



Faculty Publications

2007-10-07

Robust Multi-Modal Biometric Fusion via Multiple SVMs

Jonathan Dinerstein
jondinerstein@yahoo.com

Sabra Dinerstein
sabra.dinerstein@gmail.com

Dan A. Ventura
ventura@cs.byu.edu

Follow this and additional works at: <https://scholarsarchive.byu.edu/facpub>



Part of the [Computer Sciences Commons](#)

Original Publication Citation

Sabra Dinerstein, Jonathan Dinerstein and Dan Ventura, "Robust Multi-Modal Biometric Fusion via Multiple SVMs", Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, pp. 153-1535, 27.

BYU ScholarsArchive Citation

Dinerstein, Jonathan; Dinerstein, Sabra; and Ventura, Dan A., "Robust Multi-Modal Biometric Fusion via Multiple SVMs" (2007). *Faculty Publications*. 944.
<https://scholarsarchive.byu.edu/facpub/944>

This Peer-Reviewed Article is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Faculty Publications by an authorized administrator of BYU ScholarsArchive. For more information, please contact ellen_amatangelo@byu.edu.

Robust Multi-Modal Biometric Fusion via Multiple SVMs

Sabra Dinerstein, Jonathan Dinerstein, and Dan Ventura

Abstract—Existing learning-based multi-modal biometric fusion techniques typically employ a single static Support Vector Machine (SVM). This type of fusion improves the accuracy of biometric classification, but it also has serious limitations because it is based on the assumptions that the set of biometric classifiers to be fused is local, static, and complete. We present a novel multi-SVM approach to multi-modal biometric fusion that addresses the limitations of existing fusion techniques and show empirically that our approach retains good classification accuracy even when some of the biometric modalities are unavailable.

I. INTRODUCTION

MANY biometric modalities, including fingerprint and facial recognition, are used for verification and identification purposes. However, despite significant research, biometric matching accuracy remains low. This accuracy problem has recently been addressed through multi-modal biometric (multi-biometric) fusion, which combines the match scores that are output by individual biometric classifiers. Multi-modal biometric fusion has been shown empirically to improve the accuracy of biometrics-based verification (one-to-one comparison) and identification (one-to-many comparison) [17]. Further, parametric machine learning algorithms, including Support Vector Machines and Bayesian networks, have been shown to outperform both non-parametric learning techniques and voting schemes, when combining biometric match scores [17], [19].

Existing multi-biometric fusion techniques face a number of limitations since they are based on the assumptions that each biometric modality is local, complete, and static. These limitations are particularly pronounced when considered in the context of biometric identification, as opposed to verification. Key limitations include:

1. *Each registered person must be entered into every modality.* This may not be plausible and is very restrictive [12]. Moreover, this makes adding additional modalities to an existing system difficult or impossible.
2. *All of the classifiers must always be available.* This will not be the case if the modalities are part of a distributed system, such as when a multi-biometric system is composed of traditional biometric systems that are maintained by different groups or

organizations and are connected via the Internet.

3. *No support for “offline” biometrics.* “Offline” biometrics (such as DNA profiles) require laboratory processing to register individuals into the biometric system; the associated time and cost exacerbates limitations #1 and #2 listed above, and makes the utilization of offline biometrics impossible in existing biometric fusion systems [16].
4. *Registration changes may decrease system accuracy.* If learning is only performed when initially creating the multi-biometric system, the accuracy of the biometric fusion may degrade as individuals are later added to or removed from the system.
5. *Limited to verification.* Due to the other limitations listed above, most existing fusion techniques are explicitly designed for verification only – identification is not supported.

We propose a novel multi-biometric fusion technique that addresses the issues listed above and is suitable for both identification and verification. A mediator agent controls the fusion of the individual biometric match scores, using a “bank” of SVMs that cover all possible subsets of the biometric modalities being considered. This agent selects an appropriate SVM for fusion, based on which modality classifiers are currently available and have sensor data for the identity in question. (Our fusion technique differs from a traditional SVM ensemble – rather than combining the output of all of the SVMs [7], we apply only the SVM that best corresponds to the available modalities.) The mediator agent also controls the learning of new SVMs when modalities are added to the system or sufficient changes have been made to the data in existing modalities.

Our experiments utilize the following biometric modalities: face, fingerprint, and DNA profile data. We empirically show that our multiple SVM technique produces more accurate results than the traditional single SVM approach.

II. PREVIOUS WORK

Accurate user verification and/or identification are necessary for a wide variety of applications. Biometric classification is particularly interesting because biometric data is typically bound to a specific user (rather than being disembodied, such as in the case of a Social Security number), and is often unique enough to be used effectively in classification [11]. However, in practice, biometric classification using only a single biometric modality is typically not accurate enough [9], [16]. For example, facial recognition techniques are often sensitive to changes in

Sabra Dinerstein is with the Computer Science Department, Brigham Young University, Provo, UT 84602 USA (e-mail: sdinerstein@NPL.com).

Jonathan Dinerstein is with DreamWorks Animation, Redwood City, CA 94062 USA (e-mail: jondinerstein@yahoo.com).

Dan Ventura is with the Computer Science Department, Brigham Young University, Provo, UT 84602 USA (e-mail: ventura@cs.byu.edu).

lighting, camera angle, and distance from the camera. Additionally, not every user can provide data for the desired single biometric modality [12].

Fusion has been shown empirically to improve the accuracy of biometric classification and overcome the weakness of individual classifiers [8], [17], [13]. Additionally, in the case of a missing modality, a multi-modal biometric (multi-biometric) fusion system can still output a classification decision, by merely using one of the available modalities in a traditional manner [17]. Multi-biometric fusion is similar in spirit to bagging, stacking, and other techniques for combining complimentary classifiers. For example, in bagging, the output of two or more classifiers may be combined through voting, ideally to achieve more accurate classification results.

Multi-biometric fusion is typically applied in one of three specific steps in the classification process: fusion of the input feature vectors, fusion of the match scores output by the individual classifiers, or fusion at the decision level [17]. Fusion of the input feature vectors is not always feasible, as these input features may not be directly accessible via professional biometric collection systems. Additionally, the specific input features that are used by the different modalities may not be compatible [12]. Fusion at the match score level is applicable to general multi-biometric systems, and has been shown to be more informative than decision-level fusion [17], [18].

Voting-based fusion improves the results of using only individual biometric classifiers, and provides a simple, understandable fusion technique. The addition of quality-based weighting has further improved the results of standard fusion techniques, providing quantification of both the quality of the biometric data itself [15] and of the accuracy of the specific biometric classifier [3].

Supervised parametric learning techniques, such as SVMs and Bayesian networks, have been shown empirically to produce more accurate fusion results than either voting or non-parametric learning [2], [3], [13], [17], [19]. Of these parametric learning-based fusion techniques, the SVM appears to be the most popular choice in the literature. In short, the SVM learns to map the vector of individual biometric match scores into a joint (i.e., fused) match score or classification. Current techniques learn a single, static SVM immediately before the multi-biometric system is made available for use [18], [16]. These techniques make an implicit assumption that all modalities in the system are always available; otherwise the fusion breaks down or the accuracy degrades. Thus they also require that users provide biometric data for every modality (missing data is analogous to an unavailable modality). Additionally, the learned SVM is specific to the current biometric data that is enrolled in the system – if the data changes too much, or if we wish to add another modality to the multi-biometric system, the fusion has to be completely re-learned. Adding a new modality is further complicated by the need to gather data for this modality from all individuals previously registered in the

system.

We present a multi-SVM fusion technique that addresses these limitations and improves upon the accuracy of the single static SVM when there are missing biometric modalities. We allow for the possibility of missing biometric modalities (due either to the non-universality of biometric data [12] or to the temporary unavailability of individual biometric classifiers) by learning multiple SVMs that are trained on all possible subsets of the biometric modalities. We demonstrate our multi-SVM fusion technique on an offline biometric modality, nuclear DNA, as well as the more traditional biometric modalities: face and fingerprint.

III. BIOMETRIC CLASSIFIERS

To test our multi-SVM fusion technique, we have implemented biometric classifiers using professional biometric SDKs. We use one classifier for each of the following biometric modalities: face, fingerprint, and nuclear DNA.

A. Face Classifier

Our 2D facial recognition classifier utilizes AcSys Biometric's Face Recognition System (FRS) SDK. This system employs a neural network-based implementation, Holographic Neural Technology (HNeT) [1], which uses machine learning to improve recognition accuracy over time, as it is presented with multiple images of the same subject.

During initialization, we presented our face classifier with static images of multiple subjects. Several images were presented for each subject, and the pictures were taken from a variety of angles and camera distances. For each subject, the best image (based on head size and the system's image quality metrics) was enrolled into the database. The face classifier was then trained on all of the images in the database.

The use of multiple images (both of the same subject and of different subjects) resulted in a wide range of match scores, for both positive and negative examples. In our experiments, this facial recognition system produced extremely variable results (e.g., positive match scores in the range [-0.702, 0.868] and negative match scores scattered throughout the range [-1.167, -0.043]), which allowed us to experiment with a less than perfect biometric classifier. It should be noted that AcSys Biometric's Face Recognition System is reported to be more accurate when using a live video stream as opposed to static images, due to their neural network-based implementation, but this reduced accuracy provided interesting information for our experiments.

B. Fingerprint Classifier

We implemented a fingerprint classifier using the Identix BioEngine® SDK, which provides a minutiae-based fingerprint verification algorithm. In short, this system operates by extracting minutiae from the fingerprint ridges (such as the locations of ridge endings and bifurcations). The match score between two fingerprints is calculated by comparing these minutiae. For further information on minutiae-based fingerprint matching, see [10] and [14].

We created examples of fingerprint match scores using both live scans and static images. We used multiple scans of the same fingerprint, in addition to fingerprints from different subjects, in order to obtain a variety of match scores, including positive match scores in the range [65, 8364] and negative match scores in the range [-1, 132].

Fingerprint matching is known to be a relatively accurate biometric, even with only partial fingerprint data [9], [14]. Also, fingerprint acquisition hardware is quite affordable. Thus fingerprint matching is an excellent modality to include in any multi-biometric fusion system.

C. DNA Classifier

Our DNA profile examples are based on the United States Federal Bureau of Investigation's 13 core Short Tandem Repeat (STR) loci [4]. This is the standard used in the FBI's Combined DNA Index System (CODIS) [6]. This standard is both important and pertinent because it is admissible as identifying evidence in the legal courts of the United States and various other countries [4].

Each DNA profile is represented by a string made up of the characters, {A, T, G, C}: the profile string describes the allele values of the person's STR DNA for the 13 loci of interest. The profile string for a given person is derived through a laboratory typing process [4]. Our DNA classifier uses the Levenshtein distance metric to calculate the match score of two profile strings; these DNA profile match scores are in the range [0, 1], where a score of 1 represents a perfect match.

DNA profile data represents an extremely robust and information-dense biometric modality. However, due to the time and cost of the offline processing requirements of DNA classification, DNA has not previously been tested in biometric fusion research [16]. Thus its inclusion is an interesting aspect of our work.

IV. MULTI-SVM FUSION

Our technique centers around the use of multiple specialized SVMs that are learned by a fusion agent. As discussed earlier, previous biometric fusion techniques utilize only a single SVM, resulting in the limitations stated in the introduction. In contrast, our fusion agent learns multiple SVMs. This overcomes the limitations of previous techniques by allowing the agent to perform effective fusion even when every modality is not currently available. We describe our technique in detail below.

A. Learning Multiple SVMs

We denote the set of available biometric modalities as:

$$S = \{face, fingerprint, DNA\}.$$

Of course, the elements of S correspond to the biometric modalities chosen for inclusion in the specific system – we list the modalities that we employ in our experiments for the purpose of clarity.

Our fusion agent learns and utilizes multiple SVMs – one SVM for each possible subset of S that contains 2 or more

elements. Note that this is simply the power set of the available modalities minus those sets of cardinality < 2 (in which case no fusion can be performed). This reduced power set for the given set S is:

$$S^* = \{\{face, fingerprint\}, \{face, DNA\}, \{fingerprint, DNA\}, \{face, fingerprint, DNA\}\}. \quad (1)$$

One SVM is learned for each set in S^* . Thus the fusion agent learns $2^{|S|} - (|S| + 1)$ total SVMs, where each SVM learns to fuse a specific, unique set of biometric modalities.

We utilize LIBSVM [5] for our implementation of these SVMs. Specifically, we use a Radial Basis Function (RBF) kernel: $\exp(-\gamma|u-v|^2)$. For each SVM, we choose the appropriate γ -value and constraints-violation cost, C , at run-time, by performing k -fold (stratified) cross-validation on the current set of training examples. The γ - and C -values that produce the best accuracy percentage in cross-validation are then used to train the SVM on the entire set of training examples. In our current implementation, each SVM learns to output a classification decision, rather than a specific fused score.

B. Fusion

Biometric data is collected and processed by the individual classifiers. Each classifier outputs a match score for their specific biometric modality. Upon receiving match scores from the participating individual biometric classifiers, the fusion agent creates an attribute vector out of these individual scores, and applies the learned SVM that best corresponds to the incoming data. (This fusion process differs from a traditional SVM ensemble – rather than combining the output of all of the SVMs [7], we apply only the SVM that best corresponds to the input modalities.) The selected SVM outputs a single classification decision for the joint attribute vector. See Fig. 1 for a conceptual description of this fusion process.

The fusion agent selects an appropriate SVM based on the operational status and data completeness of each biometric classifier. Two conditions are necessary for a classifier to be included in the fusion process: the classifier must produce a match score for the identity in question, and the classifier must report that match score to the fusion agent. Currently, our implementation does not employ thresholding at the fusion level. However, some of the individual classifiers featured in our system perform thresholding on their reported matches: our fingerprint classifier returns -1 (indicating a completely negative classification) for matches that are below a certain score, and our face classifier allows the user to set the minimum score threshold at run-time.

V. EXPERIMENTS AND RESULTS

A. Data Preprocessing

Our individual biometric classifiers, as described above, do not share a common scale for match scores. Instead, each




Fig. 1: Multi-SVM Fusion. Biometric data is collected from the user, and the individual classifiers are applied; these individual match scores are sent to the fusion agent. The fusion agent selects the appropriate fusion SVM, based on the available modality match scores. The system outputs one overall classification decision.

classifier outputs scores that correspond to their own scale. Specifically, our fingerprint classifier tends to output very large scores (e.g., up to about 8000 for a good match), while our DNA classifier outputs scores in the range of $[0, 1]$. To remove the bias of the large fingerprint match scores, we scale the output of each classifier to be in the range of $[-1, 1]$, where a value of 1 represents a perfect match.

We create both positive and negative multi-biometric examples by combining the scaled match scores of the individual biometric classifiers into one attribute vector with a corresponding label. In short, single-modality examples of the same class (either positive or negative) are combined, in order to create the multi-biometric examples. For example, for the SVM that fuses $\{face, fingerprint\}$, we create examples that contain both the face and the fingerprint match scores:

- <Classification of the multi-biometric example>
- <Scaled match score for the 1st modality (face)>
- <Scaled match score for the 2nd modality (fingerprint)>

Training examples are drawn randomly with replacement, using a uniform distribution, from this set of multi-biometric examples.

B. Comparison of the traditional single SVM with our multi-SVM technique

Let us compare the multi-biometric fusion accuracy obtained when using a single static SVM to the accuracy of our multi-SVM approach. We calculate the fusion accuracy of each SVM by comparing the actual classification value of

each example to its target classification value, and determining the percentage of the examples that are correctly classified. We estimate the test accuracy of each SVM by performing k -fold cross-validation on the training examples (using the γ - and C -values that produced the highest accuracy, as described in the previous section); the highest accuracy achieved by the k -fold cross-validation is reported as the estimated test accuracy of the SVM.

We performed these fusion experiments with several different training set sizes, letting the number of training examples vary from 25 to 4500, and found that the number of training examples did not have a significant impact on the accuracy levels of the fusion algorithms, as long as the examples were sufficiently random. We performed 10 runs for each training example set size, to account for the variability in the selection of the training examples; accuracy values were averaged over all runs.

1) Fusion accuracy when there are no missing biometric modalities

Table 1 describes the accuracy of each fusion SVM when the corresponding biometric modalities are all available. Each SVM was trained and tested on examples that included exactly the specified set of biometric modalities. For example, the SVM that performs fusion on the set of modalities, $\{face, fingerprint\}$, was trained and tested on a set of examples whose feature vectors contained match scores for both the face and fingerprint modalities (and only these modalities). The other SVMs shown in Table 1 were trained and tested in a similar manner.

As expected, all of the fusion SVMs in our system produced reasonable average accuracy, as described in Table 1, verifying that the use of multiple SVMs does not reduce the accuracy of multi-modal fusion. Notice that the fusion SVM for $\{face, fingerprint\}$ has the lowest average accuracy of all of the SVMs in the system. This can be attributed to the extremely variable output of our face classifier. Also notice that in our experiments, the fusion SVMs that incorporate DNA profile data seem to be very accurate.

Table 1: Accuracy of each SVM, when all of the corresponding modalities are present. Each SVM was trained and tested on the specified set of biometric modalities. The values shown represent each SVM's average fusion accuracy over 10 runs, and the corresponding standard deviation. These accuracy values correspond to the use of 500 training examples.

	Accuracy % (No Missing Modalities)	Standard Deviation
Fusion SVM for $\{face, fingerprint\}$	97.558	0.3135
Fusion SVM for $\{face, DNA\}$	100.0	0
Fusion SVM for $\{fingerprint, DNA\}$	100.0	0
Fusion SVM for $\{face, fingerprint, DNA\}$	99.492	0.0821

2) Fusion accuracy with missing biometric modalities

Next, we compare our multi-SVM technique to previous fusion techniques that use only a single static SVM. Just as before (when there were no missing modalities), each SVM is trained using data for only the corresponding set of biometric modalities.

In our experiments, we let the fusion SVM for all of our modalities, $\{face, fingerprint, DNA\}$, represent the single static SVM approach to fusion. We simulate the single SVM approach to fusion (with missing biometric modalities) by applying the $\{face, fingerprint, DNA\}$ SVM to examples that contain subsets of the available modalities. We let the absence of a match score represent a missing modality: LIBSVM [5] implements each attribute in the feature vector as an index and value pair. Our implementation uses this index to indicate which modality the current attribute represents. For missing modalities, the corresponding index and value are simply not included in the feature vector.

In our multi-SVM approach, the fusion agent selects and applies the fusion SVM that best corresponds to the current subset of available modalities, and therefore the accuracy of each SVM in our multi-SVM technique is not affected by the missing modalities. Instead, we again estimate the test accuracy of each SVM in our multi-SVM approach, by performing k -fold cross-validation on the training examples containing the specific subset of modalities.

Average accuracy percentages (over 10 runs) for both the single SVM and multi-SVM techniques are shown in Fig. 2. As can be seen, the traditional single static SVM approach to multi-biometric fusion is highly sensitive to missing modalities. (The missing modality in each case is noted along the x -axis in Fig. 2.) The fusion accuracy of the single SVM has decreased noticeably, for each of the missing modalities. Our multi-SVM technique, on the other hand, retains high average accuracy despite the missing biometric modalities; our multi-SVM technique appears to be robust against missing biometric modalities.

Consider the massive loss of accuracy for the single SVM when fingerprint is the missing modality, as shown in Fig. 2. This severe decrease in accuracy can partially be accounted for when we consider the relative strength of the individual biometric classifiers used in our experiments. For example, we see the smallest loss of accuracy when we ignore the face modality (shown in the 3rd column of Fig. 2), because our facial recognition classifier is the least accurate of our individual classifiers. Therefore, the lack of match score data for the face modality causes the least perturbation to the classification decision that is output by the single static SVM.

Our fingerprint classifier is much more accurate than our face classifier, and in fact, appears to be the most important of our individual classifiers. Therefore, without the fingerprint data (shown in the 2nd column of Fig.2), we see the largest difference in accuracy between our multi-SVM technique and the traditional single static SVM.

Our DNA classifier is accurate, but it tends to produce match scores for positive examples that are at the very top of

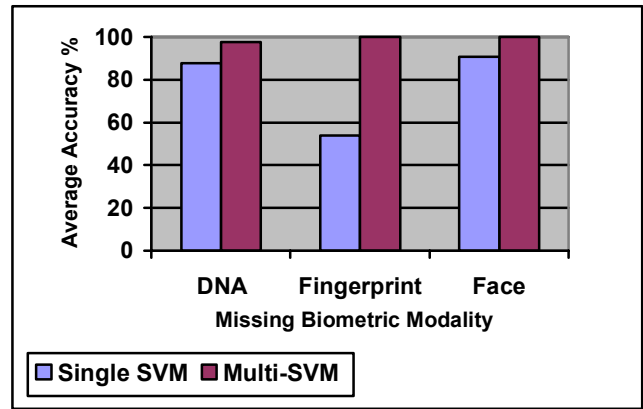


Fig. 2: Accuracy of the single, static SVM fusion technique vs. our multi-SVM technique, with missing biometric modalities. In our multi-SVM approach, the fusion agent applies the SVM that best corresponds to the available biometric modalities. As can be seen here, our multi-SVM technique produces higher average accuracy (over 10 runs, using 500 training examples) than the single static SVM, when there are missing biometric modalities. In each case, the missing biometric modality is noted along the x -axis. (These accuracy differences between the single, static SVM approach and our multi-SVM approach were determined to be statistically significant at $p < 0.001$ using a paired permutation test.)

the scaled range $[-1, 1]$, and therefore the DNA match score examples are less informative than those produced by our fingerprint classifier.

Our experimental data clearly shows that our multi-SVM fusion technique retains distinctly higher average accuracy than the single SVM fusion technique, when there are missing biometrics.

VI. CONCLUSION

Existing learning-based multi-biometric fusion techniques utilize only a single static SVM that is dependent upon both the currently enrolled biometric data and the modalities that are currently in use. This static SVM approach to fusion improves biometric matching accuracy [17], but degrades when faced with missing biometric data. Biometric modalities are known to be non-universal [12], and therefore we would like a multi-biometric fusion system to be robust against missing biometrics.

We now consider our multi-SVM approach to fusion, in the context of the limitations of the traditional single SVM approach:

1. *Each registered person must be entered into every modality.* As shown in Fig. 2, our multi-SVM fusion technique remains highly accurate, even when some of the biometric modalities are missing. If a registered person is missing any of the biometric modalities, our multi-SVM technique can still take advantage of the increased accuracy provided by multi-modal fusion (if that person supplies data for at least two modalities): our multi-SVM technique simply fuses the biometric data that is currently available. Further, this allows new biometric modalities to be added to the system, without affecting the persons that are already

registered with the system – fusion can still be performed, even without collecting data for the new modality.

2. *All of the classifiers must always be available.* Instead of being dependent upon the availability of all modalities, our multi-SVM technique takes advantage of whatever data is currently available. If an individual classifier is unavailable, its output is simply not used in the fusion. Just as with a missing modality, if the classifier is not available, our multi-SVM technique is still able to perform fusion with whatever classifiers are currently available.
3. *No support for “offline” biometrics.* Our multi-SVM technique allows for the incorporation of offline biometrics, such as DNA, that have previously been excluded from multi-biometric fusion systems [16]. The inclusion of DNA profile data in our experiments implies that the use of multiple specialized SVMs allows our multi-modal fusion system to utilize whatever data is available at the moment, rather than requiring that all of the biometric data be collected and used at the same time.
4. *Registration changes may decrease system accuracy.* In traditional static SVM implementations of biometric fusion, the single SVM is only effective as long as nothing has changed – if any of the biometric modalities are replaced or if new modalities are added, the entire learned fusion system must be replaced. In our implementation, however, much of the system can be re-used: only those SVMs that are directly affected by a modality change need to be replaced. Further, the addition of a new modality does not affect the existing SVMs. Instead, the fusion agent simply trains additional SVMs to handle the new modality combinations, as described in (1). Our implementation therefore provides the flexibility to easily add or modify biometric modalities as needed.
5. *Limited to verification.* Current learning-based fusion techniques are typically limited to verification, rather than identification. Biometric verification often assumes that all of the biometric data has been collected at the same time (typically using multiple sensors) and fed into the system immediately. Biometric identification, on the other hand, is well suited to a distributed implementation – large repositories of biometric data, such as the CODIS and Integrated Automated Fingerprint Identification System (IAFIS) databases, are typically not hosted in a single location. Combining biometric match information from multiple sources should only increase the odds of successful identification. Identification, therefore, can benefit from the use of delayed information, not just what is known at the moment. Our experiments suggest that our multi-SVM fusion technique retains high accuracy regardless of which biometric modalities are available, and

therefore our technique should be useful for biometric identification, as well as for verification.

REFERENCES

- [1] AcSys Biometric's Face Recognition System: http://www.acsysbiometrics.com/product_sdk.html
- [2] S. Ben-Yacoub, Y. Abdeljaoued, and E. Mayoraz, “Fusion of face and speech data for person identity verification”, *Technical report, IDIAP Research Institute*, number IDIAP-RR 99-03, 1999.
- [3] J. Bigun, J. Fierrez-Aguilar, J. Ortega-Garcia, and J. Gonzalez-Rodriguez, “Multimodal biometric authentication using quality signals in mobile communications”, In *Proceedings of International Conference on Image Analysis and Processing (ICIAP'03)*, 2003.
- [4] J. Butler, “Genetics and Genomics of Core Short Tandem Repeat Loci Used in Human Identity Testing”, *Journal of Forensic Science*, vol. 51, no. 2, 2006.
- [5] C. Chang and C. Lin, “LIBSVM: a library for support vector machines, 2001”. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [6] FBI's Combined DNA Index System (CODIS) Homepage: <http://www.fbi.gov/hq/lab/codis/index1.htm>
- [7] T.G. Dietterich, “Machine learning research: Four current directions”, *AI Magazine*, 118(4), pp. 97-136, 1997.
- [8] L. Hong, A. Jain, and S. Pankanti, “Can multibiometrics improve performance?” In *Proceedings of AutoID'99*, pp. 59-64, 1999.
- [9] L. Hong and A. Jain, “Integrating faces and fingerprints for personal identification”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no 12, pp. 1295-1307, 1998.
- [10] A. Jain, L. Hong, R. Bolle, “On-Line Fingerprint Verification”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, issue 4, pp. 302 - 314, 1997.
- [11] A. Jain, L. Hong, S. Pankanti, “Biometric identification”, *Communications of the ACM*, vol. 43, issue 2, pp. 90 - 98, 2000.
- [12] A. Jain and A. Ross, “Multibiometric systems”, *Communications of the ACM*, vol. 47, issue 1, pp. 34 - 40, 2004.
- [13] S.Y. Kung and M.W. Mak, “On Consistent Fusion of Multimodal Biometrics”, *ICASSP'06*, 2006, pp. 1085-1088.
- [14] D. Maio and D. Maltoni, “Direct Gray-Scale Minutiae Detection in Fingerprints”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, issue 1, p. 27 - 40, 1997.
- [15] K. Nandakumar, Y. Chen, A. Jain, S. Dass, “Quality-based Score Level Fusion in Multibiometric Systems”, In *Proceedings of the 18th International Conference on Pattern Recognition (ICPR'06)*, vol. 04, pp. 473-476, 2006.
- [16] J. Ortega-Garcia, J. Bigun, D. Reynolds, and J. Gonzalez-Rodriguez, “Authentication gets personal with biometrics”, *IEEE Signal Processing Magazine*, pp. 50-62, March 2004.
- [17] A. Ross and A. Jain, “Information fusion in biometrics”, *Pattern Recognition Letters*, vol. 24, pp. 2115-2125, 2003.
- [18] C. Sanderson and K. K. Paliwal, “Identity verification using speech and face information”, *Digital Signal Processing*, vol. 14, pp. 449-480, 2004.
- [19] P. Verlinde, G. Chollet, and M. Acheroy. “Multi-modal identity verification using expert fusion”, *Information Fusion*, vol. 1, no. 1, pp. 17-33, 2000.