

Analysis of Rachmaninoff's Piano Performances Using Inductive Logic Programming (Extended Abstract)

Matthew J. Dovey,
Programming Research Group,
Oxford University Computing Laboratory,
Oxford, OX1 3QD, Great Britain.

Introduction

The project outlined here is an attempt to use inductive logic programming ([6]) to determine various interpretative rules which the pianist, Sergei Rachmaninoff, may have used during his pianoforte performances. During the 1920's Rachmaninoff recorded a number of recitals on the Ampico Recording Piano ([4]). This method of capturing the performance not only recorded the notes, duration and tempo, but also the dynamics of the key pressure and pedalling, in a digital (in actual fact, binary) format, which easily lends itself to conversion into a computer readable representation. To complement this performance information, it was also necessary to represent the musical structure of the piece being performed, so that a general analysis could be achieved, rather than one specific to that particular piece. For simplicity, only the melodies of the pieces involved were subjected to analysis, although better results may be obtained from a full analysis of the accompanying harmonic and contrapuntal structures. The two sets of information, structural analysis and performance analysis were encoded into PROGOL scripts ([5]), which were used to attempt to determine general rules (in the form of universal predicates) underlying the data set. Two pieces were analysed in this manner: Rachmaninoff's Prelude in C Sharp Minor, Opus 3 Number 2 (Ampico Roll Number 57504) and Mendelssohn's Song Without Words Opus 67 Number 4 (Ampico Roll Number 59661).

Performance Data

The Ampico Recording and Reproducing Piano, devised by Charles Stoddart in the early 1900's, provides a digital and fairly accurate representation of the performance of many pianists of the time by recording and reproducing the dynamics of the piano keys in terms of their speed and pressure ([2]). However, there are a few caveats of which to be aware. Rather than recording the dynamics of each key individually, only two sets of dynamics were recorded: one set for the bass notes and one set for the treble notes. As the analysis for this project was only performed on the melody line of the pieces concerned, this was not a major consideration. The dynamics were recorded using a mixture of analogue and discrete techniques. Punches on the piano roll could be used to set the air-pressure in the pneumatic mechanism to one of seven different levels, giving seven different levels of volume. In addition, the roll could also control slow and fast *crescendos* (gradual increases in volume) and *decrescendos* (gradual decreases in volume), which enabled smooth articulation of the dynamics between the seven steps. There was also a control for the *una chorda* (left)

pedal, but since the mechanism used on Ampico grand pianos was that normally found on uprights, namely that of moving the hammers nearer the strings, thereby reducing the tone, this was used by the Ampico roll editors as a method to further modify the dynamics. Despite the fact that the rolls were edited, both to reduce any wrong notes and to capture the subtleties of the use of the damper pedal by extending the lengths of certain notes, the reproductions are fairly accurate to the original performances. Indeed an article by O'Connell contains a quantitative proof that Rachmaninoff's Ampico recording of the C Sharp Minor Prelude is very close to the 78 rpm disc recording ([7]). Using this data, therefore, it was possible not only to extract the timing and durations of the notes, but also their volume, which were quantised on a scale of 1 (quietest) to 14 (loudest).

Structural Data

As regards the encoding of the musical score, a number of different encoding techniques were used. On the surface level the encoding consisted of the note duration, dynamic as indicated in the score (on the same scale of 1 to 14 as for the performance dynamic), position in bar, pitch, touch (varying from *staccato*, i.e. detached, to *portamento*, i.e. smoothly linking between notes), accent (from *leggiere*, i.e. lightly to *sforzando*, i.e. heavily accented) and whether the note occurred at the beginning of, end of, or within a phrase. A more general encoding method consisted of merely noting whether the pitch of the melody was moving up, down or remaining on the same note. Such a method has been used to index musical themes with some success. A small modification was to encode high and low neighbourhoods, where a note is surrounded by two lower or higher notes, both of which have the same pitch. A more complicated attempt to encode the general structure of the music was to apply a method of classification devised by Cope for analysing compositional style ([1]). While Cope's classification can also be applied to harmony and rhythm in a hierarchical manner, here it was merely applied in a non-hierarchical way to individual notes, classifying their context within the musical phrase. In this scheme there are five classes or contexts: *Statements*, *Preparations*, *Extensions*, *Antecedents* and *Consequences* (the initials spelling SPEAC). These should not be confused with their use and meaning in logic programming. *Statements* can exist alone, and not as a result of other activity, with nothing expected except iteration. They can be prefaced by *Preparations* or lengthened by *Extensions*. *Antecedents* cause a significant musical implication in the structure and require musical resolution. They typically demand a *Consequence*. *Consequences* are often harmonically similar to *Statements* but have different musical implications in respect to their function and positioning within the musical phrase.

PROGOL Encoding

The PROGOL predicates devised take constants of the type "note" as their parameters, and defined various characteristics of each note in the piece. The predicates used can be divided into fourteen groups, the first six of which describe the performance of the note, while the remaining eight describe various properties of the note as indicated in the score. The six performance-related groups are: *rubato_hold* - describing whether the note is played longer or shorter than indicated in the score;

rubato_detach - describing whether the time interval between the note and the preceding note is longer or shorter than indicated; *rubato_overlap* - describing whether the overlap (if any) of the note with the preceding note is longer or shorter than indicated; *rubato_timing* - indicating whether the note occurs earlier or later in the bar than indicated; *dynamic_force* - describing whether the note is played louder or softer than indicated; and *dynamic_actual*, describing crescendos or decrescendos between the note and the preceding one. The nine musical score related groups are: *accent*, *touch*, *phrasing* - which were encoded directly from the relevant information mentioned in the section on “Structural Data” above; *structure* - illustrating the note’s context using the SPEAC system; *pitch* - which describes whether the note is higher or lower than the preceding note and if it is a low or high neighbourhood note (as described above); *dynamic_expected* - describing whether the note as written is louder or softer than the preceding note; *syncopation* - describing whether the note occurs on a beat, or barline; *rhythm* - describing the value of the note (e.g. crotchet, which equals one beat, or minim, which equals two beats, or quaver, which equals half a beat, etc.). Two other predicates used are: *note*, to define the type (“note”) of the constants used (which were of the form *note_1*, *note_2*, etc.), and finally *pre*, which is used to indicate the ordering of the notes in the score. Certain simplifications were introduced at this point. No predicates took integers as parameters, so comparisons are qualitative only, e.g. the *rubato_hold* predicates merely indicate that a note is shorter, longer, or over twice as long or short than indicated in the score, but does not give a more precise value as to how much longer or shorter. Also the PROGOL searches involved in the analysis only processed chains of up to four consecutive notes.

Since the first six groups of predicates (the performance-related ones) could appear in both the head and the body of the rules generated, six PROGOL scripts were generated for each piece, so that each of these six groups took it in turns to be goals, (i.e. they were declared with *modeh* declarations while the others were declared with *modeb* declarations). In each case, the negative data consisted of the set of negative predicates from the *modeh* declared group which were not properties of each note in turn. In all six scripts, the positive data from all fourteen groups were included. In this manner, the positive and negative examples come from the *modeh* declared group while the other groups form the background knowledge.

Results and Conclusions

A number of expected result came out of the analysis as well as some more interesting results, including a number of common or similar rules between the pieces, hinting at a performance style rather than predicates merely generated by style of the music being performed. Amongst the expected results is a one indicating that a note marked staccato resulted in a severe shortening of the note, followed by a corresponding lengthening of the gap before the next note. A more satisfying predicate is one characterising Rachmaninoff’s technique of accentuation described by Martyn as his “rhythmic snap” ([4]), namely that he slightly anticipates notes occurring at the beginning of a bar. Other rules found include a tendency to exaggerate the detachment between two notes if he were to linger on the second note and to stick closely to the written dynamic of a note when playing it slightly quicker than indicated.

Future Research Directions

There are a number of areas in which further research could be done, and there is a wealth of data in the form of Piano rolls (from other manufacturers as well as Ampico) and that supplied from modern performances through MIDI. A more comprehensive structural analysis than the SPEAC system outlined here could be applied, making use of other forms of musical analysis such as Schenkerian analysis ([3]); the structural analysis could incorporate the entire score rather than just the melody, possibly with some form of automation using optical character recognition to scan the score and an expert system to perform the musical analysis; the use of more powerful inductive logic tools, allowing longer chains in the search space and the identification of numeric ranges rather than the qualitative analysis performed in this project. However, even with the shortcomings of the method described here, I believe a number of interesting results were obtained, indicating that this may be a rich avenue for further investigation.

Acknowledgements

I acknowledge support given by the ESPRIT Basic Research Action 6020 on Inductive Logic Programming. I would also like to thank Stephen Muggleton for his assistance and guidance as regards inductive logic programming, and Rex Lawson and Dennis Hall for their help and advice as regards the Ampico Reproducing Piano. Finally, I would like to thank Sergei Rachmaninoff for his piano recordings.

Bibliography

- [1] David Cope, *Computers and Musical Style*, Oxford University Press 1991.
- [2] John Farmer, "The Reproducing Piano" in *Journal of the Pianola Institute*, No. 6 1993.
- [3] Allen Forte and Steven Gilbert, *Introduction to Schenkerian Analysis*, W. W. Norton 1982.
- [4] Barrie Martyn, *Rachmaninoff - Composer, Pianist, Conductor*, Scolar Press 1990.
- [5] Stephen Muggleton, *Inductive Logic Programming*, SIGART Bulletin, Vol. 5 Num. 1, 1994.
- [6] Stephen Muggleton (editor), *Inductive Logic Programming*, Academic Press 1992.
- [7] Sid O'Connell, "Rachmaninoff Performance Analysis" in *Rachmaninoff Society Journal*, No. 20 September 1994.