

CHARACTER IDENTIFICATION

Technical Field

The present invention relates to a method and apparatus for modifying a classification scheme. In particular, the modification is adapted to allow a user independent classification
5 scheme to be modified into a user dependent classification scheme.

Background Art

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgment or any form of suggestion that the prior art forms part of the common
10 general knowledge.

Handwriting recognition systems are generally classed as writer-dependent, meaning they have been trained to recognise the specific writing style of a single user, or writer-independent, meaning they have been trained to recognise the writing of any user. Writer-
15 dependent systems are usually able to produce more accurate recognition (for a specific user) than writer-independent systems, since they need only model the style of a single writer. However, they generally require the user to enter a large amount of training information to allow the user-specific training to take place. Conversely, writer-independent systems do not require user-specific training as they are generally developed
20 using a corpus of training data from a large number of users. Since these systems must cater for a wide range of stylistic variation in stroke and letter formation, they are more likely to encounter confusion between character classes.

Writer adaptation is the process of transforming a writer-independent recognition system
25 into a writer-dependent system using data provided by the user. Adaptive systems are attractive since they do not require the user to perform the lengthy process of entering training data, and can, over time, approach the recognition accuracy of a writer-dependent system. In addition to this, they are able to adapt to the changing writing style of a user over time.

30

One of the difficulties in developing an adaptive system is to ensure that the adaptations that occur lead to improved recognition. If not carefully implemented, adaptive procedures

can decrease the overall recognition rate for a specific user, for example, by the inclusion of an incorrectly labelled prototype, or adaptation using a badly formed or ambiguous pattern. As a result, many adaptive systems require guidance from the writer to perform adaptation, using interaction with the user to ensure the data used for adaptation is well-formed and correctly labelled. Examples of this are described in US Patent 5,917,924, US Patent 5,754,686, US Patent 5,544,260, and US Patent 5,550,930.

While this can reduce the occurrence of degenerative adaptation, it requires the user to be involved with performing a number of possibly lengthy and tedious operations.

Adaptive classifiers must use some kind of learning process to allow the system to adapt to a user's specific style. Supervised learning is the process whereby the behaviour of a classifier is modified based on the correct labelling of a set of samples (ie. the correct category of each example is supplied). However, this information is generally not available to an adaptive system without user intervention, since the only labelling information available is the output of the classifier itself. Alternatively, unsupervised learning techniques (also known as self-organised learning) do not require labelled samples for the learning processes, and thus are suited to adaptive systems where the correct result is not known.

Competitive learning is an unsupervised learning process that requires elements of a system to compete with each other for activation, and is described for example in C. von der Malsburg, "Self-Organisation of Orientation Sensitive Cells in the Striate Cortex". *Kybernetik*, 14:85 – 100, 1973, for the self-organisation of orientation-sensitive nerve cells. Similarly, it is also described in K. Fukushima, "Cognitron: a Self-Organising Multilayered Neural Network". *Biological Cybernetics*, 20:121 – 136, 1975 for the self-organisation of a multi-layer neural network called the cognitron. There are also neurobiological justifications for competitive learning techniques as described in J. Ambros-Ingerson, R. Granger, and G. Lynch, "Simulation of Paleocortex Performs Hierarchical Clustering". *Science*, 247: 1344-1348, 1990.

Further examples of documents describing these techniques will now be described.

V. Vuori, J. Laaksonen, E. Oja and J. Kangas, "On-line Adaptation In Recognition of Handwritten Alphanumeric Characters", *Proceedings of the Fifth International Conference on Document Analysis and Recognition, ICDAR '99*. IEEE Computer Society, Los Alamitos, CA, USA describe a user-specific adaptive system for handwritten alphanumeric characters that includes various combinations of three strategies. These include adding prototypes based on a k-NN search, inactivating prototypes which appear more harmful than useful, and prototype reshaping, based on Learning Vector Quantisation, as described in T. Kohonen, "Self Organising Maps", *Volume 30 of Springer Series in Information Sciences*. Springer-Verlag, 1997.

The classifier is based on the simple Condensed Nearest-Neighbour rule, with a semi-automatic prototype-clustering algorithm used during training to condense the prototypes. Dynamic Time-Warping using various point-to-point, point-to-line, and area metrics is used to calculate the similarity measure between input and prototypes. The approach to prototype deactivation uses supervised learning ("user-reported misclassifications are used to revise the system", as described in J. Laaksonen, V. Vuori, E. Oja and J. Kangas, "Adaptation of Prototype Sets In On-line Recognition Of Isolated Handwritten Latin Characters"). The system also makes decisions driven by inter-class confusion rather than specific class-based allograph identification. The approach is also based on a binary decision, with the confusing prototypes "removed from the set of active prototypes".

L. Schomaker, H. Teulings, E. Helsper, and G. Abbink, "Adaptive Recognition Of Online, Cursive Handwriting", *Proceedings of the Sixth International Conference on Handwriting and Drawing*. Paris, July, 4-7, 1993: Telecom, (pp. 19-21) and L. Schomaker, H. Teulings, G. Abbink, and E. Helsper "Adaptive Recognition of On-line Connected-cursive Script for use in Pen-based Notebook Computers." Handout, distributed with demonstrations presented at the IWFHR III, CEDAR, SUNY Buffalo, USA, May 25-27, 1993 describe a cursive script recognition system based on prototypical strokes clustered using a Kohonen Self Organising Map (SOM). The processes described use a writer-independent stroke-transition network that is used to recognise the written input.

If any of the top twenty possible words output by the recogniser are found to be a valid word as defined by a dictionary, the values in the stroke-transition network are

"incremented in small steps until either the target word is at the top of the output list of words, or until a maximum number of iterations is reached." The adapted system contains user-specific probabilities for individual stroke interpretations in the stroke-transition network.

5

L. Heutte, T. Paquet, A. Nosary and C. Hernoux, "Defining Writer's Invariants To Adapt the Recognition Task", *Proceedings of the Fifth International Conference on Document Analysis and Recognition*, IEEE Computer Society, 1998 uses morphological writer-specific invariants to improve the recognition of an Offline Character Recognition (OCR) system. The technique requires the detection and clustering of writer-specific invariants, which are used in combination with contextual knowledge, to disambiguate the recognition process.

A number of approaches have been proposed that use adaptive techniques for the training of classifiers such as S. Connell and A. Jain, "Learning Prototypes For On-Line Handwritten Digits", *Proceedings of the 14th International Conference on Pattern Recognition*, Brisbane, Australia, pp. 182-184, Aug. 1998, and G. Hinton, C. Williams, and M. Revow, "Adaptive elastic models for character recognition", *Advances in Neural Information Processing Systems 4*, Morgan Kaufmann, San Mateo, CA.

20

Similarly S. Connell and A.K. Jain, "Writer Adaptation of Online Handwritten Models," *Proc. 5th International Conference on Document Analysis and Recognition*, Bangalore, India, pp. 434-437, Sept. 1999 describes a system of writer adaptation that attempts to construct a user-specific writing style based on the identification of lexemes within a writer-independent model, and then uses this writer-dependent model to retrain their classifier. Since the approach is based around the training of a Hidden Markov Model (HMM) classifier, the adaptation technique is presumably performed once only for each writer (due to the large overhead required to retrain an HMM system). None of these techniques attempt to perform continuous user-specific adaptation based on user input.

25
30

US Patent 6,038,343, describes an approach to adaptation that uses the generation of writer-specific feature vectors in combination with a user-independent "universal recognition dictionary" to improve recognition results. They employ statistical regression to "predict a

writer-specific feature vector for each of multiple alternative categories from the feature vector of an input character". This method "corrects" the user-independent feature vectors by combining them with the writer-specific feature vectors, which are then used generatively to create character predictions for future user input vectors.

5

US Patent 5,917,924 uses adaptive weights that modify the likelihood of prototypes being matched as candidates in the recognition process and "only varies the weighting values in editing mode". That is, the method only makes changes to the prototype weights based on results determined from user interaction.

10

US Patent 5,754,686 describes an approach to using a user-specific dictionary to store writer-dependent models. "If recognition reliability is low, a warning is issued. In response to a warning, the user or operator can decide whether the character pattern should be registered in the user dictionary." The pattern is automatically assessed for suitability for inclusion in the user dictionary (presumably using a metric of ambiguity with existing prototypes) but no claim is made as to how the user-specific dictionary prototypes are combined with the existing writer-independent models during recognition.

15

US Patent 6,256,410 describes a standard scheme for training a writer-dependent HMM classifier, whereby user-specific training data is segmented using writer-independent models, and a set of character-based models are iteratively trained using the training data.

20

US Patent 5,319,721 describes a method for evolving a set of user-independent prototypes into a writer-dependent set. If the input stroke data matches a prototype to within a certain threshold, the user data is merged with the existing prototype to produce a writer-dependent prototype and "one or more starter prototype symbols of the same label as the input symbol" are deleted. If the stroke data does not match the existing prototypes, a new prototype is created, and again, one or more starter prototypes of the same label are deleted.

25

In US Patent 5,544,260 it is described that using "information provided during error correction for modifying character prototypes," that is, using correction strokes made by the user to update mis-recognised prototypes. Similarly, US Patent 5,550,930 describes a

30

method of storing recognition results, and, when requested by a user, displaying the results and allowing the user to select the input and corresponding symbols for classifier training.

5 US Patent 5,285,505 describes a method for "creating character prototypes for improving recognition accuracy of similarly shaped characters" by emphasising the sections of the character that are critical for discrimination and de-emphasising sections that are similar between characters. This approach is targeted specifically in resolving two-class ambiguities, such as 'g'/'y' and 'A'/'H' confusion.

10 Disclosure of the Invention

In a first broad form the present invention provides a method of modifying a classification scheme for classifying hand-written characters to thereby make the classification scheme user dependent, the classification scheme including a number of allographs, each allograph representing a respective style of a respective letter, the method including:

- 15
- a) Obtaining at least one hand-written character;
 - b) Selecting an allograph representing the character; and,
 - c) Modifying one or more of the allographs in accordance with the selection.

20 Typically, the method includes modifying the allographs to reflect that the selected allograph corresponds to a style of the respective letter used by the user.

The classification scheme can include a number of classes, each class representing a respective letter and containing one or more of the allographs, the method including:

- 25
- a) Selecting a class representing the character to thereby determine the letter represented by the character; and,
 - b) Modifying one or more of the allographs in the respective class.

Each allograph typically contains one or more prototypes, each prototype representing a variation in the respective allograph, the method including:

- 30
- a) For at least some of the prototypes, determining a selection value based on the similarity of the character and a respective prototype; and,
 - b) Selecting a respective one of the prototypes in accordance with the determined selection values.

The method generally includes selecting the class in accordance with the selected prototype.

5 The method usually includes selecting the allograph in accordance with the selected prototype.

Each prototype is typically represented as a prototype vector, each prototype vector being formed from a number of values, each value quantifying a respective feature of the respective prototype. In this case, the method of determining a selection value typically
10 includes:

- a) Determining a feature vector representing the character, the feature vector being formed from a number of values, each value quantifying a respective feature of the character; and,
- b) Determining a distance value representing the distance between the feature vector
15 and the respective prototype vector.

Each allograph preferably includes an associated weighting, in which case the method typically includes:

- a) Selecting the allograph in accordance with the weightings; and,
- 20 b) Modifying the allographs by modifying the weighting of one or more of the allographs.

Alternatively, the method can involve simply removing the allographs that are not used by the user, although it will be appreciated that this is generally less accurate.

25 The selection value is typically determined in accordance with the determined distance value and the weighting associated with the respective allograph.

The method typically includes modifying the weighting of at least the allograph containing
30 the selected prototype.

The method usually includes selecting the prototype having the lowest selection value.

The method of modifying the weightings can include at least one of:

- a) Decreasing the weighting for the allograph containing the selected prototype; and,
- b) Increasing the weighting for one or more of the allographs that do not contain the selected prototype.

5

The method generally includes increasing and/or decreasing the weighting of the allographs by a predetermined amount.

Alternatively the method can include:

- 10 a) Selecting a reference prototype for each allograph to be modified;
- b) Determining the distance value for each reference prototype;
- c) Determining an average of the determined distance values; and,
- d) Modifying the weighting of each allograph in accordance with the distance value of the respective reference prototype and the average distance value.

15

In this case, the reference prototype is usually the prototype having the lowest distance value for the respective allograph.

The method is typically performed using a processing system, the processing system including:

20

- a) A store for storing allograph data representing the allographs;
- b) A processor, the processor being adapted to:
 - i) Obtain the hand-written character;
 - ii) Select the allograph; and,
 - 25 iii) Modify the one or more allographs.

25

In a second broad form the present invention provides apparatus for modifying a classification scheme for classifying hand-written characters to thereby make the classification scheme user dependent, the apparatus including:

30

- a) A store for storing allograph data representing a number of allographs, each allograph representing a respective style of a respective letter,; and,
- b) A processor, the processor being adapted to:
 - i) Obtain at least one hand-written character;

- ii) Select an allograph representing the character; and,
- iii) Modify the allograph data, thereby modifying one or more of the allographs in accordance with the selection.

5 In this case, the allograph data typically includes a weighting associated with each allograph, the processor being adapted to modify the allograph data by modifying the weightings.

10 The store can also be adapted to store class data, the class data representing a number of classes, each class representing a respective letter and containing one or more of the allographs.

15 The store may be adapted to store prototype data representing prototypes, each allograph containing one or more prototypes, each prototype representing a variation in the respective allograph.

The processor is preferably adapted to perform the method of the first broad form of the invention.

20 In a third broad form the present invention provides a method of classifying hand-written characters using a user dependent classification scheme, the classification scheme including a number of allographs, each allograph representing a respective style of a respective letter and having a respective weighting representing the user dependency, each allograph containing one or more prototypes, each prototype representing a variation in the respective
25 allograph, the method including:

- a) Obtaining at least one hand-written character;
- b) For at least some of the prototypes, determining a selection value based on the similarity of the character and a respective prototype, and the weighting of the allograph containing the respective prototype;
- 30 c) Selecting a respective one of the prototypes in accordance with the determined selection values; and,
- d) Classifying the character in accordance with the selected prototype.

Each prototype can be represented as a prototype vector, each prototype vector being formed from a number of values, each value quantifying a respective feature of the respective prototype. In this case, the method of determining a selection value typically includes:

- 5 a) Determining a feature vector representing the character, the feature vector being formed from a number of values, each value quantifying a respective feature of the character;
- b) Determining a distance value representing the distance between the feature vector and the respective prototype vector; and,
- 10 c) Determining the selection value in accordance with the determined distance value and the weighting associated with the respective allograph.

The classification scheme can include a number of classes, each class representing a respective letter and containing one or more of the allographs. In this case, the method
15 usually includes selecting a class in accordance with the selected prototype to thereby determine the letter represented by the character.

In general the classification scheme is modified in accordance with the method of the first broad form of the invention.

20

The method generally includes using a processing system, the processing system including:

- a) A store for storing:
 - i) Allograph data representing the allographs and the associated weightings; and,
 - ii) Prototype data representing the prototypes; and,
- 25 b) A processor, the processor being adapted to:
 - i) Obtain the hand-written character;
 - ii) Determine the selection values;
 - iii) Select the prototype; and,
 - iv) Classify the character.

30

In a fourth broad form the present invention provides apparatus for classifying hand-written characters using a user dependent classification scheme, the apparatus including:

- a) A store for storing:

- 5
- i) Allograph data representing a number of allographs, each allograph representing a respective style of a respective letter, and a weighting associated with each allograph, the weighting representing the user dependency of the allograph; and,
 - ii) Prototype data representing one or more prototypes, each prototype representing a variation in the respective allograph; and,
- b) A processor, the processor being adapted to:
- i) Obtaining at least one hand-written character;
 - ii) For at least some of the prototypes, determine a selection value based on the similarity of the character and a respective prototype, and the weighting of the allograph containing the respective prototype;
 - 10 iii) Select a respective one of the prototypes in accordance with the determined selection values; and,
 - iv) Classify the character in accordance with the selected prototype.

- 15 In this case, the store can be adapted to store class data representing a number of classes, each class representing a respective letter and containing one or more of the allographs.

The processor can therefore be adapted to perform the method of the third broad form of the invention.

20

Brief Description of Figures

The present invention should become apparent from the following description, which is given by way of example only, of a preferred but non-limiting embodiment thereof, described in connection with the accompanying figure, wherein:

25

Figure 1 is an example of a processing system suitable for performing the present invention;

Figure 2 is an example of two static allographs of the letter "a";

Figure 3 is an example of two dynamic allographs of the letter "a";

30 Figure 4 is an example of two similar allographs of the letter "a" and the letter "d";

Figure 5 is a schematic diagram of an example of the structure of a classification system; and,

Figure 6 is a schematic diagram of an example of modifying the allograph weight on the classifier decision regions.

Modes for Carrying out the Invention

- 5 The following modes are described as applied to the written description and appended claims in order to provide a more precise understanding of the subject matter of the present invention.

10 An example of apparatus suitable for implementing the present invention will now be described with reference to Figure 1, which shows a processing system 10 adapted to perform handwriting recognition.

In particular, the processing system 10 generally includes at least a processor 20, a memory 21, and an input device 22, such as a graphics tablet and/or keyboard, an output device 23, 15 such as a display, coupled together via a bus 24 as shown. An external interface is also provided as shown at 25, for coupling the processing system to a store 11, such as a database.

20 In use, the processing system is adapted to receive and decode handwritten text characters using a predetermined classification scheme. Initially the classification scheme is user independent, however, the processing system 10 is also adapted to modify the scheme, thereby making the scheme user dependent, to thereby increase the accuracy of the character recognition.

25 From this, it will be appreciated that the processing system 10 may be any form of processing system such as a computer, a laptop, server, specialised hardware, or the like, which is typically adapted to perform these techniques by executing appropriate applications software stored in the memory 21.

30 In order to achieve this, the processing system 10 implements a classification scheme that utilises allographs for identifying respective styles of letter. The processing system then uses adaptive learning of the letter allographs used by a specific writer.

Letter allographs model the stylistic variation in the writing of individual character classes. That is, they represent the different ways that a character may be drawn. Allographs can be defined as static, meaning that the overall shape of the characters discriminates one allograph from another. Examples of static allographs of the letter "a" are shown in Figure 2.

Allographs can also be dynamic, where the allographs may look similar but are differentiated by dynamic stroke information (ie. the order or direction of strokes that make up the letter). Examples of dynamic allographs of the letter "a" are shown in Figure 3. In particular, in this example the basic shape of the two letters is similar, but the version 1 has been drawn with a single stroke, while the version 2 second has been drawn with two strokes. Since allographs can have both a dynamic and static meaning, it is possible for allograph-based adaptation to work in both online and offline (ie. Optical Character Recognition) systems.

A writer-independent system must model all the possible allographs that it may encounter. However, users generally use a single allograph for each letter when writing. Thus, a particular user will generally write a given letter in the same style each time.

As a result, a writer-independent system will contain a large number of allographs that are not required for recognition for a particular user. These superfluous allographs can lower the recognition accuracy of a system since they may be similar to an allograph of a different class that the writer does use.

An example of this is shown in Figure 4. In particular, the allograph on the left represents the letter 'a', while the allograph on the right represents the letter 'd'. Clearly, there is a great deal of similarity between these two allographs, with the obvious possibility of incorrect recognition.

The letter 'a' allograph shown here is rarely used in handwriting, and is not needed for the recognition of most writing styles. Writer-independent systems need to support this allograph however because a small number of writers use this style, but by allowing this allograph, the potential for the incorrect recognition of letters is increased.

Thus, it will be appreciated that a particular individual will not generally write the letters "a" and "d" in such a similar fashion, and accordingly, it is unlikely that both the allographs shown in Figure 4 would be required in a user dependent classification system.

5

Thus, the processing system is adapted to learn which allographs a writer uses, allowing other allographs for the same letter to be ignored, thereby improving the character recognition for a respective writer. This therefore allows the processing system to be able to support a wide range of writing styles with higher recognition accuracy.

10

In use, the processing system 10 therefore operates to receive a character, and compares the character to letters, and allographs stored in the database 11. This allows the processor to determine the letter and allograph to which the character corresponds.

15 This is generally achieved by comparing the character to prototypes stored in the database 11. The prototypes are grouped into respective allographs and letters. Accordingly, the similarity between the character and the prototypes is used to select the prototype that most closely matches the character, thereby allowing the character to be assigned to a respective letter and allograph group.

20

In particular, this is usually achieved by defining the character in terms of a feature vector. In this case, the feature vector includes values representing different features of the letter. The prototypes are stored as prototype vectors, allowing the distance between the feature vector and prototype vectors to be measured in a feature space, thereby allowing the most
25 similar prototype to be determined.

The processor then modifies the allographs corresponding to the selected letter, thereby indicating that the allograph is an allograph used by the respective user. This is generally achieved by having a weighting associated with each allograph, with the processor being
30 adapted to modify the weightings of the allographs. Thus for example, the weighting of the allograph used by the user can be increased/decreased, with the weightings of the other allographs being decreased/increased accordingly.

These techniques will now be described in more detail.

In particular, the following is a description of a competitive learning algorithm for handwritten character allographs. In this example, the technique will be described in terms of a simple template-matching classifier (specifically, a Euclidean-distance nearest-neighbour classifier). However, persons skilled in the art will appreciate that the procedure can be applied to any classifier type that can support the partitioning of classes into subclasses, and the subsequent weighting of those subclasses.

- 10 In addition to this, the weight updating procedure is not metric dependent and simply relies on some numeric indication of prototype similarity. Thus, the technique can be used with distance metrics such as Euclidean-distance, Dynamic Time-Warping, etc.

15 In the initial stages, the processor 20 will receive an input signal representing the character under consideration. This is usually received from the input device 22, but may also be retrieved from characters stored in the database. In any event, the processor 20 performs pattern classification by mapping the input signal, which may be pre-processed to perform some kind of normalisation, into a multi-dimensional feature space. The processor therefore assigns values to quantify respective predetermined features of the letter, such as
20 letter height, width, number of strokes, maximum stroke length, or the like. The mapping is used to generate a feature vector.

The processor 20 then applies the feature vector to a set of class models to determine which class is most likely to have generated the given features.

25

Template-matching classifiers work by storing prototypical examples of classes (prototypes), and perform classification by using some similarity metric to compare the input features with the prototypes. The class of the prototype that is most similar to the input feature vector is assumed to be the class of the input pattern. As an example, a simple
30 Euclidean-distance nearest-neighbour classifier simply stores the feature vectors of a set of training examples, and labels an input pattern with the class of the training pattern feature vector that has the minimum Euclidean-distance from the input feature vector.

For writer-dependent allograph training to occur, the writer-independent classifier must contain allograph information for each prototype. This means that each training example must be labelled by both class and allograph. To label allographs, a number of semi-automatic procedures exist. For example, L. Vuurpijl and L. Schomaker. "Finding Structure
5 In Diversity: A Hierarchical Clustering Method For the Categorization of Allographs In Handwriting," in ICDAR, pages 387-393. IEEE, August 1997 describes an agglomerative hierarchical clustering method that assists the selection and labelling of allographs for individual letter classes.

- 10 In general the structure of the classification scheme will be as shown in Figure 5. Thus, a single classifier will generally contain many letters. In this regard, it will be appreciated that the term letters will include any alphanumeric character, punctuation, or ideogram. Each letter includes one or more allographs, with each allograph representing a different letter style. Each allograph will then include one or more prototypes, each of which
15 represents variations in the respective allograph. Thus, for example, the prototypes will be used to account for variations in size, shear, angle and the like.

Thus, it is possible for many prototypes of the same class to be identified as the same allograph -- allographs model a specific writing style for a particular letter, whereas
20 individual prototypes represent instances of that writing style. For example, in a nearest-neighbour classifier, it is common to have many prototypes representing the same class, with the differences between the prototypes modelling the possible in-class variations of the features (such as translation, rotation, shearing, etc.) Similarly, many prototypes of the same allograph may model the variations in writing of that allograph. Thus a classifier
25 supports many letter classes, a class supports many allographs, and an allograph may support many prototypes.

In this procedure, classification proceeds by calculating the distance between each prototype vector and the input feature vector in the feature space. This distance represents
30 the similarity between the feature vector and the prototype vector, and hence the similarity between the character and a respective prototype. The processor 20 then determines the class of the prototype in accordance with the minimum distance as the label of the input data. Thus, the processor 20 determines the prototype vector separated from the feature

vector by the minimum distance value and uses this to determine the letter represented by the character.

In order to support the writer-specific determination of allographs, the distance value for each prototype is multiplied by a weighting factor depending on its allograph type. Thus allographs with higher weights are less likely to be matched, since the higher weighting will increase the distance metric, so *decreasing* a weight *increases* the probability that an allograph will be matched; *increasing* the weight *decreases* the probability of a match.

- Initially all weights for all allographs are initialised to the same non-zero value (for example, 1), indicating that no writer-specific preference is given to an allograph. At this point the classification scheme is independent.

Adding in the weighting, the distance value is calculated for a Euclidean-distance nearest-neighbour classifier as follows:

$$\delta(\mathbf{x}) = \arg \min_{i=1}^n (D_i(\mathbf{x}))$$

$$D_i(\mathbf{x}) = w_i \times \|\mathbf{x} - \mathbf{y}_i\|$$

20

Where: $\mathbf{y}_i$ is the feature vector of prototype i ;
 w_i is the (shared) allograph weight of prototype i .
 $p_i \in (\mathcal{R}^d, \Omega)$,
 $\mathbf{x}$ is the input feature vector such that $\mathbf{x} \in \mathcal{R}^d$,
 $\delta(\mathbf{x})$ is the classifier decision function such that $\delta(\mathbf{x}): \mathcal{R}^d \rightarrow \Omega$ is:
 Ω are the classes such that $\Omega = \{\omega_1, \dots, \omega_n\}$
 $\{p_1, \dots, p_m\}$ are the prototypes

25

The weightings are determined to reflect the allographs that the user uses. Thus the processor 20 operates to perform weight adaptation to modifying the allograph weights to give preference to allographs that the writer is more likely to use.

30

Thus, the selection of a prototype as the best match for the input signal is an indication that the writer uses the allograph of the prototype. Accordingly, in a simple approach the processor 20 performs weight adaptation by reducing the weight of the allograph containing the selected matching prototype by a small amount η (the learning-rate). Similarly, the processor 20 increases the weights of the remaining allographs in the respective class by a similar amount. Allograph weights in the unmatched classes are not modified.

Thus, for each allograph i in the matching class:

$$\begin{aligned} w_i &= w_i - \eta && \text{if } i = c \\ &= w_i + \eta && \text{otherwise} \end{aligned}$$

where: c is the allograph number of the matching prototype, and
 η is the learning rate parameter.

It will be appreciated that the classifier will not always match the correct class.

As an example, a badly formed 'a' might match the letter 'd', resulting in an update that decreases the weight of an allograph of 'd' that the writer may not use. However, subsequent correct matches of the 'd' allograph used by the writer will increase the unused allograph weight and decrease the correct allograph weight.

It follows that weight adaptation should occur slowly, to avoid inaccurate classifications giving too much weight to an incorrect allograph. However, the higher the learning rate, the more quickly that adaptation will take place, resulting in higher recognition accuracy. It is therefore necessary to find a balance for the magnitude of the learning rate.

It will be appreciated that this may be implemented by having the magnitude of the learning rate adjustable by the user. Thus, the learning parameter η , which is typically stored in the database 11, can be adjusted by the user, to allow the rate of learning to be adjusted. This allows the user to set a desired level of learning, which can then be reduced if too many inaccuracies occur.

It is also desirable to ensure that the weights do not get too large or too small. If an allograph weight is too small, the decision boundary for that allograph (and thus the class) may encroach on surrounding classes, causing incorrect class decisions to be made by the classifier. Thus, if the weighting increases by too much, it is unlikely that the feature vector
5 will ever be matched to one of the prototypes contained therein. In this instance, this can result in entire allographs being excluded from the matching procedure, even though there is a chance that these could on occasions be used.

Accordingly, upper and lower thresholds are generally provided for the weightings, with
10 the weighting being held at the threshold if the modification would push the weighting above the upper threshold, or below the lower threshold. Again, provision may be made to allow the user to adjust the thresholds if so desired.

Ensuring the weight does not get too large allows allographs to increase in priority if the
15 user changes their writing style at a later date.

The effect of modifying the allograph weight on the classifier decision regions is shown in Figure 6. In particular, in this example, as the weight for the class 3 is decreased, this pushes the original decision boundary 5 (solid line) toward the class 4 forming a new
20 decision boundary 6 (dashed line) that gives a higher probability of matching class 3 over class 4.

The simple weight adaptation procedure described above can suffer from a number of problems.
25

In particular, a number of allographs might match almost equally (ie. there is only a very small difference between the distances to two different allographs from the feature vector) indicating there is no strong evidence of one allograph being usage in preference to the other. Conversely, there may be a large difference between the allograph scores, indicating
30 a clear preference for one allograph type over another.

In the first situation, the weight updating should be small, reflecting the lack of confidence in the identification of the correct allograph. In contrast, in the second situation the weight

updating should be large to ensure unused allographs are quickly identified and de-prioritised.

Thus, the weight updating should compare the minimum distances of the allograph groups (ie. the distances of the best matching prototype for each allograph) and update the weights proportionally to these values. In this way, the allographs are made to compete with one another, since the distance results of each allograph have an impact on the weight modification made to the remaining allographs.

- 10 To implement this kind of weight updating, the processor 20 first determines the class containing the character. This is performed as described above by having the processor 20 locate the prototype that most closely resembles the character. Once this has been performed, the processor 20 operates to determine a minimum distance for each allograph group in the class. This is achieved by determining the distance between the feature vector
15 and the closest prototype vector for each allograph group.

The processor 20 then determines the mean of the minimum distances for the allographs in the class, and uses this to update each weighting. In particular, the weighting modification for a particular allograph is based on the difference between the allograph's minimum
20 distance and the mean distance. To ensure that very large distances do not produce very large changes to the mean, all distances are converted to an exponential form, as follows.

for each allograph i within the matched class

$$M_i = \arg \min_j^n (D_j)$$

25 $E_i = e^{-M_i}$

$$\bar{E} = \frac{\sum_i^n E_i}{n}$$

$$w_i = w_i + \eta \times (\bar{E} - E_i)$$

where: η is the learning rate parameter.

Using this procedure, allographs with minimum distances below the mean distance will have larger values of E_i (since $e^0 = 1$ and $\lim_{x \rightarrow \infty} e^{-x} \rightarrow 0$) resulting in negative updates to the allograph weight since the mean will be less than E_i . The greater E_i is above the mean distance, the greater the weight update.

5

Thus, poorly matched allographs will be updated with positive weight increment depending on how small E_i is compared to the mean. Further, ambiguous results (where minimum allograph distances do not differ significantly) will produce very small weight updates, since all distances will not lie far from the mean.

10

The procedure described above is automatic and therefore requires not intervention from the user. In particular, the process is unsupervised as there is no feedback from the user or another source to teach the correct identification of characters in the event that incorrect determinations are made.

15

However, it is possible to support some supervision in the weight adaptation process if additional information about the classification of the input signal is available (for example, contextual-language models). In this situation, weight updating can be delayed until all recognition processing has been performed, when the final results can be used to modify or

20

verify the raw similarity metrics of the classifier.

For example, weight updating might only be performed if the classifier decision matches the final decision produced by the recogniser. Thus, if the classifier selects the letter 'd' as the most likely match, but later processing concludes that the letter is really an 'a', the

25

weight update would not take place.

Alternatively, most recognition systems generate some kind of confidence level for letter alternatives, and this probability can be incorporated into the weight update procedure so that updates occur proportionally to the level of confidence that the letter is correctly and

30

unambiguously identified.

Furthermore, user feed back could be provided, such that if the processor 20 determines the wrong letter, the user can supply an indication of the correct letter, allowing the processor 20 to adjust the weightings of the allographs accordingly.

5 Accordingly, the above-described process provides a procedure for the competitive learning of letter allographs. This adaptive learning process is able to improve the accuracy of a handwritten character recognition system by minimising the recognition probability of letter allographs that are not used by a specific writer. A number of weight adaptation mechanisms are described.

10

It will be appreciated that in the hardware implementation described above, the weightings are stored in the database 11. Accordingly, it is possible to store the weightings in accordance with a user identifier. This allows different sets of weightings to be stored associated with different users. In turn, this allows the processing system 10 to be used to
15 classify handwriting for a number of different users. In this case, each time handwriting from a different user is analysed, the identifier will be used to select the appropriate set of weightings for both identification of characters, and for subsequent weighting modification.

This in turn allows a single central processing system 10 to be used to identify handwriting
20 from a number of users. This can be used, for example in a network based environment, such as the Internet, to allow a single processing system to analyse handwriting submitted from different users via a web-site or the like.

The invention may also be said broadly to consist in the parts, elements and features
25 referred to or indicated in the specification of the application, individually or collectively, in any or all combinations of two or more of said parts, elements or features, and where specific integers are mentioned herein which have known equivalents in the art to which the invention relates, such known equivalents are deemed to be incorporated herein as if individually set forth.

30

Although the preferred embodiment has been described in detail, it should be understood that various changes, substitutions, and alterations can be made herein by one of ordinary

skill in the art without departing from the scope of the present invention as hereinbefore described and as hereinafter claimed.

Thus for example, instead of modifying the weighting of allographs, the method can
5 involve simply removing the allographs that are not used by the user. This may be on a single case analysis basis, so when a used allograph is determined all others are eliminated from the class. Alternatively, this could be combined with the weighting modification, such that if the weighting of an allograph is modified more than a predetermined amount, the allograph is removed from the classification scheme.

Claims

- 1) A method of modifying a classification scheme for classifying hand-written characters to thereby make the classification scheme user dependent, the classification scheme including a number of allographs, each allograph representing a respective style of a respective letter and having an associated weighting, the method including:
- a) Obtaining at least one hand-written character;
 - b) Selecting an allograph representing the character; and,
 - c) Modifying the weighting of one or more of the allographs in accordance with the selection.
- 2) A method according to claim 1, the method including modifying the allographs to reflect that the selected allograph corresponds to a style of the respective letter used by the user.
- 3) A method according to claim 1 or claim 2, the classification scheme including a number of classes, each class representing a respective letter and containing one or more of the allographs, the method including:
- a) Selecting a class representing the character to thereby determine the letter represented by the character; and,
 - b) Modifying one or more of the allographs in the respective class.
- 4) A method according to any one of the claims 1 to 3, each allograph containing one or more prototypes, each prototype representing a variation in the respective allograph, the method including:
- a) For at least some of the prototypes, determining a selection value based on the similarity of the character and a respective prototype; and,
 - b) Selecting a respective one of the prototypes in accordance with the determined selection values.
- 5) A method according to claim 4, the method including selecting the class in accordance with the selected prototype.

- 6) A method according to claim 4 or claim 5, the method including selecting the allograph in accordance with the selected prototype.
- 7) A method according to any one of the claims 4 to 6, each prototype being represented as a prototype vector, each prototype vector being formed from a number of values, each value quantifying a respective feature of the respective prototype, the method of determining a selection value including:
 - a) Determining a feature vector representing the character, the feature vector being formed from a number of values, each value quantifying a respective feature of the character; and,
 - b) Determining a distance value representing the distance between the feature vector and the respective prototype vector.
- 8) A method according to any one of the claims 1 to 7, the method including:
 - a) Selecting the allograph in accordance with the weightings; and,
 - b) Modifying the allographs by modifying the weighting of one or more of the allographs.
- 9) A method according to claim 7 and claim 8, the selection value being determined in accordance with the determined distance value and the weighting associated with the respective allograph.
- 10) A method according to claim 8 or claim 9, the method including modifying the weighting of at least the allograph containing the selected prototype.
- 11) A method according to any one of the claims 8 to 10, the method including selecting the prototype having the lowest selection value.
- 12) A method according to any one of the claims 8 to 11, the method of modifying the weightings including at least one of:
 - a) Decreasing the weighting for the allograph containing the selected prototype; and,
 - b) Increasing the weighting for one or more of the allographs that do not contain the selected prototype.

- 13) A method according to claim 12, the method including increasing and/or decreasing the weighting of the allographs by a predetermined amount.
- 5 14) A method according to claim 12, when dependent on claim 7, the method including:
- a) Selecting a reference prototype for each allograph to be modified;
 - b) Determining the distance value for each reference prototype;
 - c) Determining an average of the determined distance values; and,
 - d) Modifying the weighting of each allograph in accordance with the distance value of
- 10 the respective reference prototype and the average distance value.
- 15) A method according to claim 14, the reference prototype being the prototype having the lowest distance value for the respective allograph.
- 15 16) A method according to any one of claims 1 to 14, the method being performed using a processing system, the processing system including:
- a) A store for storing allograph data representing the allographs;
 - b) A processor, the processor being adapted to:
 - i) Obtain the hand-written character;
 - 20 ii) Select the allograph; and,
 - iii) Modify the one or more allographs.
- 17) Apparatus for modifying a classification scheme for classifying hand-written characters to thereby make the classification scheme user dependent, the apparatus including:
- 25 a) A store for storing allograph data representing a number of allographs, each allograph representing a respective style of a respective letter and having an associated weighting; and,
- b) A processor, the processor being adapted to:
 - i) Obtain at least one hand-written character;
 - 30 ii) Select an allograph representing the character; and,
 - iii) Modify the allograph data, thereby modifying the weighting of one or more of the allographs in accordance with the selection.

- 18) Apparatus according to claim 17, the store being adapted to store class data, the class data representing a number of classes, each class representing a respective letter and containing one or more of the allographs.
- 5 19) Apparatus according to any one of claims 17 or claim 18, the store being adapted to store prototype data representing prototypes, each allograph containing one or more prototypes, each prototype representing a variation in the respective allograph.
- 10 20) Apparatus according to any one of claims 17 to 19, the processor being adapted to perform the method of any one of the claims 1 to 16.
- 15 21) A method of classifying hand-written characters using a user dependent classification scheme, the classification scheme including a number of allographs, each allograph representing a respective style of a respective letter and having a respective weighting representing the user dependency, each allograph containing one or more prototypes, each prototype representing a variation in the respective allograph, the method including:
- a) Obtaining at least one hand-written character;
 - 20 b) For at least some of the prototypes, determining a selection value based on the similarity of the character and a respective prototype, and the weighting of the allograph containing the respective prototype;
 - c) Selecting a respective one of the prototypes in accordance with the determined selection values; and,
 - d) Classifying the character in accordance with the selected prototype.
- 25 22) A method according to claim 21, each prototype being represented as a prototype vector, each prototype vector being formed from a number of values, each value quantifying a respective feature of the respective prototype, the method of determining a selection value including:
- 30 a) Determining a feature vector representing the character, the feature vector being formed from a number of values, each value quantifying a respective feature of the character;

- b) Determining a distance value representing the distance between the feature vector and the respective prototype vector; and,
- c) Determining the selection value in accordance with the determined distance value and the weighting associated with the respective allograph.

5

23) A method according to claim 21 or claim 22, the classification scheme including a number of classes, each class representing a respective letter and containing one or more of the allographs, the method including selecting a class in accordance with the selected prototype to thereby determine the letter represented by the character.

10

24) A method according to any one of the claims 21 to 23, the classification scheme being modified in accordance with the method of any one of the claims 1 to 15.

15

25) A method according to any one of claims 21 to 24, the method including using a processing system, the processing system including:

- a) A store for storing:
 - i) Allograph data representing the allographs and the associated weightings; and,
 - ii) Prototype data representing the prototypes; and,
- b) A processor, the processor being adapted to:
 - i) Obtain the hand-written character;
 - ii) Determine the selection values;
 - iii) Select the prototype; and,
 - iv) Classify the character.

20

26) Apparatus for classifying hand-written characters using a user dependent classification scheme, the apparatus including:

- a) A store for storing:
 - i) Allograph data representing a number of allographs, each allograph representing a respective style of a respective letter, and a weighting associated with each allograph, the weighting representing the user dependency of the allograph; and,
 - ii) Prototype data representing one or more prototypes, each prototype representing a variation in the respective allograph; and,
- b) A processor, the processor being adapted to:

30

- i) Obtaining at least one hand-written character;
 - ii) For at least some of the prototypes, determine a selection value based on the similarity of the character and a respective prototype, and the weighting of the allograph containing the respective prototype;
 - 5 iii) Select a respective one of the prototypes in accordance with the determined selection values; and,
 - iv) Classify the character in accordance with the selected prototype.
- 27) Apparatus according to claim 26, the store being adapted to store class data representing
- 10 a number of classes, each class representing a respective letter and containing one or more of the allographs.
- 28) Apparatus according to claim 26 or claim 27, the processor being adapted to perform the method of any one of the claims 21 to 25.