

The Musical Expression Project: A Challenge for Machine Learning and Knowledge Discovery

Gerhard Widmer

Dept. of Medical Cybernetics and Artificial Intelligence, University of Vienna, and
Austrian Research Institute for Artificial Intelligence, Vienna
gerhard@ai.univie.ac.at

Abstract. This paper reports on a long-term inter-disciplinary research project that aims at analysing the complex phenomenon of *expressive music performance* with machine learning and data mining methods. The goals and general research framework of the project are briefly explained, and then a number of challenges to machine learning (and also to computational music analysis) are discussed that arise from the complexity and multi-dimensionality of the musical phenomenon being studied. We also briefly report on first experiments that address some of these issues.

1 Introduction

This paper presents a long-term inter-disciplinary research project situated at the intersection of musicology and Artificial Intelligence. The goal is to develop and use machine learning and data mining methods to study the complex phenomenon of *expressive music performance* (or *musical expression*, for short). Formulating formal, quantitative models of expressive performance is one of the big open research problems in contemporary (empirical and cognitive) musicology. Our project develops a new direction in this field: we use *inductive learning techniques* to discover general and valid expression principles from (large amounts of) real performance data. The project, financed by a generous research grant by the Austrian Federal Government, started in early 1999 and is intended to last at least six years. The research is truly inter-disciplinary, involving both musicologists and AI researchers. We also expect to contribute new results to both disciplines involved, and our first experimental results show that this is realistic — for instance, in [26] both a new, general rule learning algorithm and some interesting, novel musical discoveries are presented.

In recent years, there has been an increasing number of attempts, in the field of empirical musicology, to formulate quantitative, mathematical or computational models of (aspects of) expressive performance (e.g., [1, 12, 13, 16–21]). This work has produced a wealth of detailed hypotheses and insights, but has often been based on rather limited sets of performance data (which were sometimes also produced under ‘laboratory conditions’). What distinguishes our project is the use of large amounts of ‘real-world’ data, and the application of inductive

learning methods to discover interesting and possibly novel patterns and regularities in the data. In short, we aim at performing the most data-intensive investigations ever done in musical expression research.

The purpose of the present paper is to give an overview of the project and its current state, and to discuss the challenges that this application problem presents to machine learning and knowledge discovery. In section 2, we explain the basic notions of expressive music performance. Section 3 sketches the general research framework of the project and briefly touches upon the enormous difficulties involved in data acquisition and preparation (an aspect often neglected in machine learning publications). Section 4 looks at the problem from a machine learning point of view and discusses some of the particular challenges posed by the complex nature of the target phenomenon. Section 5 briefly summarizes some interesting results obtained so far and talks about some of our ongoing research.

2 Expressive Music Performance

When played exactly as notated in the musical score, a piece of music would sound utterly mechanical and lifeless; it is both unmusical and physically impossible for a musician to perform a piece with perfectly constant tempo, even loudness, etc. What makes a piece of music come alive (and what makes some performers famous) is the art of *music interpretation*, that is, the artist's understanding of the structure and 'meaning' of a piece of music, and his/her (conscious or unconscious) expression of this understanding via *expressive performance*: a performer shapes a piece by continuously varying important parameters like tempo, dynamics (loudness), articulation, etc., speeding up at some places, slowing down at others, stressing certain notes or passages by various means, and so on. It is this shaping that can turn a lifeless piece of music into a moving experience, and that also makes both the composer's and the performer's ideas clear to the listener. What types of parameters are at a performer's disposal partly depends on the instrument being played, but the most important dimensions are *tempo* and *timing*, *dynamics* (variations in loudness), and *articulation* (basically, the way successive notes are connected).

Expressive music performance plays a central role in our current musical culture, and musicologists are showing increased interest in understanding exactly what it is that artists do when they play music. Are there explainable and quantifiable principles that govern expressive performance? To what extent and how are 'acceptable' performances determined by the (structure of the) music? What are the cognitive principles that govern the production (in the performer) and the perception (in the listener) of expressive performances? And what does this have to do with how we experience music?

Our project hopes to contribute to answering the first two of these questions. We collect precise measurements of performances by skilled musicians, and try to detect patterns and regularities (and intelligible characterizations of these) via inductive learning. As we also enable the computer to recognize structural

aspects of the music, potential relationships between expressive patterns and musical structure should emerge naturally from these investigations.

This approach is based on earlier work by the author [23,24], where it was shown that given some knowledge about musical structure, a computer can indeed learn general performance rules that produce rather sensible ‘interpretations’ of musical pieces. The central problem with these early studies was a lack of real performance data (the investigations were based largely on performances by the author himself). In our current work, we go beyond this by working with large collections of performances by skilled musicians, recorded on special instruments (pianos) that precisely measure and record each action of the performer. Ideally, we would also like to study the performance style of famous artists, on the basis of, e.g., audio CDs, but that will depend on the availability of computational methods for precise musical information extraction from audio, which is still an open problem in signal processing.

3 The Project: A High-level View

To give the reader an impression of the complexity of such a ‘real-world’ knowledge discovery project, Fig. 1 sketches the overall structure of our approach. As explained above, the basic goal is to take recordings of pieces as played by musicians, measure the ‘expressive’ aspects (e.g., tempo fluctuations) in these, and apply some machine learning algorithms to these measurements in order to induce general, predictive models of various aspects of expressive performance (e.g., a set of classification or regression rules that predict the tempo deviations a pianist is likely to apply to a given piece). These models must then be validated, e.g., by comparing them to theories in the musicological literature, by applying them to new pieces and analysing the musical quality of the resulting computer-generated performances, and, of course, by measuring their generalization accuracy on unseen data. All this is sketched in the lower half of Fig.1.

However, the story is much more complex. The problems involved in acquiring and pre-processing the data turned out to be formidable and forced us to develop a whole range of novel music analysis algorithms. And since we spent so much effort on these issues, I take the liberty of at least briefly mentioning them here.

3.1 Data Acquisition

The first problem was obtaining high-quality performances by human musicians (e.g., pianists) in machine-readable form. There are currently no signal processing algorithms that can extract the precise details of a performance from audio signals, so we cannot use sound recordings (e.g., audio CDs) as a data source. Our current source of information is the Boesendorfer SE290, a high-class concert grand piano that precisely measures every key, hammer, and pedal movement and records these measurements in a symbolic form similar to MIDI (though with higher precision). We did eventually manage to get large sets of performances that had been recorded on this instrument by a number of excellent pianists.

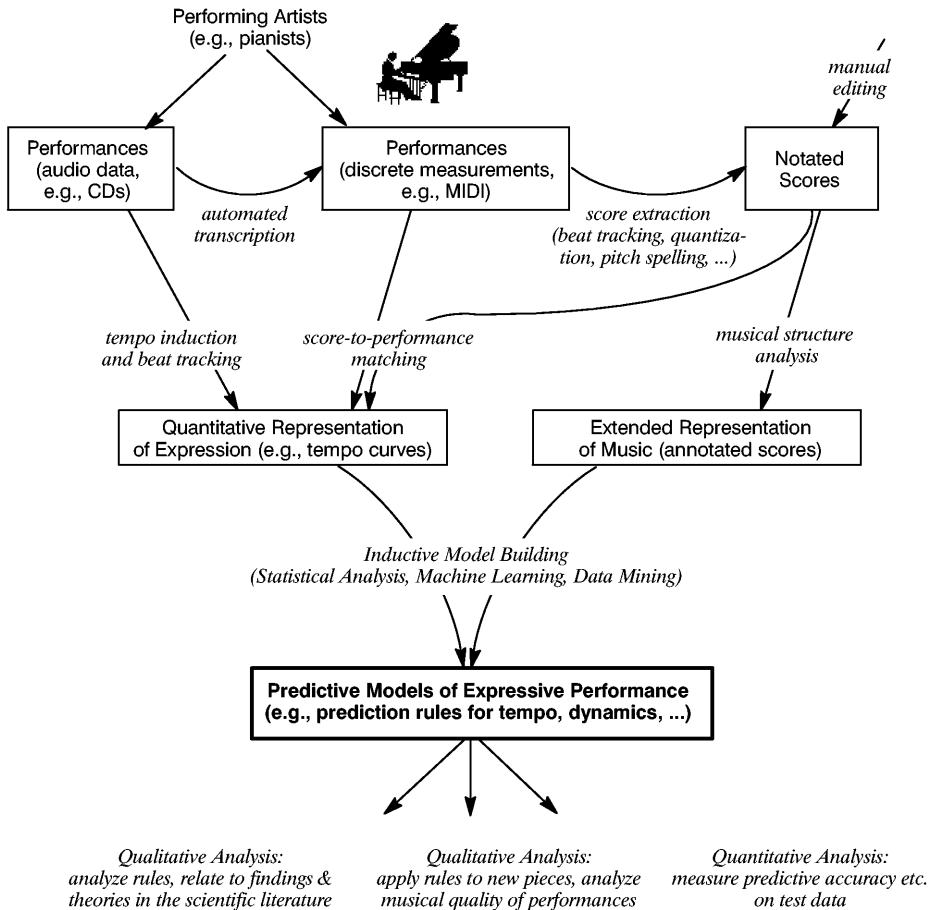


Fig. 1. The research framework: a sketch of data processing/analysis steps.

For instance, we currently have performances of 17 complete piano sonatas by W.A. Mozart as played by a highly skilled concert pianist. This data set corresponds to some 5 1/2 hours of music and contains around 150.000 notes. We also have performances, by a famous Russian pianist, of essentially the entire piano works by Frédéric Chopin (more than 9 hours of music, 300.000 notes, 2 million pedal measurements). This is a huge amount of data indeed; in fact, it is by far the largest collection of detailed performance measurements that has ever been compiled and studied in expression research.

Another line of current research, which cannot be discussed here, concerns the extraction of performance information directly from *digital audio* data, e.g., audio CDs [9] (see top left corner in Fig. 1). This will eventually allow us to also study at least certain limited aspects of expression in arbitrary recordings by famous artists.

3.2 Data Preprocessing

Preprocessing these data to make them usable for analysis and machine learning is a formidable task. What we need is not only the performances (i.e., information about how the notes were played), but also the notated music score (i.e., information about how the notes ‘should be’ played) and the exact note-to-note correspondence between the two. Manually coding musical scores consisting of tens of thousands of notes is not feasible; in order to get at the scores, we had to develop computational methods for extracting (re-constructing) the score information from the expressive performances themselves. The result is a whole range of new algorithms for music analysis problems like *beat induction* and *tempo tracking* i.e., inferring the metrical structure of the piece in the face of (sometimes rather extreme) tempo changes [10, 11], *quantization*, i.e., inferring the ‘intended’ onset times and durations of notes in the underlying score [2], and inducing the *correct enharmonic spelling* of notes (e.g., G \sharp vs. Ab) [3], which is not merely an aesthetic issue, but absolutely vital for the correct interpretation of a musical passage.

The ‘raw’ score files extracted by these algorithms from the performance data (up to now, some 150.000 lines of text) still needed to be manually corrected and further annotated. And finally, the resulting score files were matched, in a semi-automatic process, with the performance files to establish the exact note-to-note correspondence; thousands of notes were manually identified and labelled as missing or extraneous (most of these are related to ornaments like trills etc.). From this information we could then finally compute all the detailed aspects of a performer’s expressive playing (e.g., tempo changes, articulation details etc.) that serve as training data in the inductive learning process.

3.3 Enhancing the Data: Musical Structure Analysis

The next problem concerns the representation of the music. What we are searching for are systematic connections between the *structure* of the music (e.g., harmonic, metrical, and phrase structure) and patterns in the performances (e.g., a gradual rise in loudness (*crescendo*) over a given phrase). The representation of the musical pieces must therefore be extended with an explicit description of certain structural aspects. Again, a complete manual analysis of a large number of complex pieces is infeasible or at least highly impractical, so there is a need for computational methods. In the context of our project, we have developed a number of new music analysis algorithms that make explicit different structural aspects of a piece such as its segment structure [5], categories of melodic motifs and their recurrence [6], and various types of common melodic, harmonic, and rhythmic patterns, as postulated by music theorists [15]. These algorithms are of general interest to musicology, as they constitute formal computational models of aspects of musical structure understanding that had not hitherto been sufficiently formalized in music theory. The analyses computed by these tools can be used as additional descriptors in the representation of musical pieces.




Fig. 2. Frédéric Chopin, Etude Op.10 No.3, mm.1–9, as played by three different pianists: dynamics (relative to average dynamics of entire piece) and tempo (relative to average tempo) of the melody.

4 Challenges to Machine Learning

The result of all these efforts are training data as exemplified in Fig. 2, which shows the dynamics and tempo deviations extracted from performances, by three different pianists, of a well-known piece by Frédéric Chopin. For the moment, we restrict our attention to how the *melodies* of the pieces are played (and neglect more complex aspects like interactions between different voices of a piece). All the three expression dimensions that we are currently studying — tempo/timing, dynamics, and articulation — can then be represented as curves that associate a particular tempo, loudness, or relative duration value with each melody note. Fig. 2 also contains a number of annotations added by the author to highlight different structural aspects of both the piece itself, and the performances. These will be of help in the following discussion.

The first question one might ask is: is there something to be learned at all? Isn't expressive performance something intangible, something that reflects the artistic uniqueness of a performer and thus necessarily escapes any attempt at formalisation or explanation? A look at Fig. 2 reassures us: there are, of course, individual differences in the interpretations by the three pianists, but there are also very clear commonalities in the three curves. In other words, there seem to be some strong common principles at work that lead performers to do things in a similar way. And these common performance patterns must somehow be determined by the structure of the music being played.

In fact, this situation affords opportunities for at least two different types of learning. The one that better fits the 'traditional' inductive learning setting is learning to characterise and predict the *commonalities* between performances. In the simplest case, a learner that is given different performances of the same piece can be expected to find descriptions of those patterns that are common to most of the performances, and treat those situations where the individual performers differ as *noise*.¹ Characterising common performance patterns that point to some fundamental underlying principles is indeed the primary goal of our project. But it would also be interesting to try to learn about characteristic *differences* between individual artists. Here, the problem is not to find out where two performers differ — that is directly obvious from the data — but to find classes of situations in which there is a *systematic* and *explicitly characterisable difference* in behaviour. This might be a novel problem for machine learning.

A related question is *how much* we can expect to learn and formalise. Clearly, we cannot expect artists to be entirely predictable. We will have to make do with models that explain only a (possibly small) fraction of the observed phenomena. This requirement favours learners that, rather than trying to cover all of the instance space, can focus on those subspaces where something can be learned and produce models that clearly indicate when something is outside their area of expertise. Some interesting results along these lines are reported in [26].

¹ 'Real' noise (in the sense of mistakes or inaccuracies by the performer) is not much of a problem — high-class pianists are extremely precise, both in terms of motor control and in terms of their memory and capacity to reproduce particular expressive patterns over repeated performances.

Another fundamental question is: what are the target concepts? And that relates to a number of deep problems concerning representation, abstraction level, and context. At first sight, the curves in Fig. 2 are reminiscent of *time series*, which suggests the use of methods from time series analysis and forecasting. However, this is an inappropriate view. It is not so much the past states that determine how a curve is going to continue into the future; it is the structure of the underlying musical piece that partly determines what ‘shapes’ or ‘envelopes’ (in tempo, dynamics, etc.) a performer will apply to the music. The question then arises as to what exactly the scope of these ‘shapes’ is, and what the structural units in the music are to which these ‘shapes’ are applied — in other words, what is the appropriate *abstraction level*?

Actually, musical expression is a *multi-level phenomenon*. Good performances exhibit structure at several levels. Local deviations expressing detailed nuances (e.g., the stressing of a particular note) will be embedded in more extended, higher-level expressive shapes, such as a general *accelerando–ritardando* (speeding up – slowing down) over an entire phrase. For instance, the expression curves in Fig. 2 exhibit both local, note-level (see notes marked by asterisks) and more global structural patterns (e.g., a clear *crescendo–decrescendo* applied the medium-level phrase A.1 (dynamics curve), and an ever so slight *accelerando–ritardando* over phrase A in the tempo dimension). Thus, it will be necessary to learn models at different structural abstraction levels, which introduces the additional problems of discerning and separating multiple pattern levels in given training observations, and of combining learned models of different granularity at prediction time. Moreover, apart from the note and the phrase levels, there may be other, intermediate structural units relevant to explaining certain aspects of the curves. Discovering these is an intriguing musicological problem. One of our plans here is to study the utility of new *substructure discovery algorithms* [8].

Generally, the *representation* problem is a non-trivial one. There are many conceptual frameworks in which music can be described. Finding the most appropriate music-structural descriptors is a question of musicological interest. Systematic experimentation with different music-theoretic vocabularies will be necessary to identify these. In addition, the representation should capture the relevant *context* of notes and musical structures, which is a tricky issue not only because we do not know exactly how large this context should be, but also because there are also some highly *non-local* effects at work (e.g., when the recurrence of a melodic motif prompts the performer to ‘fall back’ into a previous pattern). As for the essentially *relational* nature of music, which would suggest the use of first-order logic for knowledge representation and Inductive Logic Programming for learning, it will be a matter of experimentation to study the trade-off between the increase in expressive power and the increase in search complexity implied by the use of ILP algorithms (see [14]).

Another interesting observation, which may be a source of new learning problems, is that the different target dimensions are very likely to *interact* or be *inter-dependent*. The performances in Fig.2 exhibit some clear parallels between dynamics and tempo, particularly in the case of some local deviations. For in-

stance, there seems to be a strong correlation between dynamic emphasis and individual note lengthening (see the events marked by asterisks in Fig. 2). At a higher level, one could construe a certain parallelism between the dynamics and tempo shapes of the second of the high-level phrases (B) (see the arcs in the dynamics and tempo plots), which would confirm a general hypothesis by musicologists.² In general, performers have different means of stressing musical passages, by combining timing, dynamics, and articulation in certain ways. This suggests that expressive performance might be an ideal candidate domain for *multi-task learning* [7], where multiple learning tasks are pursued in parallel using a shared representation, which presumably enables the learner to transfer information between different related problems. Moreover, we would be interested in an explicit *characterisation* of the connection between, say, timing and dynamics, if there is one. This seems to be a new type of learning problem.

And finally, there is the *evaluation* problem. How is one to evaluate and quantify the validity of a given theory in a domain where there is no unique ‘correct’ solution (there are usually many ‘acceptable’ ways of performing a piece)? The empirical evaluation methods used in machine learning (measuring classification accuracy and prediction error on unseen data, estimating true error via cross-validation etc.) do have their place here, but they need to be complemented with more music-specific methods that, while avoiding to make judgments concerning the musical or aesthetic quality of a performance, do account for musical aspects of a model’s predictions. This is a challenging research question for musicology and is beyond the scope of the present paper.

5 First Results and Ongoing Research

It is only rather recently that we have begun to perform systematic learning experiments with the huge data collections mentioned in section 3.1, so most of the above questions and challenges are still open. Our investigations so far have mostly concentrated on the *note level*, i.e., on describing and predicting how individual notes will be played, given various features of the notes and their immediate context. Here is a brief list of the most interesting results so far:

Basic learnability: In a first suite of experiments [25], we succeeded in showing that even at the level of individual notes, there is structure that can be learned. Standard inductive learners managed to predict the performer’s choices with better than chance probability. Extended feature selection experiments showed that different sets of music-theoretic descriptors are relevant for different expressive dimensions (timing, dynamics, articulation).

New rule learning algorithm: Based on experiences gathered in these initial investigations, we developed a new rule learning algorithm named PLCG that can find simple partial theories in complex data where neither high coverage nor high precision can be expected. The PLCG algorithm and some

² In fact, this parallelism becomes clearer once certain local distortions and artifacts in the expression curves (caused, e.g., by the grace notes in bars 7 and 8) are removed.

experiments with it are described in more detail elsewhere in this volume [26].

Partial note-level rule model: PLCG has discovered a number of surprisingly simple and surprisingly general and robust³ note-level expression principles [27, 28]. These rules are currently investigated more closely from a musicological perspective; some of them will probably form the nucleus of a quantitative rule-based model of note-level timing and articulation.

Learning at higher structural levels: In some limited earlier studies [23, 24], we had already found indications that learning at multiple structural levels does indeed improve the results (and the musical quality of the resulting computer-generated performances) considerably. However, the definitions of these higher musical levels and particularly the methods for combining learned theories of different granularity were very *ad hoc*, and the training material was extremely limited. We are currently developing a more principled approach.

Discovering stylistic differences: Regarding the possibility of discovering *stylistic differences* between different performers, we had obtained first indirect positive evidence in an early experiment that involved performances of the same piece by both the famous Vladimir Horowitz and a number of advanced piano students [22]. There it turned out that rules learned from Horowitz yielded a significantly higher predictive accuracy on other Horowitz data than on the student data, and vice versa. Recently, we have started new focussed investigations on this issue, with the aim of finding *characterisations* of these differences. This can be done with standard inductive rule learning algorithms, but requires the design of a different type of learning scenario. In a small initial experiment, several interesting rules were discovered that might describe characteristic differences in behaviour between the two great pianists Alfred Cortot and V. Horowitz. But the data were much too limited to permit general conclusions. We are now planning to repeat this type of experiments with a much more extended data set.

Machine learning for structural music analysis: And finally, computational music research offers many other opportunities for machine learning that are not necessarily related to the performance issue itself. There are many problems in *automated structural music analysis* for which there are as yet no reliable algorithms (e.g., harmonic analysis, phrase structure analysis, etc.) and which could benefit from inductive learning. For instance, we have developed an algorithm for finding classes of musical motifs and for elucidating the motivic structure of a piece, based on a new *clustering method*. This algorithm has been shown to be capable of reproducing motivic analyses by human musicologists of such complex pieces as Schumann's *Träumerei* and Debussy's *Syrinx* [6], and of predicting the categorizations made by human listeners [4].

³ For instance, 4 simple timing rules turn out to be sufficient for correctly predicting more than 20% of a pianist's local *ritardandi*, and these rules seem to generalize well to music of different styles.

Obviously, these are just first steps in a long research journey that should take us closer to our final goal — a quantitative, composite computational theory that explains as much as possible of the various dimensions of expressive music performance, and the interactions between them — and that will force us to address a number of novel machine learning problems on the way. This is a long-term undertaking, and we would like to extend an invitation to motivated young researchers to join our project team and work with us towards this goal.

Acknowledgements

The project is made possible by a very generous START Research Prize by the Austrian Federal Government, administered by the Austrian *Fonds zur Förderung der Wissenschaftlichen Forschung (FWF)* (project no. Y99-INF). Additional support for our research on machine learning and music is provided by the European project HPRN-CT-2000-00115 (MOSART). The Austrian Research Institute for Artificial Intelligence acknowledges basic financial support by the Austrian Federal Ministry for Education, Science, and Culture. I would like to thank my colleagues Emiliós Cambouropoulos, Simon Dixon, and Werner Goebl for their cooperation and many fruitful and enjoyable discussions.

References

1. Bresin, R. (2000). *Virtual Virtuosity: Studies in Automatic Music Performance*. Doctoral Dissertation, Royal Institute of Technology (KTH), Stockholm, Sweden.
2. Cambouropoulos, E. (2000). From MIDI to Traditional Musical Notation. In *Proceedings of the AAAI'2000 Workshop on Artificial Intelligence and Music*, 17th National Conference on Artificial Intelligence (AAAI'2000), Austin, TX. Menlo Park, CA: AAAI Press.
3. Cambouropoulos, E. (2001). Automatic Pitch Spelling: From Numbers to Sharps and Flats. In *Proceedings of the 8th Brazilian Symposium on Computer Music*, Fortaleza, Brazil.
4. Cambouropoulos, E. (2001). Melodic Cue Abstraction, Similarity, and Category Formation: A Formal Model. *Music Perception*, 18(3) (in press).
5. Cambouropoulos, E. (2001). The Local Boundary Detection Model (LBDM) and its Application in the Study of Expressive Timing. In *Proceedings of the International Computer Music Conference (ICMC'2001)*. San Francisco, CA: International Computer Music Association.
6. Cambouropoulos, E. and Widmer, G. (2000). Automatic Motivic Analysis via Melodic Clustering. *Journal of New Music Research*, 29(4) (in press).
7. Caruana, R. (1997). Multitask Learning. *Machine Learning* 28(1), 41–75.
8. De Raedt, L. & Kramer, S. (2001). The Levelwise Versionspace Algorithm and its Application to Molecular Fragment Finding. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, Seattle, WA.
9. Dixon, S. (2000). Extraction of Musical Performance Parameters from Audio Data. In *Proceedings of the First IEEE Pacific-Rim Conference on Multimedia (PCM 2000)*, Sydney, Australia.

10. Dixon, S. (2001). Automatic Extraction of Tempo and Beat from Expressive Performances. *Journal of New Music Research* (in press).
11. Dixon, S. and Cambouropoulos, E. (2000). Beat Tracking with Musical Knowledge. In *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI-2000)*, Berlin. IOS Press, Amsterdam.
12. Friberg, A. (1995). *A Quantitative Rule System for Musical Performance*. Ph.D. dissertation, Department of Speech Communication and Music Acoustics, Royal Institute of Technology (KTH), Stockholm.
13. Friberg, A., Bresin, R., Frydén, L., and Sundberg, J. (1998). Musical Punctuation on the Microlevel: Automatic Identification and Performance of Small Melodic Units. *Journal of New Music Research* 27(3), 271–292.
14. Kramer, S. (1999). *Relational Learning vs. Propositionalization. Investigations in Inductive Logic Programming and Propositional Machine Learning*. Ph.D. thesis, Technical University of Vienna.
15. Narmour, E. (1992). *The Analysis and Cognition of Melodic Complexity: The Implication-Realization Model*. Chicago, IL: University of Chicago Press.
16. Palmer, C. (1988). *Timing in Skilled Piano Performance*. Ph.D. Dissertation, Cornell University.
17. Repp, B. (1992). Diversity and Commonality in Music Performance: An Analysis of Timing Microstructure in Schumann’s ‘Träumerei’. *Journal of the Acoustical Society of America* 92(5), 2546–2568.
18. Shaffer, L.H. (1980). Analyzing Piano Performance: A Study of Concert Pianists. In G.Stelmach and J. Requin (eds.), *Tutorials in Motor Behavior*. Amsterdam: North-Holland.
19. Sundberg, J., Friberg, A., and Frydén, L. (1991). Common Secrets of Musicians and Listeners: An Analysis-by-Synthesis Study of Musical Performance. In P. Howell, R. West & I. Cross (eds.), *Representing Musical Structure*. London: Academic Press.
20. Todd, N. (1989). Towards a Cognitive Theory of Expression: The Performance and Perception of Rubato. *Contemporary Music Review*, vol. 4, pp. 405–416.
21. Todd, N. (1992). The Dynamics of Dynamics: A Model of Musical Expression. *Journal of the Acoustical Society of America* 91, pp.3540–3550.
22. Widmer, G. (1996). What Is It That Makes It a Horowitz? Empirical Musicology via Machine Learning. In *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI-96)*, Budapest. Wiley & Sons, Chichester, UK.
23. Widmer, G. (1996). Learning Expressive Performance: The Structure-Level Approach. *Journal of New Music Research* 25(2), pp. 179–205.
24. Widmer, G. (1998). Applications of Machine Learning to Music Research: Empirical Investigations into the Phenomenon of Musical Expression. In R.S. Michalski, I. Bratko and M. Kubat (eds.), *Machine Learning, Data Mining and Knowledge Discovery: Methods and Applications*. Chichester, UK: Wiley & Sons.
25. Widmer, G. (2000). Large-scale Induction of Expressive Performance Rules: First Quantitative Results. In *Proceedings of the International Computer Music Conference (ICMC'2000)*. San Francisco, CA: International Computer Music Association.
26. Widmer, G. (2001). Discovering Strong Principles of Expressive Music Performance with the PLCG Rule Learning Strategy. In *Proceedings of the 11th European Conference on Machine Learning (ECML'01)*, Freiburg. Berlin: Springer Verlag.
27. Widmer, G. (2001). Inductive Learning of General and Robust Local Expression Principles. In *Proceedings of the International Computer Music Conference (ICMC'2001)*. San Francisco, CA: International Computer Music Association.
28. Widmer, G. (2001). *Machine Discoveries: Some Simple, Robust Local Expression Principles*. Submitted.