tion feature, which was made up of a vector of 128 on/off values, one for each General MIDI patch. This feature had one network dedicated to it.

These features in particular were selected because they were easy to implement and give a general description of recordings without being optimized to the particular genre taxonomy that was used. Although there is no doubt that twenty better features could be devised, these particular features were chosen simply to show that even non-optimal features could still perform well.

Feature	Explanation
Orchestration	Which of the 128 MIDI instruments are played
Number of instruments	Total number of instruments played
Percussion prevalence	Fraction of note-ons belonging to unpitched instruments
Dominant pitch prevalence	Fraction of note-ons corresponding to the most common pitch
Dominant pitch class prevalence	Fraction of note-ons corresponding to the most common pitch class
Dominant interval	Number of semi-tones between the two most common pitch classes
Adjacent fifths	Number of consecutive pitch classes separated by perfect 5ths that represent at least 9% of the notes
Pitch class variety (common)	Number of pitch classes that represent at least 9% of the notes
Pitch class variety (rare)	Number of pitch classes played at least once
Register variety	Number of pitches played at least once
Range	Difference between highest and lowest pitches
Pitchbend fraction	Number of pitch bends divided by total number of note-ons
Dominant periodicity	Magnitude of the highest periodicity bin
Second dominant periodicity	Magnitude of the second highest periodicity bin
Combined dominant periodicities	Combined magnitude of the two highest periodicity bins
Dominant periodicity strength ratio	Ratio of the frequencies of the two highest periodicity bins
Dominant periodicity ratio	Ratio of the periodicities of the two highest periodicity bins
Number of strong periodicities	Number of periodicity bins with normalized magnitude > 0.1
Number of moderate periodicities	Number of periodicity bins with normalized magnitude > 0.01
Number relatively high periodicities	Number of periodicity bins with frequencies at least 25% as high as the highest magnitude

Table 1: Features extracted from MIDI files and fed into neural networks.

4 Details of the Experiment

The training and testing data consisted of 225 MIDI files hand classified hierarchically into three parent genres (Classical, Jazz and Pop) and nine sub-genres (Baroque, Romantic, Modern Classical, Swing, Funky Jazz, Cool Jazz, Rap, Country and Punk). The particular files that were chosen were selected so as to represent each category as broadly as possible (e.g. the Baroque category included operas, violin concertos, harpsichord sonatas, etc., not just organ fugues, for example). This significantly increased the difficulty of the task, as each sub-genre only had 20 training recordings (five recordings were reserved for testing in each run) to learn a broad range of music. This was done in order to truly test the viability of the system and its features.

The recordings were classified using an array of eight feed-forward neural networks that consisted of four networks for identifying parent genres and four networks for identifying sub-genres. Each network had a single hidden layer. This division into two groups made it possible to classify parent genres independently from sub-genres.

The input units of each network took in different groups of features (orchestration, pitch statistics, rhythm statistics or stylistic), thus making it possible to study the relative success of the different features in classifying the test data. This made it possible to compare how well different feature groups performed.

A coordination system considered the certainty score output by the networks for each sub-genre in combination with the certainty for each parent genre, and produced a final classification using weighted averages.

This particular classification system was used because it allowed the independent comparison of different groups of features as well as a comparison of how well parent genres were classified relative to sub-genres.

5 Results

A five-fold cross-validation was used to test the performance of the system. The results are shown below:

	Set 1	Set 2	Set 3	Set 4	Set 5	Average
Classical	93	80	100	93	100	**93.2**
Jazz	73	80	60	53	40	**61.2**
Pop	100	100	100	100	100	**100.0**
Average	88.7	86.7	86.7	82.0	80.0	**84.8**

Table 2: Classification success rates (in percentages) for parent genres for all five cross-validation testing runs.

	Set 1	Set 2	Set 3	Set 4	Set 5	Average
Baroque	80	40	80	80	80	**72.0**
Romantic	0	40	0	20	40	**20.0**
Modern	100	40	100	40	80	**72.0**
Swing	40	80	20	40	20	**40.0**
Funky Jz.	60	40	60	40	0	**40.0**
Cool Jz.	40	20	20	20	0	**20.0**
Rap	80	60	80	60	20	**60.0**
Country	80	100	100	100	100	**96.0**
Punk	100	100	100	100	100	**100.0**
Average	64.4	57.8	62.2	55.6	48.9	**57.8**

Table 3: Classification success rates (in percentages) for sub-genres for all five cross-validation testing runs.

Overall success rates of 84.8% were achieved for parent genres and 57.8% for sub-genres across all five train-