

Predicting Gender and Weight from Human Metrology using a Copula Model

Deng Cao, Cunjian Chen, Donald Adjeroh, Arun Ross

West Virginia University

Morgantown WV 26506, USA

{dcao, cchen10}@mix.wvu.edu

{Donald.Adjeroh, Arun.Ross}@mail.wvu.edu

Abstract

We investigate the use of human metrology for the prediction of certain soft biometrics, viz. gender and weight. In particular, we consider geometric measurements from the head, and those from the remaining parts of the human body, and analyze their potential in predicting gender and weight. For gender prediction, the proposed model results in a 0.7% misclassification rate using both body and head information, 1.0% using only body information, and 12.2% using only head information on the CAESAR 1D database consisting of 2,369 subjects. For weight prediction, the proposed model gives 0.01 mean absolute error (in the range 0 to 1) using both body and head information, 0.01 using only body information, and 0.07 using only measurements from the head. This leads to the observation that human body metrology contains enough information for reliable prediction of gender and weight. Furthermore, we investigate the efficacy of the model in practical applications, where metrology data may be missing or severely contaminated by various sources of noises. The proposed copula-based technique is observed to reduce the impact of noise on prediction performance.

1. Introduction

Automated gender classification from biometric data has been an active research topic in the computer vision literature. Most existing work has focused on assessing gender using information from the human face, be it appearance-based [19, 14] or metrology-based [4, 11]. In surveillance type of applications, researchers have considered using whole body images for gender prediction. Cao et al. [5] studied the problem of gender recognition from whole body images using a part-based representation and an ensemble learning algorithm. They reported a 75.0% accuracy for predicting gender from either the front view or back view. Guo et al. [15] used biologically-inspired features in combination with manifold learning techniques and

achieved around 80% accuracy in gender prediction from body. Collins et al. [8] investigated several popular image representation methods, such as Histogram of Gradients (HOG), spatial pyramid HOG (PHOG) and bag of words model, with the goal of finding effective human body appearance models for gender prediction. Shan et al. [23] fused gait and face cues for automated gender recognition based on canonical correlation analysis. Their work demonstrated that the two sets of measurements from gait and face are complementary, resulting in improved recognition accuracy at the feature level.

Though the above approaches show that automatic gender prediction from human body is feasible, the methods are largely based on *appearance* or *texture* information. Our work, on the other hand, utilizes a set of *geometric measurements* to predict gender in the absence of textural details. A similar line of work was undertaken by Adjeroh et al. in [2] where the problem of gender prediction from whole-body human metrology was considered. However, in [2], the associations between different body measurements were neither characterized nor utilized to improve prediction performance. This work bridges this gap by introducing a novel copula-based prediction model that exploits the association structure between different human metrological features. In addition to gender prediction, the proposed model is also used to deduce an individual's weight from metrology. Weight prediction from metrology has been previously studied by Adjeroh et al. [2] and Velardo and Dugelay [25].

In general, the use of metrology for deducing soft biometric traits has several applications. In video-based surveillance systems, it may be easier to quickly extract geometric measurements of the human body for classification, rather than primary biometric traits such as face or iris. In recognition-at-a-distance applications, primary biometric traits may not be readily available thereby necessitating the use of the dynamic geometry of the human body for identification. In applications based on the Kinect device, deducing gender or weight information from human anthropometric

measurements may be useful for enhancing perceived user experience.

The contribution of this work is two-fold. Firstly, we develop a copula-based gender and weight prediction model that accounts for associations between geometric attributes of the human body. Secondly, we evaluate the efficacy of the model on the CAESAR 1D database, both in the absence and presence of (simulated) noise. The rest of the paper is organized as follows. Section 2 discusses the basics of a prediction (estimation) model. Section 3 presents the copula model used in this work. Section 4 provides experimental evidence substantiating the efficacy of the model. Section 5 concludes the article.

2. Background

2.1. A General Prediction System

An end-to-end metrology-based prediction system has several discrete stages. In the first stage (feature extraction), a person is characterized using a collection of metrology measures known as features. These features need to be properly extracted from an individual. In the second stage (feature representation), the raw features are transformed into a new feature space, which is expected to be suitable for further analysis. This stage usually involves various types of normalizations. In the third stage (feature selection), certain techniques are used to reduce the dimension of the feature space in order to enhance the generalization capability, or to speed up the learning process. After this stage, the raw feature vector $V = (x_1, \dots, x_n)$ becomes $V' = (x'_1, \dots, x'_m)$, where $m \leq n$. In the final stage (feature prediction), V' is sent to a classifier (for classification) or a regressor (for estimation). The output is a discrete class for classification or a continuous value for regression.

In this work, we do *not* focus on feature extraction. We assume that a set of features (measurements), referred to as the feature vector, is already provided. Our goal is to analyze these features and develop a prediction model based on these features. The input to our prediction system is a set of metrological features pertaining to an individual. These features correspond to the head and the body. The output of the prediction system is the gender (classification) or weight (regression) of the individual.

2.2. The Association Structure between Features

Previous gender prediction models [2, 25, 10, 15] do not explicitly consider the association between the features in the feature vector. However, significant correlation between human metrological features have been observed in [2]. For example, a tall person is *likely* to have long arms, long feet, and long fingers. Clearly, incorporating information on the association or correlation structure between human measurements is likely to lead to improved prediction perfor-

mance. In order to better understand the interaction among different metrological features, we use a specific statistical tool. The term **association** should be considered as **statistical dependence** rather than a narrower term **correlation**, which is usually used to describe the linear relationship between random variables [20]. Thus, we do not use the Pearson correlation coefficient since the linearity between features is not guaranteed. We also avoid the Chi-square goodness of fit test, since it requires prior knowledge of the distribution of samples.¹ Instead, we use a non-parametric test which does not rely on any assumptions on the distribution of samples. The Kendall's tau rank correlation coefficient [17] is selected as our tool.

Let $(x_1, y_1), \dots, (x_n, y_n)$ be a set of joint observations from two random variables X and Y , respectively, such that all the values of x_i and y_i are unique. A pair of observations (x_i, y_i) and (x_j, y_j) are said to be **concordant** if both $x_i < x_j$ and $y_i < y_j$ or if both $x_i > x_j$ and $y_i > y_j$. They are said to be **discordant**, if $x_i > x_j$ and $y_i < y_j$ or if $x_i < x_j$ and $y_i > y_j$. If $x_i = x_j$ or $y_i = y_j$, the pair is neither concordant nor discordant. Let n_c and n_d be the number of concordant and discordant pairs, respectively. Then, the Kendall's tau coefficient is defined as follows:

$$\tau = \frac{n_c - n_d}{\frac{1}{2}n(n-1)}. \quad (1)$$

If X and Y are independent, the coefficient would be approximately zero. The coefficient is 1 for perfect agreement between X and Y and -1 for perfect disagreement. Figure 1 shows the absolute values of pairwise Kendall's tau coefficients for the 43 anthropometric measures in the CAESAR 1D database [1]. The features themselves are listed in Figure 8. The warmer color denotes a higher dependence and the colder color denotes a lower dependence. The figure shows that significant associations do exist between pairwise features. The average absolute coefficient value between features varies from 0.13 to 0.51.

We want to incorporate this association structure into our prediction model in order to boost its performance. Thus, a novel copula-based prediction system is proposed in the next section.

3. Copula-based Features

3.1. Bayesian Framework

Inspired by Kwitt et al [18], we construct the features based on the class-conditional joint probability density. Let $1, \dots, M$ be a set of classes, and $p(r)$ be the corresponding probability density function (*pdf*) for $r = 1, \dots, M$. Let $z = (z_1, \dots, z_B)$ be the B -dimensional feature vector in which the features may be dependent on each other. Assume that we are to classify a sample based on the evidence

¹A sample is a single feature vector.

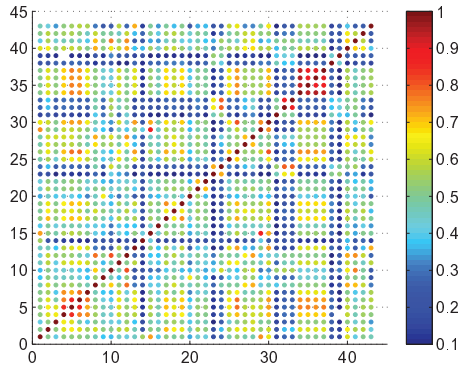


Figure 1. A display of pairwise Kendall's tau coefficients between the 43 anthropometric measures found in the CAESAR database.

provided by z . An optimal decision rule is to choose the class that is most probable given observation z . Define a function $g(z) \rightarrow 1, \dots, M$ that maps z to one of M classes. This decision rule can be formulated as a Bayes classifier [13]:

$$g(z) = \arg \max_r p(r|z). \quad (2)$$

However the posterior probability $p(r|z)$ is usually hard to obtain. So, we consider Bayes' theorem:

$$p(r|z) = \frac{p(z|r)p(r)}{p(z)}. \quad (3)$$

We may reasonably assume that each feature vector belongs to one and only one of the M classes with equal prior probability $p(r)$. Thus, by Bayes theorem, Eqn. (2) can be rewritten as a maximum likelihood (ML):

$$g(z) = \arg \max_r p(z|r). \quad (4)$$

In practice, $p(z|r)$ can be estimated from a collection of training samples z^1, \dots, z^n from class r . Also, a classifier such as Support Vector Machine (SVM) can be used in lieu of Eqn (4) as done in this work.

3.2. Copula Modeling

Our next objective is to compute the class-conditional likelihood $p(z|r)$. In this work, we choose the copula model to construct $p(z|r)$, which offers two advantages. Firstly, the copula representation does not require explicit mathematical relations between features, which are usually unknown in practice. Instead, it relies on the study of marginal distributions of components in z , which are substantially easier to obtain in practice. Secondly, the copula construction does not constrain the choice of marginal distributions, so the model can be adapted to different feature spaces.

Consider B uniform random variables, u_1, \dots, u_B , where $u_i \in [0, 1]$ for $i = 1, \dots, B$. Let $u = (u_1, \dots, u_B)$. A copula is defined as follows:

$$C(u_1, \dots, u_B) = Pr(U_1 \leq u_1, \dots, U_B \leq u_B). \quad (5)$$

Thus, by Sklar's theorem [12], given a B -dimensional random vector $z = (z_1, \dots, z_B)$, there exists a B -dimensional copula C such that:

$$C(F_1(z_1), \dots, F_B(z_B)) = F(z_1, \dots, z_B), \quad (6)$$

where $F_i(z_i), i = 1, \dots, B$ are marginal cumulative distribution functions (cdf's): $F_i(z_i) = Pr(Z_i \leq z_i)$. If F_1, \dots, F_B are given and they are continuous and non decreasing, we have [20]:

$$C(u) = F(F_1^{-1}(u_1), \dots, F_B^{-1}(u_B)), \quad (7)$$

where $F_i^{-1}(u_i)$ denotes the inverse cdf of F_i . Eqn. (7) is an important property of copulas that allows us to utilize the information about the marginals. Figure 2 and Figure 3 show the mapping from original data (Sitting Acromial Height and Ankle Circumference) to the unit square copula scale using a kernel estimator of the cumulative distribution function.




Figure 2. The scatter plot with marginal histograms of Sitting Acromial Height (x) and Ankle Circumference (y).

If the joint density f of F exists, it can be written as the product of the copula density c and the marginal densities [20, 18]:

$$\begin{aligned} f(z) &= \frac{\partial^B C(u)}{\partial z_1 \dots \partial z_B} = \frac{\partial^B C(u)}{\partial u_1 \dots \partial u_B} \prod_{i=1}^B \frac{\partial u_i}{\partial z_i} \\ &= c(u) \prod_{i=1}^B p_i(z_i). \end{aligned} \quad (8)$$

Assuming that the class labels are provided in the training data (in our case, gender labels), the class-conditional




Figure 3. Transformed data in Figure 2 to the copula scale using the estimated cdf .

copula density $c(u|r)$ can be easily calculated from Eqn (8) and has the form:

$$c(u|r) = \frac{f(z|r)}{\prod_{i=1}^B p_i(z_i|r)}. \quad (9)$$

Assuming that we have enough training samples for each class, in Eqn (9), $f(z|r)$ and $p_i(z_i|r)$ for all i can be estimated from the corresponding training samples for each r . Thus, the association structure between B dependent features for a given class r is converted to a single variable $c(u|r)$. We denote $c(u|r)$ as our new feature and refer to this as the CFeature (C for copula). In our proposed algorithm, we do not consider a single association matrix. Rather, we deduce multiple association matrices corresponding to pairs of features. Consequently, we have multiple CFeatures corresponding to these pairs. To distinguish the set of original measurement features from the set of CFeatures, the former is referred to as MFeatures (M for metrology or measurement). In practice, computing $c(u|r)$ involves the estimation of the copula parameter matrix Λ^r , which is defined by the copula *type* that is chosen. There are several copula types possible: Gaussian copula, student t copula and various Archimedean copulas (e.g., Clayton, Frank, and Gumbel). For example, an Archimedean copula with Clayton generator is defined in Eqn (10), where ψ is the generator, ψ^{-1} is the generator inverse, and $\theta \in (0, \infty)$ is the only parameter. Since we consider pairwise associations, note that $C(u)$ has *two* components.

$$\begin{aligned} C(u) &= \psi_\theta [\psi^{-1}(u_1), \psi^{-1}(u_2)], \\ \psi(t) &= (1+t)^{-1/\theta}, \\ \psi^{-1}(t) &= t^{-\theta} - 1. \end{aligned} \quad (10)$$

3.3. Estimation of Parameters

The copula parameter matrix Λ^r in the copula model should be correctly estimated so that the probability of mis-

classification error is minimized. We use the local maximum likelihood estimation (MLE) method [18] to estimate this for each CFeature. Note there are two types of parameters: marginal parameters and copula parameters. To estimate these parameters, we can use a two-step procedure called as the Inference Functions for Margins (IFMs) method [2]. Consider a pair of distinct features denoted by subscripts, i and j . For a given set of independent and identically distributed (iid) training samples, $z^1 = (z_i^1, z_j^1), \dots, z^n = (z_i^n, z_j^n)$, $i, j \in \{1, 2, \dots, B\}$, the marginal parameters $\hat{\theta}_i$ and $\hat{\theta}_j$ are first estimated. In the second step, we transform z_i and z_j to their corresponding cumulative marginals F_i and F_j using the probability integral transform (PIT) [21] and estimate the copula parameter Λ using MLE [16]:

$$\hat{\Lambda}^r = \arg \max_{\Lambda^r} \sum_{t=1}^n \log c \left(F(z_i^t; \hat{\theta}_i), F(z_j^t; \hat{\theta}_j) | \Lambda^r \right). \quad (11)$$

Although Sklar's theorem shows that a copula function always exists, Eqn (11) does not always have an explicit expression [18]. In practice, $\hat{\Lambda}$ can be obtained from the training data using Matlab's `copulafit` routine.

Our approach can be summarized in the following steps:

1. Input: Consider a set of training samples (feature vectors). Each sample is a B -dimensional feature vector (MFeatures).
2. Estimation: We first select K pairwise features from each training sample. Note that the maximum value of K is $\binom{B}{2}$. These pairs are used to construct K bivariate copulas. The copula parameter matrices are estimated for each class using the training set (see Eqn (11)). We represent this as $\hat{\Lambda}^r = \hat{\lambda}_1^r, \dots, \hat{\lambda}_K^r$, $r = 1, \dots, M$. (See below the procedure for feature selection to reduce dimensionality).
3. Transformation: Given a test sample that has to be classified, we transform it into $M - 1$ K -dimensional CFeature vectors $c(u|r)$, $r = 1, \dots, M - 1$ using the estimated parameter matrices (see Eqn (9)).
4. Classification: We concatenate these $(M - 1)$ CFeatures with the original MFeatures and input the combined feature vector (MFeatures + CFeatures) to a SVM (that has been trained using training samples) for classification.

If a test vector $z \in r$, the copula density $c(u|r)$ should be high, otherwise it should be low. For gender prediction, since there are only two classes, we have one CFeature vector. We choose pairwise associations instead of higher dimensional associations due to the consideration of the curse of dimensionality and error propagation. In practice, an

“outlier” feature should not be associated with too many other features. To further reduce the redundancy and error propagation, a d-prime method [4] is used for feature selection:

$$d'_k = \frac{\mu_k^m - \mu_k^f}{\sqrt{[(\sigma_k^m)^2 + (\sigma_k^f)^2] / 2}}, \quad (12)$$

where (μ_k^m, μ_k^f) and (σ_k^m, σ_k^f) are the mean values and standard deviations of the distributions of the k -th CFeature given male (m) and female (f) classes, respectively. The d-prime value should be high when two distributions are well separated. Only those CFeatures that are well separated between the two classes are selected. We set an empirical value $d' = 0.2$ for all CFeatures in our experiments. Note that the same procedure is used to select the CFeatures for weight prediction.

4. Experiments

The Civilian American and European Surface Anthropometry Resource (CAESAR) 1D database [1] contains 1119 male and 1250 female subjects, and 43 measurements for each subject, after removing missing data. It also contains two major attributes which we considered as ground truths: gender and weight. 500 randomly selected males and 500 randomly selected females are included in the training set and the rest are included in the test set. For statistical validation, the experiment is repeated in a Leave- T -Out manner 50 times with replacement. Here T is the size of the test set. The lib-SVM classifier [7] with RBF kernel (Eqn 13) is used for classification purposes. For gender prediction, the parameters we used are $C = 2000$ for soft margin [9], and $\sigma = 0.0001$ for the width of the basis function. For weight prediction, we used a nu-SVR regression scheme with the default parameter setting from lib-SVM ($C = 1$, $\sigma = 1/d$), where d is the number of features.

$$Ker(v_1, v_2) = \exp\left(-\frac{\|v_1 - v_2\|^2}{2\gamma^2}\right), \quad (13)$$

where v_1 and v_2 are the feature vectors, and γ is the width of the basis function.

We separate the 43 measurements into two clusters: body cluster and head cluster. The reason for such a separation is to investigate if the measurements in the body cluster and head cluster can be independently used to predict soft biometrics, and to determine their variability in prediction performance. In practice, one of them may not be available. For instance, a webcam surveillance system may not be able to capture the face (head) information, prompting the use of body information for prediction.

The features in the two clusters are non-overlapping (see Figure 8). That is, in each cluster, there is no measurement

containing information or partial information from the other cluster. For the head cluster, the sitting height and sitting eye height are further combined into a single measure calculated as (**sitting height - sitting eye height**). As a result, we have 35 measurements in the body cluster and 6 measurements in the head cluster.

We first observe how the prediction performance without CFeatures is affected by noise. Recalling that our features have been normalized in the range $[0, 1]$, we use a threshold $R = 3 \cdot \text{StdDev}$ to describe the Gaussian noise; this means 99.7% of the noise value will be in the range $(-R, R)$. We use 7 different R values: 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6. In this representation, $R = 0$ means no noise is added and $R = 0.6$ means a large Gaussian noise, which could be up to 60% of the maximum feature value or more, is added.

Figure 4 shows the results for gender prediction using MFeatures. Here, performance is measured using the misclassification rate.

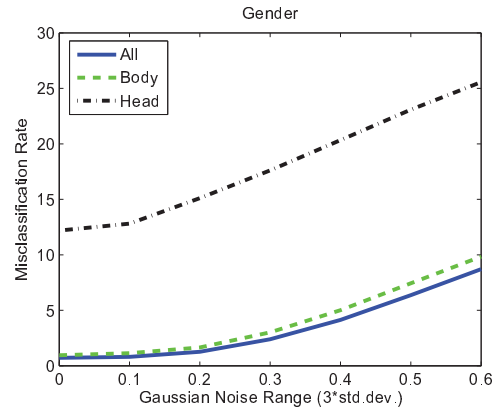


Figure 4. Misclassification rate (%) for gender prediction at various noise levels using MFeatures.

Figure 5 shows the results for weight prediction using MFeatures. The performance is measured using the mean absolute error.

Now we consider the pairwise CFeatures. In this paper, an Archimedean copula with Clayton generator is used [Eqn (10)]. Compared to the more complex student t copula, using Archimedean copula is computationally more efficient with almost the same classification performance. Compared to the Gaussian copula, the Archimedean copula yields more stable results, since the normality assumption is not always satisfied. Based on the feature selection process described previously, the approximate number of CFeatures selected for gender and weight prediction is described in Table 1. The exact number varies depending upon the training set used.

The results for gender prediction using the copula model are shown in Figure 6. Corresponding results for weight

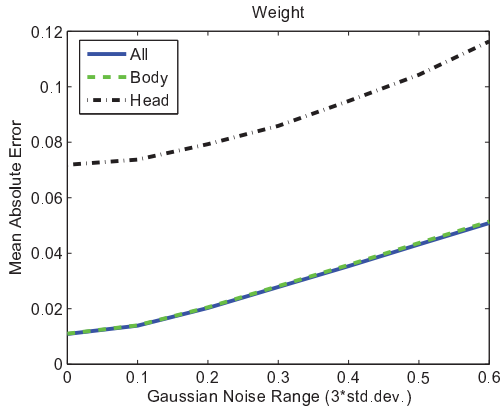


Figure 5. MAE for weight prediction at various noise levels using MFeatures.

prediction are shown in Figure 7. The results show that the CFeatures play an important role in boosting the performance, especially when the noise is severe.

Table 1. The approximate number of CFeatures selected by the proposed algorithm in our experiments. The number of MFeatures considered for “All”, “Body” and “Head” are 43, 35 and 6, respectively.

	All	Body	Head
Gender	169	165	3
Weight	220	194	2

To address the effect of the number of MFeatures in the copula model, we manually divide the 43 MFeatures into 3 categories by their measurability ranks (Figure 8). The rank 1 features are usually 1D measures and are larger compared to other features, such as stature and shoulder breadth, thus they are relatively easy to be extracted in practice. The rank 2 features are usually 2D measures such as chest or head circumference. The rank 3 features are usually hard to be extracted, such as hand circumference, triceps skinfold, or foot length. There are 25 features in rank 1, 10 in rank 2 and 8 in rank 3. We first randomly select 5 MFeatures from rank 1, and use the copula model to generate the corresponding CFeatures which are then input to the SVM. Next, we randomly select 10 features from the set of rank 1 features. We repeat this (i.e., increment number of features in steps of 5) until all 25 rank 1 features are selected. The next 5 features are then chosen from rank 2, until all 10 rank 2 features are selected. Each step is repeated 100 times with replacement for cross validation. In this experiment, we choose the first 500 males and 500 females in the database for training, and used the rest for testing. We do not use rank 3 features in this experiment. Note that for every set of MFeatures selected, the corresponding set of CFeatures are computed using the proposed algorithm. Figure 9 shows the results

under maximum noise ($R = 0.6$). We see that the performance of the copula model is affected by the number of features, especially for weight prediction. Although the copula model generally leads to a better performance, it works best when the number of features is moderate (between 10 to 25).

Feature ID	Feature Name	Measurable Rank
1	'Acromial Height, Sitting (mm)'	1
2	'Ankle Circumference (mm)'	3
3	'Spine-to-Shoulder (mm)'	1
4	'Spine-to-Elbow (mm)'	1
5	'Arm Length (Spine to Wrist) (mm)'	1
6	'Arm Length (Shoulder to Wrist) (mm)'	1
7	'Arm Length (Shoulder to Elbow) (mm)'	1
8	'Armscye Circumference (Scye Circ Over Acromion) (mm)'	2
9	<i>'Bizygomatic Breadth (mm)'</i>	<i>1</i>
10	'Chest Circumference (mm)'	2
11	'Buttock-Knee Length (mm)'	1
12	'Chest Girth at Scye (Chest Circumference at Scye)(mm)'	1
13	'Crotch Height (mm)'	1
14	'Elbow Height, Sitting (mm)'	2
15	<i>'Eye Height, Sitting (mm)'</i>	<i>1</i>
16	<i>'Face Length (mm)'</i>	<i>1</i>
17	'Foot Length (mm)'	3
18	'Hand Circumference (mm)'	3
19	'Hand Length (mm)'	3
20	<i>'Head Breadth (mm)'</i>	<i>1</i>
21	<i>'Head Circumference (mm)'</i>	<i>2</i>
22	<i>'Head Length (mm)'</i>	<i>1</i>
23	'Hip Breadth, Sitting (mm)'	1
24	'Hip Circumference, Maximum (mm)'	2
25	'Hip Circ Max Height (mm)'	2
26	'Knee Height (mm)'	1
27	'Neck Base Circumference (mm)'	2
28	'Shoulder Breadth (mm)'	1
29	<i>'Sitting Height (mm)'</i>	<i>1</i>
30	'Stature (mm)'	1
31	'Subscapular Skinfold (mm)'	3
32	'Thigh Circumference (mm)'	2
33	'Thigh Circumference Max Sitting (mm)'	2
34	'Thumb Tip Reach (mm)'	1
35	'TTR 1 (mm)'	1
36	'TTR 2 (mm)'	1
37	'TTR 3 (mm)'	1
38	'Triceps Skinfold (mm)'	3
39	'Total Crotch Length (Crotch Length) (mm)'	3
40	'Vertical Trunk Circumference (mm)'	3
41	'Waist Circumference, Pref (mm)'	2
42	'Waist Front Length (mm)'	1
43	'Waist Height, Preferred (mm)'	1

Figure 8. The measurements represented in red italics belong to the head cluster. The measurements represented in black belong to the body cluster. The others, represented in blue, do not specifically belong to the body or the head, but are overall measurements. The last column indicates the relative difficulty in the automated acquisition of the measurements (1 – easy; 2 – moderate; 3 – most difficult).

5. Discussions and Conclusions

Due to the effectiveness of using copula to characterize the interaction between variables, our model combines the characteristics of copula and SVM to boost the gender and weight prediction performance on contaminated data.

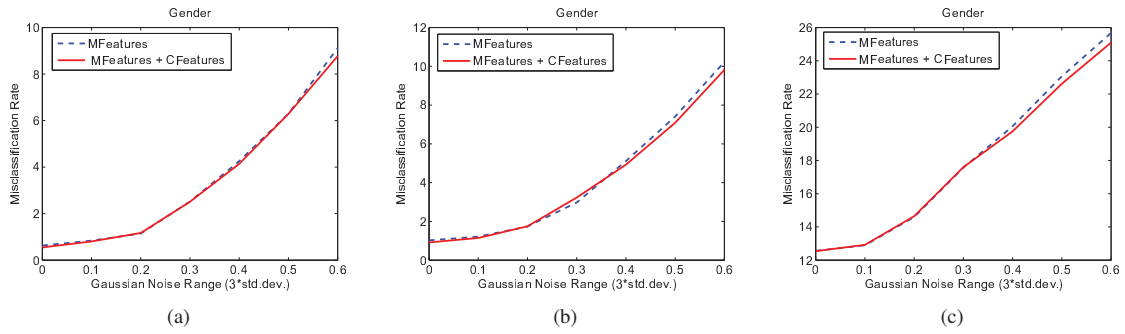


Figure 6. Gender prediction using: (a) all features; (b) body features only; (c) head features only.

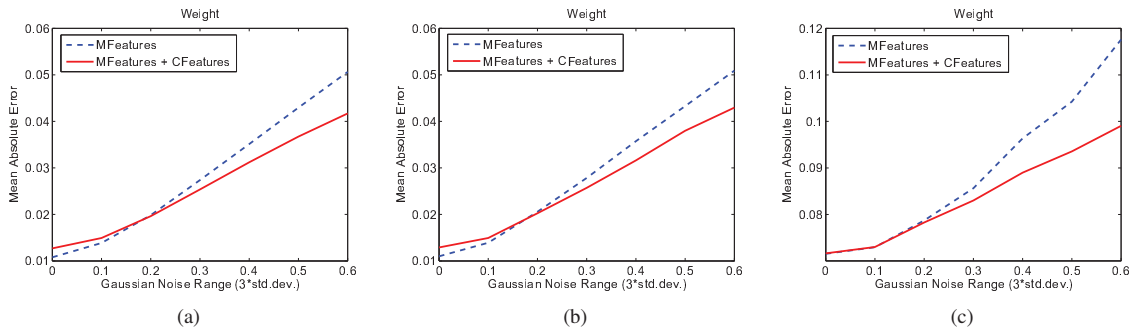


Figure 7. Weight prediction using: (a) all features; (b) body features only; (c) head features only.

Table 2. Comparison of recent studies on gender prediction

Data	#Subjects	Method	Perf.	Ref.
FERET	2,409	Raw Pixels+Adaboost	7.0%	[3]
FERET	994	SIFT+Bayesian	16.3%	[24]
LFW	7,443	LBP+Adaboost	5.6%	[22]
MUCT	276	Facial Metrology	13.2%	[4]
CAESAR	2,369	Metrology+Linear Reg.	9.0%	[2]
CAESAR	2,369	Metrology+SVM	0.6%	Ours

The results show that the body measurements (without head measurements) can provide comparable performance when using all the measures, including head measurements. The head cluster has lower performance, which is reasonable, since it only contains 6 features. When predicting with no noise or errors, we obtained a 0.63% misclassification rate for gender using all measures, a 0.93% error rate using only the body cluster, and a 12.6% error rate using only the head cluster. The above results are based on MFeatures only. Table 2 shows a comparison with other recent studies on gender prediction. Compared to previous work, precisely measured full body metrology (with no error or noise) provides a significant performance improvement in gender prediction.

For weight prediction, we have a 0.0108 mean absolute error (MAE) using all measures, a 0.0110 MAE using only body cluster and a 0.0719 MAE using head cluster only.

To our knowledge, very little prior work has been done on automated weight prediction. One related work by Velardo and Dugelay [25] showed a 5%~10% error rate on weight estimation, using 7 measurements from the NHANES [6] database. Adjeroth et al’s work [2] showed a 0.028 MAE using all measures. However, the impact of noise was not systematically considered in the previous work.

The benefit of using the copula model is rather evident. Our study shows that the impact of Gaussian noise is moderate, even under severe contamination. Note that for the head cluster, the performance drops faster than the others: gender prediction drops by 13.11% from the no-noise case to the highest noise case. Compared to the head cluster, the body cluster shows a 9.04% drop in gender prediction, while all-measures shows a 8.48% drop in gender prediction. This implies that the extra information provided by the association structure compensates for the effect of noise in the measurements.

Ignoring the minor fluctuations due to randomness, the results show that using the CFeatures has a positive impact, especially when the test set is severely contaminated. Among the three clusters, the impact of CFeatures is most obvious when using the head cluster (which has the least number of features). In our study, the Gaussian noise is “evenly” assigned to all test features. Under this condition, one can argue that the class-conditional probability density

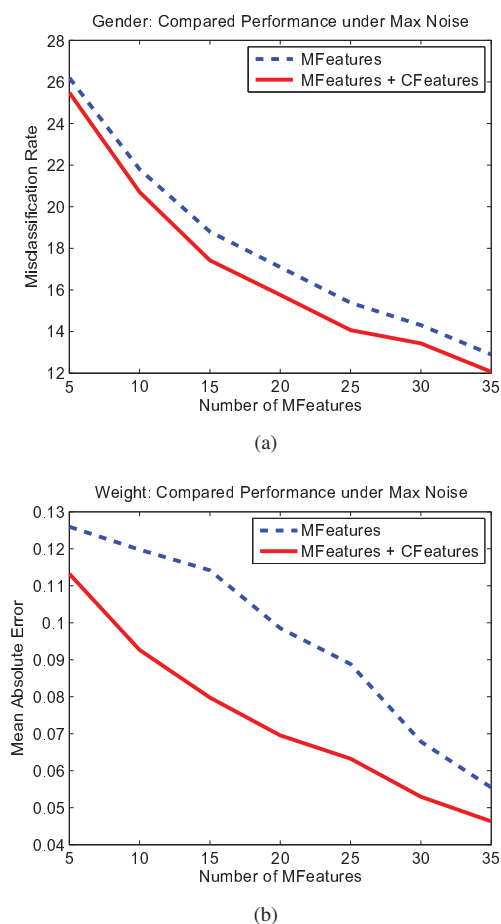


Figure 9. The performance using the copula model (MFeatures + CFeatures) under maximum noise: (a) gender; (b) weight.

is unlikely to result in incorrect class prediction unless both features are significantly shifted in the same “direction” into an adjacent class due to noise. This property makes CFeatures less sensitive to noise than MFeatures. If the error bound for certain measures are known, we could pair features having low error bounds with features having high error bounds, so that the impact of error can be further mitigated. This issue warrants a more careful further study.

Work sponsored by the Office of Naval Research (ONR) and National Institute of Justice (NIJ).

References

- [1] Civilian American and European Surface Anthropometry Resource, <http://store.sae.org/caesar/>.
- [2] D. Adjeroh, D. Cao, M. Piccirilli, and A. Ross. Predictability and correlation in human metrology. In *WIFS*, 2010.
- [3] S. Baluja and H. A. Rowley. Boosting sex identification performance. *IJCV*, 71(1):111–119, 2007.
- [4] D. Cao, C. Chen, M. Piccirilli, D. Adjeroh, T. Bourlai, and A. Ross. Can facial metrology predict gender? In *IJCB*, 2011.
- [5] L. Cao, M. Dikmen, Y. Fu, and T. S. Huang. Gender recognition from body. In *ACM MM*, pages 725–728, 2008.
- [6] Center for Disease Control. National Health And Nutrition Examination Survey. 1999–2005.
- [7] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Trans. on Intelligent Systems and Technology*, 2:1–27, 2011.
- [8] M. Collins, J. Zhang, P. Miller, and H. Wang. Full body image feature representations for gender profiling. In *ICCV workshop*, pages 1235–1242, 2009.
- [9] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, 1995.
- [10] A. Dantcheva, C. Velardo, A. D’Angelo, and J.-L. Dugelay. Bag of soft biometrics for person identification - new trends and challenges. *Multimedia Tools Appl.*, 51(2):739–777, 2011.
- [11] J. Fellous. Gender discrimination and prediction on the basis of facial metric information. *Vision Res.*, 37(14):1961–73, 1997 Jul.
- [12] E. W. Frees and E. A. Valdez. Understanding relationships using copulas. *North American Actuarial Journal*, 2(1):1–25, 1998.
- [13] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. San Mateo, CA: Morgan Kaufmann, 1990.
- [14] A. B. A. Graf and F. A. Wichmann. Gender classification of human faces. In *BMCV*, pages 491–500, 2002.
- [15] G. Guo, G. Mu, and Y. Fu. Gender from body: A biologically-inspired approach with manifold learning. In *ACCV*, 2009.
- [16] R. V. Hogg, J. W. McKean, and A. T. Craig. *Introduction to mathematical statistics*. Prentice Hall, 6 edition, 2004.
- [17] M. Kendall. A new measure of rank correlation. *Biometrika*, 30(12):81:89, 1938.
- [18] R. Kwitt, P. Meerwald, and A. Uhl. Efficient texture image retrieval using copulas in a Bayesian framework. *IEEE Trans. on Image Processing*, 20(7):2063–2077, 2011.
- [19] E. Makinen and R. Raisamo. Evaluation of gender classification methods with automatically detected and aligned faces. *IEEE Trans. on PAMI*, 30(3):541–547, 2008.
- [20] D. D. Mari and S. Kotz. *Correlation and Dependence*. Imperial College Press, 2001.
- [21] M. Rosenblatt. Remarks on multivariate transformation. *Ann. Math Stat.*, 23(3):470–472, 1952.
- [22] C. Shan. Gender classification on real-life faces. In *ACIVS*, pages 323–331, 2010.
- [23] C. Shan, S. Gong, and P. W. McOwan. Fusing gait and face cues for human gender recognition. *Neurocomput.*, 71(10-12):1931–1938, 2008.
- [24] M. Toews and T. Arbel. Detection, localization, and sex classification of faces from arbitrary viewpoints and under occlusion. *IEEE Trans. on PAMI*, 31(9):1567–1581, 2009.
- [25] C. Velardo and J. Dugelay. Weight estimation from visual body appearance. In *BTAS*, 2010.