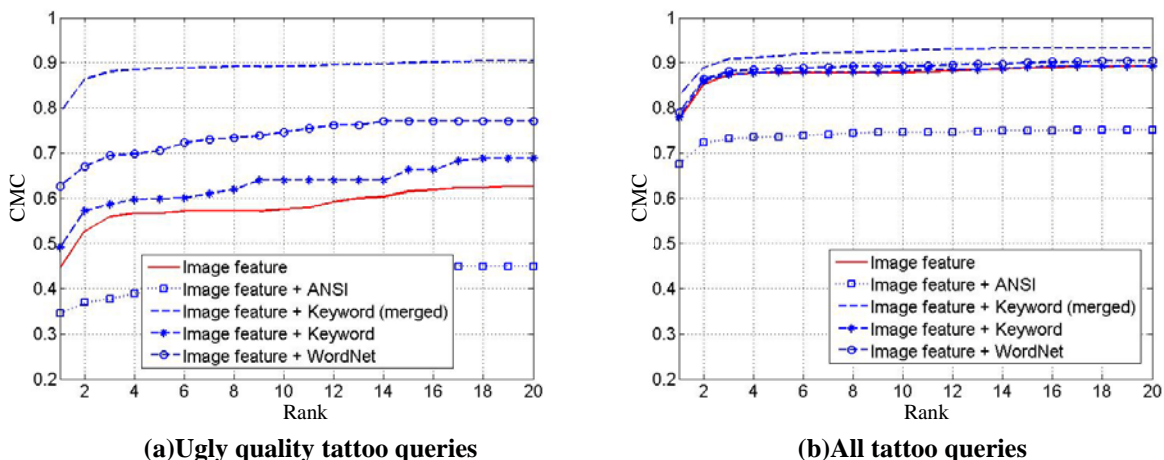


Figure 7. Retrieval performances for (a) 252 ugly quality queries and (b) all the 1,000 tattoo queries with/without metadata information against the database of 21,000 images.

4.2.1 Metadata Generated by Free Keyword Annotations

We treat the keyword annotations as free text and apply the standard text retrieval methods to compute the similarity score for metadata. More specifically, we use the *tf-idf* weighting scheme for text retrieval and the *Lemur* text search engine [16] to efficiently compute the matching scores between free keyword annotations. Given the similarity $S_W(I_q, I)$ based on the weighted keypoint matching, and the similarity $S_T(I_q, I)$ based on keyword matching, the combined similarity score is computed as $S(I_q, I) = S_W(I_q, I) + w \cdot S_T(I_q, I)$, where the weight parameter w is empirically tuned to optimize the retrieval performance.




Figure 8. Examples of inconsistent assignment of ANSI/NIST classes to near-duplicate tattoo pairs. While (a), (b), and (c) show near duplicate images of the same tattoo, they have been annotated differently by the subjects in our experiment based on ANSI/NIST classes (shown under each image).



Figure 9. Comparison of retrieval results with and without free keyword annotation. The first number under each image is the ranking position for the correct retrieval based on image feature alone and the second number (in parenthesis) is the ranking position for the correct retrieval based on image features together with merged free keywords.

The plot in Figure 7 labeled as *Image Feature + Keyword (merged)* shows that the retrieval results of combining the free-keyword-based matching with image matching. There is a significant improvement in retrieval performance for both ugly quality queries (~27%) and all the tattoo queries (~10%). This indicates that the free keyword annotation is much more effective than the ANSI/NIST classes for retrieving near duplicate tattoo images. This is because, unlike the classes in ANSI/NIST standard that are often ambiguous in terms of labeling tattoos, most human subjects appear to be consistent in choosing keywords for describing the similar visual content.

One potential problem with the above experiment is that the free keyword annotations for query images are created by the same subjects who created the annotations for the gallery images. In an operational system, we may expect different subjects to perform keyword annotation for query images than for gallery images, which could degrade the retrieval performance. In fact, for the 21,000 annotated tattoo images, we observe that, on average, less than 50% of the keywords are shared by two different subjects. To accommodate this scenario, we changed the design of the metadata experiment as follows: we used the free keyword annotations for query images by one subject, and the annotations for gallery images by a different subject. The retrieval results for ugly quality queries and all the 1,000 queries are shown in Figure 7 with the legend *Image Feature + Keyword*. It is not surprising that now there is a significant drop in retrieval accuracy compared to the case when both query images and gallery images are annotated by the same subjects. On the other hand, compared to using image features alone, we still observe a significant improvement (~7%) for ugly quality queries, and a marginal but consistent improvement (~1%) for all the 1,000 tattoo queries. Figure 9 shows examples of retrieval results based on combination of free keyword

annotations and image features, where the images in (a)-(c) are successful retrievals and images in (d)-(e) are failure cases.

An analysis of failure cases shows that subjects in our experiments assigned different free keywords to describe similar tattoos. For example, the image in Figure 9(d) was annotated as “face” and “skull” by two different subjects. To address this problem, we expanded the annotation keywords using *WordNet* [18]. The underlying assumption is that different keywords used to describe similar tattoo images are likely to share the same semantic concept, and as a result, the concept expansion from *WordNet* may be able to bridge this gap. *WordNet* is a large lexical database where nouns, verbs, adjectives, and adverbs are grouped into sets of cognitive synonyms, called *synsets*. *Synsets* interlink different conceptual-semantic and lexical relations. In our study, we use the *hypernym* hierarchy in *WordNet* for keyword expansion. A hyponym shares a type-of relationship with its *hypernym*. For example, the *hypernym* of “dog” is “canine”. We choose the *hypernym* relation because two words sharing the same concept are likely to share a common *hypernym*s in *WordNet*. Among the 2,019 different free keywords used by the subjects in annotating 21,000 tattoos, 1,737 keywords are found in *WordNet* and were expanded with the corresponding *hypernym* hierarchy.

The plot in Figure 7 labeled as *Image feature + WordNet* shows the retrieval results using *WordNet* expansion for both ugly quality queries and all the 1,000 queries. For both cases, we observe up to 8% improvement by using the *WordNet* expansion. The *WordNet* expansion clearly helps bridge the gap due to differences in free keyword annotations. For example, for the query tattoo in Figure 9(a), the correct retrieved image is found at rank 12 by fusion of the weighted keypoint matching and free keyword matching scores. By expanding the free keywords with *WordNet*, the correct retrieved image is found at rank 8 and the matching score is improved from 5 to 8.6. The *WordNet* expansion fails (see Figures 9(d), (e)) when the *gap* between free keyword annotations by different subjects is too large. For example, the keyword annotation for tattoo in Figure 9(e) is “Symbol” while the keyword annotation for its true mate image in the database is “Cross”.

5. Summary

The use of soft biometrics in forensics has been recognized as a valuable tool for solving crimes. We have focused on one such soft biometric, namely tattoo images, which are routinely collected by law enforcement agencies and used in apprehending criminals and identifying suspects. The current practice of matching and retrieval of tattoos is based on ANSI/NIST classes, and it is prone to significant errors due to limited vocabulary and subjective nature of labeling. To improve the performance and robustness of keyword-based tattoo matching, we introduced a content-based image retrieval (CBIR) system, called Tattoo-ID. It automatically extracts features from a query image and retrieves near-duplicate tattoo images from a database. We present two modifications to Tattoo-ID that further improve the retrieval accuracy, particularly for queries with low quality, called ugly tattoos. The modifications involve (i) robust similarity measure, and (ii) metadata utilization in the form of free keyword annotation in conjunction with *WordNet*. The best retrieval performance, as measured by top-20 retrieval on 1,000 tattoo queries and a database of 21,000 tattoos, is 94%. For the same 1,000 queries against a database of 100,000 images, the top-20 accuracy without the metadata is 90.5%. One limitation of the proposed algorithm is that it depends on manual annotations of tattoo images. We plan to overcome this limitation by exploiting supervised and semi-supervised learning algorithms to automatically annotate tattoo images with free keywords.

Acknowledgment

This research was partially supported by WCU (World Class University) program funded by the Ministry of Education, Science and Technology through the National Research Foundation of Korea (R31-10008). All correspondences should be directed to Anil K. Jain.

References

- [1] A. K. Jain, S. C. Dass, and K. Nandakumar, "Can soft biometric traits assist user recognition?," In Proc. SPIE Conf. on Biometric Technology for Human Identification, 2004
- [2] U.S. Department of Homeland Security, US-VISIT, <http://www.dhs.gov/files/programs/usv.shtm>.
- [3] Press Release. The Federal Bureau of Investigation, <http://www.fbi.gov/pressrel/pressrel08/ngicontract021208.htm>
- [4] J.-P. Beauthier, P. Lefevre, and E. D. Valck, "Autopsy and Identification Techniques", in Nils-Axel Mörner (Ed.), The Tsunami Threat-Research and Technology, InTech, 2011
- [5] ANSI/NIST-ITL 1-2011, "Data Format for the Interchange of Fingerprint, Facial, & Scar Mark & Tattoo (SMT)", http://www.nist.gov/itl/iad/ig/ansi_standard.cfm
- [6] Tattoo Facts and Statistics, http://www.vanishingtattoo.com/tattoo_facts.htm, Oct. 2006
- [7] J.-E. Lee, A. K. Jain, and R. Jin, "Scars, Marks and Tattoos (SMT): Soft Biometric for Suspect and Victim Identification", Proc. Biometric Symposium, Biometric Consortium Conference, 2008.
- [8] A. K. Jain, J.-E. Lee, R. Jin, and N. Gregg, "Content-based image retrieval: An application to tattoo images," In Proc. ICIP, 2009. pp. 2745-2748
- [9] J.-E. Lee, R. Jin and A. K. Jain, "Unsupervised Ensemble Ranking: Application to Large-Scale Image Retrieval," In Proc. ICPR, 2010. pp. 3902-3096
- [10] The CBS Interactive Business Network. "MorphoTrak acquires innovative tattoo matching technology from Michigan State University", http://findarticles.com/p/articles/mi_m0EIN/is_20100119/ai_n48674730/
- [11] S.T. Acton and A. Rossi, "Matching and retrieval of tattoo images: active contour CBIR and global image features," Proc. IEEE Southwest Symposium on Image Analysis and Interpretation, 2008.
- [12] R. Datta, D. Joshi, J. Li and J. Wang, "Image Retrieval: Ideas, Influences, and Trends of the New Age", ACM Computing Surveys, Vol. 40. pp. 1-60, 2008.
- [13] D. Lowe, "Distinctive image features from scale invariant keypoints," Int. J. Comp. Vision, Vol. 60. pp. 91-110, 1999.
- [14] K. Mikolajczyk and C. Schmid, "A Performance Evaluation of Local Descriptors", IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 27, pp. 1615-1630, 2005
- [15] ESP Game. <http://www.gwap.com/gwap/gamesPreview/espgame/>.
- [16] The Lemur Project, <http://www.lemurproject.org/>
- [17] George A. Miller, "WordNet: A Lexical Database for English", Comm. of ACM, Vol. 38, pp. 39-41, 1995.