Support Vector Machines for Mathematical Symbol Recognition

Christopher Malon[1], Seiichi Uchida[2], and Masakazu Suzuki[1]

[1] Engineering Division, Faculty of Mathematics, Kyushu University
6–10–1 Hakozaki, Higashi-ku, Fukuoka, 812–8581 Japan
[2] Faculty of Information Science and Electrical Engineering, Kyushu University
6–10–1 Hakozaki, Higashi-ku, Fukuoka, 812–8581 Japan

Abstract. Mathematical formulas challenge an OCR system with a range of similar-looking characters whose bold, calligraphic, and italic varieties must be recognized distinctly, though the fonts to be used in an article are not known in advance. We describe the use of support vector machines (SVM) to learn and predict about 300 classes of styled characters and symbols.

1 Introduction

Optical character recognition problems were considered very early in the development of support vector machines, with promising results [1]. However, the problem of OCR for mathematical documents is substantially more difficult than standard OCR problems for three principal reasons:

1. Although a variety of fonts is used in mathematical literature, when reading any single paper, it is important to keep appearances of italic, bold, roman, calligraphic, typewriter, and blackboard bold letters distinguished.
2. A rich set of symbols is used, and distinctions between letters may be more subtle than within the character set of a typical human language.
3. The symbols are not arranged in a simple one–dimensional pattern. Subscripts, superscripts, fractional relationships, and accents occur, and may be nested [2].

The *Infty Project[7]* in the Suzuki Laboratory at Kyushu University is developing *Infty Reader* software [5] to perform OCR of scanned mathematical journal articles, and produce output in languages that allow symbol relationships to be properly encoded, including TEXand MathML. Although Infty Reader nominally achieved 99 percent accuracy of single–letter recognition before this investigation (October 2005), its failure to distinguish certain common symbols would be bothersome to any serious user.

The Infty Project defined entities for about 600 characters and symbols used in mathematical research, and created a ground truth database identifying their appearances in page–by–page scans of hundreds of journal articles. Many character pairs could be distinguished in different styles by simple clustering techniques

applied to directional features measured in a mesh grid. Runtime accuracy exceeded 99% , but hundreds of letter pairs remained consistently problematic. We aim to improve the accuracy of single–character recognition through the use of support vector machines.

2   Test Data

The Infty character set comprises 1,571 Roman and Greek letters, numbers, and mathematical symbols, divided into 46 categories according to purpose and style. Further details of these characters appear in [8].

The Infty Project has selected journal articles representing diverse fields of higher mathematics, taken from a thirty year period. These articles were scanned page by page at 600 dpi to produce bitmap image files. The Infty OCR engine extracted the symbols from each page and recognized each symbol as a character from the Infty character set. College and graduate mathematics students manually inspected and corrected the results.

The results of this process appear in a "ground truth" database. Namely, for each character on a scanned page, a bitmap framed by the bounding box of that character is taken from the page. This bitmap is tagged with the correct identification of the symbol. [3] "Link information" describing the character's relationship to others on the page (subscripts, superscripts, limits, portions of a fraction, *etc.*) is also available in the database, but it is not utilized in the present study.

In fact, the Infty project has produced two databases of this kind. One, called InftyCDB-1, is freely available for research purposes upon request, and is summarized in [8]. The other is used internally by the Infty Reader OCR engine. We use the latter database in this experiment, because it has more data, and because it makes it easier to compare our results with those of the actual Infty Reader. Our data sample consists of 284,739 character symbols extracted from 363 journal articles. There are 608 different characters represented.

At random, we divide the 363 articles into three parts consisting of 121 articles each. The data from the corresponding articles is marked as "training", "selection", or "testing" accordingly. To make sure we had enough data to train and evaluate our classifiers, we examined only the characters with at least ten samples in training, selection, and testing portions of the database. This left 297 characters, pictured in Figure 1.

---

[3] Some bitmaps would not be identifiable solely on the basis of this bitmap. For example, a hyphen could not be distinguished from an underscore, without knowing its relationship to the baseline on the page when it was scanned. The original position on the page is part of the database, but this information was discarded prior to our experiment.

Big Symbol	Calligraphic
Greek Upright	Arrow
Latin Upright	Latin Italic
Greek Italic	Relational Operators
Binary Operators	Other Symbols
German Upright	Blackboard Bold
Accents	Punctuation
Symbol Fragments	Ligature Italic
Brackets	Ligature Upright

Fig. 1. Symbols with 10 training, selection, and testing samples.

3 Directional Features

Given an instance of a symbol, let w be its width and h be its height. Our feature vectors consist of the aspect ratio $(\frac{h}{w})$, followed by 160 floating–point coordinates of mesh directional feature data.

This mesh data is divided into "tall", "square", and "short" blocks of 48, 64, and 48 coordinates respectively. When the aspect ratio of a character exceeds 1.3, the tall block contains directional feature data computed from a 3×4 mesh; otherwise it contains zero–valued entries. When the aspect ratio of a character is between $\frac{1}{1.7}$ and 1.7, the square block contains directional feature data from a 4×4 mesh; otherwise it contains zero–valued entries. When the aspect ratio of a character is less than $\frac{1}{1.3}$, the short block contains directional features computed from a 4×3 mesh; otherwise it contains zero–valued entries. Thus, for any symbol, one or two (but never three) of the blocks are assigned nonzero entries.

We describe roughly the algorithm for associating directional feature data to an $m \times n$ mesh block. Divide the original bitmap horizontally into m equally sized lengths, and vertically into n equally sized lengths. Assign a "chunk" of four coordinates of the block to each of the $m \times n$ grid positions; initially, their values are zero. These four coordinates represent the horizontal and vertical directions, and two diagonals.

The contribution of part of the outline's direction to the mesh features is determined from its position in the bitmap, using a partition of unity. Given a positive integer r, our r–fold partition of unity consists of functions $p_i^r : [0,1] \rightarrow [0,1]$, $i = 0, \ldots, r-1$, with the property that p_i^r is supported on $[\frac{i-1}{r}, \frac{i+2}{r}]$.

Discard every isolated black pixel from the original bitmap. In the remaining bitmap, trace every outline between white and black pixels, following its chain code description. When visiting the pixel in location (x, y) during this trace, identify the direction (horizontal, vertical, diagonal one, or diagonal two) where the next pixel in the outline will be. For every i, $0 \le i < m$, and every j, $0 \le j < n$, add $p_i^m(\frac{x}{w}) \cdot p_j^n(\frac{y}{h})$ to the coordinate of the (i, j) chunk representing that direction.

After completing the trace of each outline component, divide all the values by the perimeter of the bounding box. This result gives the values to be entered in the corresponding block of the feature vector.

4   Naive classifier

Typically, a support vector machine learns a binary classification. There are various techniques for putting SVM's together to distinguish multiple classes; a comparison of some popular methods (1–vs–1, 1–vs–all, and the Directed Acyclic Graph) may be found in [4]. Except for the 1–vs–all method, these methods require the construction of $O(n^2)$ classifiers to solve an n–class classification problem. Because the Infty character set includes more than 1,500 entities, this seemed unnecessarily burdensome. Therefore, we try to extract an easier part of the classification problem that can be solved without SVM.

Taking the data assigned to the "training" portion of the database, we compute the mean feature vectors for the instances of each symbol. We create a naive classifier that assigns an input to the class whose mean feature vector is nearest, by Euclidean distance.

We run this naive classifier on the "selection" portion of the database, to produce a confusion matrix. The (i, j) entry of this matrix counts the number of samples in which a character truly belonging to class i was assigned to class j by this rule. The 297 by 297 confusion matrix we produced had 947 nonzero off-diagonal entries, an average of 3.2 misrecognitions per character.

We consider some of the misrecognitions to be too difficult for any classifier to resolve on the basis of our mesh of directional features. Particularly, we do not expect bold and non–bold variants of the same character to be distinguishable. Also, we do not expect upper and lower case variants of the letters C, O, P, S, V, W, X, and Z to be distinguishable in the same style, or in styles that are

identical except for boldness. Disregarding misrecognitions of these two kinds, 896 other nonzero off–diagonal entries remain in the confusion matrix.

For 62 of the 297 characters with ten training, selection, and testing samples, the naive classifier recognized less than half of the selection samples correctly. These characters are displayed in Figure 2 (a). In comparison, ninety percent

(a)   (b)
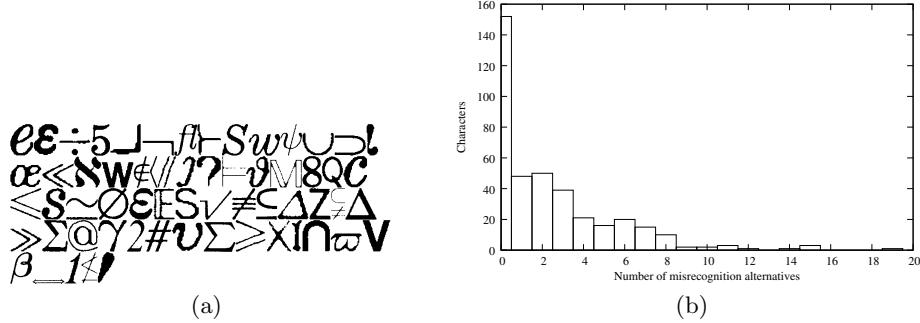
Fig. 2. The naive classifier: (a) Characters the naive classifier usually fails to recognize. (b) Histogram of distinct misrecognitions of an input character by the naive classifier.

accuracy is achieved for 197 of the 297 symbols, 95 percent accuracy for 163 symbols, and 99 percent accuracy for 123 symbols.

Although the confusion matrix is relatively sparse, certain troublesome characters have many misrecognition results, as can be seen in Figure 2 (b). For 95 of the 297 characters, at least four distinct characters occur as misrecognition results. Eleven letters (plain '1', '4', 'E', 'I', 'l', 'r', 's', 't', ' " ', and italic 'γ' and 'ψ') had ten or more distinct characters appear as misrecognition results.

At runtime, the naive classifier will be used to assign each letter to a cluster of possible candidates, consisting of the recognition result and the other candidates most likely to have produced that recognition result (as determined by our confusion matrix). The harder problem of distinguishing between the letters in each of these clusters will be assigned to support vector machines.

5   Linear SVM

Within each cluster, we will use the 1–to–1 approach to multiclass classification. This requires first creating a binary SVM for each pair of classes in the cluster.

Because they are simple and can be computed quickly, we begin our experiment with SVM's that use the linear kernel:

$$K(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x} \cdot \boldsymbol{y}. \tag{1}$$

The naive classifier, when restricted to two classes, can be thought of as the linear classifier determined by the hyperplane equidistant from the two cluster

centers. The support vector method enables us to search for hyperplanes in the original feature space that perform better on the training data.

There are no kernel parameter choices needed to create a linear SVM, but it is necessary to choose a value for the soft margin (C) in advance. Then, given training data with feature vectors $\boldsymbol{x_i}$ assigned to class $y_i \in \{-1, 1\}$ for $i = 1, \ldots, l$, the support vector machines solve

$$\min_{\boldsymbol{w}, b, \boldsymbol{\xi}} \quad \tfrac{1}{2} K(\boldsymbol{w}, \boldsymbol{w}) + C \sum_{i=1}^{l} \xi_i \tag{2}$$
$$\text{subject to } y_i(K(\boldsymbol{w}, \boldsymbol{x_i}) + b) \geq 1 - \xi_i$$
$$\xi_i \geq 0$$

where $\boldsymbol{\xi}$ is an l–dimensional vector, and $\boldsymbol{w}$ is a vector in the same feature space as the $\boldsymbol{x_i}$ (see, *e.g.*, [3]). The values $\boldsymbol{w}$ and b determine a hyperplane in the original feature space, giving a linear classifier. *A priori*, one does not know which value of soft margin will yield the classifier with the best generalization ability. We optimize this choice for best performance on the selection portion of our data, as follows.

Our basic parameter search method, here and in the following sections, is a grid search method that generalizes to any number of dimensions. For each candidate parameter assignment, we train an SVM with those parameters on the training portion of our data. Then we measure its performance on the instances of two classes that appear in the selection data. The score of the parameter is the minimum of the accuracy on the first class's input and the accuracy on the second class's input. *Hereafter, "accuracy" by itself, in the context of a binary classification problem, will refer to this score.*

Often, grid searches require a search interval to be specified for each dimension. Our approach requires only an initial parameter choice, and then grows the search region outward, until performance stops improving. Initial choices may matter under this grid search algorithm, if the algorithm terminates before reaching a selection of parameters that produces globally optimal results. This possibility seems unlikely as long as the resulting SVM performs better than random guessing in each case. The linear SVM problem has only the soft margin C as a parameter, and we set it initially to be 1024.

Table 1 displays the accuracy achieved by the linear SVM selected, on the testing data for pairs of symbols that the naive classifier sometimes confused.

We compared the chosen linear SVM classifier's performance on the letters where the naive classifier did not reach 100% accuracy, to the performance of the naive classifier. The 896 misrecognitions of the naive classifier comprise 795 unordered pairs of symbols. For nine of these pairs, both the naive classifier and the linear SVM always misassigned one of the two characters. Figure 3 (a) compares the performance of the two methods on the remaining 786 pairs. Of the 786 pairs, 34 did not perform as well under the linear SVM as with the naive classifier. The exact same performance was achieved on 95 pairs, and improvement occurred on 657 pairs. The histogram does not report the 24 symbols with more than a three–fold improvement in accuracy. Thirteen of these symbols

Accuracy >	Number of pairs	Accuracy >	Number of pairs
Total	795	.9	750
0	783	.95	742
.5	774	.97	720
.6	770	.99	684
.7	767	.995	650
.8	759	.999	609

Table 1. Linear SVM performance.

received zero accuracy from the naive classifier, for an infinite improvement in performance.

(a)   (b)
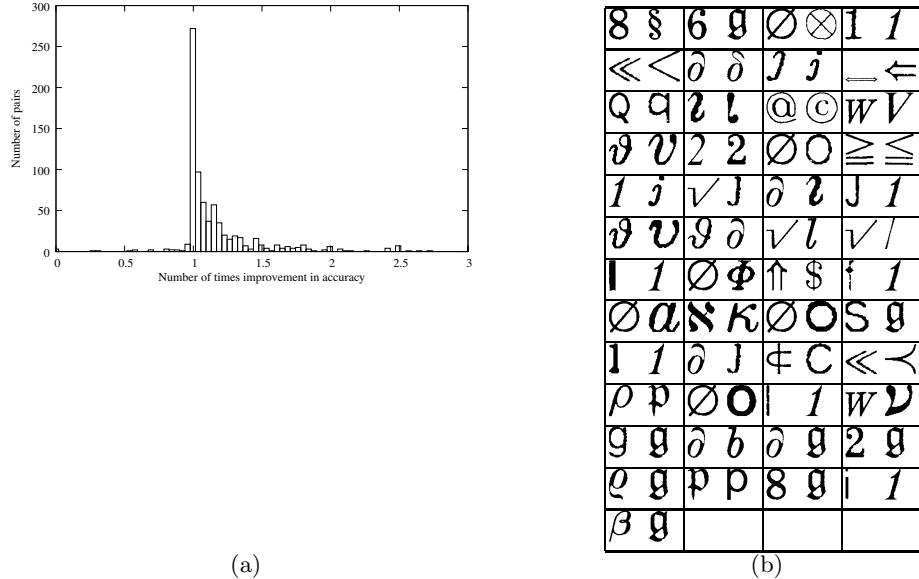
Fig. 3. Linear SVM improvement compared to naive classifier: (a) Histogram. (b) Pairs with two–fold improvement.

Figure 3 (b) illustrates the cases where the linear SVM achieved double the accuracy of the naive classifier.

6 Gaussian SVM

Just by using linear kernel support vector machines, our symbol recognition rates dramatically improved, but the use of a linear kernel severely limits the potential

benefit of a support vector machine. The use of a Gaussian (radial) kernel

$$K(\boldsymbol{x}, \boldsymbol{y}) = e^{-\gamma \|\boldsymbol{x} - \boldsymbol{y}\|^2} \tag{3}$$

in the SVM problem (2) effectively transforms the input feature space into an infinite–dimensional one, where the search for an optimal separating hyperplane is carried out. Classifiers of this form may perform better on classes whose feature data is not linearly separable in the original feature space. However, the addition of the parameter γ in the kernel definition makes the parameter search two–dimensional, adding computational expense to the selection of a classifier.

According to a result of Keerthi and Lin [6], given a soft margin C, the sequence of Gaussian SVM classifiers with kernel parameter γ and soft margin $\frac{C}{2\gamma}$ converges pointwise, as $\gamma \to 0$, to the linear SVM classifier with soft margin C. Thus, if our parameter search is wide enough, we should achieve higher accuracy with the Gaussian kernel than with the linear one.

We constructed Gaussian–kernel SVM classifiers for the 75 pairs of letters that the linear kernel failed to distinguish with 97% accuracy. A comparison of the performance of the chosen classifiers for each kernel type is given in Figure 4 (a). In Figure 4 (b), we display the eight pairs on which the Gaussian SVM

(a)   (b)
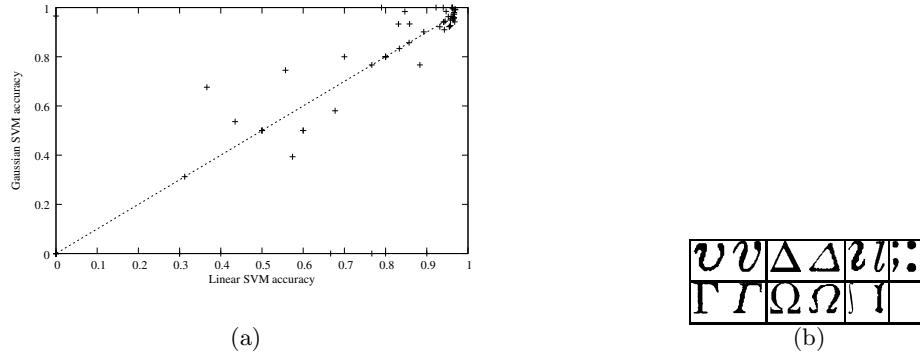
Fig. 4. Comparison of linear and Gaussian SVM: (a) Accuracies. (b) Pairs with 10% improvement from linear to Gaussian kernel.

performed with at least 10% higher accuracy than the linear SVM. The 31 pairs where Gaussian SVM accuracy falls below 80% are shown in Figure 5.

7 Conclusion

Even with the simplest kernel, the support vector method is strong enough to achieve good generalization accuracy on an optical character recognition problem that causes difficulty for simpler classification methods. We believe that our

Fig. 5. Pairs with under 80% accuracy by Gaussian SVM.

SVM results may be the best classification possible on the basis of the mesh of directional features we are using.

To distinguish the characters that confuse our SVM classifier, we plan to add new features. For example, by counting the number of connected components in a symbol, we could distinguish many variants of the greater–than sign ($>$). We also plan to record the convexity or concavity of a symbol as traced along its outline, to distinguish various nearly vertical characters. These features will be the topic for a future paper.

To our surprise, the SVM's we constructed with the Gaussian kernel did not show significantly stronger performance on the testing data. We attribute this phenomenon to the simple nature of our mesh of directional features. We plan to repeat this comparison after attaching a greater variety of features to our data.

References

[1] CORTES, C., AND VAPNIK, V. Support-vector networks. *Mach. Learn. 20*, 3 (1995), 273–297.

[2] ETO, Y., AND SUZUKI, M. Mathematical formula recognition using virtual link network. In *Sixth International Conference on Document Analysis and Recognition (ICDAR)* (2001), IEEE Computer Society Press, pp. 430–437.

[3] HSU, C.-W., CHANG, C.-C., AND LIN, C.-J. A practical guide to support vector classification. `http://www.csie.ntu.edu.tw/%7Ecjlin/papers/guide/guide.pdf`, July 2003.

[4] HSU, C.-W., AND LIN, C.-J. A comparison of methods for multi-class support vector machines. *IEEE Transactions on Neural Networks 13* (2002), 415–425.

[5] The infty project. `http://infty.math.kyushu-u.ac.jp`.

[6] KEERTHI, S. S., AND LIN, C.-J. Asymptotic behaviors of support vector machines with gaussian kernel. *Neural Comput. 15*, 7 (2003), 1667–1689.

[7] SUZUKI, M., TAMARI, F., FUKUDA, R., UCHIDA, S., AND KANAHORI, T. Infty: an integrated ocr system for mathematical documents. In *DocEng '03: Proceedings of the 2003 ACM symposium on Document engineering* (New York, NY, USA, 2003), ACM Press, pp. 95–104.

[8] SUZUKI, M., UCHIDA, S., AND NOMURA, A. A ground-truthed mathematical character and symbol image database. In *ICDAR '05: Proceedings of the Eighth International Conference on Document Analysis and Recognition (ICDAR'05)* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 675–679.