

Selecting a Small Set of Optimal Gestures from an Extensive Lexicon

Jacob Grosek

Air Force Research Laboratories
Albuquerque, NM 87116

J. Nathan Kutz

Department of Applied Mathematics
University of Washington, Seattle, WA 98195-3925

ABSTRACT

Finding the best set of gestures to use for a given computer recognition problem is an essential part of optimizing the recognition performance while being mindful to those who may articulate the gestures. An objective function, called the *ellipsoidal distance ratio metric* (EDRM), for determining the best gestures from a larger lexicon library is presented, along with a numerical method for incorporating subjective preferences. In particular, we demonstrate an efficient algorithm that chooses the best n gestures from a lexicon of m gestures where typically $n \ll m$ using a weighting of both subjective and objective measures.

General Terms:

Gesture Recognition, Computer Vision

Keywords:

Best Gestures, Variable Selection, Optimal Gesture Lexicon

1. INTRODUCTION

Research in computer vision continues to be of great technological importance given the potentially vast impact in a wide range of applications in recognition software as well as human-computer interactions. Computer vision broadly includes mathematical methods and algorithms for acquiring, processing, analyzing, and understanding images, often which are high-dimensional, in order to produce accurate decisions and classifications about what is observed. One increasingly important branch of the computer vision field concerns gesture recognition, where computers are trained to recognize hand signals, facial expressions, and/or eye movements in order to better interface and interact with humans [1, 2]. For many software applications that require gesture detection, it is often the case that only a few gestures are needed for controlling the software or program from among a nearly endless number of gestures that a person can articulate. Thus given a large lexicon of gestures and an application that requires only a small subset, one would like to know which are the best gestures to choose for the given application. More succinctly, we develop an algorithm that chooses the best n gestures from a lexicon of m gestures ($n \ll m$). There are many factors, both objective and subjective, that determine which gestures are most appropriate for a particular application, i.e. best gestures. Certain gestures may be easier for the computer to recognize, but may not necessarily be comfortable or

suitable for humans to articulate. Ergonomics, or the ease of articulating the various gestures, and relations to physical signs and gestures that are already in use in the culture or that are appropriate for the given application constitute *subjective* measures for claiming some gestures may be better than others. The *objective* reasons for ranking the quality of gestures comes purely from the computer's ability to distinguish and recognize gestures in a statistical sense. Both the subjective and objective reasons for determining the n best gestures from a lexicon with m elements need to be considered for designing gesture-based, robust software interfaces.

As an example, consider replacing a standard computer mouse with a set of articulated gestures using the computer's built-in camera. The gestures would mimic (i) mouse movement, (ii) a one-click button, (iii) a two-click button, (iv) a scroll-up function and (v) a scroll-down function. Thus the required lexicon set for these basic mouse functionalities gives $n = 5$. The potential lexicon library for executing these functional behaviors is tremendously large ($m \gg 5$). Limiting our lexicon to the sign language alphabet gives $m = 26$. Thus we would want to identify the five best sign language gestures that give the most robust statistical classification. Such applications are already finding their way into the consumer electronics markets.

Objectively, the best lexicons of size n are determined by having the computer attempt to recognize every combination of n out of m gestures, and then choosing the best gesture set as the one with the highest overall successful recognition rate during a training process [3]. This can be a time-consuming and combinatorially challenging problem. In this paper a new metric, called the *ellipsoidal distance ratio metric* (EDRM), is introduced that provides an excellent indication as to which gestures will be easily recognizable. The EDRM is applied to the feature space of the gestures, and therefore doesn't require the computer to complete the entire recognition process in order to gain some notion as to which are the best gestures in the entire lexicon of available gestures [4]. Thus a robust and efficient algorithm is developed to extract the best n of m gestures.

2. GESTURE REPRESENTATION AND FEATURES

In computer vision, the gesture recognition process starts when a raw image is imported into the computer. The gesture recognition process can be broken down into the following steps:

- (i) *gesture detection, and segmentation*, which is the process by which the gesture is found and isolated in the image frame,
- (ii) *pre-processing*, which normalizes like-gestures to similar sizes, shapes, colors, positions, and orientations,

- (iii) *feature selection*, which determines important characteristics and aspects of gestures that will simultaneously distinguish between different gesture classes and will highlight like-gesture classes (e.g. principal component analysis), and
- (iv) *statistical learning and classification*, which trains the computer to identify a gesture's articulation by statistical means and predictive functions that draw upon feature data.

Not all recognition schemes use all of these steps, or in this order. However, the general procedure is still valid [2].

2.1 Feature Space

Most gesture recognition algorithms eventually represent the gestures as a sequence of features, which can be thought of as points in a feature space (See Fig. 1). This feature space provides a great deal of information about how similar or different the computer views the gestures. However, it is very common that most gesture recognition algorithms extract many features in order to better their recognition performance, and this corresponds to having high-dimensional feature spaces, which are usually impossible to visualize. The high-dimensional nature of feature spaces invokes many of the issues that accompany the *curse of dimensionality* [5]. Figure 1 provides a visual example of what a feature space may look like; in this case, it is in 3 dimensions, which corresponds to having three features extracted from every image. A common way to extract features is through, for instance, a principal component analysis, but there are many other techniques available. Note that gesture class **0** in Fig. 1 is well-separated from the other four classes in feature space. Also, the feature points of this class are well-clustered such that there does not appear to be any images with feature points that are outliers to the other images of this gesture class. Well-clustered feature points are a result of extracting consistent feature values from like-gestures of the same class.

Gesture classes **a** and **e** illustrate a common occurrence with many gestures that look similar in that these classes overlap with one another in feature space. Obviously, similar gestures are likely to produce similar feature values, and therefore will lie near one another in feature space. Also, note that classes **a** and **e** are not well-clustered, and have many outliers to their apparent centers. A linear separatrix, defined by $y = (2/5)x + 9/5$, can be found by statistical learning techniques in order to distinguish these two classes in some optimal sense. However, it is evident that there isn't a perfect way of separating these classes without likely *over-fitting* the data. Gesture classes **u** and **v** of Fig. 1 exemplify another possibility in the recognition process in that the classes are separated from one another in feature space, but the boundary between the two class regions is defined by some nonlinear function. In this case, a nonlinear separatrix, defined by $y = (2/5)x + \cos(4x)/4 + \cos(4z)/4 + 6$, distinguishes the two classes in some optimal sense. One can easily imagine that without enough data points, assuming they are even available, and/or without any *a priori* idea as to how these classes ought to be separated, an algorithm may struggle to learn the nonlinear separatrix, although this may be possible with, for instance, adaptive boosting methods [3]. It is unlikely for gesture recognition that one will be able to get any useful *a priori* information as to how the features of any class will appear in feature space relative to another gesture class, especially for high-dimensional data.

Therefore, it would be ideal if all gesture classes were comprised of well-clustered feature points, and if they were all well-separated from one another into distinct regions of the feature space. One of the best ways of controlling how the gestures are rendered in feature space is by choosing excellent features in the first place. A quality feature selection algorithm will find features that remain

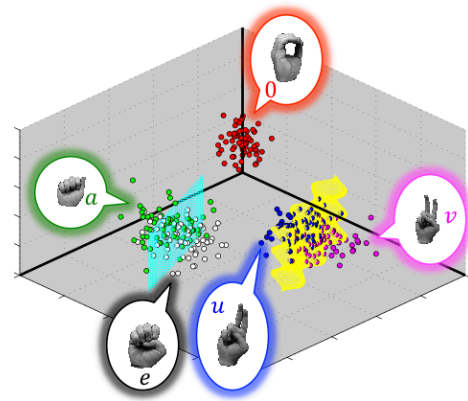


Fig. 1: This figure exemplifies how gestures may appear in feature space. Gesture classes **a** (green) and **e** (white), which overlap in feature space, are separated by a linear separatrix (cyan) that attempts to distinguish the classes in some optimal sense. Gesture classes **u** (blue) and **v** (magenta), which do not overlap in feature space, are separated by a nonlinear separatrix (yellow). Gesture class **0** is well-separated from the other classes.

consistent for like-gestures and yet are very distinct for different gestures. In this paper, feature selection will be accomplished by either principal component analysis (PCA) [6, 7] or by generalized projections (GPs) [8, 9].

As a final note, Fig. 1 already highlights some of the key ideas of this work. In particular, if $n = 3$ gestures were required for an application from the $m = 5$ lexicon shown, then one would undoubtedly (and based upon simple intuition alone) choose **0** along with one of either **a** or **e** and one of either **u** or **v**. Thus all $n = 3$ selected gestures would be well-separated for recognition and decision making in the given application.

2.2 Best Features

A common practice in the recognition field is to extract as many features as possible and allow the statistical learning algorithm to diminish the relevance of the features, usually by controlling weightings, which do not strongly effect the decision making process of the classification. However, in order to determine the best gestures, this practice may be detrimental because weak features are either inconsistent (noisy) for like-gestures of the same class or they are consistent across different classes. This compromises the chances for like-gestures to be well-clustered in feature space and yet well-separated from different gesture classes. Additionally, by including more features, the dimension of the feature space is increased, making it harder for both the human and computer to understand how the different gesture classes are separated in feature space. Thus, for the search of the best lexicons of size n , it is better to filter out the weak features from the analysis.

Currently, there are some tests designed to evaluate the efficacy of a given feature on the entire recognition process [4]. Some examples include the Fisher score [10], the Generalized Fisher score [11], mutual information [12], ReliefF [13], the Laplacian score [14], the Hilbert Schmidt Independence Criterion [15], the Trace Ratio Criterion [16], the Multi-Layer Perceptron Sensitivity Method [17], and Principle Feature Analysis [18]. Another variable selection routine that can help determine which features have the greatest impact on the recognition performance is what termed here as the *feature selection weakness* (FSW). This measure assesses how close

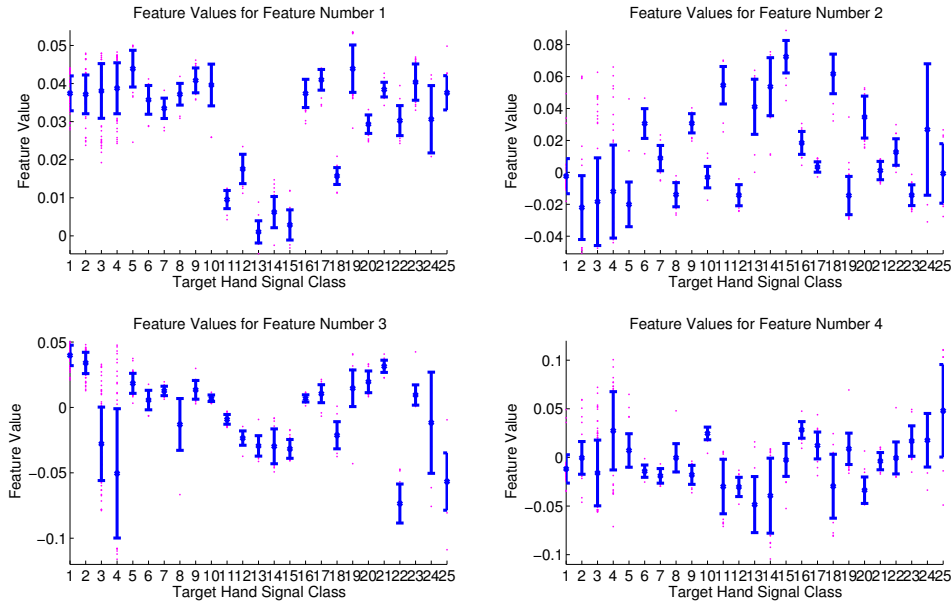


Fig. 2: The values for the first 4 PCA features of a recognition problem with 25 gesture classes. The average and one standard deviation spreads (blue) of the feature values (magenta) for each class are plotted. Note the smaller spreads in the feature values of feature 1 as compared to feature 4, which has larger error bars. Also, note how, for feature 1, some classes {11 – 15, 18} are well-distinguished from the rest of the classes. But, for feature 4, the classes are not as well-separated. These attributes are reflected in the FSW values shown in Fig. 3.

points of the same class/type/variety are to one another against how far points of different classes/types/varieties are separated from one another. This is similar to the Fisher score. Like the Fisher score, FSW treats the features independently and not in combination.

Assume that there are L features extracted from each image and M gestures in total, there being N_m images of the m^{th} gesture, where $m \in \{1, 2, \dots, M\}$. Therefore, the total number of images in the dataset is given by $\sum_{m=1}^M N_m$. Let X^ℓ be a data structure which contains all of the feature values for feature ℓ , where $\ell \in \{1, 2, \dots, L\}$, such that X^ℓ has M columns, and the m^{th} column has N_m rows. The element in the i^{th} row and m^{th} column of this data structure will be denoted by X_{im}^ℓ . The average among the feature values for feature ℓ and for gesture class m is given by $\mu_m^\ell = \frac{1}{N_m} \sum_{i=1}^{N_m} X_{im}^\ell$. The mean of these averages across all gesture classes is given by $\bar{\mu}^\ell = \frac{1}{M} \sum_{m=1}^M \mu_m^\ell$. The variance among the features values for feature ℓ and for gesture class m is given by $(\sigma^2)_m^\ell = \frac{1}{N_m - 1} \sum_{i=1}^{N_m} (X_{im}^\ell - \mu_m^\ell)^2$. With this notation, the FSW for the ℓ^{th} feature can be defined as

$$\text{FSW}(\ell) = \frac{\frac{1}{M} \sum_{m=1}^M [(\sigma^2)_m^\ell]}{\frac{1}{M-1} \sum_{m=1}^M (\mu_m^\ell - \bar{\mu}^\ell)^2}. \quad (1)$$

Upon closer inspection, the numerator of the FSW value is the mean, across all classes, of the variances in each class. And the denominator of the FSW value is the variance, across all classes, of the means in each class. If the variance in the feature values within any given class is small ($(\sigma^2)_m^\ell \sim \text{small}$), then it means that feature ℓ produces consistent values for each class, which was already noted to be a beneficial attribute for the recognition process. When there is large variation between the average feature values in each class (μ_m^ℓ), then these classes are well-separated by feature ℓ . Figures 2 and 3 show how FSW values relate to the actual feature value

data collected from a feature selection algorithm. If a given feature (ℓ) receives a large FSW value, then this is a weak feature. The contrapositive is that the better features have smaller FSW values.

3. GESTURE SEPARATION MEASURES

One major problem in the recognition field is *over-fitting*. Over-fitting occurs when a statistical learning algorithm learns its training set too well, and thus is able to distinguish between the training set gesture classes with excellent accuracies, but at the cost of not being able to maintain those accuracies in a generalized setting. Over-fitting can occur by simply not having a diverse enough training set, but also because of the complexity and over-sophistication of the learning algorithm. One might think of finding an incredibly complicated separatrix, like the one in Fig. 4 (B), that perfectly distinguishes two classes in the training set, but does not do as good of a job distinguishing the test data which may actually have overlapping features from the two classes.

For most recognition applications, one cannot determine how the different classes will render themselves in feature space. And since many sophisticated statistical learning algorithms need some *a priori* information in order for the algorithm to learn a nonlinear separation measure, one is often left to guess-work or very specific problems and applications where *a priori* information is known.

A good statistical learning algorithm may also be able to learn a nonlinear separation between classes if there are enough training data points available, especially near and well-distributed around the nonlinear separation. However, not every application can provide a sufficient number of training data points, nor would it be easy to guarantee that those points would be near and well-distributed around the critical areas of the feature space. This problem is further exacerbated by having high-dimensional data, which multiplies the number of data points needed in order to properly learn

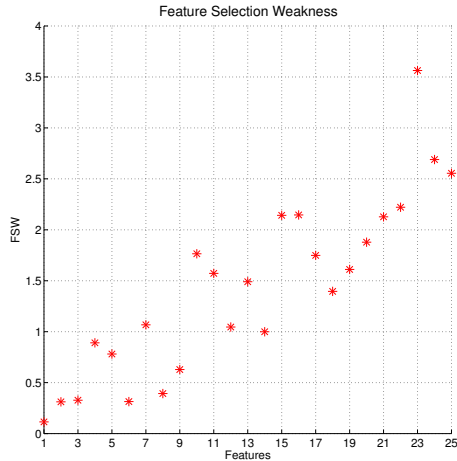


Fig. 3: FSW values plotted for the 25 features using PCA feature selection. The best features are extracted first, and then the features tend to get progressively weaker as more are extracted. The first 4 FSW values come from the feature values plotted in Fig. 2.

complicated class separations as the dimension increases. This is another reason for leaving out weak features if possible. In some recognition applications, especially when there are only a few gestures needed, it may be acceptable to have gesture classes that are sparse in feature space, as long as the classes are separated from one another. However, in general, sparse classes tend to be harder to separate than well-clustered classes [4]. Also, as more gestures are needed for the given application, it becomes more likely that a sparse gesture class will overlap in feature space with another class, making them difficult to distinguish.

3.1 Ellipsoidal Distance Ratio

In order to create an objective metric that captures the desirable feature space characteristics for the recognition process of computers, it makes sense to reward gesture classes that are distributed in ellipsoidal-like shapes in feature space. When a gesture's features lie in an ellipsoidal distribution in feature space, they're likely to be both well-clustered and easily separated from other classes. The *ellipsoidal distance ratio metric* (EDRM), to be introduced presently, is one possible metric for gesture class separation that takes advantage of rewarding optimal class feature space distributions. The distinctions between the terms ellipse, ellipsoid, and hyper-ellipsoid are ignored, and the general term ellipsoid is used.

To calculate the EDRM between two classes we find the ellipsoid in the appropriate dimension which best encompasses the feature points for each gesture class, even, of course, if those points are not in a true ellipsoidal distribution. This is accomplished by finding the positive definite matrix E for each class of the general ellipsoid equation: $(\vec{x} - \vec{x}_0)^* E^{-1} (\vec{x} - \vec{x}_0) = 1$, where $(*)$ is the Hermitian transpose, \vec{x}_0 is the center of the ellipsoid, \vec{x} defines a point on the surface of the ellipsoid, and the eigenvectors and eigenvalues of the matrix E are the principal directions of the ellipsoid and the squares of the semi-axis lengths, respectively. The centers of the ellipsoids are the centers of the feature points for each class.

Assume that the feature data for a given gesture m is in a data structure Y^m such that Y^m has L columns, which corresponds to the number of features extracted from the images, and the l^{th} column has N_m rows, which corresponds to the number of images in ges-

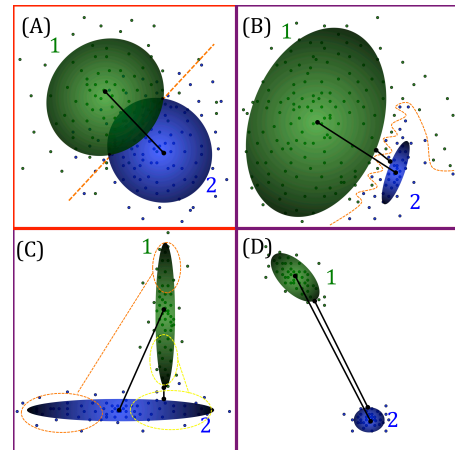


Fig. 4: Four examples that illustrate why ellipsoidal distance ratios (EDRs) are good measures of class separation. The EDR is the shortest distance between two ellipsoids divided by the distance between the ellipsoid centers (black lines). When two classes overlap (A), so do their corresponding ellipsoids; thus, the EDR is zero. Even when two classes have non-ellipsoidal distributions of feature points and are separated by a nonlinear separatrix (B), the EDR still provides a reliable measure metric for closeness. The EDR measure isn't easily fooled when many points from the two classes are well separated, but some are not (C) - compare the distances between features points in the upper half of class 1 and in the left half of class 2 to the distances using the lower half of class 1 and the right half of class 2.

ture class m . Note that the number of features extracted L is also the dimension of the feature space. By the properties of the *singular value decomposition* (SVD), the ellipsoidal matrix for gesture m is $E_m = (Y^m)^* \cdot Y^m$. This works out because the SVD states that any matrix $Z \in \mathbb{R}^{q \times r}$ can be decomposed such that $Z = U \Sigma V^*$, where $U \in \mathbb{R}^{q \times q}$ is unitary, $\Sigma \in \mathbb{R}^{q \times r}$ is diagonal, and $V \in \mathbb{R}^{r \times r}$ is unitary. This implies that $Z \cdot Z^* = U \Sigma^2 U^*$, which is the eigendecomposition of $Z \cdot Z^*$. So if the ellipsoid matrix $E = Z \cdot Z^*$, then the principal directions of the ellipsoid are found in the columns of U and the corresponding squares of the semi-axis lengths are along the main diagonal of Σ^2 . Letting $Z = (Y^m)^*$, the ellipsoidal matrix E is guaranteed to have the proper dimensions ($L \times L$) and proper attributes that best fit the data of feature m into an ellipsoid. At this point one has the option of rescaling the size of the ellipsoid for each gesture class. This somewhat arbitrary rescaling can have significant effects on the EDRM. Clearly, it is best to choose a consistent rescaling scheme across all the gesture classes in order to maintain fair comparisons of the class sizes and distances in feature space. It is suggested that each ellipsoid's surface be a certain percentage of the distance from the center of the ellipsoid to the most extreme outlier feature point in that gesture class. In this way, it is guaranteed that at least some feature points from each class are on the outside of the ellipsoid of its own class. This method acknowledges the reality that there are almost always impure articulations (outliers) of any gesture that may not fully represent the class. The actual percentage used in this convention can be determined based on how pure the gesture images are for each class. For instance, if all the images used in the dataset are of well-articulated gestures, then one might choose to have the ellipsoid encompass most of the feature points; this is assuming there are few to no outliers, and so a high percentage value is used, e.g. 90% – 100%. Whereas, a noisier sampling of the gestures might assume the existence of outliers

in each class, and so a lower percentage of 60% – 80% is used. In what follows, the convention will be set at 65%.

The *ellipsoidal distance ratio* is the shortest distance between two ellipsoids compared against the distance between the ellipsoid centers. Measuring the straight-line distance between ellipsoid centers is trivial. However, the shortest distance between two ellipsoids needs to be formulated as an optimization problem. Suppose there are two classes, **a** and **b**, with ellipsoid surface points \tilde{x}^a and \tilde{x}^b , centers \tilde{x}_0^a and \tilde{x}_0^b , and ellipsoid matrices E_a and E_b , respectively. The optimization problem can then be stated as follows

$$\begin{aligned} \arg \min & \quad \sqrt{(\tilde{x}^a - \tilde{x}^b)^* \cdot (\tilde{x}^a - \tilde{x}^b)} \\ \text{subject to} & \quad (\tilde{x}^j - \tilde{x}_0^j)^* E_j^{-1} (\tilde{x}^j - \tilde{x}_0^j) = 1 \quad j = \mathbf{a}, \mathbf{b} \end{aligned} \quad (2)$$

This problem effectively says that one wants the shortest distance between the points \tilde{x}^a and \tilde{x}^b , while constraining \tilde{x}^a to be on the surface of the ellipsoid of gesture **a** and constraining \tilde{x}^b to be on the surface of the ellipsoid of gesture **b**. Without any loss of purpose or accuracy, the optimization problem can be better formulated for numerical and algorithmic stability issues such that

$$\begin{aligned} \arg \min & \quad \left[(\tilde{x}^a - \tilde{x}^b)^* \cdot (\tilde{x}^a - \tilde{x}^b) \right] / 2 \\ \text{subject to} & \quad \left[(\tilde{x}^j - \tilde{x}_0^j)^* E_j^{-1} (\tilde{x}^j - \tilde{x}_0^j) - 1 \right]^2 = 0 \quad j = \mathbf{a}, \mathbf{b} \end{aligned} \quad (3)$$

Since, in high-dimensions, these ellipsoids cannot be easily visualized, one needs to take extra precautions to ensure that the ellipsoids do not overlap in feature space. Given the optimal points \tilde{x}^a and \tilde{x}^b , one check that can be made is to ensure that

$$\left(\tilde{x}^b - \tilde{x}_0^a \right)^* E_a^{-1} \left(\tilde{x}^b - \tilde{x}_0^a \right) > 1 \quad \left(\tilde{x}^a - \tilde{x}_0^b \right)^* E_b^{-1} \left(\tilde{x}^a - \tilde{x}_0^b \right) > 1. \quad (4)$$

These inequalities verify that the optimal point on ellipsoid **b** is not inside of the ellipsoid for gesture **a**, and vice-versa. Also, due to numerical inaccuracies it can be good to verify that

$$\left| 1 - \left(\tilde{x}^j - \tilde{x}_0^j \right)^* E_j^{-1} \left(\tilde{x}^j - \tilde{x}_0^j \right) \right| < \epsilon \quad j = \mathbf{a}, \mathbf{b} \quad (5)$$

where ϵ is some small tolerance for numerical inaccuracy, which may need to grow with the number of dimensions of the problem. These inequalities ensure that the optimal point for ellipsoid **a** is truly on the surface of ellipsoid **a**, and the same for ellipsoid **b**. Once these points on the two gesture class ellipsoids that minimize the distance between the ellipsoids have been found, one can calculate the *ellipsoidal distance ratio metric* as follows

$$\text{EDRM}_{\mathbf{a}, \mathbf{b}} = \frac{\sqrt{(\tilde{x}^a - \tilde{x}^b)^* \cdot (\tilde{x}^a - \tilde{x}^b)}}{\sqrt{(\tilde{x}_0^a - \tilde{x}_0^b)^* \cdot (\tilde{x}_0^a - \tilde{x}_0^b)}}. \quad (6)$$

The EDRM values are always between 0 and 1. Larger EDRM values indicate that the two classes that are being compared are further separated and better clustered in feature space. Also note that the EDRM normalizes the scale of the ellipsoids in that larger ellipsoids must be more separated from one another in order to have the same EDRM score as smaller ellipsoids. If the volumes of both ellipsoids go to zero, because of the clustering of feature points, then the EDRM approaches its ideal value of 1 because the shortest distance between the ellipsoid surfaces (numerator) becomes the same as the distance between the ellipsoid centers (denominator).

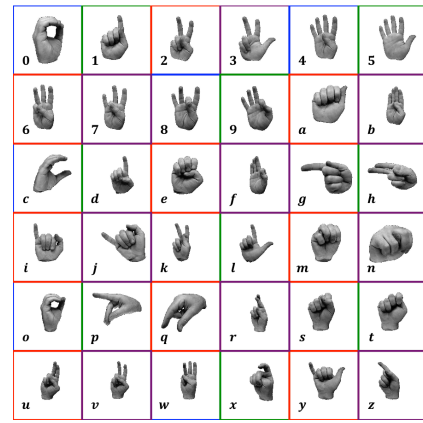


Fig. 5: The ASL static hand gesture dataset from Massey University [19] is a lexicon of 36 gesture classes, labeled 0 – 9, then a – z.

Figure 4 illustrates why the EDRM is a good metric for gesture separation. First note that when feature points from the two given classes significantly overlap (A), the EDR value will always be zero because the class ellipsoids will also overlap and so the distance between the ellipsoid surfaces is zero. In the case that one or both feature point regions for the two classes have non-ellipsoidal distributions (B), the best-fitting ellipsoids to the class regions are determined. Further, in the case that these regions are still fairly well-separated, a EDRM score can be calculated that indicates how well-separated the regions are. In some instances there can be sub-regions of the two classes that are well-separated in feature space, and yet other sub-regions of the two classes in feature space that are nearby each other (C). In this scenario, the EDRM will produce a practical value that accounts for the different sub-regions and their respective separations. The EDRM is neither misled into being too high nor too low by the positioning of the sub-regions. Ideally, the two classes are well-clustered and well-separated (D).

3.2 Subjective Separation Measures

For many applications, the comfort and ease of articulating a gesture can be an important factor for determining which gestures are more appropriate for use. Moreover, some gestures better suite an application because of the cultural norms and meanings of the gesture already in use in society. In order to incorporate these ergonomic and vernacular reasons into the best gestures decision making process, one would need to subjectively rate all the gestures in the entire lexicon of available gestures. One method of creating a subjective measure for ranking the suitability of the gestures that is comparable with the EDRM, is by first rating the ergonomic and vernacular quality of each gesture on a scale of [0, 1]. A rating of 0 would mean *not suitable* for the given application, and a rating of 1 would mean *most suitable* for the task. The fact that this rating scale is the same as the range of EDRM scores is not coincidental. Note that the EDRM scores are comparisons between two classes, therefore the subjective ratings must also be a comparisons between two classes in order to be able to integrate this subjective measure with the EDRM. This is accomplished by creating a *subjective measure* (SM) for every class pairing that is the average of the numerical ratings that each class in the pair was given.

$$\text{SM}_{\mathbf{a}, \mathbf{b}} = \frac{\text{Rating}(\mathbf{a}) + \text{Rating}(\mathbf{b})}{2} \in [0, 1] \quad (7)$$

Table 1. : Lexicon of 4 gestures, **a** – **d**, receive subjective ratings (in parentheses) which are used for subjective measures (SM) between class pairings.

	a (0.9)	b (0.2)	c (0.7)	d (0.5)
a (0.9)	-	0.55	0.80	0.70
b (0.2)	0.55	-	0.45	0.35
c (0.7)	0.80	0.45	-	0.60
d (0.5)	0.70	0.35	0.60	-

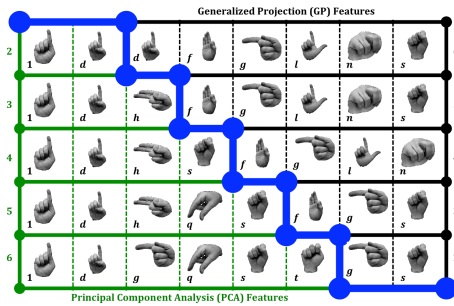


Fig. 6: Best lexicons of sizes $n = 2 - 6$ using the PCA (green) and GP (black) feature extraction methods. The lexicons of size n are read across each row; the blue line separates the two feature selection methods. Here $\alpha = 0$ and subjective considerations are ignored in the total measure (TM).

Table 1 is a four class lexicon example of how subjective measures are calculated. Note that we subjectively generated the values of this table. Ultimately, one might refine this process by having a large sample of people vote on their favorite gestures, thus generating average scores of gesture likability across a population. The subjective measure and the EDRM for each class pairing can be combined into a *total measure* (TM) by a weighting factor $\alpha \in [0, 1]$, which determines the relative influence that the subjective measure will have on the total measure.

$$TM_{a,b} = \alpha SM_{a,b} + (1 - \alpha) EDR_{a,b} \in [0, 1] \quad (8)$$

If any gesture must be included in the best lexicon of size n , then the problem becomes a matter of finding the best lexicon of size $n - 1$, where the n^{th} gesture is already determined. In the case that one of p ($p \ll n$) gestures must be included, in the best lexicon of size n , then this constraint can be enforced in the programming of the algorithmic search process. Finally, if there is a gesture or multiple gestures that should never be chosen to be a part of the best gestures set, then these gestures can simply be removed from the entire lexicon of available gestures and from consideration.

3.3 Best Gestures

Once one has calculated the total measure (TM) for every gesture class pairing in the entire lexicon of available gestures, determining the best gesture subset is a fairly straight-forward process. Assume there are a total of m available gestures in our lexicon library, and one desires to know what is the best lexicon of size n , where $n < m$. The best lexicon of size 2 is simply the gesture class pairing with the largest TM value. In order to calculate the best lexicon of size $n \geq 3$, one must first find all the unique combinations of n gestures in the set of m possible gestures. Then, the sum of the TMs from each pairing in every unique combination of n gestures

is calculated. The unique grouping of n gestures with the highest total sum TM value constitutes the best lexicon of size n .

Note that this problem is *NP hard* because there are $\binom{m}{n}$ unique combinations of size n , and $\binom{n}{2}$ pairings within each combination of n . Therefore, determining the best gesture set of size $n \geq 3$ requires at least $\binom{m}{n} \binom{n}{2}$ operations, which grows unsustainably with m , and the problem becomes intractable. However, the process is still manageable when n is small relative to m .

Note that the subjective metric (SM) is set up in such a way that the best lexicon of size n will always include the best lexicon of size $n - 1$; thus, one only needs to find the best gesture to add to the $n - 1$ already found. This makes the problem much more tractable. However, the objective EDRM does allow for situations where the best lexicon of size n does not include the best lexicon of size $n - 1$. This can occur for a variety of reasons including cases where there are two similar gestures, which of course are nearby in feature space, and one of them is better separated from a few other gestures. But, as the number of gestures included in the best gesture set increases, the other gesture of the two like-gestures becomes the better choice. Another example is when two different gestures are very well-separated from one another in feature space; however, one of the two gestures is fairly close to the other available gestures in the entire lexicon. In this case, the best lexicon of 2 is determined by the two aforementioned well-separated gestures, but the best lexicons of $n \geq 3$ will not include the gesture that is close to the rest of the available gestures. In any case, assuming that the best lexicon of size n includes the best lexicon of size $n - 1$ simplifies the search process, but at the cost of potentially not finding the true best lexicon of size n .

4. EXAMPLE AND RESULTS

In order to test this process for determining the best gesture sets, we will apply this method to publicly available hand gesture sets.

4.1 Experiments and Data

It is imperative for the purposes of testing and reproducibility to work with a public domain, archival dataset of static hand gestures. A variety of gesture articulation datasets are available, but many have different purposes than the present. Specifically, some are for demonstrating recognition of hand gestures from noisy backgrounds. Others have only a limited lexicon for performing such tasks, i.e. 10 or less hand gestures. In our example, we would like a hand gesture set with many examples of each gesture along with a large range of articulated gestures. One data set serving the purposes of this work is available: the Massey University ASL static hand gesture dataset [19] of 36 hand gestures, which has 70 renditions of each gesture class, except for class **t** which has 65 renditions (See Fig. 5). The hands have already been isolated within each image and segmented with black backgrounds. The raw images are pre-processed by converting to grayscale, centering the hand within the frame, down-sampling the images to 32×32 (without changing the aspect ratio of the hand), and normalizing the brightest pixel intensity [20]. As has been stated, feature selection is done by the PCA method and the GP method, both only using 10 features. Statistical learning will be done by the LDA method [21], with a one-vs-the-rest classification style [20]. The training set consists of 20 randomly chosen images from each gesture class, and the test data consists of the remaining 50 images (45 for class **t**) from each class. Note that the results were cross validated by performing this random selection of 20 and 50 images repeatedly (more than 1000 times). The results are highly consistent under cross-validation,

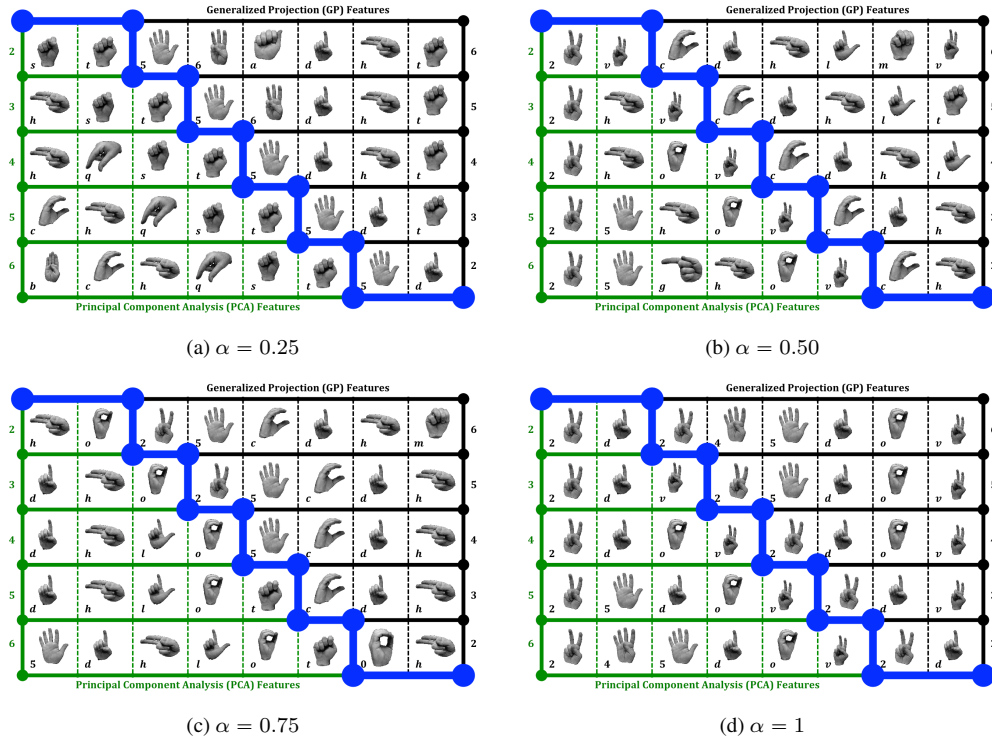


Fig. 7: Best lexicons of sizes $n = 2 - 6$ using the PCA (green) and GP (black) feature extraction methods. The lexicons of size n are read across each row; the blue line separates the two feature selection methods. The relative influence of subjective considerations are weighting by the factor α . When $\alpha = 1$, the total measure (TM) comes solely from the subjective measure (SM), and is not swayed by the EDRM.

showing only slight changes in performance from one realization to the next, i.e. the variability is negligible. This validates the methodology and shows it to be highly robust under experimentation.

4.2 Best Lexicon Results

For the case where there are no subjective considerations to be incorporated into the decision making process ($\alpha = 0$), the computer will rely solely on the objective EDRM to determine the best lexicon of size n . Figure 6 illustrates the results of this experiment using both the PCA (left) and GP (right) feature selection methods. Because these methods produce profoundly different features, the best gesture sets that are determined are also different.

Of course, the true best gesture sets, without subjective considerations, are those which produce the highest recognition rates from the given algorithms. Recognition accuracy can take on two forms: (1) the *within class success rate*, meaning when images are correctly labeled within their respective class, and (2) the *out-of-class success rate*, meaning when images are correctly not labeled to belong to classes to which they do not belong. By counting the number of times each image is correctly labeled within its own class and not labeled to belong to another class, the *within class* and *out-of-class* success rates can be calculated and then averaged together to create an overall average recognition rate. By finding this overall average recognition rate for every possible combination of n gestures, and picking the maximal rate, one thus discovers the corresponding true best gesture set of size n .

Using this prescribed method, all of the best lexicons of size n in Fig. 6 that were determined by the EDRM have been confirmed to also be the true best lexicons. For all of the different cases presented in Fig. 6, there are multiple lexicons of size n that tied for being the best gesture set according to the average recognition rates (only one is shown), which isn't surprising for small lexicons.

One caveat to mention is that the LDA statistical learning method relies on Fisher linear discriminants, which attempt to group like-classes while separating different classes in a projection from feature space onto a line. Since the EDRM rewards pairs of classes that are both well-clustered and well-separated in feature space, the LDA method is more likely to produce good recognition rates on the same classes that bode well under the EDRM. Other statistical learning algorithms may produce slightly different best lexicons, but are still very likely to rate the best lexicons found using the EDRM very highly, for the reasons mentioned in Section 3.

One can incorporate subjective considerations into the TM metric as was described in Section 3.2. Table 2 lists an example of possible subjective rankings given to each gesture for the purposes of this experiment. Using these rankings, the SM for each class pairing is calculated and incorporated into the TM according to the weighting $\alpha > 0$. Figure 7 displays the best lexicons of sizes 2-6 for (a) $\alpha = 0.25$, (b) $\alpha = 0.50$, (c) $\alpha = 0.75$, and (d) $\alpha = 1$.

5. CONCLUSIONS AND OUTLOOK

Finding a limited number of best gestures in a large lexicon (library) of gestures for a particular software application is of grow-

Table 2. : The subjective rankings given to each gesture in the Massey University lexicon are listed, from which the measures are calculated.

0 → 0.97	1 → 0.10	2 → 1.00	3 → 0.30	4 → 0.95	5 → 0.96
6 → 0.90	7 → 0.45	8 → 0.55	9 → 0.50	a → 0.85	b → 0.85
c → 0.95	d → 0.99	e → 0.85	f → 0.60	g → 0.90	h → 0.95
i → 0.70	j → 0.20	k → 0.10	l → 0.95	m → 0.90	n → 0.85
o → 0.97	p → 0.75	q → 0.70	r → 0.45	s → 0.95	t → 0.90
u → 0.95	v → 0.98	w → 0.90	x → 0.55	y → 0.30	z → 0.30

ing important for consumer electronics and computer vision applications. To have broad technological appeal, such optimal selection must not only be accomplished through an objective algorithm, but must include consumer preference, or subjective constraints, as well. The manner in which feature data from the gesture images lie in feature space provides excellent clues as to how well a general recognition algorithm will perform. As is expected, well-clustered features of the same class that are also well-separated from the features of other classes are ideal for consistent and accurate gesture recognition by computer algorithms. However, many features are not well separated, and more sophisticated techniques, like those outlined here, play a key role in providing an objective and subjective measure of the best gesture selection process. Specifically, it is best to remove weaker features and select the best features when attempting to find the best gesture set of size n . The *feature selection weakness* (FSW) is one measure well-suited for determining which features will be well-clustered and well-separated in feature space. The *ellipsoidal distance ratio metric* (EDRM) is a measure devised to reward optimal feature clusterings in feature space by assigning a value to each gesture class pairing that accounts for their respective separation compared to their individual clustering sizes. This measure has been shown to match the true average recognition rate results that decide which combination of n gestures is best, which are obtained by actually completing entire recognition process. The EDRM value only needs the feature data to provide a well-educated guess as to what is the best lexicon of size n . One major advantage to using the EDRM is the ease of incorporating subjective considerations. A simple $[0, 1]$ rating scale can be used to rank individual gestures, and the subjective strength of class pairings is calculated by averages, which keeps subjective measure in the EDRM range. These subjective rankings come from ergonomic and vernacular considerations that are mindful of the one who articulates the gestures, and account for ease of articulation, comfort, cultural norms, and suitability to the given application.

6. REFERENCES

[1] R. Cipolla, A. Pentland, *Computer Vision for Human-Machine Interaction*, Cambridge University Press, 1998.
 [2] V. Pavlovic, R. Sharma, T. Huang, *Visual interpretation of hand gestures for human-computer interaction: A review*, IEEE Trans. Patt. Anal. Machine Intell. 19 (1997) 677–695.
 [3] X. Wu, V. Kumar, J. R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. McLachlan, A. Ng, B. Liu, P. Yu, Z. Zhou, M. Steinbach, D. Hand, D. Steinberg, *Top 10 algorithms in data mining*, Knowledge and Information Systems 14 (1) (2008) 1–37.
 [4] F. J., L. J., *A Selective Overview of Variable Selection in High Dimensional Feature Space*, Stat. Sinica 20 (2010) 101–148.
 [5] D. Donoho, *High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality*, in: Math Challenges of the 21st Century, Stanford University, Los Angeles, CA, 2000.

[6] H. Birk, T. Moeslund, C. Madsen, *Real-Time Recognition of Hand Alphabet Gestures Using Principal Component Analysis*, in: 10th Scandinavian Conf. Image Analysis, 1997.
 [7] M. Turk, A. Pentland, *Eigenfaces for Recognition*, Journal of Cognitive Neuroscience 3 (1) (1991) 71–86.
 [8] G. C. Feng, P. C. Yuen, *Variance projection function and its application to eye detection for human face recognition*, Pattern Recognition Letters 19 (1998) 899–906.
 [9] Z. Zhou, X. Geng, *Projection functions for eye detection*, Journal of Pattern Recognition 37 (2004) 1049–1056.
 [10] S. Mika, G. Ratsch, J. Weston, B. Schölkopf, K. Müller, *Fisher Discriminant Analysis with Kernels*, in: Proceedings of the IEEE Workshop on Neural Networks for Signal Processing, 1999, pp. 41–48.
 [11] Q. Gu, Z. Li, J. Han, *Generalized Fisher Score for Feature Selection*, in: Proceedings of the International Conference on Uncertainty in Artificial Intelligence, 2011.
 [12] D. Koller, M. Sahami, *Toward Optimal feature Selection*, Technical Report, Stanford InfoLab (1996).
 [13] M. Robnik-Šikonja, I. Kononenko, *Theoretical and Empirical Analysis of ReliefF and RReliefF*, Machine Learning 53 (1-2) (2003) 23–69.
 [14] X. He, D. Cai, P. Niyogi, *Laplacian Score for Feature Selection*, in: Advances in Neural Information Processing Systems 18, MIT Press, Cambridge, MA, 2006, pp. 507–514.
 [15] L. Song, A. Smola, A. Gretton, K. Borgwardt, J. Bedo, *Supervised feature selection via dependence estimation*, in: Proc. 24th Intern. Conf. Machine Learning, 2007, pp. 823–830.
 [16] F. Nie, S. Xiang, Y. Jia, C. Zhang, S. Yan, *Trace Ratio Criterion for Feature Selection*, in: Proc. 23rd Assoc. Adv. Artificial Intell. Conf., 2008, pp. 671–676.
 [17] P. Lisboa, A. Mehri-Dehnavi, *Sensitivity Methods for Variable Selection Using the MLP*, in: Inte. Work. Neural Net. Ident. Cont. Rob. Signal/Image, 1996, pp. 330–338.
 [18] Y. Lu, I. Cohen, X. Zhou, Q. Tian, *Feature Selection Using Principal Feature Analysis*, in: Proceedings of the 15th Annual International Association of Computing Machinery Multimedia Conference, 2007, pp. 301–304.
 [19] A. Barczak, N. Reyes, M. Abastillas, A. Piccio, T. Susnjak, *A New 2D Static Hand Gesture Colour Image Dataset for ASL Gestures*, Research Letters in the Information and Mathematical Sciences 15 (2011) 12–20.
 [20] J. Grosek, P. Shi, J. Kutz, *Enhanced Gesture Recognition Performance through Improved Pre-Processing*, International Journal of Computer Applications 62 (9) (2013) 1–8.
 [21] S. Mika, G. Ratsch, J. Weston, B. Schölkopf, K. Müller, *Fisher Discriminant Analysis with Kernels*, in: Proc. IEEE Work. Neural Net. Signal Proces., 1999, pp. 41–48.