

Similarity Knowledge Formalisation for Audio Engineering

Christian Severin Sauer, Thomas Roth-Berghofer,
Nino Auricchio, and Sam Proctor

School of Computing and Technology, University of West London,
St Mary's Road, London W5 5RF, United Kingdom,
{christian.sauer|thomas.roth-berghofer|
nino.auricchio|sam.proctor}@uwl.ac.uk

Abstract. This paper describes an approach to externalisation of the tacit knowledge used by experienced audio engineers to affectively describe emotions evoked by a sound or piece of music. We formalised the adjectives describing the timbre of a sound as well as their relationships. The main problems are the vagueness of emotions and the variation in the emotions the same single percept can trigger in different people. We demonstrate how similarity knowledge can be used to process fuzzy and incomplete queries to emulate the vagueness and differentiation associated with the emotions triggered by a sound percept. We capture the experience of audio engineers by mapping the formalised vocabulary of timbre-describing adjectives to their workflows, which describe the actions to change the spectral shaping of a sound and its emotional effect.

Keywords: Case-based reasoning, audio engineering, similarity measures, knowledge formalisation

1 Introduction

The formalisation of affective, emotional statements or descriptive adjectives of an emotion is still a problem [13, 8]. This problem is often encountered by applications dealing with art, as art is deeply linked to emotions and perception of such. In the case of music, a variety of approaches already exist to formalise emotional annotations of music, see for example [23].

There are already a variety of approaches to automated composition of expressive music [19, 27] and the expressive performance of music [4]. Both tasks include the problem of formalising emotions. This formalisation is needed as information relating to the intended emotional effect of a composition and/or performance must be encoded and then integrated in the systems knowledge. However, formalising emotions is not easy due to the different perception and emotional links individual humans might have to a percept. Next to composition and performance, a third very important task in professional music production is the mastering of a recording of a sound or song. *“Mastering is the last creative*

step in the audio production process, the bridge between mixing and replication-your last chance to enhance sound or repair problems in an acoustically-designed room-an audio microscope. Mastering Engineers lend an objective experienced ear to your work; we are familiar with what can go wrong technically and aesthetically. Sometimes all we do is-nothing! The simple act of approval means the mix is ready for pressing. Other times we may help you work on that problem song you just couldn't get right in the mix, or add the final touch that makes a record finished and playable on a wide variety of systems.” [16].

Mastering is the process of applying a set of spectral modifications to sounds in order to achieve a change in timbre or more specifically the emotional effect of the perception of the sound on a listener. This process is goal oriented, with the goal being a desired change in the emotional effect of a sound. The vocabulary describing this effect-change consists of terms that describe the emotion desired to be triggered or altered, i.e., increasing or decreasing an emotional effect. We find terms like ‘make it sound more warm’ or ‘make it sound less harsh’ and onomatopoeia in the language of audio engineers. The experience of audio engineers is in the linkage between these emotional descriptors and in the choice and application of spectral modifications used to achieve the desired change of the sound. The emotional effect of a timbre is also linked to the context in which it occurs. The modelling and (re-)use of such context embedded timbres is part of our future work.

This paper introduces our work on a systematic approach that allows an audio engineer during the mastering stage of a music production to apply descriptive adjectives for the automatic selection of workflows using presets of spectral modifications that deliver the intended alteration of the sounds emotional effect. A preset can be described as a selection of frequency descriptors with definite values for said frequencies. A preset can further contain information on defined effects such as reverb or delay and the values to be applied to these effects.

We provide a dynamic recommendation of workflow steps, consisting of the application of presets to the sound being mastered. To achieve this we map workflow descriptions such as the temporal ordering of preset use, next to the selection of presets, to the descriptive adjectives and descriptors of amounts of an effect contained in the vocabulary used by audio engineers. This approach seems worthwhile as audio engineers label almost any changes applied to an audio product during the mastering stage of its production with descriptive adjectives. The adjectives are almost always based on the change of a sound’s timbre that is to be achieved by the application of said changes. It is not uncommon to encounter a description such as “*make it sound a bit more warm and punchy*” [5]. (See also “*Gentle - Opposite of edgy. The harmonics-high and upper midst-are not exaggerated, or may be even weak. Grungy - Lots of harmonic I.M. distortion. Harsh - To much upper midrange, usually around 3 kHz. Or, good transient response, as if the sound is hitting you hard*” [5].) The interlinking of such descriptive terms, which we call *timbre descriptors*, to defined settings of frequency reshaping is mainly based upon years of experience of a sound engineer. We try to make this experience available in music production software.

The paper is structured as follows: We interlink our approach with the current state-of-the-art in the field of artificial music composition and performance in the following section. Based upon the identified problems we introduce the challenges we expect for our approach in Section 3. We then introduce our approach to externalising and formalising the tacit knowledge of experienced audio engineers. A summary and outlook on future work then concludes the paper.

2 Related Work

A variety of approaches to formalise emotional annotations and/or descriptive terms that either describe the mood of the music or the way it is to be played [20, 11] already exists. Such approaches for now deal with either playing music in a certain defined way to convey an emotion [9] or to select songs or sounds that are associated with a mood or emotional state [29]. For automated composing, the question of integrating a formal description of the mood the composed music should match is already well researched [20, 4].

Emotions or emotional perceptions are not easy to be a) defined and b) quantised/formalised [15, 10, 13]. Another problem we were facing was that we tried to quantify and cluster descriptive adjectives based on very vague data given by the individual descriptions of the emotional effect a sound has on a person describing this effect. The difficulties of capturing a sound's timbre are [12]: *“It is timbre’s ‘strangeness’ and, even more, its ‘multiplicity’ that make it impossible to measure timbre along a single continuum, in contrast to pitch (low to high), duration (short to long), or loudness (soft to loud). The vocabulary used to describe the timbres of musical instrument sounds indicates the multidimensional aspect of timbre. For example, ‘attack quality,’ ‘brightness,’ and ‘clarity’ are terms frequently used to describe musical sounds.”* The vagueness of the data is based on said variation in the individuals perceptions when they either should describe an emotional effect or perceive something that is annotated with a particular emotion but have a complete different idea of the actual emotion this percept triggers [24, 12, 14]. An experiment which tried to establish if there is a common understanding of how humans perceive and describe timbres showed that although some agreement among musicians upon basic descriptive adjectives was reached, the more complicated a sound or impression got, the more variance occurred in the assessment of the sound’s perception [10].

A way to circumvent the lack of quantifiable measures and vagueness is to allow for vagueness and a certain amount of ambiguity within the techniques used for formalising and retrieving presets based on descriptive adjectives. The vagueness accompanying the formalisation of emotional descriptors can be described by similarity knowledge in Case-Based Reasoning (CBR) systems [1, 17] for instance CBR can handle vague queries and returns similar results if no exact match can be made [30]. The ability of CBR to handle said vagueness has already been used to guide the emotional component of automatic composition as well as performance of music, see e.g., [9, 3, 25, 20].

3 Challenges of Formalising Descriptive Adjectives and Preset Mapping

When describing the timbre of a sound, tacit knowledge is present in the implicit emotional descriptions using descriptive adjectives on the timbre. Externalising this tacit knowledge about a timbre or emotional effect of a sound was the main challenge of the knowledge formalisation task at hand. The first challenge was to find common ground, a basic selection of timbres and their emotional descriptors. The questions raised by this challenge were:

1. Do annotations of timbre with descriptive adjectives/terms vary between different people?
2. Are there significant clusters, distances, patterns in the classification of the adjectives/terms used to describe timbre?

Related work showed that the first question can be answered with a solid *yes*. This is especially true with respect to the vocabulary we try to establish as it is not free of redundancy and ambiguities yet. Also true the terms often overlap and have different meanings in different contexts, which is common for emotions and their formalisation (see, e.g., [2]. As Donnadiu, Porcello, Darke [10, 12, 26]) In this paper, we rely on interviews with experienced audio engineers for knowledge gathering. We used their expertise to establish a first set of timbre descriptors most frequently encountered in the domain of audio mastering. Based on this we mapped the changes of timbre to the application of workflows consisting of a sequence of presets applied to the sound. Therefore, we had to measure the effect/change of timbre the application of a preset has on a sound.

4 Similarity Knowledge Formalisation

We chose CBR as a suitable methodology for our task as it is able to handle vagueness. CBR already has been used in music composition and expressive performance to date as well as for handling otherwise difficult to formalise knowledge [3, 25]. For the purpose of modelling and testing the similarity knowledge in a CBR system we employed the myCBR Workbench and SDK¹.

4.1 Knowledge and Data present in the Audio Mastering Domain

In our domain we face three sets of artefacts, presets, descriptive adjectives, and workflows. The first set has been already described. The descriptive adjectives are present in audio engineering literature and day to day practise of audio engineers. Workflows describe step-by-step best practises of the application of one or more presets. Application of these workflows aims at reshaping an audio product, resulting in a shift of the specific emotion being evoked by the audio product. Again this knowledge is partly available from literature but mainly only present as tacit knowledge of experienced audio engineers.

¹ <http://mycbr-project.net>, Version 3.0

4.2 Advantages of CBR in our Domain of Interest

CBR is able to make use of the customer’s language, in our case descriptive adjectives and likely vague terms describing the amount of an effect desired. This means that we can use fuzzy descriptive adjectives like ‘muddy’ or ‘bright’ to define queries or problem descriptions, to retrieve cases holding the workflow descriptions to achieve this effect as their solution part. CBR is able to retrieve cases based on only sparse problem descriptions. CBR heavily relies on similarities which are, as introduced in section 3, comparatively easy to elicit within our domain of interest. Additionally CBR allows for queries that combine retrieval and filtering in the way of queries like: ‘This should sound really airy, but not harsh’. Furthermore by being able to retrieve a number of cases linear to their similarity CBR can offer a choice to the audio engineer with respect to the possible workflows to apply to achieve the desired change of the timbre. This is a particular advantage again with reference to the overall vagueness of the audio mastering process. In the following we detail our formalisation of audio mastering knowledge into the four knowledge containers of CBR [28].

4.3 Case Structure and Attributes

Our cases consist of a problem description part, specifying the present timbre of a sound and the desired change in the timbre and an indicator for the amount of this change. The solution part of our cases is, for now, given by a workflow description of one or more presets to be applied to the sound and their order of application. The basic mapping of workflows to the three descriptive adjectives, given by the timbre descriptor describing the present sound, the timbre descriptor describing the timbre the sound should change to and the amount descriptor specifying the amount of change that should effect the timbre, suits the approach of structural CBR [6] very well. Table 1 shows an example case.

Table 1. Example case with problem descriptors (left) and solution (right)

Problem attribute	Value	Workflow (solution)
Input Timbre:	Bassy	Apply preset 4
Target Timbre:	Hollow	then use filter 7 with
Amount of change:	+50	with 80 percent treble.

For now we rely on a percentage for the amount descriptor, ranging from 0 to 100, whilst 0 means no change to the timbre of the input sound at all and 100 translates to the total conversion of the timbre of the input sound to the timbre specified by the timbre descriptor of the cases problem description. We have not mapped these percentages of effect to verbal amount descriptors such as ‘a bit’, ‘a lot less’, ‘much more’ here.

Our case structure reflects the problem description, consisting of an Input timbre as the timbre descriptor characterising the input sound on which the

modulating workflow is to be used. The problem is further specified by a target timbre which is a timbre descriptor characterising the way the sounds timbre should change (make it sound more/less ‘timbre descriptor’). Thus our approach is goal driven by the way that an initial timbre given is to be changed to a desired target timbre. Nevertheless this change of timbre can also be achieved without an initially given timbre as in audioengineering timbre changing workflows applied to a sound always have the same effect on a timbre regardless of its original state. The third part of our problem description is the Amount of change desired to take effect, ranging from 0 (no effect) to 100 per cent (total conversion of the timbre), that can hold negative’. The second part of our case structure is the solution description, i.e., a workflow description of how to obtain the timbre change to the extent desired. This workflow description can be stored as an URL, a text or a sequence of presets selected from a database.

In addition to the timbre descriptor and amount descriptor as problem description and a workflow for the desired timbre change, possible further attributes for a case can be seen in onomatopoeia describing the sound itself and a description of the sounding situation, e.g. ‘opera house’, ‘marching’ or ‘club’ [10].

4.4 Vocabulary

The vocabulary we use was identified to consist of timbre descriptors and the names or id’s of pre-sets that provide this timbre after being applied to an audio product. The vocabulary further consists of a set of amount descriptors as for example ‘a bit’, ‘much’, ‘a touch’. Additionally we also incorporate workflow descriptions into the vocabulary to provide the workflow suggestions as the solution parts of our case structure.

Our initial approach of establishing a vocabulary was limited to general music settings, with regard to the domain complexity and the described problems known with the formalisation of timbre describing adjectives or even onomatopoeia. By omitting specific genres, such as rock or jazz, we aimed at keeping the vocabulary as ‘flat’ or simple as possible. We did so to prevent our effort from being too specific (to a genre) and to be reusable in a more general way for audio mastering. From our expert interviews we elicited 34 timbre descriptors. We have not yet established amount descriptors but aim to split the interval of 0 to 100 per cent of timbre change to, for example, 20 amount descriptors to provide a 5 per cent granularity for the desired impact of a timbre changes described in the problem description of a case. For practical reasons we aim to discretise the interval to a number of amount descriptors ranging from ‘Not’ ‘None’ or ‘Should not sound’ to ‘Totally’ ‘Convert to’ ‘Fully’ rather than use numerical values.

4.5 Similarity Measures

To establish the similarity of attributes and between cases in our system, we needed to establish local as well as global similarity measures. Local similarity functions describe the similarity between single attributes of query and case.

A global similarity function, e.g., a weighted sum, then aggregates the individual local similarity values into one overall similarity value for a case. For the formalisation of the similarity between timbre descriptors we considered two options. The first option was given by employing Multidimensional Scaling (MDS). MDS was intended to achieve a dissimilarity matrix describing the dissimilarity of descriptive adjectives. The computation of a dissimilarity matrix is identical to one of the main approaches used in CBR to formalise similarities and thus offers a way to capture the similarities of timbres, if there are any to discover as patterns. By using MDS we hoped to establish if there are patterns within the terms regarding their similarity of use when describing a timbre.

The second option was to establish a taxonomy of descriptive adjectives. Such a taxonomy can be seen as a comparable to a taxonomy of colours, classifying the emotion a sound triggers by the use of timbre descriptors. An example of a parent and two child nodes in such a taxonomy would look like the following: [Treble (parent) - [Topy (child) [Bright (child)]]]. The taxonomy would be used to establish the similarity of two emotional-descriptors by their position within the taxonomy and their distance and also store adaptation knowledge as we detail in the following subsection. Both options are very close to the common data structures used within CBR to formalise similarities and thus offer easy approaches to formalise the similarity of descriptive adjectives.

We decided to apply the second option of building a taxonomy of timbre descriptors based upon the elicitation of their similarity from the tacit knowledge of experienced audio engineers. The taxonomy consists now of 32 nodes beginning with the most abstract 'timbre descriptor' and expanding down to its leafs with the most concrete descriptors of timbres. See 1 for a part of the initial taxonomy describing timbres in the higher frequency ranges. The weights of the nodes reflect the similarity of the timbres to each other ranging from 1.0 total match or in other words having the same emotional effect to -1.0 totally dissimilar or in other words negating a timbre.

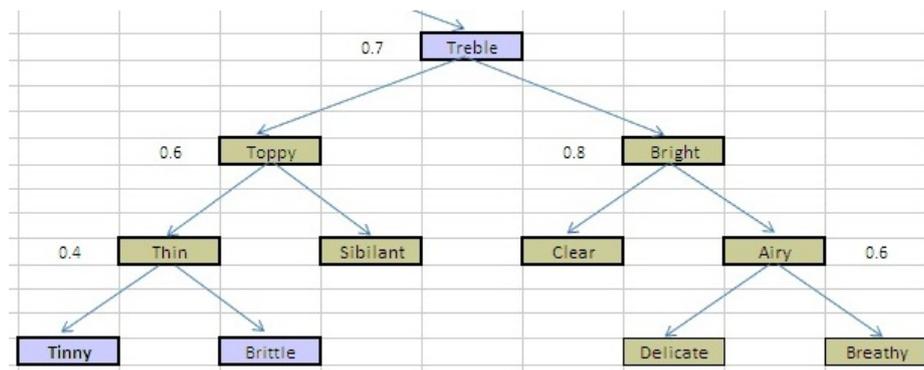


Fig. 1. Excerpt from the taxonomy of timbres in the high frequency range

The initial taxonomy was then refined and modelled using myCBR Workbench. We modelled the three local similarity measures for the present timbre, the (target) timbre a sound should change to, and the amount of this change (Figures 2 and 3).

	Muffled	High mid	Low	Sibilant	Timbre	Dark	Thin
Muffled	1.0	0.0	0.8	0.6	0.0	0.4	0.5
High mid	0.0	1.0	0.1	0.3	0.0	0.6	0.0
Low	0.8	0.1	1.0	0.0	0.0	0.1	0.2
Sibilant	0.6	0.3	0.0	1.0	0.0	0.8	0.6
Timbre	0.0	0.0	0.0	0.0	1.0	0.0	0.0
Dark	0.4	0.6	0.1	0.8	0.0	1.0	0.5
Thin	0.5	0.0	0.2	0.6	0.0	0.5	1.0
Clear	0.0	0.0	0.7	0.1	0.0	0.3	0.7
Treble	0.1	0.1	0.2	0.0	0.0	0.0	0.0
Bright	0.0	0.0	0.5	0.6	0.0	0.2	0.5
Harsh	0.2	0.2	0.0	0.7	0.0	0.0	0.2
Honky	0.4	0.3	0.1	0.0	0.0	0.7	0.3
Hollow	0.0	0.1	0.4	0.6	0.0	0.0	0.0
Airy	0.5	0.1	0.0	0.2	0.0	0.9	0.0

Fig. 2. Similarity measure for timbre in table view (*myCBR* Workbench)

Beyond the formalisation of the basic timbre descriptors we further grouped the timbre descriptors into more abstract groups describing families of timbres and even more abstract the frequency ranges were these families of timbre descriptors are most commonly used. This approach aims at being able to include additional information in our similarity measures. Next to the timbre descriptors we yet have to, as described, provide a similarity measure for the amount descriptors. A possible future addition could be seen in another taxonomy describing the context of an audio signal being manipulated. Such a context could be provided by the instrument that is used to generate the audio signal. So again we could establish a taxonomy of instruments which would begin with abstract families of instruments, like ‘brass’ or ‘strings’ and get more specific in the deeper levels of the taxonomy distinguishing individual instruments of a family, like for example: [Organ (parent) - [Hammond (child)] [Pipe (child)]]

For the global similarity measure we initially use a non-weighted sum of the local similarities. For more complex cases, that include a set of more than input and target timbre descriptors and their amount descriptors and may also be extended to include the context description of a sound, like the instrument, we plan to provide weighted sums of the respective local similarities.

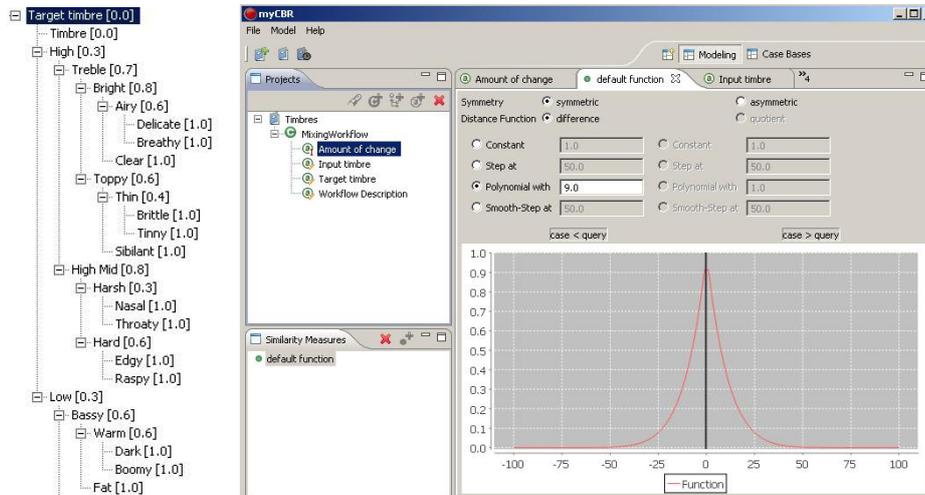


Fig. 3. On the left, similarity measure for the target timbre modelled as taxonomy. On the right, simple percentage-based similarity measure for the amount of effect (*myCBR* Workbench)

4.6 Adaptation Knowledge

The purpose of adaptation knowledge in CBR is to adapt the solution of the most similar case to the current problem. A basic example of this is adaptation by replacement. By storing adapted and tested (verified) cases the CBR system gains new knowledge. Such adaptations are also desirable for our system. We therefore plan to integrate adaptation knowledge in a number of ways. One way to obtain, formalise and use adaptation knowledge is to use of taxonomies in similarity measures [7]. Our system uses taxonomies for the descriptive adjectives and, in the future, also for the sound context, i.e., instrumental families. Adaptation knowledge is stored in the parent-child relations formalised in the taxonomies. For example the taxonomy of descriptive adjectives can be used to provide replacements for invalid or unwanted adjectives in the following way: Assuming that within the taxonomy there are nodes of the following kind: [Treble - [Toppy] [Hard]] (Figure 1). If ‘Toppy’ was defined within a query case’s problem description but is not available within any case from the case base, the taxonomy could be used to select the most similar adjective, ‘Hard’ instead and thus use either ‘Hard’ or fall back to the parent node ‘Treble’ as the next more abstract timbre descriptor to replace ‘Toppy’.

5 Summary and Outlook

In this paper we presented our approach to externalising tacit knowledge encoded in a special vocabulary used by experienced audio engineers to effectively

describe emotional effects and timbres of audio products. We have described our approach to formalise the timbre descriptors and to map them to frequency reshaping workflows of pre-sets application. We discussed the known problems associated with attempting to formalise and quantise emotions, in general, and adjectives describing timbre in music, in particular. Based on those findings we introduced CBR as an approach to amend the problems of vagueness of terms and the variance of emotions invoked by the same sound in different humans. We then introduced our approach to use CBR's ability to process fuzzy and incomplete queries and the ability to choose between grades of similarity of retrieved results to emulate the vagueness. We detailed especially on the approaches we used to formalise the knowledge into the four knowledge containers of CBR [28].

As the very next step of developing our approach further a more complex way to adapt a case in our domain is to be enabled by the fact that there do exist timbres that are 'opposing' each other. By 'opposing' we mean that there can be two timbres, like 'Airy' and 'Boxy' that cancel each other out if they are applied to the same sound. As we elicited the knowledge from the audio engineers they pointed out that applying such opposing timbres is a common practise while mastering an audio product. We thus asked them to provide us, next to the similarity of two timbres, also with the 'oppositeness' of them. We formalised this oppositeness in a 'negative similarity measure', ranging from 0, i.e., not opposite at all, to -1, i.e., total opposition, thus describing two timbres that cancel each other out. The values between 0 and -1 describe the ability of two timbres to soften the effect of the other. So, for example, if we look at the timbre 'Nasal' the timbre 'Dark' has an oppositeness of -0.2, so applying 'Dark' to a 'Nasal' sound reduces the sounds 'Nasal' timbre by roughly 20 per cent. The way we intend to use this oppositeness as adaptation knowledge is by providing rules off the following nature: Assume a query case asking for a shift from a Nasal timbre to a Harsh timbre with 40 per cent effect strength. The best case in the case base only provides the workflow for a change for a nasal timbre to a harsh timbre with 20 per cent effect strength. The remaining 20 per cent of shifting the nasal timbre to the harsh timbre could be accomplished by applying a 20 per cent opposite timbre, like the dark timbre, thereby reducing the nasal timbre by another 20 per cent. So we cancel out 20 per cent of the nasal timbre by applying another workflow to add the -0.2 opposite dark timbre. The resulting rules thus will be of the form general form: If effect strength not reached/exceeded: Find a case similar or opposite timbre to apply with regard to the missing amount of effect strength. A particular example there would be: If shift from nasal to harsh with x per cent not reached: Apply shift from nasal to dark with x-best case applied effect strength.

The current basic system we developed will be used by audi engineering students and audio engineers at the school of technology at the University of West London to gather feedback on the useability of the system. We further aim to get feedback from the studnets and engineers as well as artists on the verstaility and exactness of our knowledge formalisation approaches implemented in the basic system.

For the future of our approach we aim at further adding a more detailed way of user group modelling into our system. We do so, as we have established that it is a major difference if our system will interact with artists from different genres and/or users of different level of experiences, i.e., novice to expert sound engineers. It is also of importance to establish what a user might have as a goal overall, because mixing, composing and mastering are three different contexts in which the retrieval of presets would differ significantly in a later version of our system. We also aim to research about the importance of the dialogue between audio engineers among themselves and audio engineers and students (see, e.g., [26]). The tacit knowledge conveyed within these dialogues is also of concern to our approach, besides the basic approach of mapping timbre descriptors to workflow selections. This concern is introduced by the fact that workflow knowledge on how to change the timbre of a sound is often encoded within the dialogues occurring during a in a mastering session. As we are aiming for an extension of our system we have to consider the possibilities to extract workflow information from dialogues, which is a current research area in CBR (see, e.g., [21, 18, 22]).

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* 1(7) (Mar 1994)
2. Arbib, M., Fellous, J.: Emotions: from brain to robot. *Trends in cognitive sciences* 8(12), 554–561 (2004)
3. Arcos, J., Grachten, M., de Mántaras, R.: Extracting performers' behaviors to annotate cases in a cbr system for musical tempo transformations. *Case-Based Reasoning Research and Development* pp. 1066–1066 (2003)
4. Arcos, J., De Mántaras, R., Serra, X.: Saxex: A case-based reasoning system for generating expressive musical performances*. *Journal of New Music Research* 27(3), 194–210 (1998)
5. Bartlett, J.: *Practical recording techniques*. Focal Press (2002)
6. Bergmann, R., Göker, M.: Developing industrial case-based reasoning applications using the INRECA methodology. In: *Workshop at the International Joint Conference on Artificial Intelligence, IJCAI - Automating the Construction of Case Based Reasoners*, Stockholm (1999)
7. Bergmann, R., Stahl, A.: Similarity Measures for Object-Oriented Case Representations. In: Smyth, B., Cunningham, P. (eds.) *Advances in Case-Based Reasoning, Proceedings of the 4th European Workshop on Case-Based Reasoning, EWCBR 98*, Dublin, Ireland. pp. 25–36. Springer-Verlag, Berlin (1998)
8. Broekens, J., DeGroot, D.: Emotional agents need formal models of emotion. In: *Proc. of the 16th Belgian-Dutch Conference on Artificial Intelligence*. pp. 195–202 (2004)
9. Canamero, D., Arcos, J., de Mántaras, R.: Imitating human performances to automatically generate expressive jazz ballads. In: *Proceedings of the AISB99 Symposium on Imitation in Animals and Artifacts*. pp. 115–20 (1999)
10. Darke, G.: Assessment of timbre using verbal attributes. In: *Conference on Interdisciplinary Musicology*. Montreal, Quebec (2005)

11. DE MANTARAS, R.: Towards artificial creativity: Examples of some applications of ai to music performance. 50 Anos de la Inteligencia Artificial p. 43 (2006)
12. Donnadieu, S.: Mental representation of the timbre of complex sounds. Analysis, Synthesis, and Perception of Musical Sounds pp. 272–319 (2007)
13. Fellous, J.: From human emotions to robot emotions. Architectures for Modeling Emotion: Cross-Disciplinary Foundations, American Association for Artificial Intelligence pp. 39–46 (2004)
14. Halpern, A., Zatorre, R., Bouffard, M., Johnson, J.: Behavioral and neural correlates of perceived and imagined musical timbre. *Neuropsychologia* 42(9), 1281–1292 (2004)
15. Hudlicka, E.: What are we modeling when we model emotion. In: Proceedings of the AAAI spring symposium—Emotion, personality, and social behavior (2008)
16. Katz, B., Katz, R.: Mastering audio: the art and the science. Focal Press (2007)
17. Lenz, M., Bartsch-Spörl, B., Burkhard, H.D., Wess, S. (eds.): Case-Based Reasoning Technology: From Foundations to Applications, Lecture Notes in Artificial Intelligence, vol. LNAI 1400. Springer-Verlag, Berlin (1998)
18. Madhusudan, T., Zhao, J., Marshall, B.: A case-based reasoning framework for workflow model management. *Data & Knowledge Engineering* 50(1), 87–115 (2004)
19. de Mantaras, R.: Making music with AI: Some examples. In: Proceeding of the 2006 conference on Rob Milne: A Tribute to a Pioneering AI Scientist, Entrepreneur and Mountaineer. pp. 90–100 (2006)
20. de Mantaras, R., Arcos, J.: AI and music: From composition to expressive performance. *AI magazine* 23(3), 43 (2002)
21. Minor, M., Bergmann, R., Görg, S., Walter, K.: Towards case-based adaptation of workflows. *Case-Based Reasoning. Research and Development* pp. 421–435 (2010)
22. Minor, M., Tartakovski, A., Bergmann, R.: Representation and structure-based similarity assessment for agile workflows. *Case-Based Reasoning Research and Development* pp. 224–238 (2007)
23. Oliveira, A., Cardoso, A.: A computer system to control affective content in music production. In: Portuguese Conference on Artificial Intelligence. pp. 69–80 (2007)
24. Pitt, M.: Evidence for a central representation of instrument timbre. *Attention, Perception, & Psychophysics* 57(1), 43–55 (1995)
25. Plaza, E., Arcos, J.: Constructive adaptation. *Advances in Case-Based Reasoning* pp. 306–320 (2002)
26. Porcello, T.: Speaking of sound. *Social Studies of Science* 34(5), 733–758 (2004)
27. Ribeiro, P., Pereira, F.C., Ferrand, M., Cardoso, A.: Case-based melody generation with MuzaCazUza. In: Convention of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB'01) (2001)
28. Richter, M.M.: Introduction. In: Lenz, M., Bartsch-Spörl, B., Burkhard, H.D., Wess, S. (eds.) Case-Based Reasoning Technology – From Foundations to Applications. LNAI 1400, Springer-Verlag, Berlin (1998)
29. Typke, R., Wiering, F., Veltkamp, R.C.: A survey of music information retrieval systems. In: ISMIR. pp. 153–160 (2005)
30. Watson, I.: Case-based reasoning is a methodology not a technology. *Knowledge-Based Systems* 12(5), 303–308 (1999)