

Chapter 2

Gait Biometric Recognition

The gait biometric has demonstrated potential promises as an alternative or complementary identifier for use in human recognition systems. However, there is no single measure that encompasses the full set of complex dynamics reflecting what we consider to be the human gait. Instead, important aspects of gait can be measured using one or more of several analysis techniques. Among these techniques are *visual approaches* involving cameras, which can capture differing angles of gait from a distance, and *sensor approaches*, which collect information about gait while in contact with the subject being analyzed. In this chapter, we explore the ways in which these varying approaches have previously been applied to achieve gait biometric recognition, while also highlighting important possible areas of concern in their usage with respect to practicality, privacy, and security. This chapter provides a foundation of knowledge with respect to *gait* and *machine learning*, which will be built upon through the remainder of the book as we demonstrate, via the application of powerful machine learning techniques and the levels of gait recognition performance we might hope to achieve using a sensor-based approach for demonstration.

2.1 Introduction to Machine Learning

Data representations of the gait biometric tend to be large and contain distinguishing characteristics that might not be obvious to even an expert; however, the field of *machine learning* provides multiple concepts for dealing with such complexity. In this section, we explore some of these concepts.

2.1.1 Machine Learning Paradigm

Building intelligent systems that can automatically learn from data or make predictions on data is essential for many types of modern applications. At the core of

the design of such systems is the idea of “learning,” [38] which consists of getting the system to accomplish one or more tasks using some intelligent decisions based on data and patterns (direct experience) or some kind of observations, rules, data experienced in the past. For instance, a learning system can be embedded in a robot, allowing it to learn how to move through its environment while avoiding collisions with other objects in its presence. In healthcare monitoring, using a learning system to understand and classify human physical activities such as sitting, lying, walking, running, climbing, etc., can help in assessing the biomechanical and physiological variables of a patient on a long-term basis. Other practical examples justifying the need for building intelligent systems that are capable of learning include speech recognition, image and text retrieval, commercial software, and web page classification, to name a few. Learning systems can be built using artificial intelligence techniques such as data mining, learning automata [38], and machine learning [35], to name a few. The latter is considered to be a future driver of innovation [34].

Machine learning [35] is a paradigm that focuses on devising computer programs (often called *agents*) that are able to learn from present and/or past experience or adapt to changes in environment without human user assistance. Its fundamental goal is to *generalize* beyond the examples in a training set. These learning algorithms must satisfy certain requirements in terms of *practicality* and *efficiency*, depending on their learning style and the way they model the problem [11]. They must also rely on the so-called *learning models*, themselves distinguishable by the type of learning that they induce, which include supervised learning, unsupervised learning, reinforcement learning, and deep learning, to name a few. From a design perspective, these learning options differ from each other in the way they capture the most important learning features, usually, under some simplifying assumptions. These learning features are concerned with determining which patterns are subject to learning, finding how these patterns are generated and assigned to the learning algorithm, and determining the expected learning goal(s).

2.1.2 Machine Learning Design Cycle

The use of a machine learning approach relies on the application of a machine learning design cycle, which consists of a *dataset and feature selection*, a *learning algorithm*, and an *evaluation module*, as shown in Fig. 2.1.

Dataset: This is the input to the machine learning method, which may or may not require a *preprocessing* step (also called preparation step). The preprocessing step is meant to avoid incompleteness or inconsistency in data, for instance, due to missing attributes or attribute values, discrepancy in coding or naming conventions, errors, noisy data, or outliers, to name a few. It is also meant to check for *data biases*, i.e., the data on which the conclusions rely should be similar to the ones that were used for analysis.

The data preprocessing step can be realized using any of the following methods or a combination of them:




Fig. 2.1 Machine Learning Design Cycle. This figure demonstrates a typical machine learning design cycle

- *Data cleaning* [15]—This process is meant to identify incomplete, inaccurate, and unreasonable data and to improve the quality of data by correcting the detected errors, omissions, outliers (geographic, statistical, temporal or environmental). Its goal is to ensure some conformance to a set of predefined rules. Data cleaning methods include filling out the missing attribute values, for instance, by predicting them using a learning algorithm such as decision tree, binning the attribute values, i.e., transforming continuous values to a finite set of discrete values, clustering the attribute values, i.e., grouping the attribute values in clusters, removing redundancies, performing format checks, completeness checks, reasonableness checks, and validation checks, to name a few.
- *Data normalization* (also called *rescaling*)—This consists of scaling the attribute values so that they belong to a specific prescribed range, e.g.: [0,1] or [-1,1]. Typically, there is no hard requirement for data normalization as pre-processing is more dependent on the data than on the learning algorithm. For instance, for neural networks or deep learning-based learning algorithms where a gradient method is used along with the same learning rate across all weights, all input dimensions can be normalized so that they share a similar distribution.
- *Data aggregation*—This consists of presenting the attribute values in a summarized format; for instance, in the form of their minimum value, maximum value, or average value. Typically, a data cube is constructed for the analysis of the data at various granularities.
- *Data abstraction/new data construction*—This is meant to merge, add, or replace new attributes/or attribute values.
- *Data reduction*—This consists of using partitioning, compression, and discretization techniques to reduce the volume of data while guaranteeing similar analytical results and without compromising the integrity of the original data.

Features Selection: As the size of the dataset can turn out to be very large, it is essential to investigate its quality in terms of irrelevant and redundant information, prior to applying the learning algorithm. Feature selection is a method used to discover and remove such information (as much as possible), with the goal of constructing a small subset of highly predictive features so as to avoid *overfitting* the training data. Common feature selection methods include:

- *Wrapper methods* [6]—which rely on the use of a random statistical resampling technique to determine how accurate the feature subsets will be. These methods typically require that the learning algorithm be repeated several times using a different set of features; hence, they may be computationally expensive;
- *Filter methods* [6]—which are used to identify and remove the undesirable features of the data without the need for a learning algorithm. These undesirable features can be identified by examining the consistency in the data or by finding some correlations between them, or by other means, usually based on heuristic search methods that are designed according to the general characteristics of the dataset. As an example, features can be ranked based on some ranking criteria and the top ranking features can be selected.

In a nutshell, a feature selection method typically involves a generation procedure that produces some subsets of features of the dataset, which are then fed into an *evaluation function*¹ that iteratively works for determining the goodness value of each of these subsets by comparing it against the previous best such value based on a predefined stopping criterion. Next, a validation step is required in order to check the validity of any generated subset. In most cases, the efficiency of the feature selection method [17] is measured by either the level at which it allows the chosen learning algorithm to operate in a fast and effective way, or at which it allows for a simple and comprehensive representation of the future classification and its accuracy. A comprehensive taxonomy of feature selection algorithms is provided in [48].

Model Selection: This step refers to choosing a model (or a set of models) from a set of available ones that has the best possible inductive bias given a finite set of training data² and the above-mentioned small subset of highly predictive features derived from the *Feature Selection* step. A learning algorithm should also be selected. Model selection methods often include the use of data mining approaches and a search strategy, along with a set of criteria used to compare the available models. An initial data analysis and visualization method may also be required in order to make good guesses about the data distribution form and the shape of the desired learning function. A comprehensive survey of model selection methods is given in [25, 36]. Typically, a fraction of the training data (referred to as validation data) is set aside and the remaining training data is used to train each candidate

¹More description of the evaluation function is provided in the Subsection entitled “Learning Models”.

²The training set is typically used to train the classifier, whereas the test set is used to estimate the error rate of the trained classifier.

model. Errors are then evaluated on the held-out data and the selected model is the one with the smallest held-out error. Indeed, the chosen model is expected to work best on the training data, and may be based on either the same learning model with different complexities (hyperparameters) or different learning models. Methods to select the validation data include:

- *Holdout method*—the dataset is broken into two subsets: training set and testing set. A function (so-called approximator) is invoked on the training set only in order to predict the output value for the data in the testing set. The resulting mean test set error is then used for evaluating the model.
- *K-fold cross-validation*—the dataset is broken into K disjoint subsets (typically $K = 10$) and the holdout method is invoked K times, each time using one of the K subsets as testing set and the remaining $K-1$ subsets as training set.
- *Leave-one-out cross-validation (LOOCV)*—This is a special case of K-fold Cross-Validation where K is chosen as the total number of examples. For instance, for a dataset with N examples, LOOCV will perform N experiments, each of which uses $N-1$ examples for training purpose and the remaining example for testing purpose.
- *Random subsampling cross-validation*—where: (a) the dataset is randomly split into two subsets of sizes defined by the user (training and test sets); (b) the model is then fit using the training set, then (c) evaluated using the testing set, and (d) the procedure (a)–(c) is repeated multiple times and the mean estimated error is determined.
- *Bootstrapping*—method for which the dataset, say of size N , is sampled with replacement to create a new dataset of the same size before splitting it up into testing set and training set.
- *Other methods*—Model selection can also be thought of from a Bayesian decision-theoretic perspective. For instance, through the design of a statistical analysis technique where the goal is to study how different aspects of the observed units of the dataset affect some outcome of interest. In this regard, information criteria methods [31] can be used as a tool for optimal model selection among a set of available models. The Akaike information criterion (AIC) [2] is an example of an information criterion method that can be used to measure the goodness of fit or uncertainty for the considered range of data values. The minimum description length [31] is another method for model selection, which relies on the principle that any regularity in the data can be exploited to describe it, to name a few.

Learning Process: In this step, the chosen learning algorithm works with a given or modified dataset $\mathbf{D}$ of examples—which, for instance, represent some past experiences. It then tries to either describe that dataset in some meaningful way or develop a suitable response for the future data. As an example, in a healthcare environment, an illustration of this scenario is when the learning algorithm is provided a list of patients' records and their corresponding diagnoses; then it tries to learn from this list in order to discover the dependencies between diseases or predict how often the disease will affect future patients with similar type of records.

Typically, a training set of examples, say $(\mathbf{x}_k, \mathbf{y}_k)$ in $\mathbf{D} = \{d_l = (x_l, y_l), \dots, d_n = (x_n, y_n)\}$ is provided by the learner, where $\mathbf{x}_k = (x_{k,1}, \dots, x_{k,p})$ is an observed input and $\mathbf{y}_k$ is the output corresponding to it. The learner then outputs a classifier. The goal of the learning process is to *generalize* beyond the examples in the training set; for instance, to determine whether for future examples $\mathbf{x}_u$ the classifier³ will yield the correct output $\mathbf{y}_u$ or will accurately predict it. Since there may be many such functions $f: \mathbf{X} \rightarrow \mathbf{Y}$ to learn, the question that arises is how to select the best possible function f that will represent the mapping between $\mathbf{x}_k$ and $\mathbf{y}_k$? Usually, the learner is expected to embody some assumptions and knowledge beyond the data that it has on hand, for instance, with regard to the type of function to be considered; say for instance $f(x) = ax + b$ for each pair of parameters a and b . This equation defines the learning model M . Here, the parameters a and b should be set in such a way that the misfit between the model M and the observed data $\mathbf{D}$ (so-called error or evaluation function) is optimized or the data $\mathbf{D}$ is explained in the most meaningful way. Examples of evaluation functions include the *mean square error* and *average misclassification error* [30], to name a few. It should be noticed that the above-mentioned *overfitting* problem occurs when the knowledge and data that the learner has in hand are not sufficient to completely determine the correct classifier.

The choice of the above function f to formally represent a learner is paramount to the choice of the set of classifiers that can be learned (also called *hypothesis space*). This imposes the constraint that only classifiers belonging to this space are considered valid. Next, to differentiate good classifiers from bad ones, the aforementioned *evaluation function* used by the learning algorithm may be different from the one defined in the evaluation step (as shown in Fig. 2.1)—i.e., one that the classifier is expected to optimize. Next, a method to search for the classifier with the highest score among that of all candidate classifiers is needed, which will determine how efficient the learner is. This method relies on the choice of a suitable *optimization* technique. Often, off-the-shell optimizers are first experimented with, and then customized ones are designed later on to replace them. The learning process is therefore an optimization problem that may be difficult to solve because the right choice of the learner and evaluation functions matters.

In a nutshell, a machine learning technique involves three main components: a formal learner representation, an evaluation function, and an optimization technique to find the classifier with the highest possible score:

- *Learner representation*: One can distinguish between *parametric* (or fixed-size) learners—which are useful when enough training samples are provided—and *nonparametric* (or variable size) learners—whose learning efficiency typically depends on whether a sufficient training set is acquired or not, which is difficult to judge due to the *curse of dimensionality*; a factor encountered when the dimensionality (i.e., the number of features) of the examples increases [30].

³A classifier is defined as a machine learning system that takes, as input, a vector of continuous or discrete feature values and generates a single discrete value (so-called class) as output.

- *Evaluation function*: This is used by the chosen learning algorithm to characterize the goodness of the classifier. This function can be based on performance and/or statistical measures such as accuracy, error rate, precision and recall, squared error, likelihood, posterior probability, information gain, K-L divergence, cost, utility, and margin, to name a few.
- *Optimization technique*: The output is the classifier that has the highest possible score. Continuous constrained optimization such as linear or quadratic programming, unconstrained continuous optimization such as gradient descent, quasi-Newton methods, and conjugate gradient, in addition to combinatorial optimization techniques such as greedy search, branch-and-bound, and beam search techniques, to name a few, can be used for such purpose.

It should be noticed that choosing or defining each of the above three components is a difficult task. Indeed, the fact that a function can be represented does not mean that it can automatically be learned. As an example, a decision tree learner cannot operate on trees that have more leaves than the training set of examples. As another example, in a scenario where the *Features Selection* step has resulted in the construction of the best possible set of features, but the classifiers obtained from the learner are still not accurate enough, more examples and/or raw features must be gathered or a better learning algorithm must be designed. These examples show that several tradeoffs should be considered when designing a classifier or when using an existing one [4, 53].

Various classification algorithms are available in the literature [56]. These may be grouped based on (1) learning style—i.e., the way that the algorithm models the problem based on its interaction with the input data, (2) similarity in terms of function used, and (3) other criterion.

- *Learning style-based classifiers*: These include three main types of learning schemes: supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning.
 - *Supervised learning*: In this learning model, the training data includes the input and desired responses. The goal is to construct a predictor model that can yield reasonable predictions for the responses to new input data. For some training sets, the correct responses (targets) are known and are given as inputs to the learner during the learning process. In this case, constructing the proper test, training, and validation sets is crucial, and the goal is to obtain correct results when new data are given as input without a priori knowledge of the target.
 - *Unsupervised learning*: In this learning model, the input data is not labeled and the learner is not provided with known results during the training phase. The goal is to have the learner predict the learning relations that are present between the data components.
 - *Semi-supervised learning*: In this learning model, the training data is a mixture of labeled and un-labeled samples and the learner has to learn the structures of the input data and make predictions for new input data.

- *Reinforcement learning*: In this learning model, samples of the input vector $\mathbf{x}$ (but not $\mathbf{y}$ the corresponding response) are provided to the learner. Instead of $\mathbf{y}$, the learner gets a feedback (reinforcement) from a critic about how good the response was. The goal is to select responses that will lead to the best possible reinforcement.
- *Similarity-based classifiers*: These include *instance-based learning* methods such as *K-Nearest Neighbors* and *Support Vector Machines*; hyperplanes methods such as naïve Bayes and logistic regression; graphical models such as Bayesian networks and conditional random fields [5]; set of rules-based methods such as proposition rules, logic programs, association rule learning; other methods include *artificial neural networks*, decision trees, regularization methods, Kernel methods, clustering methods, deep learning, *dimensionality reduction*, ensemble methods, and learning vector quantization [11], to name a few.

Evaluation: As stated earlier, the learned model has to be applied to new data in order to predict the responses $\mathbf{y}$ for new input data $\mathbf{x}$ using the learned function $f(\mathbf{x})$. This is referred to as *Evaluation step*, which is meant to determine the *effectiveness* and *efficiency* of the learned $f(\mathbf{x})$ (classifier or learning algorithm) on various collected datasets [30] based on a predefined *evaluation* metric and a *confidence* estimation method. The goal of the evaluation step is to: (1) compare a new learning algorithm against other classifiers on a given domain or a set of benchmark domains, or (2) characterize some generic classifiers on some benchmark domains, or (3) compare various classifiers on a specific domain. Performance metrics of interest and a framework for the evaluation of learning algorithms and classifiers are provided in [30]. To test the learner’s performance, the basic experimental setup consists of: (1) taking the dataset $\mathbf{D}$ and dividing it into training dataset and testing dataset (for instance, using the above-described model selection methods), (2) using the training dataset and the chosen learning algorithm to train the learner, and (3) testing the learner on the testing dataset. The results obtained with the testing dataset can then be used for comparing other learners with different models and learning algorithms.

2.2 General Principles of Designing Gait Biometric-Based Systems

The machine learning concepts from the previous section outlined some of the potential building blocks for creating a gait recognition tool. To make use of these concepts and exploit the gait biometric for practical purposes, we must first design a *biometric system* that is suitable for the requirements of the intended usage. The first question that should be asked when designing a biometric system is whether the system will be used for *identification* or *verification* purposes. If a system is

configured for verification, it will be provided a claimed identity information alongside a subject's gait biometric information and will be responsible for deciding whether that identity matches the provided biometric information; such systems are typically used for user authentication. On the other hand, if a system is configured for identification, it will only be provided a gait biometric sample and will be responsible for assigning an *identity* to the sample; such systems are typically used in forensics analysis, surveillance, and catering (i.e., they provide a custom experience based on an individual's preferences). The choice of system together with its intended function may impact the design decisions regarding the system's sensitivity to errors.

When designing a gait biometric system for the purpose of verification, also referred to as the verification mode of operation, there are two key metrics that must be addressed: the *False Acceptance Rate* (FAR) and *False Rejection Rate* (FRR). The FAR (obtained via Eq. 2.2) is a measure of the rate at which the verification attempts are incorrectly accepted, reflecting the number of samples incorrectly accepted (N_{FA}) as a proportion of the total number of samples accepted (N_{SA}); a high FAR would indicate that the system is vulnerable to intrusions. Conversely, the FRR given in Eq. (2.2) is a measure of the rate at which the verification attempts are incorrectly rejected, reflecting the number of samples incorrectly rejected (N_{FR}) as a proportion of the total number of samples rejected (N_{SR}); a high FRR would indicate that the system is prone to locking out its intended users, and as such has a low usability. In the verification mode, the biometric system's machine learning algorithms will typically produce a percentage output reflecting a certainty as to whether the subject being verified matches a provided sample, and, from this, a threshold can be chosen at which a given subject will be accepted or rejected by the system.

$$FAR = \frac{N_{FA}}{N_{SA}} \times 100 \quad (2.1)$$

$$FRR = \frac{N_{FR}}{N_{SR}} \times 100 \quad (2.2)$$

In general, the goal of the biometric system in verification mode will be to simultaneously reduce both error rates. This goal can be accomplished by designing the biometric system in a way that minimizes a third metric referred to as the *Equal Error Rate* (EER). The EER identifies the single rate at which a biometric system's FAR is equal to its FRR, when given some input dataset. This metric can be visually observed using a *Detection Error Tradeoff* (DET) curve, as shown in Fig. 2.2.

Designing a gait biometric system using the identification mode of operation presents a slightly different set of concerns. This type of system may be designed using one of the following two approaches: a "closed set" approach whereby all subjects being identified are known to the system or an "open set" approach whereby subjects unknown to the system might be presented for identification. If only a small group of known users will ever be put up for identification, then the closed set approach might simplify the design of a gait-based biometric system. Moreover,




Fig. 2.2 Example DET Curve. This figure demonstrates an example of a DET Curve with an EER around 10 %

under a closed set system, the design will be faster and easier to evaluate, since it is assessed using only a single metric, measuring the percentage of correct identification matches, referred to as the identification rate. However, in many real-world identification scenarios, no assumptions can be made about the identity of the people being examined and, consequently, the open set identification system lends itself to more practical applications than its closed counterpart. This open set identification mode is measured according to two primary metrics: the *false alarm rate*, i.e. the percentage of identification instances for which an unknown user is identified as a known user or known user identified incorrectly, and the *detection and identification rate*, i.e. the percentage rate at which known users are correctly identified. In a system using the open set identification mode, the tradeoff between the two aforementioned rates is typically represented using a *Receiver Operating Characteristics* (ROC) curve, as demonstrated in Fig. 2.3. Using this information, it is left to the designer to determine whether the biometric system should be configured to tolerate a higher false alarm rate or lower detection and identification rate.

Having decided upon the mode of operation for a gait-based biometric system, the next consideration should involve the method of capture. The gait biometric is typically measured using one of the following two different approaches. The first approach involves capturing a *video* of a person’s walking motion, while the second approach involves *sensors* that are in direct contact with the subject in motion. Each approach carries a unique set of capabilities and limitations that may influence their




Fig. 2.3 Example ROC Curve. This figure demonstrates an example of an ROC curve used to assess the error tradeoff for a biometric system using an open set identification mode

incorporation into a security system. A video imaging-based capture of the gait biometric may be appropriate in situations where subjects are likely to be uncooperative or must otherwise have their information captured from a distance; it may, however invoke *privacy concerns*. Conversely, the sensor-based approach may allay privacy concerns, but requires a reasonable degree of *cooperation* on behalf of the subject. At other times, the choice of gait capture may be limited by the resources available. Yet, when possible, it may be beneficial to use more than one form of gait biometric captures, as this has been shown to the performance in previous research [13]. The example in Fig. 2.4 demonstrates the three most commonly studied methods of gait capture for biometric identification and verification, including two sensor-based approaches and one image-based approach. The chosen method of gait capture directly affects the format of data samples produced, and therefore will have a key role in the selection of machine learning algorithms most suited for the biometric system design.

At a high level, the *recognition system* from Fig. 2.4 possesses the “intelligence” of the gait biometric system. The design for this component follows a similar pattern in all biometric systems and its interaction with users can be separated into two phases: an *enrollment* phase and a *challenge* (or matching) phase. The diagram in Fig. 2.5 expands upon this component to demonstrate how it facilitates the interactions with the users. In designing a gait biometric system to satisfy this pattern, it is important to consider how the training data will be collected for enrollment, how many




Fig. 2.4 Gait Biometric Capture Techniques. This figure demonstrates an integration of the three most commonly studied gait biometric capture techniques into an example biometric system. In this system the *recognition system* captures sample data from the subject then compares it with previously stored samples or templates (*gait signatures*); the information is then passed on for *access control* or to a *security or surveillance authority* depending on the mode of operation




Fig. 2.5 Gait Biometric System Phases. This figure demonstrates how the two phases of user engagement fit into a gait biometric system design. During the *enrollment* phase enrollees would be expected to walk for a certain number of cycles until an adequate amount of information about their gait has been acquired. Later, during the *challenge* phase, the system will use this enrolled data to identify or verify a user while attempting to filter out imposters

samples might be required to sufficiently train the system, and how to ensure security of the stored data. Choices should also be made at this point in the design as to whether the system must produce real-time results, delayed results, or will only be used after

the fact; in some cases it may be possible to get more accurate results if real-time data analysis is not a requirement.

Gait data samples tend to be large and contain a high degree of variability, making potentially important patterns within the data difficult to assess visually. Discovering a process that is able to isolate and exploit these patterns to best assist in the distinguishing of one individual from others is at the core of gait recognition-based research. Previous research has suggested that *machine learning* techniques, algorithms that specialize in the extraction and prediction of patterns, are well suited for accomplishing this objective. These techniques have been of particular interest in the areas of feature extraction and classification, but could also potentially be extended to be used in the normalization of data via a form of *regression analysis* (this will be explored in Chap. 5). Before designing a gait-based biometric system, it is important to understand the various forms of machine learning and data analysis techniques that may be included and how they might fit together within a biometric system.

Gait monitoring system inputs typically run as a continuous feed of data, oblivious to the periodic cycles we associate with gait. In video approaches, capturing a cycle might mean starting and stopping the sample extraction cycle when certain movements are detected with respect to a static background, while in sensor approaches, it might mean picking up on force spikes and drop-offs that would be associated with a footstep. Consequently, the first action that will need to be performed in any gait-based biometric system is the *extraction of the sample* over the desired periodic interval. From there, some systems may wish to account for variations in scale, such as the varying distances for which subjects may appear in an image, by performing normalization. Finally, to reduce large sample sizes down to more manageable sizes for classification, a third data preprocessing step, known as *feature extraction*, may be required prior to pushing the sample to the classifier (the final component in all biometric systems). To illustrate the incorporation of these data analysis and machine learning-based components into a gait-based biometric system design, Fig. 2.6 presents the relationship between each as a simple flow diagram.

A final consideration that may arise in the designing of a gait-based biometric system relates to the lack of knowledge about the data that will be collected. If there is no history of biometric systems developed around a chosen data format, then it will be difficult to judge the most suitable values for the machine learning parameters that underpin the system. To optimize the biometric system prior to evaluation, it is generally suggested that the tested data be separated into two mutually exclusive datasets: a *development dataset* and an *evaluation dataset*. Then, using the development dataset, machine learning parameters can be optimized, with the optimal solution being incorporated into the biometric system for an unbiased evaluation via the evaluation dataset; a process we will cover later in this book through our demonstrative experiment.

Now that we have demonstrated the principles on which the design of a gait-based biometric system rely, we can explore some of the ways in which gait




Fig. 2.6 Gait Biometric System Components. This figure demonstrates the breakdown of components that add “intelligence” to a gait-based biometric system in the order in which they would appear in the processing of a data sample. Each of these components may be represented by a number of different data analysis or machine learning techniques, which will be discussed in further detail in later chapters

biometrics have been implemented. The next section provides an overview of previous research as it pertains to the different gait biometrics approaches, together with a brief discussion of *security* and *privacy*-connected issues that relate to such systems.

2.3 Authentication Using the Gait Biometric

Biometric recognition using gait presents a number of advantages over the use of more traditional biometric traits: it is generally considered unobtrusive, as it can be measured in a way that does not require a person to alter his or her typical behavior; it does not require a person to present any more information than is already available to a casual observer; and studies have suggested it is very difficult to imitate [21]. In his research, Cattin [13] (Cattin, 2002) makes special mention of the ability of the gait biometric to perform a *living person test*. The test is described as the ability to determine whether the owner of a trait being observed is alive and physically present or not. Traditional biometrics, such as fingerprints, often fail this test as the traits observed can be faked with present-day technology. The security of the gait biometric lies in the incredible difficulty required to spoof it, thus the living person test is considered intrinsic to the method. Yet, as was mentioned in the previous section, there is no single approach that captures the gait in its entirety. Moreover, while it may be considered less intrusive than other biometrics, it may also be found to be invasive of privacy due to the ease with which it can be used to track individuals without them knowing they are being monitored. In this section, we explore the various approaches to perform authentication via gait biometrics and the impacts that this biometric may have in terms of privacy and security.

2.3.1 Privacy and Security Implications of Gait Biometrics

As gait biometrics slowly find their way into real-world applications, it is worth noting that they carry with them a number of concerns common to *all biometrics* with respect to security and privacy, as well as several concerns more specific to the gait itself. Perhaps most importantly from a security standpoint, gait biometric has not yet achieved the *reliability* that would be required as a standalone biometric, and therefore it is typically used together with more reliable biometrics as part of a *multimodal* biometric system [24]. It has also been shown through the work of Gafurov and Snekenes [23] that, in the absence of information about enrolled gait users, the gait biometric cannot easily be spoofed; however, with the knowledge that an attacker has a similar gait signature to an enrolled subject, spoofing becomes substantially easier. Such an attack is just one of many that could occur if the attacker were to compromise the non-sensor components that make up the biometric system. Rath et al. [42] identified eight such potential areas in which a biometric system might be compromised. Of these risks, one of the greatest areas for concern relates to the system's *database*, which if compromised would allow an attacker to view, create, delete, and/or modify the records at will. Another area of concern would be the links between the systems, which if compromised may be subject to replay attacks to gain access to a previous user's template or a bypass of the biometric system altogether in the case of the link granting the recognition decision. To address these concerns, various technical security implementations could be implemented including equipment tampering alarms, encryption of the database, and encryption between links. Moreover, due to the sensitivity of the biometric data, it has been suggested that it only be stored in *nonreversible* template representations rather than as raw data to reduce the utility of the information were it to be compromised [14]. Alternatively, attacks may come in forms that no amount of technical security could ever be expected to prevent, via *social engineering*; an example of such an attack might come as a bribe offered to an employee in return, for access to a system. To address such a case, an organization might wish to conform to the *ISO27001* and *ISO27002* standards for information security [12].

Security in a gait-based biometric system may also be affected by undesirable behavior even without being compromised. One such case would come as a result of this biometric's fluidity and likelihood to change due to environmental factors as well as over the lifetime of an enrolled subject; system training cannot account for all such factors. Consequently, any gait-based security mechanism is likely to eventually start rejecting samples for an enrolled user simply because they do no longer possess the gait dynamics that were present during training; this may restrict the gait from ever being used as a standalone principal biometric. Furthermore, the strength of security in the system will reflect its mode of operation, with identification mode being inherently less secure than verification mode [57]. Nevertheless, a secure incorporation of gait biometrics into a system provides key advantages in enhancing the overall security and making the system less prone to plausible deniability on behalf of a malicious insider.

From a privacy perspective, the gait biometric carries with it some fairly substantial concerns. To begin, the gait biometric like most other biometrics can be considered to be relatively *static* over its intended period of use, and therefore, were it to be leaked, an attacker would gain access to a key piece of information that could not easily be changed or reset like a password or keycard; however, given the dynamic nature and difficulty in reproducing the gait, it is not entirely apparent how useful such information would be to an attacker. A greater potential privacy concern with respect to the gait is the ease with which it could be used to enroll and track individuals *without* their consent or even knowing they are being tracked; such a scenario could play out through an extensive network of video cameras or floor sensors and would violate an individual's right to anonymity. This application of gait biometrics has attracted the greatest amount of attention from privacy advocates, yet, as prominent gait biometric researchers, Boulgouris et al. [9] have pointed out that such a scenario remains technically infeasible for the foreseeable future. A greater immediate privacy concern might come as a result of the use of gait tracking technologies, not for biometric recognition purposes, but as a *surveillance* technique for detecting abnormal behaviors; such behaviors would be flagged as suspicious and may bring scrutiny upon the individual that triggered them. In one recent study, it was demonstrated how gait analysis could detect whether an individual was carrying excess load in addition to its body weight, and evidence suggested it may be possible [46]. If such a tracking system were to be deployed in public, it could be prone to very high false match rates, with undeserved scrutiny being given to individuals suffering from injury, fatigue, or psychological conditions. In a similar fashion, gait analysis could have the dual purpose of identifying the *medical conditions*; for instance, gait analysis has proven particularly valuable in the study of Parkinson's disease [45]. Yet, again, incorporating such technologies into public surveillance systems could compromise the medical privacy rights.

A strong public response should be expected in the event that a gait biometric tracking system was ever modified to take on any of the aforementioned dual-purpose surveillance objectives. Consequently, before deploying a gait biometric recognition system for real-world use, sufficient effort should be made to allay relevant public privacy fears.

2.3.2 Gait Biometric Approaches

There are a variety of characteristics that define the human gait and a variety of techniques used to extract these gait characteristics. In a summary of research in the field of gait biometrics, Derawi et al. [19] have suggested there are 24 different components that, together, can uniquely identify an individual's gait. However, there is no single device able to capture the full complex set of motions that form the human gait. Consequently, the approaches used to accomplish gait recognition relate directly to the instrumentation needed to extract the gait data, and fall into

three categories: the *machine vision* approach, the *wearable sensor* approach, and the *floor sensor* approach. The following subsections examine these approaches and describe the research that has been done to date with respect to each.

2.3.2.1 The Machine Vision Approach

The machine vision (MV) approach to gait biometrics typically involves capturing gait information from a distance using video recorder technology. This is the most common approach to gait biometric recognition referenced in the current literature [21], having benefited from the availability of large public datasets such as the NIST MV gait database [40]. Recognition via MV has been accomplished using two different techniques: *model-free* and *model-based* recognition algorithms. The model-free technique, often referred to as the *silhouette-based* technique (Fig. 2.7), involves deriving a human silhouette by separating out a moving person from the static background in each video frame. Using this technique, classifiers are developed around the observed motion of the silhouette. In contrast, the less commonly used model-based technique (Fig. 2.8) involves imposing a model for human movement [39]; this is often accomplished by extracting features, such as limbs and joints, from captured images and mapping them onto the structural components of human models for recognition [16].

Over the past decade gait recognition using MV has been investigated by a number of researchers using a variety of methods with promising results. In 2003, Wang et al. [55] developed a silhouette-based technique that used the feature space dimensionality-reducing principal component analysis (PCA) together with the nearest neighbor and Euclidean nearest neighbor classification algorithms. This approach achieved identification rates in the 70–90 % range, varying on the dataset and acceptance criteria used. Later, in 2005, Boulgouris et al. [10] proposed a novel silhouette-based system for gait recognition using linear time normalization (LTN) on gait cycles. The system demonstrated an 8–20 % improvement in its identification rates when compared with existing methodologies at the time. Another study that same year by Lu et al. [33] achieved an identification rate of 92.5 % using a genetic fuzzy support vector machine (GFSVM) classifier; this



Fig. 2.7 Gait Silhouette Sequence. This figure presents an example of a gait silhouette sequence taken from the publically available OU-ISIR gait samples [29]

Projection of Human Body Model



Modeled Gait




Fig. 2.8 Dynamic Gait Modeling. This figure demonstrates the simple modeling of gait in two and three dimensions. The more advanced *three-dimensional modeling* may require several cameras from different angles but can counter covariate factors like differences in scale that might otherwise lead to invalid results using a silhouette-based MV approach

result improved upon the results of the nearest neighbor and standard support vector machine (SVM) tested against the same dataset.

In 2006 Cheng et al. [16] introduced a gait recognition system that used a hidden Markov model (HMM) and, unlike previous systems, was designed to perform recognition on subjects walking down different paths. It accomplished this by first recognizing the walking direction, then applying an appropriate identifier to the determined path; identification rates achieved by this system varied in the 80–90 % range across differing datasets and acceptance criteria. Another silhouette-based study in 2006, by Liu and Sarkar [32], used a generic population HMM (pHMM) to normalize gait dynamics, then used PCA to reduce the feature space and a linear discriminant analysis (LDA) classifier to perform classification; the use of an HMM for normalization was unique to this study and demonstrated how normalization could be used to improve upon the recognition performance. In 2009, a study by Venkat and De Wilde [54] took a different approach and attempted to reduce the computational intensity of silhouette-based recognition techniques by examining *sub-gaits*, defined as smaller localized frames, rather than entire gait images. This technique yielded an identification rate range of about 75–90 % across various datasets and acceptance criteria. However, when vision or motion obstructing factors such as carrying condition and particularly clothing condition were considered, the identification rate dropped to as low as 29 %.

Few researchers to date have studied the effects of the various factors that can obstruct the human gait recognition; however, a real-world gait recognition system would most likely need to be designed to handle such events. To address the issue

of gait variability and obstructions, also known as *covariate factors*, Bouchrika and Nixon [8] proposed a model-based system to extract human joint positions and model gait motion using elliptic Fourier descriptors. Their 2006 study successfully extracted 92.5 % of the heel strikes observed in a dataset containing both visible joints and joints occluded by clothing, demonstrating that the key features of motion analysis could still be tracked in the presence of covariate factors. A follow-up study in 2008 [7] examined the effects of footwear, clothing, carrying condition and walking speed, and, using the model-based approach to joint extraction together with a K-nearest neighbors (KNN) classifier, achieved an identification rate of 73.4 % against a database containing variations of these covariate factors.

Two further studies attempted to mitigate the weaknesses of the MV approach to gait recognition by fusing it with a *secondary* biometric factor. In 2002, Cattin [3] developed a system that fused the data from a video sensor recognition system with a force plate-based footstep recognition system to recognize an individual walking in a monitored room. The results of his study were promising with a verification EER of 1.6 %. In 2006, Zhou and Bhanu [60] developed a different multifactor biometric technique that combined an MV gait recognition system with a facial recognition system. This system had the benefit of only requiring a single video sensor for capturing both biometrics and achieved an identification rate of 91.3 %.

All studies discussed so far have dealt with recognition using images captured from a video sensor; yet one unique study, which, however, was not based on visual data but best falls under the MV category, proposed *audio*-based footstep recognition. In 2006, Itai and Yasukawa [28] took a wavelet transform technique, widely used in feature extraction for speech recognition, and applied it to feature extraction for audible footsteps. Using this technique, an identification rate of 80 % was achieved, suggesting that this could be an exciting new realm for study in the field of gait recognition.

2.3.2.2 The Wearable Sensor Approach

Biometric recognition using wearable sensors (WS) is a new approach to gait biometrics that aims to use sensors attached to the human body (Fig. 2.9) to perform recognition. Much of the early research into gait-aware wearable sensors came from medical studies that focused on their usefulness for detecting pathological conditions [3]. Research into the usefulness of the WS approach for biometric recognition has been, at least partly, held back due to the lack of large publically available datasets [19]. However, the WS approach to biometrics presents a number of advantages, including the ability to perform *continuous authentication*, which would not always be possible with sensors fixed to a physical location. Over the past 10 years, a number of studies, using a variety of techniques, have investigated the feasibility of the WS approach.

In 2006, Gafurov et al. [22] developed a WS biometric recognition system using an accelerometer sensor attached to the lower leg. This study first collected test




Fig. 2.9 Nonintrusive Wearable Sensors for Gait Monitoring. This figure demonstrates some of the nonintrusive devices that can be used to monitor gait via embedded wearable sensors; stronger gait recognition performance may be achievable by combining multiple types of sensor

subject data then uploaded it to a computer, rather than performing real-time classification. The experiment achieved a verification EER of 5 % when recognizing individuals using a histogram similarity classification technique. Another study in 2006, by Huang et al. [27], presented a recognition system based on sensors embedded in a shoe. The sensors used included a pressure sensor, tilt angle sensor, gyroscope, bend sensor, and accelerometer. This system was developed to transmit gait data from the shoe to a computer in real time, and used the PCA feature reduction technique together with an SVM classifier to accomplish a 98 % identification rate on a small sample dataset. A follow-up study by Huang et al. [26] (Huang, Chen, Huang, & Xu, 2007) in 2007 applied a Cascade Neural Network with a node-decoupled extended Kalman filtering (CNN-NDEKF) classifier to the shoe-based WS system and achieved a 97 % identification rate.

One major weakness of the WS approach is the potential inconvenience or discomfort that may be caused by attaching sensors to the human body. For this reason, WS research has tended to focus on one of two unobtrusive WS techniques: *shoe-based* monitoring techniques, like [27] and [26] described in the previous paragraph, and *phone-based* monitoring techniques. Both techniques make use of equipment that is already a part of daily life and require no alterations to typical behavior. In the past few years, the increasing computational power and wider use of smart phones, and increasingly *smart* watches, has sparked a number of studies into feasibility of mobile-based monitoring techniques.

A study in 2009, by Spranger and Zazula [49], described a biometric recognition system that worked with a feature set consisting of cumulants of accelerometer data captured by a mobile phone attached to a person's hip. The system achieved a 93.1 % identification rate on a small dataset using PCA for feature dimensionality reduction and an SVM classifier. A separate study in 2009 by Fitzgerald [20] demonstrated a system, designed for possible future use in mobile phones, that worked with accelerometer and gyroscope data captured by a Nintendo Wii controller. Gait cycles captured by the system were normalized with respect to time. User recognition for the system was examined using KNN, naive Bayes, and quadratic discriminant analysis (QDA) classifiers, with the KNN classifier performing best, achieving an identification accuracy of about 95 %. In 2010, Derawi et al. [18] collected data from accelerometers attached to a belt on the legs of 60 volunteers, generating a much larger dataset than used in the other WS studies described in this chapter. This study focused on cycle length as a metric and, using a cross cyclical rotation metric (CRM), achieved a verification EER of 5.7 % for person recognition. Although the device used by the study was not a mobile phone, the application of this system for use in mobile phones was noted as an important area for future research. In 2011 Nickel et al. used an HMM classifier on accelerometer data from commercially available mobile phones to perform person recognition and achieved a verification EER of about 10 %. The study worked with a relatively large dataset of 48 subjects and was particularly promising for the field of WS-based authentication, because it proved that even a standard, commercially available mobile phone could now be used for gait recognition purposes.

2.3.2.3 The Floor Sensor Approach

The floor sensor (FS) approach to gait biometrics involves recognizing people based on the signals they generate as they walk over sensor-monitored flooring. Data captured by floor sensors typically falls into two categories: *binary image frames* (Fig. 2.11) of the foot while it is in contact with the ground, and single dimensional *force distribution plots* (Fig. 2.10), which describe the force exerted by the foot over time. Most FS technology was developed for the study of biomechanical processes; particularly for improving performance in athletics and discovering the effects of pathological conditions such as diabetes [47]. The first studies using FS technology for gait recognition began in the late 1990s [43]. Over the past 10 years, a small but increasing number of studies have examined the FS approach to gait biometrics. Some, like [13] by Cattin, described in the previous machine vision section, combined the FS recognition approach with another gait recognition approach, such as the MV approach, to improve the recognition accuracy; however, most other studies have focused on using the FS approach for *single factor* recognition.

Open research into using footsteps as a biometric dates back to a 1997 study by Addlesee et al. [1]. In this study, load cell floor sensors were used to capture partial ground reaction force (GRF) data for 15 volunteers and, using an HMM classifier, a




Fig. 2.10 Footstep GRF Force Time Series. This graph demonstrates the force-time series representation of a footstep commonly used for FS-based recognition

91 % footstep identification rate was achieved. Three years later, Orr and Abowd [40] outfitted a floor tile with a set of force sensors to capture the GRF profile for 15 volunteers. In the study, ten features were extracted and normalized, then passed to a Euclidean nearest neighbor (ENN) classifier for recognition; the result was a 93 % identification rate. Another study in 2005, by Suutala and Rönning [51] used a floor sensor called ElectroMechanical Film (EMFi) to capture the GRF for ten volunteers. The primary focus of this study was to compare various classifiers, combine such classifiers, and examine the effects of rejecting unreliable data samples from the classifier training. In this study, it was found that the SVM and multilayer perceptron (MLP) neural network classifiers performed best; the strongest corresponding identification accuracy on their most complicated dataset was around 92 %, which increased to 95 % when the most unreliable 9 % of sample set was rejected. A later study in 2007 by Moustakis et al. [37] captured the GRF for a larger dataset of 40 volunteers. This study used a feature extraction technique built on wavelet packet decomposition (WPD) to detect the transient characteristics and distinguishing features of the GRF, then applied an SVM classifier to the feature set; the result was a 98.3 % identification rate.

In 2009, Ye et al. [58] presented a unique technique for FS-based biometrics: instead of performing recognition on a footstep, like most previous studies, they developed a system that could recognize a person by upper-body movements




Fig. 2.11 Binary Footstep Frame. This diagram demonstrates a single frame of a footstep on a pressure sensitive sensor array. The darker region represents locations where the footstep is detected as being in contact with the floor. The series of frames produced by a single footstep can be used for FS-based footstep recognition

performed while standing on a force plate. The study obtained the center of pressure for the foot, and monitored its movement as instructed actions were completed by a person on a force plate. This study achieved its best results using a neural network classifier, with a verification EER in the 1-12 % range. Two other studies, one in 2008 [50] and another in 2011 [59], also took a different approach to FS-based gait biometrics, opting to perform gait recognition using binary images of footsteps rather than GRF force signatures. The 2008 study by Suutala et al. [50] examined the shape and pattern of individual footstep image frames, as well as the displacement between the feet during footsteps; it achieved a maximum identification rate of 84 %, using a Gaussian process classifier. The 2011 study by Yun [59] focused primarily on subjects in bare feet and also extracted features from individual footstep frames together with footstep displacement. Using an MLP classifier trained with the extracted features, this study achieved a 96 % identification rate.

Research focused on the FS approach to gait biometrics, like the WS approach, has been disadvantaged due to the lack of any *publically available dataset*. This lack of a publically available dataset makes it difficult to compare results across studies. To address this issue, one group at the University of Wales, Swansea, set out to develop such a dataset. In 2007, Rodríguez et al. [43] published a study of an FS-based recognition system and introduced a dataset of footstep force signatures

covering 41 persons and over 3000 footsteps; the intended purpose was to verify the data and make the dataset available at some future point. In their process, they presented a holistic and geometric feature set, and, using an SVM classifier, they achieved a verification EER of 11.5 % for person recognition. A later study [44] in 2008 dealt with a dataset expanded to 55 persons and more than 3500 footsteps; using the same classifier as the previous study, but, additionally normalizing and optimizing the feature set, a verification EER of 13 % was achieved, with the small increase in error rate attributed to the larger dataset. The database in the 2008 study was said to have been made publically available, but, at the time of writing, no longer appears on the project website [52]; nevertheless, the project web site has indicated that a larger dataset is currently being packaged for future release.

2.4 Summary

Like any biometric system, the gait biometric system is primarily a pattern recognition system. As such, it relies heavily on machine learning for human identification and recognition. This chapter gave a brief overview of machine learning and explored the various factors that define biometric recognition via gait. A high-level overview of the complexities that go into designing a system for the purpose of gait-based biometric recognition was provided, and it was shown that there is no single measure to fully encapsulate the many dynamics that make up the human gait; rather multiple methods, measured through a variety of capture devices, can be used to perform gait recognition. The choice of sensor was highlighted as being particularly important to the design of such a system, and the various methods for achieving gait recognition were expanded up on in the rest of the chapter, where it was shown that gait recognition could be broadly categorized into three different approaches: the MV approach, the WS approach, and the FS approach. Significant research was shown to have gone into each of the three different approaches, in some cases with promising results. Finally, this chapter dealt with some of the key security and privacy concerns that would need to be considered before deploying any gait biometric system to a practical setting.

Looking back on the research relating to the three different gait recognition approaches, one thing that stands out is the relative strength of the WS and FS approaches in comparison with the MV performance. This was particularly true for the approaches that measured footsteps, with nearly every study achieving an error rate of less than 10 %. It is possible that these results are just a reflection of the MV databases being larger and more similar to a real-world environment, with the others being restricted to a laboratory environment; however, it might also be the case that the sensors in contact with the body during a detected gait motion provide more reliability than those capturing a video from a distance. It should be noted that both sensor approaches, when configured to measure the forces of a footstep, produce a signal representation referred as the *Ground Reaction Force* (GRF). In

the next chapter, we explore this measure in greater detail and discuss some of the machine learning strategies that have previously been utilized to perform GRF-based gait recognition.

References

1. Addlesee, Michael D., Alan Jones, Finnbar Livesey, and Ferdinando Samaria. 1997. The ORL active floor [sensor system]. *IEEE Personal Communications* 4(5): 35–41.
2. Akaike, Hirotugu. 1983. Information measures and model selection. *Bulletin of the International Statistical Institute* 50: 277–290.
3. Bamberg, Stacy J. Morris, Ari Y. Benbasat, Donna Moxley Scarborough, David E. Krebs, and Joseph A. Paradiso. 2008. Gait analysis using a shoe integrated wireless sensor system. *IEEE Transactions on Information Technology in Biomedicine* 12 (4): 413–423.
4. Bauer, Eric, and Ron Kohavi. 1999. An empirical comparison of voting classification algorithms: bagging, boosting, and variants. *Machine Learning* 36(1–2): 105–139.
5. Bernardo, José M., and Adrian F.M. Smith. 2008. *Bayesian theory*. New York, USA: Wiley.
6. Bishop, Christopher M. 2006. *Pattern recognition and machine learning*, 1st ed. New York, USA: Springer.
7. Bouchrika, Imed, and Mark S. Nixon. 2008. Exploratory factor analysis of gait recognition. *FG '08. 8th IEEE international conference on automatic face & gesture recognition*, Amsterdam, 1–6.
8. Bouchrika, Imed, and Mark S. Nixon. 2007. Model-based feature extraction for gait analysis and recognition. In *Mirage: Computer vision / computer graphics collaboration techniques and applications*, 150–160. Rocquencourt: Inria.
9. Boulgouris, Nikolaos V., Dimitrios Hatzinakos, and Konstantinos N. Plataniotis. 2005. Gait recognition: A challenging signal processing technology for biometric identification. *IEEE Signal Processing Magazine* 22(6): 78–90.
10. Boulgouris, Nikolaos V., Konstantinos N. Plataniotis, and Dimitrios Hatzinakos. 2006. Gait recognition using linear time normalization. *Pattern Recognition* 39(5): 969–979.
11. Brownlee, Jason. 2013. Taxonomy of machine learning techniques. [Online]. <http://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/>.
12. Calder, Alan, and Steve G. Watkins. 2010. *Information security risk management for ISO27001/ISO27002*. United Kingdom: IT Governance Publishing.
13. Cattin, Philippe C. 2002. Biometric authentication system using human gait. Ph.D. thesis, Swiss Federal Institute of Technology, Zurich, Switzerland.
14. Cavoukian, Ann. 1999. Consumer biometric applications: A discussion paper. Information and Privacy Commissioner, Ontario, Toronto.
15. Chapman, Arthur D. 2005. *Principles and methods of data cleaning—primary species and species-occurrence data*. Copenhagen: Global Biodiversity Information Facility.
16. Cheng, Ming Hsu, Meng Fen Ho, and Chung Lin Huan. 2008. Gait analysis for human identification through manifold learning and HMM. *Pattern Recognition* 41(8): 2541–2553.
17. Dash, Manoranjan, and Huan Liu. 1997. Feature selection for classification. *Intelligent Data Analysis* 1(3): 131–156.
18. Derawi, Mohammad Omar, Patrick Bours, and Kjetil Holien. 2010. Improved cycle detection for accelerometer based gait authentication. In *Sixth international conference on intelligent information hiding and multimedia signal processing*, Darmstadt, 312–317.
19. Derawi, Mohammad Omar, Davrondzhon Gafurov, and Patrick Bours. 2012. Towards continuous authentication based on gait using wearable motion recording sensors. In *Continuous authentication using biometrics: Data, models, and metrics*, ed. Issa Traoré and Ahmed Awad E. Ahmed, Chap. 8, 170–192. IGI Global.

20. Fitzgerald, Michael Nowlan. 2009. Human identification via gait recognition using accelerometer gyro force. http://www.cs.yale.edu/homes/mfn3/pub/mfn_gait_id.pdf.
21. Gafurov, Davrondzhon. 2007. A survey of biometric gait recognition: Approaches, security and challenges. In *Annual Norwegian computer science conference*, Bergen, 19–31.
22. Gafurov, Davrondzhon, Kirs Helkala, and Søndrol Torkjel. 2006. Biometric gait authentication using accelerometer sensor. *Journal of Computers* 1(7): 51–58.
23. Gafurov, Davrondzhon, and Einar Snekkenes. 2009. Gait recognition using wearable motion recording sensors. *EURASIP Journal on Advances in Signal Processing* 2009(1): 1–16.
24. Goudelis, Georgios, Anastasios Tefas, and Ioanis Pitas. 2010. Intelligent multimedia analysis for emerging biometrics, Chap. 5. In *Intelligent multimedia analysis for security applications*, 97–125. Berlin Heidelberg: Springer.
25. Guyon, Isabelle, Amir Saffari, Dror Gideon, and Gavin Cawley. 2010. Model selection: Beyond the bayesian/frequentist divide. *Journal of Machine Learning Research* 11: 61–87.
26. Huang, Bufu, Meng Chen, Panfeng Huang, and Yangsheng Xu. 2007. Gait modeling for human identification. In *IEEE international conference on robotics and automation*, Roma, 4833–4838.
27. Huang, Bufu, Meng Chen, Weizhong Ye, and Yangsheng Xu. 2006. Intelligent shoes for human identification. In *IEEE international conference on robotics and biomimetics*, Kunming, 601–606.
28. Itai, Akitoshi, and Hiroshi Yasukawa. 2006. Personal identification using footstep based on wavelets. In *International symposium on intelligent signal processing and communication systems*, Totoori, 383–386.
29. Iwama, Haruyuki, Mayu Okumura, Yasushi Makiyara, and Yasushi Yagi. 2012. The OU-ISIR gait database comprising the large population dataset and performance evaluation of gait recognition. *IEEE Transaction on Information Forensics and Security* 7(5), 1511–1521.
30. Japkowicz, Nathalie, and Mohak Shah. 2011. *Evaluating learning algorithms: A classification perspective*. New York, NY, USA: Cambridge University Press.
31. Kadane, Joseph B., and Nicole A. Lazar. 2004. Methods and criteria for model selection. *Journal of the American Statistical Association* 99(465): 279–290.
32. Liu, Zongyi, and Sudeep Sarkar. 2006. Improved gait recognition by gait dynamics normalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28(6): 863–876.
33. Lu, Jiwen, and Erhu Zhang. 2007. Gait recognition for human identification based on ICA and fuzzy SVM through multiple views fusion. *Pattern Recognition Letters* 28(16): 2401–2411.
34. Manyika, James et al. 2011. Bigdata: The next frontier for innovation, competition, and productivity. Technical Report, McKinsey Global Institute.
35. Mitchell, Tom. 1997. *Machine learning*. New York, USA: McGraw Hill.
36. Model Selection. 2015. <http://bactra.org/notebooks/model-selection.html>.
37. Moustakidis, Serafeim P., John B. Theocharis, and Giannis Giakas. 2008. Subject recognition based on ground reaction force measurements of gait signals. *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics* 38(6): 1476–1485.
38. Narendra, Kumpati S. and Mandayam A.L. Thathachar. 1989. *Learning automata: An introduction*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc.
39. Nixon, Mark. 2008. Gait biometrics. *Biometric Technology Today* 16(7–8): 8–9.
40. Nixon, Mark S., and John N. Carter. 2006. Automatic recognition by gait. *Proceedings of the IEEE* 94(11): 2013–2024.
41. Orr, Robert J., and Gregory D. Abowd. 2000. The smart floor: A mechanism for natural user identification and tracking. In *CHI '00 conference on human factors in computer systems*, The Hague, 275–276.
42. Ratha, Nalini K., Jonathan H. Connell, and Ruud M. Bolle. 2001. Enhancing security and privacy in biometrics-based authentication systems. *IBM Systems Journal* 40(3): 614–634.
43. Rodríguez, Rubén Vera, Nicholas W.D. Evans, Richard P. Lewis, Benoit Fauve, and John S. D. Mason. 2007. An experimental study on the feasibility of footsteps as a biometric. In *15th European signal processing conference (EUSIPCO 2007)*, Poznan, 748–752.

44. Rodríguez, Rubén Vera, John S.D. Mason, and Nicholas W.D. Evans. 2008. Footstep recognition for a smart home environment. *International Journal of Smart Home* 2(2): 95–110.
45. Salarian, Arash et al. 2004. Gait assessment in Parkinson's disease: Toward an ambulatory system for long-term monitoring. *IEEE Transactions on Biomedical Engineering* 51(8): 1434–1443.
46. Samudin, Norshuhada, Wan Noorshahida Mohd Isa, Tomás H. Maul, and Weng Kin Lai. 2009. Analysis of gait features between loaded and normal gait. In *Fifth international conference on signal image technology and internet based systems*, Marrakech, 172–179.
47. Sazonov, Edward S., Timothy Bumpus, Stacey Zeigler, and Samantha Marocco. 2005. Classification of plantar pressure and heel acceleration patterns using neural networks. In *IEEE international joint conference on neural networks*, vol. 5, Montreal, 3007–3010.
48. Sewell, Martin. 2007. Feature selection. <http://machine-learning.martinsewell.com/feature-selection/>.
49. Spranger, Sabastijan, and Damjan Zazula. 2009. Gait identification using cumulants of accelerometer data. In *2nd WSEAS international conference on sensors, and signals and visualization, imaging and simulation and materials science*, Stevens Point, 94–99.
50. Suutala, Jaakko, Kaori Fujinami, and Juha Röning. 2008. Gaussian process person identifier based on simple floor sensors. In *Smart sensing and context third European conference, EuroSSC*, Zurich, 58–68.
51. Suutala, Jaakko, and Juha Röning. 2008. Methods for person identification on a pressure-sensitive floor: Experiments with multiple classifiers and reject option. *Information Fusion Journal, Special Issue on Applications of Ensemble Methods* 9(1): 21–40.
52. Swansea Footstep Recognition Dataset. [Online]. <http://eeswan.swan.ac.uk>.
53. Tang, Jiexiong, Chenwei Deng, and Guang-Bin Huang. May 2015. Extreme Learning Machine for Multilayer Perceptron. *IEEE Transactions on Neural Networks and Learning Systems* PP(99): 1. DOI:10.1109/TNNLS.2015.2424995.
54. Venkat, Ibrahim, and Philippe De Wilde. 2011. Robust gait recognition by learning and exploiting sub-gait characteristics. *International Journal of Computer Vision* 91(1): 7–23.
55. Wang, Liang, Tieniu Tan, Huazhong Ning, and Weiming Hu. 2003. Silhouette analysis-based gait recognition for human identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25(12): 1505–1518.
56. Witten, Ian H., Eibe Frank, and Mark A. Hall. 2011. *Data mining practical machine learning tools and techniques*, 3rd ed. San Mateo, CA, USA: Morgan Kaufmann.
57. Woodward, John D. et al. 2001. Biometrics: A technical primer, Chap. 2. In *Army biometric applications: Identifying and addressing sociocultural concerns*, Rand, 67–86.
58. Ye, Hong, Syoji Kobashi, Yutaka Hata, Kazuhiko Taniguchi, and Kazunari Asari. 2009. Biometric system by foot pressure change based on neural network. In *39th international symposium on multiple-valued logic*, Naha, 18–23.
59. Yun, Jaeseok. 2011. User identification using gait patterns on UbiFloorII. *Sensors* 11: 2611–2639.
60. Zhou, Xiaoli, and Bir Bhanu. 2006. Feature fusion of face and gait for human recognition at a distance in video. In *Proceedings of IEEE international conference pattern recognition*, Hong Kong, 529–532.

Machine Learning Techniques for Gait Biometric
Recognition

Using the Ground Reaction Force

Mason, J.E.; Traore, I.; Woungang, I.

2016, XXXIV, 223 p. 76 illus., 3 illus. in color., Hardcover

ISBN: 978-3-319-29086-7