ELEKTROTEHNIŠKI VESTNIK 81(3): 74–80, 2014
ORIGINAL SCIENTIFIC PAPER

A case study on multi-modal biometrics in the cloud

Ž. Emeršič[1], J. Bule[1], J. Žganec-Gros[2], V. Štruc[3] and P. Peer[1]

[1]*University of Ljubljana, Faculty of Computer and Information Science, Tržaška 25, 1000 Ljubljana, Slovenia*
[2]*Alpineon d.o.o., Ulica Iga Grudna 15, SI-1000 Ljubljana, Slovenia*
[3]*University of Ljubljana, Faculty of Electrical Engineering, Tržaška cesta 25, 1000 Ljubljana, Slovenia*

Abstract. Cloud computing is particularly interesting for the area of biometric recognition, where scalability, availability and accessibility are important aspects. In this paper we try to evaluate different strategies for combining existing uni-modal (cloud-based) biometric experts into a multi-biometric cloud-service. We analyze several fusion strategies from the perspective of performance gains, training complexity and resource consumption and discuss the results of our analysis. The experimental evaluation is conducted based on two biometric cloud-services developed in the scope of the Competence Centere CLASS, a face recognition service and a fingerprint recognition service, which are also briefly described in the paper. The presented results are important to researchers and developers working in the area of biometric services for the cloud looking for easy solutions for improving the quality of their services.

Keywords: Multi-modal biometrics, biometric-cloud services, fusion strategies, cloud computing

Multi-modalna avtentikacija uporabnikov v oblaku

Storitve biometričnega razpoznavanja oseb v oblaku odpirajo nove možnosti uporabe biometričnih tehnologij in hkrati zagotavljajo tudi vrsto dobrodošlih karakteristik. Razširljivost oblačnih storitev omogoča sprotno prilagajanje obsegu njihove uporabe, zanseljivost storitev se zaradi boljše razpoložljivosti in dostopnosti izboljša, hkrati pa se zagotovi tudi možnost uporabe s širokim spektrom naprav in uporabnikov. V članku se posvetimo problemu združevanja eno-modalnih biometrijskih storitev v več-modalno biometrično oblačno storitev, ki uporabniku zagotavlja višji nivo varnosti. V ta namen ovrednotimo različne strategije k fuziji eno-modalnih biometričnih sistemov in jih analiziramo z vidika učinkovitosti, kompleksnosti učenja in porabe virov. V članku predstavimo rezultate naše analize in jih ustrezno komentiramo. Analiza je opravljena na podlagi dveh samostojnih biometričnih storitev v oblaku - sistema za razpoznavanje obrazov in sistema za razpoznavanje prstnih odtisov.

1 INTRODUCTION

Biometric technology, capable of recognizing people based on their physiological and/or behavioral traits, is nowadays being used increasingly and deployed ever more widely. This development raises issues related to the accessibility and scalability of the existing biometric technology. To address these issues a scalable technology capable of operating on large amounts of data with sufficient storage and processing power, needs to be developed. A straight forward solution to the presented problem is the implementation of biometric technology for the cloud, where the cloud platform ensures appropriate scalability, sufficient amount of storage, and parallel processing capabilities. This solution also enables wide support for various devices [1], where the computationally heavy part is processed in the cloud, and only simple computations are left for a client device to process. The aforementioned scalability enables timely processing even when the number of client devices and verification requests change heavily and rapidly.

Another aspect of a successful biometric technology is its verification performance. In terms of performance, biometric techniques relying on a single biometric trait can only be improved to a certain extent whether they are cloud-based or not. Any additional improvements of uni-modal biometric systems can be either too costly or not even possible with the given state of technology. In such cases the use of multi-modal biometric systems may represent the solution in which the performance may be improved with no (or almost no) additional costs. Multi-modality in general is the ability of the system to base its output on more than one input modality. In the case of multi-modal biometric systems this implies that several biometric features are used for the process of verification.

In this paper, we address the problem of building cloud-based multi-biometrics system from existing uni-modal biometrics systems. This is important when building new services on top of existing commercial services, where the integration of several biometric traits ensures added value for the developed service. In this paper, we rely on implementations of face and fingerprint recognition systems developed in the scope of the KC

Received 16 December 2013
Accepted 9 May 2014

Table 1. Overview of typical fusion strategies.

Fusion level	Short description	Representative references
The signal or sensor level	Several impressions of the same biometric trait are captured and then an enhanced composite biometric sample that is better suited for recognition is created	[2], [3], [4], [5]
The feature level	Involves combining evidence of several biometric feature vectors of the same individual obtained from several information sources	[2], [6], [7]
The matching score level	Enables easy combination of matching scores of different experts and is, therefore, the most commonly used approach in multi-biometric systems	[2], [8], [9], [10], [11], [12], [13]
The decision level	Is sensible when the uni-modal experts provide access only to the final classification result [2]. Different techniques can be considered at this level, e.g., the AND- and OR-rules, majority voting, weighted majority voting and others	[2], [14], [15], [16]

CLASS project and assume that we do not have access to the implementations of the recognition services as this would be the case with commercial services as well. Thus, we assume that no modifications can be conducted on the existing experts and that we can rely only on data available as outputs from the two biometric experts. We study and analyze different combination (i.e., fusion) strategies and present our findings.

The rest of the paper is structured as follows. In Section 2 we briefly review the field of biometric fusion and present the basics of the uni-modal biometric experts used in this case study. In Section 3 we introduce the fusion strategies considered in the paper and present their assessment in Section 4. We conclude the paper with some final comments in Section 5.

2 Background

2.1 Biometric fusion

The problem of biometric fusion has been widely discussed and analyzed in the scientific literature, including: [2], [13], [17], [4], [18], [19], [14]. In general, the process of fusion can be conducted at different levels of biometric verification, as shown in Table 1.

In this case study we assume that we do not have access to the implementations of the biometric services and that therefore not all fusion strategies are feasible. Typically, existing (commercial) biometric services only provide APIs for assessing either the verification result (accept or reject) or the matching score that measures the similarity between the input biometric sample and the template of the claimed identity. This means that only strategies from the last two levels of Table 1 can be considered for buiding a multi-modal biometric service from existing uni-modal ones.

2.2 The face expert

The following is a brief summary of the face expert used in the uni-modal face recognition cloud service (see Fig. 1 (a)):

- **Face detection, localization and pre-processing.** Facial detection is conducted based on the Viola-Jones object detector [20] and face landmark localization is done with PSEF correlation filters. This is followed by geometrical normalization, re-scaling of the facial crops to a fixed size of 128×128 pixels, and photometric normalization technique using the Tan and Triggs technique [21].
- **Feature representation.** Gabor magnitude features and LBP histograms are used as the feature representation for the face expert. To improve recognition performance, a vector of the first few DCT coefficients of the normalized facial image is also used in addition to the Gabor-LBP features.
- **Verification.** Verification is based on simple similarity measurements using the Bhattacharyya distance for the Gabor-LBP histogram sequence and the whitened cosine similarity measure for the DCT coefficients. The two similarity scores are then stacked together with various image-quality measures (see [22] for details - Q-stack) and the newly combined feature vector is ultimately subjected to an AdaBoost classifier to obtain the final matching score based on which verification is conducted.

2.3 The fingerprint expert

The following is a brief summary of the fingerprint expert used in the uni-modal fingerprint recognition cloud service (see Fig. 1 (b)) [23]:

- **Segmentation and image enhancement.** Fingerprint scans are first segmented to separate the fingerprint pattern from the background. The resulting fingerprint patterns are then enhanced through

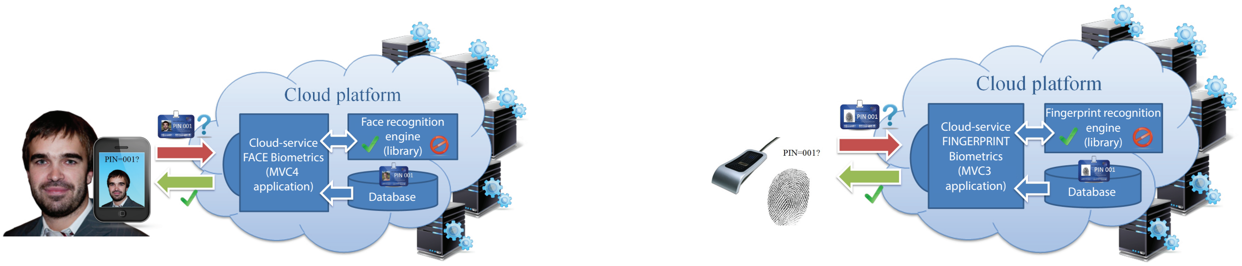
PIN=001?
PIN 001
Cloud platform
Cloud-service FACE Biometrics (MVC4 application)
Face recognition engine (library)
Database
PIN=001?
PIN 001
Cloud platform
Cloud-service FINGERPRINT Biometrics (MVC3 application)
Fingerprint recognition engine (library)
Database

(a) Cloud-implementation of face recognition service

(b) Cloud-implementation of fingerprint recognition service

Figure 1. Illustration of basic architecture of the biometric cloud-services

binarization and ridge profiling [24].

- **Minutiae extraction.** The minutiae pattern is extracted from the profiled binary image by thinning of the ridge structures, removal of structure imperfections from the thinned image, and the final process of minutiae extraction. For each detected minutia its type (bifurcation or ending), spatial coordinates (x, y) and the orientation of the ridge containing the minutia are stored as the templates for each given identity [24].
- **Matching and verification.** A minutiae-based matching algorithm is used in this stage to compare the template computed from the query fingerprint sample and the template corresponding to the claimed identity. Two fingerprints are declared a match when a sufficient number of minutiae match.

2.4 *Cloud-based biometric recognition*

For the aims of this case study, we build a multimodal cloud-based biometric service from the two experts presented in the previous two sections. Here, we perform most of the processing in the cloud exploiting the existing functionality of the two services, while the biometric fusion is implemented on the client side. Face and fingerprint images are first acquired via a camera and fingerprint scanner, respectively. The images are then sent from the client through a REST API to the two biometric cloud services. Both services consist of a middle-layer application, a biometric recognition engine and a database (see Fig. 1). In the case of face recognition, the middle-layer application forwards the received face image to the face recognition engine. In case of fingerprint recognition process, fingerprint images are forwarded from the middle-layer to the fingerprint recognition engine. Both engines then compute the verification results based on comparisons with template data stored in databases in the cloud and return the results to the client, where biometric fusion is ultimately performed.

3 FUSION STRATEGIES

As emphasized in one of the previous sections, the existing two biometric experts are capable of returning either the matching score of the given verification attempt or the verification decision (i.e., legitimate or illegitimate verification attempt). Fusion strategies at two different levels can, therefore, be considered for combining the two cloud-based experts, namely, the matching score and the decision level. In the remainder, we use the following notation to formalize the fusion strategies: we use w_1 to denote the class of legitimate verification attempts, w_2 to denote illegitimate verification attempt and δ to denote a similarity score.

3.1 *Decision-level fusion rules*

At the decision the biometric experts are first queried for the classification result $w_k^{(j)}$ (for $j = 1, 2, ..., J$ and $k \in \{1, 2\}$) and the results are then combined to arrive at the combined decision:

$$\psi : \{w_k^{(1)}, w_k^{(2)}, ..., w_k^{(J)}\} \rightarrow w_k^{fused}, \text{ where } k \in \{1, 2\} \tag{1}$$

where $w_k^{(j)}$ denotes the classification result of the $j-$th expert, w_k^f is the combined classification result and $k \in \{1, 2\}$.

For this case study we consider two of the most common options for choosing the fusion function ψ, namely, the AND- and OR-rules [16]. In the context of the biometric experts at our disposal the two rules, which assume that the class labels w_1 and w_2 are binary encoded, i.e., $w_1 = 1$ and $w_2 = 0$, are defined as:

$$\psi_{AND}(w_k^{(1)}, w_k^{(2)}) = w_k^f = w_k^{(1)} \wedge w_k^{(2)}, \text{ and} \tag{2}$$

$$\psi_{OR}(w_k^{(1)}, w_k^{(2)}) = w_k^{fused} = w_k^{(1)} \vee w_k^{(2)}, \tag{3}$$

where the indices $^{(1)}$ and $^{(2)}$ denote for the face and fingerprint experts, respectively, and $k \in \{1, 2\}$.

The decision level fusion strategies are in general easy to implement and provide a straightforward way of combining independent biometric experts. However, client applications exploiting these strategies cannot freely choosing the operating point of the multi-modal biometric system, which is imposed in a sense by the operating points of the uni-modal experts.

3.2 *Matching-score-level fusion rules*

Cloud-based biometric services can typically be queried for a similarity/matching score rather than the verification decision. The client application can then implement the verification procedure using a desired value for the decision threshold θ. Such an operating mode is common in most biometric services and gives the client applications the option of choosing their own operating points.

This mode of operation gives raise to the second possibility for combining the two uni-modal biometric experts, namely, strategies at the matching score level. The general form for combining experts at this level can be written as follows:

$$\phi : \{\delta^{(1)}, \delta^{(2)}, ..., \delta^{(J)}\} \rightarrow \delta^f, \qquad (4)$$

where ϕ is the fusion function and $\delta^f \in \mathbb{R}$ stands for the combined similarity score that can be used for verification. It is important to stress that the decision threshold for the combined scores needs to be recalculated for all operating points.

For our analysis presented in the next section, we implemented two fixed matching-level fusion rules: the weighted sum-rule and the weighted product-rule, they are defined as follows:

$$\phi_{SUM}(\delta^{(1)}, \delta^{(2)}) = \delta^f = \tau\delta^{(1)} + (1-\tau)\delta^{(2)}, \qquad (5)$$

$$\phi_{PRO}(\delta^{(1)}, \delta^{(2)}) = \delta^f = (\delta^{(1)})^{\tau} \cdot (\delta^{(2)})^{(1-\tau)}, \qquad (6)$$

where $\tau \in [0, 1] \in \mathbb{R}$ and stands for the weighting factor balancing the relative importance of the face and fingerprint scores.

4 Experiments

4.1 *Database and experimental protocol*

For our experiments we constructed a bimodal chimeric database from the XM2VTS and FVC2002 databases [25], [26]. A chimeric database represents a database, in which biometric modalities from different databases are combined and assigned common identities. Since the biometric samples in the initial databases are not taken from the same identities this procedure creates artificial (chimeric) subjects. Note that such a procedure is reasonable due to the fact that biometric modalities are independent one from another (e.g., a facial image says nothing about the fingerprint of the subject and vice versa) [27]. The constructed chimeric database consisted of facial imagery and fingerprint data of 100 subjects with each subject having a total of 8 biometric samples for each modality.

In our experiments all 800 samples from the chimeric database were matched against each other, forming a square 800 × 800 similarity score matrix. From this matrix, 6400 similarity scores corresponded to legitimate verification attempts and 633600 corresponded to illegitimate verification attempts. Note that prior to our experiments the matching scores were normalized using min-max score normalization [28].

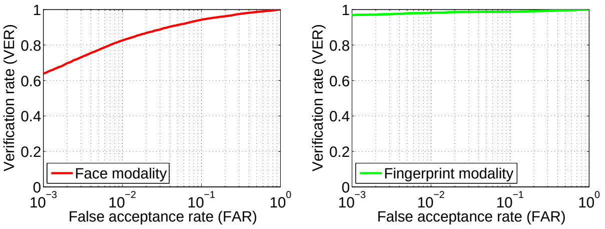

Figure 2. ROC curves of the experiments: face recognition (left), fingerprint recognition (right)

Table 2. Quantitative comparison of the biometric modalities

Procedure	EER	VER@0.1FAR
Face	0.0720	0.6394
Fingerprint	0.0163	0.9691

To evaluate the performance of the fusion techniques, we used Receiver Operating Characteristic (ROC) curves, which plot the verification rate (VER) against the false acceptance rate (FAR) at various values of the decision threshold. However, to better highlight the difference among the fusion procedures at the lower values of the false acceptance rate, we used a log scale for the x-axis of the ROC curves. In addition to the performance curves, we also computed quantitative performance measures for each of the experiments. For each series of experiments we calculated one (or several) of the following performance metrics:

- the so-called *equal error rate (EER)*, which is defined with the ROC curve operating point, where the false acceptance error rate (FAR) and the false rejection error rate (FRR) (i.e., 1-VER) take the same value,
- the verification rate at the false acceptance rate of 0.1% ($VER@0.1FAR$),
- the half total error rate, defined as HTER = 0.5(FAR+FRR),

4.2 *Analysis of fusion strategies*

A prerequisite for the evaluation of different fusion strategies is the establishment of the baseline performance of the two uni-modal biometric experts. The results of the first series of experiments, aimed towards this goal, are presented in the form of ROC curves in Fig. 2 and with quantitative performance metrics in Table 2.

The fingerprint recognition system performs, as expected, much better than the face recognition system. At the equal error rate, the face expert results in an error of around 7%, while the fingerprint expert ensures an error rate of a little more than 1.5%.

Table 3. Quantitative comparison of the decision-level fusion rules

Procedure	HTER	FAR	FRR
AND-rule	0.0440	0.0011	0.0869
OR-rule	0.0440	0.0862	0.0014

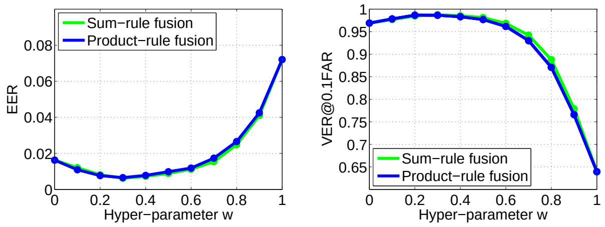

Figure 3. EERs, VER@0.1FARs and VER@0.01FARs for the sum- and product-fusion rules for different values of w

In our second series of verification experiments we evaluate the feasibility of fusion strategies applied at the decision level. In this setting, no similarity scores are sent to the client application, instead the cloud-services are asked to make a decision regarding the validity of the identity claim. The ratio between the FAR and the FRR (the operating point) of the cloud-recognition-service cannot be changed and is determined by the settings on the service-side. In our case, the operating point of the cloud-services is set to the Equal Error Rate (EER).

Two decision-level fusion schemes are implemented for the experiments as described in Section 3. The results of the experiments on the training data are shown in Table 3 in the form of various performance metrics. Note that it is not possible to generate ROC curves for this series of experiments, since no similarity scores are available.

Both tested fusion strategies result in a similar HTER with the difference that the AND-rule favors small FARs, while the OR-rule favors small FRRs. When compared to the performance of the single experts, the decision level fusion rules outperformed the face expert but performed worse than the fingerprint expert. As a general observation we can say that fusion strategies applied at the decision level are better suited for adjusting the operating point of the multi-modal biometric system than for improving its overall recognition performance.

In our third series of verification experiments we examined the performance of the two matching-score level fusion techniques on our chimeric database. Here, it was necessary to find appropriate values for the open hype-parameter w for the sum- and product-fusion rules. To this end, we gradually increased the value of w from 0 to 1 with a step size of 0.1 and observed the EER and VER@0.1FAR for the different values. The results of this series of experiments are shown in Fig. 3 Note that both the sum- and product-fusion rules result in the best performance at the value of $w = 0.3$, which is fixed for both fusion rules for all following experiments.

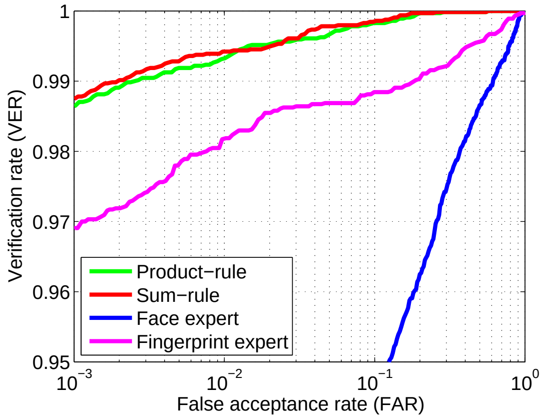

Figure 4. ROC curves for the fusion rules (training data)

Table 4. Quantitative comparison of the fusion rules with learned parameter w ($w = 0.3$ for both techniques) - training data

Procedure	EER	VER@0.1FAR
Product-rule	0.0066	0.9866
Sum-rule	0.0063	0.9875

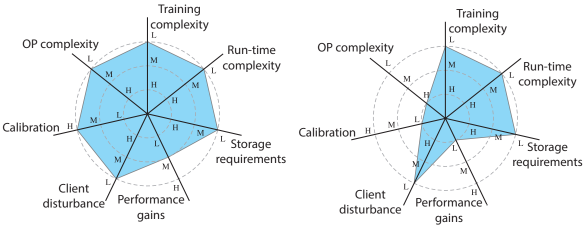

(a) Matching score-level fusion

(b) Decision-level fusion

Figure 5. Kiviat graphs of the fusion techniques generated based on the selected evaluation criteria

To compare the performance of the sum- and product-fusion rules with the selected value of the hyper-parameters to that of the single experts, we generated ROC curves for the conducted experiments. The performance curves are shown in Fig. 4 and the corresponding performance measures in Table 4. Note that he sum- and product-fusion rules perform significantly better than the uni-modal biometric experts. The EER, drops by more than 50% with both fusion rules when compared to the better performing fingerprint expert.

4.3 Strategies revisited

The experimental results presented in the previous sections show that different fusion strategies result in different verification performance on our chimeric database. However, from an application development point-of-view, other criteria next to pure verification performance, are of importance as well. To evaluate the fusion strategies based on other (non-performance related) criteria, a grade (low - L, medium - M, or high - H) was assigned to each strategy in seven different categories. The grades were given according to the perception of the authors and served as the basis for constructing the Kiviat graphs shown in Fig. 5. In the generated graphs a larger area represents a better fusion strategy with respect to the selected criteria. Note again that these grades are extremely subjective and reflect the opinion/perception of the authors.

Looking at the Kiviat graphs, we conclude that the fixed fusion rules turned out to be suited best for combining different cloud implementations of biometric experts into a multi-biometric system as they provide a good trade-off between almost all the criteria, especially when compared to decision-level fusion. Based on our analysis it is difficult to say if the sum- or product-rule should be preferred, since both approaches exhibited similar characteristics.

5 Conclusion

We have presented a case study on multi-modal biometrics in the cloud. We have shown that various fusion strategies can be exploited to combine existing implementations of biometric experts and that the strategies can lead to significant performance gains of the uni-modal systems. The main finding of the paper is the fact that existing biometric cloud-services can be combined with minimal effort and ensure enhanced security for potential client applications.

Acknowledgements

The work presented in this paper was supported in parts by the national research program P2-0250(C) Metrology and Biometric Systems, by the national research program P2-0214 Computer Vision and by the EU, European Regional Fund, in scope of the framework of the Operational Programme for Strengthening Regional Development Potentials for the Period 2007-2013, contract No. 3211-10-000467 (KC Class) and the European Union's Seventh Framework Programme (FP7-SEC-2011.20.6) under grant agreement number 285582 (RESPECT). The authors additionally appreciate the support of COST Actions IC1106 and IC1206.

References

[1] P. Peer, J. Bule, J. Žganec Gros, and V. Štruc, "Building cloud-based biometric services," *Informatica*, vol. 37, no. 1, pp. 115–122, 2013.

[2] A. Ross, K. Nandakumar, and A. Jain, *Handbook of Multibiometrics*. New York, USA: Springer Science+Business Media, LLC, 2006.

[3] X. Xia and L. O'Gorman, "Innovations in fingerprint capture devices," *Pattern Recognition*, vol. 36, no. 2, pp. 361–369, 2003.

[4] A. B. Khalifa and N. B. Amara, "Bimodal biometric verification with different fusion levels," in *International Multi-Conference on Systems, Signals and Devices*, vol. 1, 2009, pp. 1–6.

[5] A. Noore, R. Singh, and M. Vatsa, "Robust memory-efficient data level information fusion of multi-modal biometric images," *Information Fusion*, vol. 8, no. 4, pp. 337 – 346, 2007.

[6] A. Ross and R. Govindarajan, "Feature level fusion using hand and face biometrics," in *SPIE Conference on Biometric Technology for Human Identification*, vol. 5779, 2005, pp. 196ďż˝–204.

[7] A. Nagar, K. Nandakumar, and A. Jain, "Multibiometric cryptosystems based on feature-level fusion," *IEEE TIFS*, vol. 7, no. 1, pp. 255–268, 2012.

[8] Q. Tao and R. Veldhuis, "Robust biometric score fusion by naive likelihood ratio via receiver operating characteristics," *IEEE TIFS*, vol. 8, no. 2, pp. 305–313, 2013.

[9] N. Poh and J. Kittler, "A unified framework for biometric expert fusion incorporating quality measures," *IEEE TPAMI*, vol. 34, no. 1, pp. 3–18, 2012.

[10] N. Poh, D. Windridge, V. Mottl, A. Tatarchuk, and A. Eliseyev, "Addressing missing values in kernel-based multimodal biometric fusion using neutral point substitution," *IEEE TIFS*, vol. 5, no. 3, pp. 461–469, 2010.

[11] N. Poh, J. Kittler, and T. Bourlai, "Quality-based score normalization with device qualitative information for multimodal biometric fusion," *IEEE TSMC, Part A: Systems and Humans*, vol. 40, no. 3, pp. 539–554, 2010.

[12] M. Vatsa, R. Singh, A. Noore, and A. Ross, "On the dynamic selection of biometric fusion algorithms," *IEEE TIFS*, vol. 5, no. 3, pp. 470–479, 2010.

[13] K. Nandakumar, Y. Chen, S. Dass, and A. Jain, "Likelihood ratio-based biometric score fusions," *IEEE TPAMI*, vol. 30, no. 2, pp. 342ďż˝–347, 2008.

[14] J. Kittler, M. Hatef, R. Duin, and J. Matas, "On combining classifiers," *IEEE TPAMI*, vol. 20, no. 3, pp. 226ďż˝–239, 1998.

[15] L. Lam and C. Suen, "Application of majority voting to pattern recognition: An analysis of its behavior and performance," *IEEE TSMC, Part A: Systems and Humans*, vol. 27, no. 5, pp. 553–568, 1997.

[16] Q. Tao and R. Veldhuis, "Threshold-optimized decision-level fusion and its application to biometrics," *PR*, vol. 42, no. 5, pp. 823 – 836, 2009.

[17] L. Kuncheva, *Combining Pattern Classifiers: Methods and Algorithms*. Hoboken, New Yersey: Wiley-Interscience, 2004.

[18] ——, "A theoretical study on six classifier fusion strategies," *IEEE TPAMI*, vol. 24, no. 2, pp. 281–286, 2002.

[19] J. Kittler and F. Alkoot, "Experimental evaluation of expert fusion strategies," *PRL*, vol. 20, pp. 1361ďż˝–1369, 1999.

[20] P. Viola and M. Jones, "Robust real-time face detection," *IJCV*, vol. 57, no. 2, pp. 137ďż˝–154, 2004.

[21] X. Tan and B. Triggs, "Enhanced local texture sets for face recognition under difficult lighting conditions," *IEEE TIP*, vol. 19, no. 6, pp. 1635–1650, 2010.

[22] K. Kryszczuk and A. Drygajlo, "Improving biometric verification with class-independant quality information," *IET Signal Processing*, vol. 3, no. 4, pp. 310–321, 2009.

[23] U. Klopčič and P. Peer, "Fingerprint-based verification system : a research prototype," in *Proceedings of IWSSIP 2010*, A. Conci and F. Leta, Eds., vol. 1, 2010, pp. 150–153.

[24] J. Fierrez-Aguilar, J. Ortega-Garcia, J. Gonzalez-Rodriguez, and J. Bigun, "Kernel-based multimodal biometric verification using quality signals," 2004.

[25] K. Messer, J. Matas, J. Kittler, J. Luettin, and G. Maitre, "Xm2vtsdb: The extended m2vts database," in *Proceedings of AVBPA*, vol. 1, 1999, pp. 72–77.

[26] D. Maio, D. Maltoni, R. Cappelli, J. Wayman, and A. Jain, "Fvc 2002: Second fingerprint verification competition," in *Proceedings of ICPR*, vol. 1, 2002, pp. 811–814.

[27] N. Poh and S. Bengio, "Using chimeric users to construct fusion classifiers in biometric authentication tasks: an investigation," in *Proceedings of IEEE ICASSP*, vol. 1, 2006, pp. 1077–1080.

[28] A. Jain, K. Nandakumar, and A. Ross, "Score normalization in multimodal biometric systems," *PR*, vol. 38, no. 12, pp. 2270–2285, 2005.

Žiga Emeršič received his BSc from the Faculty of Computer and Information Science (University of Ljubljana) in 2013. His research interests include computer vision and fusion strategies.

Jernej Bule received his BSc from the Faculty of Electrical Engineering and Computer Science (University of Maribor) in 2011. He is currently a researcher at Faculty of Computer and Information Science (University of Ljubljana). His research interests include cloud computing, computer vision and biometrics.

Jerneja Žganec-Gros received her PhD from the Faculty of Electrical Engineering (University of Ljubljana) in 1997. She is currently CEO of Alpineon d.o.o. and focuses on research on speech and image technologies.

Vitomir Štruc received his PhD from the Faculty of Electrical Engineering (University of Ljubljana) in 2010. His research interests include pattern recognition, computer vision and biometrics.

Peter Peer received his PhD from the Faculty of Computer and Information Science, University of Ljubljana, in 2003. He is currently an Assistant Professor at the same institution. His research interests are computer vision, biometry and cloud computing.