

A Meta-Analysis of Face Recognition Covariates

Yui Man Lui, David Bolme, Bruce A. Draper, J. Ross Beveridge, Geoff Givens, P. Jonathon Phillips

Abstract—This paper presents a meta-analysis for covariates that affect performance of face recognition algorithms. Our review of the literature found six covariates for which multiple studies reported effects on face recognition performance. These are: age of the person, elapsed time between images, gender of the person, the person's expression, the resolution of the face images, and the race of the person. The results presented are drawn from 25 studies conducted over the past 12 years. There is near complete agreement between all of the studies that older people are easier to recognize than younger people, and recognition performance begins to degrade when images are taken more than a year apart. While individual studies find men or women easier to recognize, there is no consistent gender effect. There is universal agreement that changing expression hurts recognition performance. If forced to compare different expressions, there is still insufficient evidence to conclude that any particular expression is better than another. Higher resolution images improve performance for many modern algorithms. Finally, given the studies summarized here, no clear conclusions can be drawn about whether one racial group is harder or easier to recognize than another.

I. INTRODUCTION

Since 2003, the authors have published a series of papers quantifying the impact of covariates on the performance of face recognition algorithms [19], [17], [18], [8], [7]. The motivation for these studies is both scientific and practical. The scientific reason is to better understand algorithms and phenomenology underlying face image formation as expressed through recognition performance. The practical reason is to provide guidance to individuals responsible for fielding face recognition systems. The studies consistently show that changes in covariate values often affect the performance of a face recognition algorithm. The question taken up in this paper is, over multiple studies, which effects are consistent and which are not. For example, is there an emerging consensus about the way age, or gender, or image resolution influences face recognition performance.

Our approach is similar to that taken by Phillips and Newton in their iris recognition meta-analysis [27]. Here we present a quantitative summary of papers in the face recognition literature that report performance broken out by covariates. Face recognition covariates, in general, represent some aspect of the person, the image or some other aspect of the intended face recognition application. Covariates associated with people include such things as age, gender and race. They also include more transitive aspects of the person's appearance such as whether they are smiling or

wearing glasses. Properties of the image includes such things as image resolution, the number of pixels between the eyes, or whether the image was compressed, and if so, by how much. Other aspects of the intended application may include such factors as whether or not images are taken outdoors, or how much time has elapsed between the time two images being compared were taken.

The rule for including a covariate in this analysis was that results appear in more than one study, and ideally from more than one group of researchers. After an extensive review of the literature, the decision was made to restrict our meta-analysis to results from experiments on frontal, still, visible light images of faces. The following covariates meet the above criteria for inclusion in our study: age, gender, and race of the subject; expression; elapsed time between the two images being compared; and size of the face in the image.

Some covariates not included in our analysis warrant special mention. Arguably, the two covariates with the greatest impact on face recognition performance are pose and illumination. As a result, these two covariates have received significant attention. The Yale [16] and PIE [34] databases were created to systematically study pose, illumination, and their interactions, and have been widely used by researchers. As a result, there is already a substantial literature on pose and illumination.

There are also a number of interesting covariates that are worth briefly mentioning, but are not included in our analysis because they are only reported in one study. Boulton et al. [12] looked at how the time of day effects recognition performance in outdoor settings, and determined that recognition is better in the morning and afternoons than in either the daybreak or evening twilight in outdoor environment. The Foto-Fahndung report [3] noted that face recognition systems perform better in the middle of a day at the Mainz railway station. Image quality measures are also not analyzed, because virtually every study employs a unique quality measure. For example, Hsu and Martin [22] look at whether human subjective judgments of image quality correlate to machine recognition performance (they do), while Adler and Dembinsky [5] look at the correlation between a black-box vendor-supplied quality measure and performance. Abdel-Mottalib and Mahoor [4] study the effect of a machine quality measure based on kurtosis, while Weber [37] evaluates a quality measure based on "sharpness" and Beveridge et al. [8] measure quality via edge density. While there is clearly a relationship between image quality and recognition performance, structured comparisons are not possible without a common measure of image quality.

Yui Man Lui, David Bolme, Bruce A. Draper, and J. Ross Beveridge are with the Department of Computer Science, Geoff Givens is with the Department of Statistics, Colorado State University, Fort Collins, CO 80523, and P. Jonathon Phillips is with the National Institute of Standards and Technology, Gaithersburg, MD 20899.

A. Quick Summary of Findings

For the age covariate we found that virtually all studies report that older people are easier to recognize than younger ones. This is consistent with the intuition on the phenomenology of faces. Because faces change over time, it should not be surprising that older faces are either harder or easier to recognize than younger faces. Even though this effect is consistent over multiple studies, more research is needed to determine why older faces are easier to recognize.

Similarly, all studies agree that recognition performance degrades as the time between the query and target images increases. As expected, images taken on the same day consistently have a higher recognition rate than images taken on different days. Furthermore, it is expected that the appearance of the face changes as a person ages and therefore the time difference makes recognition more difficult. While this effect is clearly present in most studies, it is only significant when measured in years, not days or months.

Gender, on the other hand, is a point of contention. Studies disagree about whether men or women are easier to recognize, although most agree that any effect of gender on recognition performance is small. Gender is also an example of a covariate that interacts with other covariates. The effect of gender, for example, decreases as the age of the subject increases.

Regarding expression, there is consensus that performance improves when a person's expression is the same in both the query and target images. Less clear is whether, if forced to enroll only a single image, it is better to request the person to smile or instead to adopt a neutral expression.

Image resolution effects also do not always agree, but in this case the chronology of the studies takes on particular significance. Older studies, often using algorithms that match whole face images in dimensionality reduced subspaces showed a sweet spot in terms of resolution and actually in some cases showed that higher resolution images could reduce recognition performance. However, more recent studies with more modern algorithms appear to remove this barrier, and instead largely suggest that increasing resolution increases recognition performance.

Multiple studies show race matters in so much as recognition performance improves for individuals belonging to one race versus another. However, based upon current studies it is difficult to tease apart how much this has to do with the specific race versus the relative proportion of different races in a particular data set. There is a suggestion in the current studies that recognition may be easier when a person belongs to a race that appears less frequently in the system's database of enrolled people.

II. METHODOLOGY

A three stage methodology is used here. The first stage consists of deciding what are the initial selection criteria for including papers in this study and a strategy for searching the literature for papers. For the first stage the selection criteria were papers that reported experimental results broken out by a subject, image, or quality covariate. The search

utilized IEEEExplore, Google Scholar, major computer vision conferences, biometric and face conferences and workshops, and references in papers.

The second stage consists of refining the selection criteria and deciding how to summarize findings. At this stage the selection criteria was restricted to papers reporting on frontal, still, visible light images of faces. Some papers had to be dismissed because they did not specifically relate a change in a commonly recognized measure of face recognition performance to a clearly identified covariate. It was also the case that for covariates such as age, gender and expression, the vast majority of papers we reviewed addressed how to classify face images by age, gender or expression, rather than relating these attributes to face recognition. At the conclusion of the second stage we had tabulated just under 100 distinct effects drawn from 46 studies.

In the third stage, each of these findings was further reviewed in order to form some reasonable quantitative summary of the finding. As the reader will see in the tabulated results presented below, ultimately we had to be able to relate a change in a covariate, for example gender, to a percent change in a commonly used measure of face recognition performance. Examples of such measures include rank 1 identification rate or verification rate at a fixed false accept rate. Of the nearly 100 candidates coming out of stage two, for only about half were we able to take this final step.

The tabulated results coming out of stage three are presented in two forms below. There are abbreviated tables that indicate succinctly the findings of different studies with little auxiliary information. These are included in the main body of the paper and are sufficient to see the major trends in our analysis. For each major table, there is an expanded version included in a supplemental document ¹ that provides supplemental details such as the data set used in the study, the algorithm used in the study, and what if any controls were used to account for other covariates.

The covariate results are summarized in six tables. Each row of a table represents a single research finding, and the rows are typically ordered chronologically so that the most recent, and therefore presumably most relevant, findings appear near the top. The columns to the left, filled with asterisks, show the reported effect of a covariate on recognition. These columns share a common header, indicating the direction of the effect.

The magnitude of the effect is reported as a percent change in a common performance measure. For example, reported effects for verification studies are measured as the change in verification rate at a fixed false reject rate. For rank based identification studies, the magnitude reflects the change in the rank one recognition rate. In both cases, magnitudes are coarsely quantized into ranges.

An asterisk may appear in more than one column, indicating an interaction with some other covariate leading to a range of observed performance. The right most column cites the paper in which the finding is reported.

¹<http://www.cs.colostate.edu/~vision/publications/btas09lui.pdf>

TABLE III
GENDER EFFECT ON RECOGNITION PERFORMANCE.

Male Easier				Female Easier				Interactions with other covariates	Source
+17.5%	15%	10%	5%	0%	5%	10%	15%		
		*		*	*			Outdoor/Indoor	[9]
		*		*	*			Outdoor/Indoor	[9]
				*	*				[8]
				*	*				[7]
		*		*					[7]
		*		*					[18]
				*					[17]
				*					[17]
*	*	*		*	*			Age	[17]
*	*	*		*	*			Age	[19]
*	*	*		*	*			Age	[30]
*	*	*		*	*			Age	[30]
				*	*			Expression, Lighting	[20]

illumination, camera position or clothing, than with aging. We believe that when same-day pairs are not considered, the effect of aging over a few months or even one or two years is small. This conclusion is also supported by a study by Bartlett [6] that suggests it is easier to match same-day images with different expressions than to match different-day images with the same expression (see Section VI).

V. GENDER

Many studies have looked at the effects of gender on face recognition as shown in Table III, but something less than a consensus emerges. There is a slight result in favor of men being easier to recognize; men are easier to recognize than women 7 to 5, with 6 reporting no effect. One plausible inference might be that gender has little or no effect on face recognition, since 11 studies have reported either no effect or a range of effects which includes no effect. Also, no study has reported as large an effect as reported in some of the studies of age and time.

Gender effects, when reported, often interact with other covariates. Note the addition in Table III of a column indicating interactions between gender and other covariates. For example, Phillips et al. [30] reported that for the Identix [1] system, the gender effect could range from men being 10% easier to women being 10% easier, depending on the age of the subject (see row 14 of Table III). Other studies have shown interactions between gender and setting (indoor vs. outdoor) and among gender, expression and illumination. In all but one finding with an interaction, the interaction not only changed the magnitude of the gender effect, it changes its direction in terms of whether men or women are easier to recognize.

VI. EXPRESSION

Most work on expression is focused on recognizing a person's expression, with fewer studies relating expression to recognition performance. Our summary of these studies

TABLE IV
EXPRESSION EFFECT ON FACE RECOGNITION PERFORMANCE.

Same Expression				Different Expression				Source	
+17.5%	15%	10%	5%	0%	5%	10%	15%		+17.5%
*									[7]
	*								[7]
		*							[7]
			*						[17]
			*						[17]
	*			*					[17]
				*					[19]
Neutral vs. Smiling				Smiling vs. Neutral					
			*						[7]
			*						[7]
		*							[7]
Regular vs. Alternate				Alternate vs. Regular					
			*						[10]
			*						[10]
			*						[10]

appears in Table IV. These studies have focused on two subtly different questions: (1) does a change in expression matter, and (2) if it does matter, what expression should be used when a single image of a person is to be enrolled in a database?

Studies on whether changes in expression matter are reported in the top part of Table IV. Not surprisingly, there is relative consensus on this question: 6 out of 7 findings report that it is easier to match subjects across images when the expressions are the same. For 5 of the 7 studies same versus different expressions improved recognition performance from between 5% to 10%.

It should be noted that half of these studies were conducted on the FERET data, in which people were asked to adopt a neutral expression for one image and an "alternate" expression for the other. The other half of the studies were conducted on the FRGC dataset, in which people were asked to adopt a neutral expression for one image and a smiling expression for the other. Nonetheless, when both parts of an image pair share a common expression, whatever it may be, recognition is easier.

The bottom half of Table IV is divided into two parts: one for the study by Blackburn et al. [10], the other for the study by Beveridge et al. [7]. Both these studies included results for cross expression recognition. Thus, both studies considered the situation where a new, query, image is compared to an enrolled, target, image and the expressions do not match.

In the study by Blackburn et al. [10], carried out in 2000, recognition performance was compared between two cases: 1) a query image with a regular expression was compared to a target image with an alternate expression, and 2) a query image with an alternate expression was compared to a target image with a regular expression. Blackburn et al. found that performance improved for cross-expression image

pairs if the query image had the alternate expression and the enrolled image the regular expression. This is true across three algorithms, although the effect is small (about 5%).

In the study by Beveridge et al. [7], carried out on data and algorithms from FRGC 2006, the two cases compared were: 1) a query image with a neutral expression was compared to a target image with an smiling expression, and 2) a query image with smiling expression was compared to a target image with a neutral expression. The authors found that when the choice was restricted to smiling versus neutral, enrolling a smiling face is slightly better. However, the effect is small and algorithm dependent. In another study, Yacoob and Davis [40]³ report that extreme expressions, whether smiling or angry, produce better recognition performance than neutral expressions.

Overall it is better to compare matching expressions than non-matching expressions, but from a performance standpoint, what expression should be enrolled when the query cannot be controlled is unknown. This finding is of interest in relation to the common practice of requiring a neutral expression on face images being used in an official capacity.

For example, the Canadian government has decreed that smiles are not allowed in passport photos [2]. Many U.S. states are adopting or considering similar regulations for driver's license photos. This bias toward neutral expressions makes sense if one assumes that neutral expressions are more likely to be present in query images against which these enrolled images are to be compared. It is not supported by evidence of any intrinsic superiority of one expression over another.

VII. IMAGE RESOLUTION

Image resolution is an important covariate that can greatly affect an algorithm's performance. It is also a covariate over which system engineers typically have a great deal of control. Modern cameras make comparatively high resolution images available with essentially no overhead. However, there are still computational incentives to downsample images in order to make storage and particularly comparison of images more efficient. Thus understanding the performance trade offs related to resolution is critical.

Table V presents results grouped into three sections. The top section of the table describes the effect of image resolution while both query and target images are varied based on the distance between the eyes. The middle section of the table illustrates the effect of image resolution while the resolution of the target images is fixed. The bottom section of the table depicts the effect of image resolution in terms of image size rather than pixels between the eyes.

There is general agreement that recognition performance degrades when image resolution is below 32×32 [11], [36], [13], [38], [23] as evident by the results shown in the bottom section of Table V. This statement suggests an operational lower bound on image resolution. Note that enhancement

³This study is noteworthy, but was not tabulated in our results because it reported results in terms of discrimination power that we could not easily map to explicit changes in recognition performance.

techniques, like super-resolution [38], may be applied prior to image matching when the resolution drops below this lower bound.

Still focusing on the bottom section of Table V, the optimal recognition performance can be found in image size between 32×32 and 64×64 . Note the results for this section are derived for variants of PCA and LDA (see supplemental materials for details). Low resolution imagery may be adequate for these older techniques. In addition, Keil et al. [23] reported that a 37×37 face chip has the highest class separability for Fisher linear discriminant analysis, nonparametric discriminant analysis, and mutual information algorithms on the FRGC subset. Beveridge et al. [7] also reported that PCA benefits from using lower resolution shown in the fourth row on the top section of Table V.

What about more advanced algorithms? The FaceIt system [1] is one of the earliest commercial face recognition products. Like PCA and LDA algorithms, the FaceIt system derived no benefit from higher resolution imagery when tested in the FRVT 2000 evaluation [10] shown in the first row of the middle section of Table V. More specifically, according to the FRVT 2006 evaluation report [31], the FaceIt system performs gracefully with lower resolution imagery when the eye distance is varied between 75 and 400 pixels in both controlled and uncontrolled settings. This result is shown in the top section of Table V, row 14. Similar findings are also supported by Wheeler et al. [38] where the optimal recognition result is found in 37 pixels between the eyes using the FaceIt system on the GE video and the FERET dataset. Finally, O'Toole et al. [28] noted a similar finding; as shown in the top section of Table V, row 12, the FaceIt system does not favor very high resolution imagery.

On the other hand, some modern algorithms take considerable advantage of high resolution imagery. As the top section of Table V reveals, most of the results summarized indicate considerable improvement in performance is possible when using higher resolution images. Recently, O'Toole et al. [28] have found that NevenVision⁴ performs exceptionally well on very high resolution images as shown in the eleventh row on the top section of Table V. The fact that NevenVision performs better in higher resolution can be further supported by the FRVT 2006 evaluation [31] where NevenVision consistently yields higher recognition performance in both controlled and uncontrolled settings using high resolution imagery, see the fifteenth row on the top section of Table V. Furthermore, Beveridge et al. [7] also noted that algorithms provided by CMU [39] and NJIT [25] achieve better recognition results when high resolution imagery is used as shown in the fifth and sixth rows on the top section of Table V, respectively.

The choice of image resolution is clearly algorithm dependent which is demonstrated through a wide range of algorithms like LDA [11], PCA [36], FaceIt [1], CMU [39], NJIT [25], and NevenVision. When there is an interaction

⁴Neven Vision, Inc was acquired by Google, Inc in 2006.

between indoor and outdoor, image resolution is also interacted with these environment settings. Beveridge et al. [8], [9] reported that higher resolution imagery is favored when the query image is taken outdoor whereas lower resolution may be adequate when the query is acquired in indoor. As the first three rows on the top section of Table V depicts, the environment effect could range from 17.5% or more easier when high resolution imagery is used in an outdoor setting to 10% easier when low resolution imagery is used on indoor images.

Face recognition is advancing quickly and will find more applications in uncontrolled environments. Therefore, we can safely surmise that the use of high resolution imagery will increase. As Table V shows, some modern algorithms have already taken advantages of the additional detail present in high resolution imagery.

VIII. RACE

As shown in Table VI, there have been several studies which attempt to determine if members of one racial group are easier or harder for machines to recognize than members of another. Unfortunately, the findings from these studies are difficult to interpret because all of the studies confound race with sampling effects. To be specific, all three of the studies in Table VI draw their subjects from racially mixed populations in which the groups are unevenly represented. In the FERET, FRGC and FRVT2006 datasets, Caucasians are always the most heavily sampled group, followed by East Asians, by approximately 3 to 1. All other groups are insufficiently represented to report.

So, while Table VI shows that for these studies, Caucasians are harder to recognize than East Asians, this could be explained either as an intrinsic property of the two races, or alternatively, as a by product of the relative proportion of the data set constituted by each. If that latter explanation is the correct one, the presumably reversing the relative proportion of Caucasians and East Asians in the study would correspondingly suggest Caucasians are easier to recognize than East Asians. What is needed to resolve this ambiguity is a study comparing performance across two uniform populations of different racial groups, but to the best of our knowledge this has not yet been done.

IX. CONCLUSION

The goal of this review was to quantitatively summarize what is known from the literature about the effects of six covariates on face recognition: age, time, gender, expression, resolution and race. In so doing, we collect findings reported over different face recognition algorithms and data sets using different statistical methods. The goal was to look for consensus; if a result holds over different algorithms, data sets and studies, it is probably a universal property of face recognition. Conflicting findings, on the other hand, suggest that an effect is algorithm or data specific, or interacts with another unknown factor. Using this methodology, we identify four universal covariates: (1) older people are easier to recognize than younger ones; (2) recognition gets harder

TABLE V
IMAGE RESOLUTION EFFECT ON RECOGNITION PERFORMANCE.

High Resolution					Low Resolution						
+17.5%	15%	10%	5%	0%	5%	10%	15%	+17.5%	Pixels between the eyes	Source	
*	*	*	*	*	*	*			202 ~ 134	[9]	
*	*	*	*	*	*	*			202 ~ 134	[9]	
*	*	*	*	*	*	*			202 ~ 134	[9]	
								*	180 ~ 120	[7]	
									180 ~ 120	[7]	
*									180 ~ 120	[7]	
*					*				210 ~ 137	[8]	
				*				*	190 ~ 110	[28]	
				*				*	190 ~ 110	[28]	
			*					*	190 ~ 110	[28]	
*	*	*	*	*	*				190 ~ 110	[28]	
*	*	*	*	*	*	*	*	*	190 ~ 110	[28]	
				*					400 ~ 75	[31]	
				*	*	*	*	*	400 ~ 75	[31]	
				*	*	*	*	*	400 ~ 75	[31]	
*	*	*	*	*	*				400 ~ 75	[31]	
*	*	*	*	*	*				400 ~ 75	[31]	
*	*	*	*	*	*				400 ~ 75	[31]	
*	*	*	*	*	*				48 ~ 17	[38]	

High Resolution					Low Resolution						
+17.5%	15%	10%	5%	0%	5%	10%	15%	+17.5%	Pixels between the eyes	Source	
				*					60 ~ 15	[10]	
*									60 ~ 15	[10]	
									60 ~ 15	[10]	

High Resolution					Low Resolution						
+17.5%	15%	10%	5%	0%	5%	10%	15%	+17.5%	Image size	Source	
			*	*					128x128 vs 32x32	[11]	
			*	*					32x32 vs 16x16	[11]	
			*	*					144x108 vs 64x48	[36]	
			*	*					64x48 vs 32x24	[36]	
			*	*					144x108 vs 64x48	[36]	
		*		*					64x48 vs 32x24	[36]	
		*		*					64x64 vs 32x32	[13]	
*	*	*	*	*					64x64 vs 16x16	[13]	
*	*	*	*	*					64x64 vs 8x8	[13]	

as the time between the query and target images grows, but this is only significant if the time difference is measured in years; (3) recognition is easier if the subject adopts the same expression in the query and target images; and (4) high resolution imagery yields better recognition performance for many modern algorithms. None of these results are particularly surprising.

Of perhaps greater interest are the effects we did not find. There is no universal gender effect. Although individual studies report that men are easier to recognize than women or vice-versa, there is no consensus. Gender effects appear to be algorithm and/or dataset dependent rather than universal. Expression is another interesting case. While it is universally easier to match images with the same expression, there is

TABLE VI
RACE EFFECT ON RECOGNITION PERFORMANCE.

Caucasian Easier		East Asian Easier			Interactions with other covariates	Source
+17.5%	15%	10%	5%	0%		
					Outdoor/Indoor	[9]
					Outdoor/Indoor	[9]
			*	*		[8]
			*	*		[7]
			*	*		[7]
			*	*		[17]
			*	*		[17]
			*	*		[17]
			*	*	Dist. Measure	[19]
			*	*		[15]
			*	*		[15]

no empirical support for the accepted dictum that subjects should be enrolled with neutral expressions. Finally, although there are several studies suggesting a race effect, all of these studies are confounded with sampling imbalances between racial groups. As a result, there is no strong evidence to believe that one race group is easier to recognize than another; the more probable explanation is that whichever group is most heavily sampled in a population becomes the most difficult racial group to recognize.

Not surprisingly, this review suggests the need for additional studies. Studies that look at the effect of expressions during enrollment are needed, particularly in light of the current efforts to regulate passport and other identification photos. Studies that separate sampling effects from ethnic effects would also be useful. The most important studies, however, may be about other covariates. Even taking age, time and expression into account, it is still difficult to predict which images will fail to match. More research is needed.

ACKNOWLEDGMENTS

The work was funded in part by the Technical Support Working Group (TSWG) under Task T-1840C. Jonathon Phillips was supported by the Department of Homeland Security, Director of National Intelligence, and Federal Bureau of Investigation. The identification of any commercial product or trade name does not imply endorsement or recommendation by Colorado State University or the National Institute of Standards and Technology.

REFERENCES

- [1] Faceit. <http://www.11id.com>.
- [2] Passport canada. <http://www.ppt.gc.ca/cdn/photos.aspx?lang=eng>.
- [3] Face recognition as a search tool "foto-fahndung", 2007.
- [4] M. Abdel-Mottaleb and M. Mahoor. Application notes-Algorithms for Assessing the Quality of Facial Images. *IEEE Computational Intelligence Magazine*, 2(2):10–17, 2007.
- [5] A. Adler and T. Dembinsky. Human vs. automatic measurement of biometric sample quality. In *CCECE*, 2006.
- [6] M. Bartlett. *Face image analysis by unsupervised learning and redundancy reduction*. PhD thesis, University of California, San Diego, 1998.

- [7] J. Beveridge, G. Givens, P. Phillips, and B. Draper. Factors that influence algorithm performance in the Face Recognition Grand Challenge. *Computer Vision and Image Understanding*, 2009.
- [8] J. Beveridge, G. Givens, P. Phillips, B. Draper, and Y. M. Lui. Focus on quality, predicting frvt 2006 performance. In *FGR, Amsterdam, The Netherlands*, 2008.
- [9] J. R. Beveridge, G. Givens, P. J. Phillips, B. Draper, D. S. Bolme, and Y. M. Lui. Frvt 2006: Quo vidas face quality. *Image and Vision Computing*, (to appear).
- [10] D. Blackburn and M. B. J. Phillips. Frvt 2000 - evaluation report. *Technical Report*, 2001. <http://www.frvt.org>.
- [11] B. Boom, G. Beumer, L. Spreuwers, and R. Veldhuis. The effect of image resolution on the performance of a face recognition system. In *Proceedings of the Ninth International Conference on Control, Automation, Robotics and Vision*, pages 409–414, 2006.
- [12] T. E. Boulton and R. Scheirer, Walter J. and Woodworth. Faad: face at a distance. In *SPIE Defense and Security Symposium : Biometric Technology for Human Identification*, volume 6944, 2008.
- [13] J. Czyz and L. Vandendorpe. Evaluation of LDA-based face verification with respect to available computational resources. In *Int'l Workshop on Pattern Recognition in Information Systems*, 2002.
- [14] P. J. Flynn, K. W. Bowyer, and P. J. Phillips. Assessment of time dependency in face recognition: An initial study. In *Audio- and Video-Based Biometric Person Authentication*, 2003.
- [15] N. Furl, P. Phillips, and A. O'Toole. Face recognition algorithms and the other-race effect: computational mechanisms for a developmental contact hypothesis. *Cognitive Science*, 26:797–815, 2002.
- [16] A. Georgiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660, 2001.
- [17] G. Givens, J. Beveridge, B. Draper, P. Grother, and P. Phillips. How features of the human face affect recognition : A statistical comparison of three face recognition algorithms. In *CVPR*, 2004.
- [18] G. Givens, J. Beveridge, B. Draper, and P. Phillips. Repeated Measures GLMM Estimation of Subject-Related and False Positive Threshold Effects on Human Face Verification Performance. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 40–40, 2005.
- [19] G. Givens, J. R. Beveridge, B. A. Draper, and D. Bolme. A statistical assessment of subject factors in the pca recognition of human faces. In *CVPR Workshop on Statistical Analysis in Computer Vision*, 2003.
- [20] R. Gross, J. Shi, and J. Cohn. Quo vadis face recognition? In *CVPR Workshop on Empirical Evaluation Methods in Computer Vision*, 2001.
- [21] W. H. Ho, P. A. Watters, and D. Verity. Are younger people more difficult to identify or just a peer-to-peer effect. In W. G. Kropatsch, M. Kampel, and A. Hanbury, editors, *CAIP*, volume 4673, pages 351–359. Springer, 2007.
- [22] R.-L. Hsu, J. Shah, and B. Martin. Quality assessment of facial images. In *Biometric Consortium Conference, 2006 Biometrics Symposium*, 2006.
- [23] M. S. Keil, A. Lapedriza, D. Masip, and J. Vitria. Preferred spatial frequencies for human face processing are associated with optimal class discrimination in the machine. *PLoS ONE*, 3(7), 2008.
- [24] H. Ling, S. Soatto, N. Ramanathan, and D. Jacobs. A study of face recognition as people age. In *IEEE Conf. ICCV*, 2007.
- [25] C. Liu. Capitalize on dimensionality increasing techniques for improving face recognition grand challenge performance. *IEEE transactions on pattern analysis and machine intelligence*, 28(5):725–737, 2006.
- [26] Z. Liu and S. Sankar. Outdoor recognition at a distance by fusing gait and face. *Image and Vision Computing*, 25(6), 2007.
- [27] E. Newton and P. Phillips. Meta-analysis of third-party evaluations of iris recognition. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 39(1):4–11, Jan. 2009.
- [28] A. O'Toole, P. Phillips, and A. Narvekar. Humans versus algorithms: Comparisons from the face recognition vendor test 2006. In *FGR, Amsterdam, The Netherlands*, 2008.
- [29] U. Park, Y. Tong, and A. K. Jain. Face recognition with temporal invariance: A 3d aging model. In *FG*, pages 1–7. IEEE, 2008.
- [30] J. Phillips, P. Grother, R. Micheals, D. Blackburn, E. Tabassi, and M. Bone. Frvt 2002 - evaluation report. *Technical Report*, 2002. <http://www.frvt.org>.
- [31] P. Phillips, W. Scruggs, A. O'Toole, P. Flynn, K. Bowyer, C. Schott, and M. Sharpe. FRVT 2006 and ICE 2006 large-scale results. *National Institute of Standards and Technology, NISTIR*, 7408, 2007.
- [32] N. Ramanathan and R. Chellappa. Face verification across age progression. *IEEE Transactions on Image Processing*, 15(11):3349, 2006.

- [33] K. Ricanek Jr and E. Boone. The effect of normal adult aging on standard PCA face recognition accuracy rates. In *2005 IEEE International Joint Conference on Neural Networks, 2005. IJCNN'05. Proceedings*, volume 4, 2005.
- [34] T. Sim, S. Baker, and M. Bsat. The cmu pose, illumination, and expression database. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1):1615 – 1618, December 2003.
- [35] R. Singh, M. Vatsa, A. Noore, and S. K. Singh. Age transformation for improving face recognition performance. In A. Ghosh, R. K. De, and S. K. Pal, editors, *Pattern Recognition and Machine Intelligence, Second International Conference*, volume 4815 of *Lecture Notes in Computer Science*, pages 576–583. Springer, 2007.
- [36] J. Wang, C. Zhang, and H. Shum. Face image resolution versus face recognition performance based on two global methods. In *Proceedings of Asia Conference on Computer Vision (ACCV'04)*, 2004.
- [37] F. Weber. Some quality measures for face images and their relationship to recognition performance. In *Biometric Quality Workshop. National Institute of Standards and Technology, Maryland, USA*, 2006.
- [38] F. W. Wheeler, X. Liu, and P. Tu. Multi-frame super-resolution for face recognition. In *BTAS*, 2007.
- [39] C. Xie, M. Savvides, and B. Kumar. Redundant class-dependence feature analysis based on correlation filters using frgc2.0 data. *Computer Vision and Pattern Recognition, 2005 IEEE Computer Society Conference on*, 3:153–153, June 2005.
- [40] Y. Yacoob and L. Davis. Smiling faces are better for face recognition. In *In Proc. International Conference Face Recognition and Gesture Analysis*, pages 59–64, 2002.