A System for Aiding Discovery in Musical Analysis

EDILSON FERNEDA *, CARLOS ALAN PERES DA SILVA **,
LUCIÊNIO DE MACÊDO TEIXEIRA **, HÉLIO DE MENEZES SILVA *

Universidade Federal da Paraíba
** Departamento de Sistemas e Computação / ** Departamento de Artes*
Av. Aprígio Veloso, s/n — 58.109-970 Campina Grande - PB - Brazil
e_mail: edilson@dsc.ufpb.br / peres@brufpb2.bitnet

Abstract

We start by proposing a computer aided scientific discovery system. This system may be seen as a knowledge acquisition environment. We present a knowledge representation (the Semi-Empirical Theories) usable to formulate, to experiment, and to divulge that knowledge, and a protocol (MOSCA) for cooperation between *rational* agents. The protocol is geared to the acquisition and evolution of knowledge. The objective of this system, rather than producing an exact knowledge, is yielding a knowledge which may present a high level argumentation on its validity and may also be improved via a dialog protocol. As an application, we aim at making the machine to behave rationally when performing Musical Analysis, which involves the four technical fields of Music: *composition, execution, theory, and sound digital processing*.

1. Introduction

In this paper we propose the specification of a Rational-Agent machine and its application to the field of Musical Analysis. A Rational Agent is an autonomous intelligent system that appears to the user as having reasoning abilities, because it is capable of common sense reasoning (such as those that we exert in our daily lives) and of handling intentions, beliefs, and knowledge that is tolerated to be, to some extent, evolutionary, incomplete, imprecise, and erroneous. This project is characterized mainly by making use of expertise in varied domains of the Cognitive Sciences: Artificial Intelligence, Informatics, Music, Psychology, Didactics, etc.

From an Artificial Intelligence vantage point, this work may be seen as lying in the confluence of the streams of Knowledge Acquisition and Machine Learning. According to the definition proposed in (Aussenac-Gilles, Krivine & Sallantin, 1992):

> *The domain of Knowledge Acquisition for Knowledge Based Systems (KBS's) is characterized by the identification and management of the processes necessary to the elaboration (conception, evaluation, and evolution) of a KBS from heterogeneous sources of knowledge (documented, human, and experimental). The result expected from our approach is to furnish the future system with the knowledge that will be the foundation of its expertise. The conductor of the process of knowledge acquisition is the knowledge engineer: he orchestrates the intervention of different processes, actors, and agents.*

While Knowledge Acquisition uses the machine as a mere tool for helping the knowledge engineer to elicit the expert's knowledge, Machine Learning studies the set of mechanisms that gives the machine the faculty of building the knowledge base by analyzing data, explanations, criticisms to solutions, etc. Several works (e.g. Barboux, 1990) have shown the necessity of making Machine Learning and Knowledge Acquisition to synergetically work together for modeling the control component of the learning process.

Learning proficiently is not enough if the expert is left without proper ammunition to efficiently check and validate the information acquired by the machine. This information has a too large volume and the expert has no tool capable of helping him to efficiently criticize the choices done, particularly the choice of the description language (which implies the choice of the learning tool) and of the selected sample of examples.

A System for Aided Discovery (SAID) is a synergetic combination of Machine Learning and Knowledge Acquisition. It follows principles (summarized in Wielinga, Boose, Gaines Screiber & van Someren, 1990), such as those for data acquisition, for abstraction based on information about a conceptual model, for particularization of this model, etc. The study of expert systems has shown a pervasive dichotomy between deep knowledge (characterized by having theoretical justification and by being found in scientific books, articles, etc.) and shallow knowledge (which is characterized by being situational, empirical, and not found in conventional scientific writings, though massively used by true experts, being the very cause of their expertise). As it was well put in (Sallantin & Haiech, 1993), a SAID discovers this shallow knowledge by taking advantage of both some deep (theoretical) knowledge made available and a set of incomplete, partially erroneous data. The knowledge base is assumed to be revisable (by error correction) and evolutionary (by making the knowledge more precise, more broad, more deep, more structured, more understandable,... or, in short, by improving the knowledge, in any sense).

In this paper we see scientific discovery as being the result of examining and revising a modeling process over which both theoretic models and experimental data intervene. During the modeling process, discovery is seen as that which was not yet learned by the on going modeling.

A first effort for conceptualizing an artificial apprentice generated a conceptual framework: the Semi-Empirical Theories (SET), introduced in (Sallantin, Szczeciniarz, Barboux, Lagrange & Renaud, 1991; Sallantin, Quinqueton, Barboux & Aubert, 1991). That effort established the elementary concepts which allow building (modeling) an apprentice's knowledge and studying its evolution. SET, however, being focused on the structures and mechanisms of an apprentice, neglected a fundamental learning aspect: the environment for the interaction between the apprentice and the external world. Therefore, a learning environment was proposed, as well as a description of a learning protocol. (Ferneda, 1992) shows how this protocol can be merged with the SET framework.

Since concepts formulated by an apprentice can be erroneous, an intervening agent should be able to determine counter-examples (and also examples likely to be in the frontier or beyond the frontier of the apprentice's current knowledge), testing and exercising him to the limits of his capabilities, hopefully embarrassing him by exposing his deficiencies, therefore stimulating him to revise and improve his knowledge. The goal is not to have an apprentice capable of acquiring a perfect (exact and complete) knowledge, but rather to have an apprentice capable of acquiring a knowledge which will be considered as quite plausible (because the apprentice can yield a high-level argumentation of its plausibility) and may also be corrected/improved via a dialog protocol.

The objective of a SAID applied to Musical Analysis is the analysis of Music's horizontality (melody, theme, scale, ...) and verticality (harmonic structure, instrumental coloring). An immediate application would be a study comparing the works of a same composer, or the works of a set of composers.

2. Rational agents

The scientific community, in spite of having tried really hard, has not yet come to a consensual definition of *intelligence*. There are, however, active entities which display behavior conventionally considered as being intelligent. These entities will hereafter be named *agents*.

Researches for the conceptualization and conception of artificial systems (or agents) capable of exhibiting behavior accepted as intelligent must, therefore, take into consideration the several characteristics presumed as necessary for a conduct to be classified as being intelligent. Among these attributes, we are here particularly

interested in the one of rationality (Newell,1982). The notion of rationality, more specific than the one of intelligence, is related to the treatment of a well delineated class of problems.

A *Rational Agent* is defined as being any (human or artificial) system capable of producing and controlling its own knowledge in a certain domain, in such way that the system will be able to proficiently perform some classes of complexes tasks (such as deciding, classifying, diagnosing, predicting, simulating, restricting, conceiving, and planning) conventionally considered as requiring intelligence for being well accomplished.

J. P. Müller (Müller, 1987) showed the possibility of constructing systems that: *(i)* are able to interpret symbolic structures; *(ii)* are conscious of their limitations; *(iii)* act in logical accordance with their beliefs; *(iv)* are able to adapt their actions to the changes in their knowledge. These systems, therefore, are capable of improving their representation of the external world and of better interacting with this world. This capability of constructing and evolving their representation of the world may be added to the learning aptitude of an intelligent agent.

Next, we will describe the behavior of an apprentice agent (*apprentice*, for short) whose knowledge is the result of communicating with other agents. This agent builds and controls the evolution of its knowledge. It has reasoning mechanisms such as those of the Semi-Empirical Theories (section 3), and its learning environment is based on the MOSCA protocol (section 4).

3. Semi-Empirical Theories

Semi-Empirical Theories (SET) are a language-independent knowledge conceptualization introduced by J. Sallantin. SET defines how the knowledge is formulated, experimented, and divulged.

Figure 1 depicts a taxonomy of the terms used for expressing knowledge in SET. This taxonomy is based on the work of T. Addis (Addis, 1988), which revised C. S. Pierce's work (Pierce, 1934) on modeling knowledge. The taxonomy includes: *(i) data* representing the knowledge; *(ii) mechanisms* for creating the data (by *abduction*), organizing them (by *induction*) and propagating restrictions on them (by *deductions*); *(iii) methods* related to the interactions with an external agent that plays the role of criticizing or the role of proposing a statement to be proved. The methods examine the adequacy of information such as *being a lemma*, *being an objection*, *being a proof*, *being a conjecture*, etc.

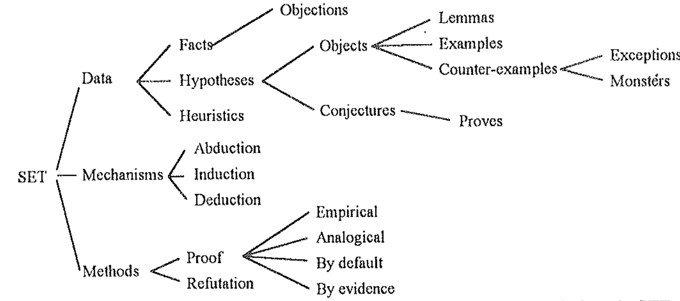

Figure 1: ***Terms intervening in the knowledge formalization and evolution via SET.***

4. A Protocol for learning

Formal theories for learning (Boucheron, 1992) define a minimal learning environment which is made of an *apprentice* communicating with an *oracle*. From the point of view of problem solving, the protocol controlling the dialog between these two actors may be summarized as follows: the oracle sends pairs <problem, solution> to the apprentice, each pair being named a *sample*, the problems having been already solved and their solutions known by the oracle; upon receiving each pair, the apprentice stores the information received and, if it does not perfectly match the *currently learned hypothesis* (the knowledge base), the apprentice searches the *hypothesis space* (the set of all hypothesis which may be formed in the light of all the information, old and new) looking for a hypothesis that, when measured by a *learning criterion*, will be considered better than the previous hypothesis and all the other candidate ones. Therefore, the presented model sees a learning problem as composed of *(i)* a hypothesis space, *(ii)* a learning criterion measuring how a hypothesis fits the set of samples, and *(iii)* a strategy for traversing the hypothesis space.

Two kinds of noise add to the inherent complexity of searching for a hypothesis: *(i)* the pair <problem, solution> may have been erroneously described, and *(ii)* the language adopted for describing those pairs may be too coarse, leading to the consequence that it, not perceiving-and-representing the difference between two problem specifications, may present the apprentice with a unique problem having two distinct solutions. Real world applications can not completely escape the existence and negative effects of noise.

This minimal learning environment is implemented as follows: While the machine plays just the role of the apprentice, the expert plays the role of the oracle and may also play some other roles, as we will see. The expert, by choosing the way of structuring and representing knowledge, determines the type of apprentice deemed more adequate to the problem at hand. Well, determining the type of the apprentice is determining the hypothesis space on which the apprentice can operate. This way, the expert is the one who takes all the crucial decisions: *(i)* He, by choosing the type of the apprentice (and, therefore, the format of the hypothesis, of the samples, and of the problems to be solved), decides the underlying theoretical framework for learning in the application domain; and *(ii)* The expert selects the examples to be offered to the apprentice. When playing his role as an *oracle*, the expert has available a first manner of pressing the apprentice, imposing a knowledge on him.

After having received a first set of examples and counter-examples and built a learned hypothesis, the apprentice leaves his learning mode and enters his probation mode, ready for solving problems proposed by the expert. Three difficulties appear: *(i)* examining merely the solutions produced by the apprentice is not enough for evaluating the learned hypothesis; *(ii)* we are trying to achieve scientific discovery, therefore the expert does not exactly know how to characterize whether or not a hypothesis is a good one, deserving to be maintained; and *(iii)* hypotheses is usually too large, too unstructured and too complex to be directly utilized by the expert. For these reasons, we provided the apprentice with a high-level argumentation mechanism whose importance has acknowledged by some researchers (Fisher, Lemke, Mastaglio & Morch, 1991). The expert will judge the goodness of a hypothesis by judging the argumentations presented by the apprentice as an justification of the solutions he found for the posed problems. When playing his role as an examiner, the expert has available a second manner of pressing the apprentice, making him to revise the learned hypothesis.

This informal presentation of the MOSCA[1] (Reitz, 1992) depicts five distinct roles: *(i)* the *apprentice*, yielding a learned hypothesis which fits well the sample of examples and counter-examples previously made available to him; *(ii)* the *oracle*, yielding unrefutable *<problem,solution>* pairs; *(iii)* the *Client*, which submits problems to the apprentice and expects to receive solutions from him; *(iv)* the *probe*, yielding refutable *<problem,solution>* pairs, making him to present the due argumentations; and *(v)* the *master*, who analyses the apprentice's argumentations and then offers useful criticisms to him. The learning environment is summarized in Figure 2. Additional explanations follow:

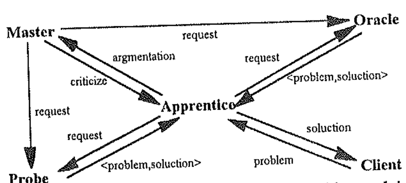

Figure 2: ***The MOSCA learning protocol, seen from a problem solving vantage point.***

[1] **MOSCA** = Master + Oracle + Sonda (Probe, in Portuguese) + Client+ Apprentice.

The *apprentice* asks a set of *<problem,solution>* pairs from the oracle. The pairs are then stored (it should be defined whether or not data may be eventually erased/changed) by the apprentice and compose the *sample* to be used for learning.. Every change in the sample makes the apprentice to revise the hypothesis so far learned. This hypothesis is extracted from a hypothesis space and satisfies a learning criterion. Of course, any learned hypothesis finds the right solution for all the problems in the sample, and finds none of the wrong solutions for the counter-examples in the sample.

The apprentice may request a *<problem,solution>* pair, or even a set of those pairs, from the oracle. There two ways for making this request: *(i)* When the choice of the problem is left to the oracle (who may do it following or not nay plan, such a previously defined teaching plan), the request is made just by sending a signal; *(ii)* when the apprentice desires to learn to solve a certain problem, he makes his request by precisely specifying this problem (of course this must be allowed only in a controlled way, or else the apprentice would take the time just unproductively interrogating the oracle).

Similarly, the apprentice receives *<problem,solution>* pairs from the probe. These pairs may be intentionally erroneous. The apprentice compares his solution against the one proposed and displays to the master an argumentation justifying the solution. There are two forms of argumentations: explanations (when the solutions agree one with the other) and objection (otherwise).

For each argumentation presented, the apprentice receives a criticism from the master. Whenever possible, every negative criticism makes the apprentice to present an alternative argumentation. When no alternative argumentation is possible, then the apprentice either weakens the learned hypothesis, in such way that it no more produces the pair <problem, solution>, or he consults the oracle, aiming at revising the learned hypothesis and, consequently, revising the argumentation.

The master, by either sending a signal or a problem specification, controls the probe's production of pairs *<problem, solution>*.

Whenever the apprentice weakens his learned hypothesis, thus becoming silent about certain problems, the master finds a way of forcing the apprentice to re-strength his learned hypothesis: the master sends (again either via a signal or via a problem specification) to the oracle a request for generating an adequate set of pairs.

A *Client* submits a problem to the apprentice and expects a solution from him. Whether or not an argumentation on the solution is sent to the master, depends on how the apprentice was defined.

5. Discovery in musical analysis

The musical thinking, when generating a certain work, spans aspects pertaining to a knowledge branch difficult to approach. This has lead us to conceptualizations emphasizing terms, such as "gift" and "inspiration", which were quite used, in the past, to impose an end to any discussions on music creation. Our days, however, such extremist position fully satisfies neither the artist, nor the scientist, as both of them, knowing how sounds effect the men, have been studying the organization of sounds, encompassing aspects spanning from its atomic form, pitch up to its final manifestation, the composition.

Perception of this organization in the composition, and perception of the observance of clearly defined forms and rules regulating the creative process, indicate the existence of a systematization in the musical thinking. This systematization is common to all composers in a given context (style-and-school, local-and-time, etc.) and may be represented by means of a knowledge base constructed by a musician and which can constantly have its rules corrected, added, modified, substituted, or eliminated (in short, enriched in any form) at the discretion of the musician.

As it was shown by R. B. Dannenberg (1993), the musical thinking does not follow a linear pattern as it does in exact sciences. The reason for this is that many complex factors interact in the musical thinking: creativity, emotions, intuition, and the proper vibratory nature of sound. However, during the creative process the musical thinking maintains its relationship with form, structure, and harmony, determining a musical logic which allows classifying a work as belonging to a style. Those characteristics strongly associate themselves with those peculiar to given composers.

According to H. J. Koellreutter (Zagonel & la Chiamulera, 1985), the composition process follows four steps: consciousness of the idea, formal conception, choice of the musical signs, and structuring. From the above, (Kugel, 1990; Roads, 1985; Widmer, 1992) perceived that the reasoning of a composer may be simulated by a knowledge based system zeroing not in the musical work proper, but in the process which he used for generating his work. This system, besides the computational aspects involved in analyzing a musical work in terms of information processing, should handle knowledge on domains influencing musical conception: *(i) Aesthetics*, involving concepts from Physics (study of the vibratory nature of the sound), Psychology (study of sound as a psychic phenomenon), and Sociology (study of the ideological aspects of the creator), and Statistics; *(ii) Music*, with its rules for harmony, melody, etc.; and *(iii) History of Arts*, approaching the peculiarities of each style in a given period of History. All these knowledge areas interact for analyzing a composer's work and also for further understanding men's reasoning process.

The study of musical analysis has several facets and is intimately bound to Aesthetics (the study of the conditions and the effects of the artistic activity), which can not be dissociated from History of Arts. We can say that the analysis of a musical work is really much more that a study for deciphering a language, or a simple reading of the formal aspects of the work. We adopted the view that knowledge on Aesthetics and on History of Arts should come to the aid of knowledge on Musical Theory. This way, we define a musical work as being the result of knowledge on the domains of Aesthetics, of Music, and of History of Arts.

We aim at making the machine to behave rationally when performing Musical Analysis, which involves the four technical fields of Music: *composition, execution, theory, and sound digital processing*. We investigate, therefore, the conception and development of a system capable of aiding musical analysis.

6. Conclusions

We presented a learning environment whose protocol, identifying the set of communications needed for controlling an agent, is an extension of the classic protocol and permits analyzing the process of revising the knowledge acquired by the apprentice. This environment was studied within an conceptual framework, the Semi-Empirical Theories, supporting the expression of both the reasoning and the structure of the apprentice. Under the light of both theoretical and practical Machine Learning current results, integrating high level argumentation techniques into the learning system was deemed necessary to get the expert's validation, approval, and confidence in the acquired knowledge.

Our experience has shown that most of the currently available learning tools do not fully satisfy the expert's expectations and needs. More than interested in a system that merely has the capability of learning correctly, experts are looking for systems able to go beyond that by understandably and convincingly explaining what and how they have learned. One of the major reasons for this is that the explanations may hint what should be changed in order to improve the knowledge base. All of this grows in importance when the expert uses the machine as an aid for scientific discovery (of course this involves modeling a phenomenon).

If the user wants to teach the system, than he, while playing the role of the oracle, he should select the relevant problems. While playing the role of the master, he should use his deep knowledge to criticize the currently learned knowledge in order to identify lemmas already validated the proofs. While playing the role of the probe, he should produce relevant examples and counter-examples. Finally, while playing the role of the apprentice, he should choose methods and heuristics to be used for advancing a new learning step in the light of the last information stimulus received from any of the other agents.

Our approach was corroborated just by few and microscopic experiments (Ferneda, 1992), thus it still calls for more numerous and larger scale experiments in order to fully establish itself as a really useful learning framework. In spite of this, the way our proposed system solves learning problems may be seen as an advancement at least as a methodologically and didactically relevant concept, in as much as it sees the learning problem in a way understandable and profitable to both the application domain expert and the Artificial Intelligence researcher. It should be noticed that we assured the possibility of refuting the learned knowledge. For accomplishing that, it is necessary that all heuristics involved in the knowledge acquisition may be reevaluated. This justifies our approach in adopting the Semi-Empirical Theories and the MOSCA protocol for representing and evolving the knowledge, respectively.

Musical Analysis seeks to explain a musical work by making use of bodies of knowledge such as Aesthetics, Music, and History of Arts. It is our firm and well-founded belief that the environment here presented will show itself to be a satisfactory aid for the task of performing Musical Analysis, since the environment makes room for those bodies of knowledge (what is fundamental for the conception of a musical work) and allows the construction of a theory that permits the characterization of the work through dialogues with a human, application-domain expert agent.

References

T. R. Addis (1988). Knowledge organization for abduction, *Interdisciplinary Information Technology Conference*, Bradford University (England).

N. Aussenac-Gilles, J.-P. Krivine, J. Sallantin (1992). L'acquisition des connaissances pour les systèmes à base de connaissance, *Revue d'Intelligence Artificielle*, Vol. 6, nº 1-2, Hermès, Paris.

C. Barboux (1990). "Contrôle par objection d'une théorie incomplète", *Doctorate Thesis*, Université de Montpellier (France).

S. Boucheron (1992). *Théorie de l'apprentissage: de l'approche formelle aux enjeux cognitifs*, Hermès, Paris.

R. B. Dannenberg (1993). Music representation issues, techniques, and systems, *Computer Music Journal*, Vol. 17, nº 3, pp. 20-30, MIT Press.

E. Ferneda (1992). *Conception d'un Agent Rationnel et examen de son raisonnement en géométrie*, Doctorate Thesis, Université de Montpellier (France).

G. Fisher, A. C. Lemke, T. Mastaglio, A. I. Morch (1991) The role of critiquing in cooperative problem solving, *ACM Transaction on Information System*, Vol. 9, nº 3, pp. 123-151.

P. Kugel (1990). Myhill's Thesis: There's more than computing in musical thinking, *Computer Music Journal*, Vol. 14, nº 3, pp. 12-25, MIT Press.

J.-P. Müller (1987). *Contribution à l'étude d'un agent rationnel : spécification en logique intensionnelle et implantation*, Doctorate Thesis, Institut National Polytechnique de Grenoble (France).

A. Newell (1982). The knowledge level, *Artificial Intelligence*, nº 18, pp. 87-127.

C. S. Pierce (1934) Scientific method, in *Collected Papers of Charles Saunders Pierce*, P. Weiss (Ed), Harvard University Press.

Ph. Reitz (1992). *Contribution à l'étude des environnements d'apprentissage. Conceptualisation, Spécifications et Prototypage*, Doctorate Thesis, Université de Montpellier (France).

C. Roads (1985). "Research in Music and Artificial Intelligence", *ACM Computing Surveys*, Vol. 17, n. 2, pp. 163-190.

J. Sallantin, J.-J. Szczeciniarz, C. Barboux, M.-S. Lagrange, M. Renaud (1991). Théories semi-empiriques: conceptualisation et illustrations, *Revue d'Intelligence Artificielle*, Vol. 5, nº 1, pp. 9-67, Hermès, Paris.

J. Sallantin, J. Quinqueton, C. Barboux, J.-P. Aubert (1991). Théories semiempiriques: éléments de formalisation, *Revue d'Intelligence Artificielle*, Vol. 5, nº 1, pp. 69-92, Hermès, Paris.

J. Sallantin, J. Haiech (1993). L'aide à la découverte scientifique: évaluation sur l'investigation des séquences génériques, *Revue d'Intelligence Artificielle*, Hermès, Paris.

G. Widmer (1992). Qualitative perception modeling and intelligent musical learning, *Computer Music Journal*, Vol. 16, nº 2, pp. 51-68, MIT Press.

B. Wielinga, J. Boose, B. Gaines, G. Screiber, M. van Someren (Eds) (1990). Current trends on knowledge acquisition, *Proceedings of the European Knowledge Acquisition Workshop*.

B. Zagonel, S. M. la Chiamulera (orgs) (1985). *Introdução à Estética e à Composição Musical Contemporânea - H. J. Koellreutter*, Editora Movimento, Porto Alegre (Brazil).