![frontiers in Digital Humanities]

Technology Enhanced Learning: The Role of Ontologies for Feedback in Music Performance

Matthew John Yee-King*, Thomas Wilmering, Maria Teresa Llano Rodriguez, Maria Krivenski and Mark d'Inverno

Department of Computing, Goldsmiths University of London, London, United Kingdom

In this paper, we present an analysis of feedback as it occurs in classroom-based and technology supported music instrument learning. Feedback is key to learning in music education and we have developed technology based on ideas from social media and audio annotation which aims to make feedback more effective. The analysis here aims to enhance our understanding of technology-mediated feedback. The result of this analysis is three ontologies describing feedback and feedback systems. First, we developed the *teacher's ontology* using a qualitative, observational approach to describe the types of feedback that music instrument tutors give to their students. We used this ontology to inform the design of an online music annotation platform for music students. Second, we develop the *grounded ontology* using a grounded theory approach, based on 2,000 annotations made by students and tutors using the annotation platform. We compare the grounded and teacher's ontologies by examining structural, semantic and expressive features. Through this comparison, we find that the grounded ontology includes elements of the teacher's ontology as well as elements relating to practical and social aspects of the annotation platform, while the teacher's ontology contains more domain knowledge. Third, we formalize the transactional capabilities of the platform into the third ontology, the *platform ontology*, which we have written in the OWL language, and show how this allows us to develop several practical use cases, including the use of semantic web capabilities in music education contexts.

Keywords: ontology, feedback, computer-assisted instruction, music education, formal modeling

1. INTRODUCTION

Feedback is at the heart of learning a musical instrument; It provides vital information about performance to students. Feedback can come from different sources, both the physical and the social world. For example, the learner listening in real time to their playing, the learner listening later to a recording, and in the social world, from other learners and from expert instrumentalists and teachers who can provide nuanced, empathetic and constructive feedback about a range of different aspects of developing musicianship.

When done well, feedback has a strong, positive effect on learning and achievement. However, according to research by Hattie and Timperley (2007) which measured this effect across multiple studies using meta-analysis techniques, "the type of feedback and the way it is given can be differentially effective," We believe that education technology can be used to address this problem if it supports the provision of effective types of feedback in effective ways.

In particular, we are interested in how we can apply social media and audio annotation technology in music education contexts. Social media presents opportunities for gathering feedback from groups of people, leading to learners receiving more, persistent feedback from the social world. It is well known that feedback received on social media is not always useful; therefore we combine social media with audio annotation tools, which allow people to place annotations at particular points within a recording of music, encouraging them to leave more useful and specific feedback Brenton et al. (2014).

To create and evaluate this technology, we feel it is necessary to gain a systematic understanding of the nature and mechanisms of feedback in music instrument learning. The work presented in this paper contributes to that goal by developing and analysing three ontologies that describe the nature and mechanisms of feedback in face to face and technology supported music education.

Before presenting the ontologies, we will briefly introduce the different types of ontologies we have developed, and we will explain how they can be of use to music teachers and technologists.

1.1. Formal and Informal Ontologies

Firstly, we would like to differentiate between formal and informal ontologies. We view informal ontologies as hierarchies of categories describing the various aspects of something. In our case, the something is feedback about music instrument playing. In work presented here, we derive two such ontologies, one from lesson observations and the expert knowledge of a music teacher, the other from a grounded theory analysis of observations of students and teachers feeding back on each other's music instrument performances in online music tutorials. We present these ontologies as hierarchical diagrams, which should be accessible to music teachers and technologists.

We view formal ontologies as descriptions of operations that can be carried out on information of a specific type, where the information is in a machine-readable form, and the ontology is in a form that can be operationalized as software. In work presented here, the information consists of audio recordings of people playing musical instruments and text-based annotations connected to particular regions of those recordings. The operations are the available actions within an online audio annotation platform we have developed. We created our formal ontology through a process that is similar to a software development process, and we present it here using the Web Ontology Language (OWL).

Music teachers could make use of informal ontologies about music instrument playing in several ways. The ontologies could inform the design of individual lessons or complete curricula, by mapping lessons to the categories in the ontologies. The ontologies could inform assessment methods, presenting a student with a clear overview of the areas in which they will be assessed and how they inter-relate. Teachers could then evaluate the lesson plans in their coverage of the areas for assessment. Trainee teachers could use the ontologies to develop their awareness of the subject, ensuring they have appropriate teaching techniques to address the various areas of the ontologies.

The ontologies could inform the design of learning technology that aims to support music instrument learning, by providing common ground between system designers and teachers. Thus, it could inform a set of system requirements such as 'the system must allow a teacher to feedback about the timbre of individual notes'.

With a formal, computationally operationalizable ontology, the development of learning support technology becomes more straightforward, as it involves the implementation of the required processes in software. Artificial agents could use the ontology to reason about the information, for example, generating automatic feedback, or using the social network to locate appropriate people to provide feedback.

We can draw a parallel here with clinical decision support systems in medicine, which have been shown to make significant improvements to clinical practice (Kawamoto et al., 2005). The problem the patient has is implicitly embedded in a set of symptoms, which can be observed and fed into a (sometimes computerized) decision support system in response to prompts, resulting in a set of recommendations for treatment. The expert clinician carries out the observation and makes the decision about which recommendation to follow, but they are not required to maintain in memory the entire set of possible interactions between symptoms, causes and treatments.

Likewise, in a typical music lesson, the music student's performance on an instrument contains implicit information about problems with their playing technique; the music teacher listens to and watches the performance and then suggests remedies. A music instrument diagnosis system might help the tutor unlock the implicit information and provide more useful guidance to the student.

1.2. Aims and Contributions

Building on the discussion above, this paper works toward a more systematic understanding of the nature of feedback in face to face and technology supported music instrument tuition. It does so by making the following contributions:

1. Descriptions of three different approaches to creating ontologies.
2. Descriptions and a comparison of two informal ontologies specifying acts of feedback in music education.
3. A description and use case evaluation of a formal ontology suitable for describing an audio annotation platform, informed by analysis of the two informal ontologies.

1.3. Structure of the Paper

In order to achieve these contributions the paper is structured as follows: Section 2 provides a review of relevant related work on existing methods for creating ontologies, methods for their comparison, and descriptions of existing ontologies for the domains of music and music education. Section 3 presents the three approaches used in this paper for the creation of ontologies in the domain of feedback in music education, along with the three ontologies resulting from them. In Section 4 we carry out an analysis of the ontologies by comparing two of them, which provide different perspectives of the feedback process. In Section

5 we present use cases with the third ontology (which defines the way feedback occurs in an audio annotation platform). We conclude and provide some future work remarks in Section 6.

Throughout the paper we use Turtle syntax[1] for RDF and OWL listings, and SPARQL[2] for RDF queries. Namespace prefixes are defined as listed in **Table A1**.

2. RELATED WORK

In this section, we discuss previous work about the creation of ontologies, which supports the ontology development work reported in section 3. Following that, we discuss work comparing ontologies, in support of our ontology comparison found in section 4. We end this section by discussing previous, relevant ontologies.

2.1. Ontology Specification Languages

The W3C describes the Semantic Web as "a vision for the future of the Web in which information is given explicit meaning, making it easier for machines to automatically process and integrate information available on the Web" (McGuinness and Van Harmelen, 2004). The W3C maintains specifications for a stack of technologies with which it is possible to implement the Semantic Web. A key element of the Semantic Web is the Web of Linked Data which is constructed from linked statements consisting of *subject, predicate, object* triples, each identified with a uniform resource identifier.

The Resource Description Framework (RDF) is a data model used for making these statements, which can be thought of as a distributed database. As with other databases, querying is a key feature of the Web of Linked Data, but RDF does not provide querying capabilities. Instead, specialized languages exist which provide querying functionality, for example, the Web Ontology Language (OWL).

OWL is a formal ontology specification language for the semantic web which allows machines to operate upon and to reason about the Web of Linked Data (Antoniou and van Harmelen, 2009). Therefore, RDF represents epistemology, or information and OWL specifications represent ontology or modes of operation upon that information.

In this paper, we use the OWL language to specify an ontology describing operations that are possible on a social, audio annotation platform.

2.2. Creating Ontologies

Efficient methodologies for ontology engineering and design have been a much-discussed topic over the years (Gruber, 1993; Uschold, 1996; Gómez-Pérez and Benjamins, 1999; López et al., 2002; Léger et al., 2003; Suárez-Figueroa et al., 2008; Sure et al., 2009). To guide the design of an ontology (Gruber, 1993) proposed a set of five design criteria for ontologies, which has been extended by Gómez-Pérez and Benjamins (1999). The criteria have been developed to maximize the efficiency of an

ontology and to optimize the ontology for knowledge-sharing, i.e., to make it possible for different agents to commit to it. Examples of the criteria are clarity, completeness, coherence and extensibility. We considered these criteria in the development of the platform ontology.

In Uschold (1996) the authors point out that a methodology for building an ontology acts as a guideline, and there are several ways how to tackle the problem. The methodology best employed is largely dependent on the level of formality required for the ontology, and the outcome of the preliminary stages.

Sure et al. (2009) propose a procedure for ontology development that is similar to a software engineering process. We used this process to develop the platform ontology presented in Section 3.4. The procedure begins with a requirements gathering stage that feeds into prototyping of a semi-formal ontology, finally leading into a refinement phase where the formal ontology is developed. The refinement of the ontology is a cyclic process including evaluation and application. Refinement can be based on experience gathered from testing software agents using the ontology or representing domain-specific knowledge using its concepts. The ontology can be updated and adjusted according to the needs emerging from the application.

In the analysis of the online discussion of music practice in subsection 3.3, we apply a grounded theory methodology, as first described by Glaser and Strauss (1967). We describe our application of this method in more detail later, but Charmaz (1996) declares one of the distinguishing characteristics of grounded theory to be "the creation of analytic codes and categories developed from data," so these codes and categories are a key output of a grounded theory analysis. The grounded theory process is iterative, where the set of codes and categories are expanded and back tested against the previously analyzed part of the data, The theory that emerges comes along with evidence from the data and is, therefore, grounded in the data Burck (2005).

Researchers widely use grounded theory in the social sciences for qualitative data analysis, but some have criticized overly dogmatic use which completely ignores existing theory. More recently, researchers have proposed new grounded theory models that can integrate with existing theory, such as Goldkuhl and Cronholm (2010)'s multi-grounded theory.

2.3. Comparing Ontologies

As specified above, there are different methodologies for creating ontologies. However, these act as guidelines rather than as formal procedures. This lack of rigor results in different conceptualizations of the same domains. Consequently, much work has been carried out on methods to compare and merge ontologies. These methods usually fall into two categories (or a combination of them): those that compare ontologies by their structure and those that compare ontologies based on their semantics. In general terms, the first type of methods compare concepts based on how they are distributed into classes, subclasses and individuals, while the second type compares concepts based on their semantic similarity.

An approach for ontology comparison at the structural level has been to map the concepts of distinct ontologies into

---

[1]Terse RDF Triple Language: https://www.w3.org/TR/turtle/
[2]SPARQL Protocol and RDF Query Language: https://www.w3.org/TR/sparql11-query/

a single shared ontology. The aim here is to find relations between concepts by measuring the path (i.e., the distance) between them in the hierarchical structure of the final shared ontology (Bright et al., 1994). Other approaches suggest potential mappings between concept names using different measures of linguistic similarity to find relations between two ontologies (McGuinness et al., 2000; Noy, 2004) or different versions of the same ontology (Klein et al., 2002). Analysing the hyponymy relation (i.e., the is-a relation), commonly used when building ontologies, has also been used as a mean for comparison. Here, concept features are projected through the hierarchy and the subtype-supertype relation is exploited to search for matching points (Fankhauser and Neuhold, 1993; Corbett, 2003). Comparing meta-data about the ontologies also provides a basis for comparing the structures of different ontologies as carried out in Krogstie (2004). Specifically, the number and types of entities and relationships between them provide insights about the density and level of abstraction of ontologies.

At the semantic level, Rodriguez and Egenhofer (2003) and Bai (2013) suggest an integrated assessment by comparing concepts using synonym sets, distinguishing features (specifically their functions, parts and attributes), and semantic neighborhoods, to obtain a measure of semantic similarity between concepts of different ontologies. A similar strategy was followed in Håkansson et al. (2010), where a knowledge base is used to reason about the contents of the ontologies to find suitable concept matches. Corpus statistical information is another approach to semantic comparison; in particular, information content (Resnik, 1995), a technique that measures the level of informativeness of nodes in a taxonomy by calculating their probability of occurrence, is used in Cho et al. (2007) to compute semantic similarity between concepts from different ontologies. Starting from a common root concept, the similarity is given by finding the matching point between the concepts that presents the higher score for information content.

Combined methodologies for comparing ontologies have also been widely employed. In Hovy (2002) and Ehrig et al. (2005) general characteristics of form and content of ontologies are categorized into abstract taxonomies, which are then used to identify similarities and differences between ontologies. Approaches such as Jiang and Conrath (1997) combine the edge-based distance approach with information content analysis to calculate a semantic distance between concepts in a taxonomy, while in Maedche (2001) ontologies are seeing as sign systems that contain different levels of communication that are used to perform the comparison of ontologies. Specifically, they define a syntactic level by calculating the Levenshtein (1966) distance between concept names, and a semantic level by looking at how similar their domain and range concepts are.

As can be observed, a number of these methods work based on a prior integration of the ontologies, usually achieved by defining a top-level common starting point (Rodriguez and Egenhofer, 2003). It is not the scope of this paper to merge the ontologies; however, we follow some of the underlying principles in this section to perform the comparison of the music teacher and grounded ontologies. More detail will be given in Section 4.

2.4. Existing Ontologies for Music and Music Education

Computational musicologists have developed the Music Ontology specifically for the representation and publishing of music-related information on the Semantic Web. It is a formal framework defining concepts and relationships within the music domain that makes it possible to express editorial, cultural and acoustic information. It is built on top of other ontologies, namely the Timeline Ontology[3], the Event Ontology[4] and the Functional Requirements for Bibliographic Records ontology (FRBR) (Davis and Newman, 2005). The basis for the Timeline Ontology are the concepts of *Interval* and *Instant* as defined in OWL-Time (Cox and Little, 2017). The Timeline Ontology defines the *Timeline* as an additional concept thus making it possible to represent intervals and instants on a specific timeline, e.g., the recording date of a music piece as an interval on the universal timeline, or a particular event within a musical recording where the duration of the audio signal itself is the relative timeline. The Event Ontology defines an event in the context of music as a relatively broad concept. These are physical events associated with a music production process which itself can be associated with a specific place and time, people involved, or the musical instruments involved. FRBR is used mainly for the Group 1 entities work (an abstract, distinct, artistic creation), manifestation (physical embodiment, e.g., a digital recording), and item (a single exemplar of a manifestation). Based on the Music Ontology several extensions for specific musical domains have been developed, such as the Chord Ontology[5], the Temperament Ontology[6], the Audio Feature Ontology (Allik et al., 2016) and the Audio Effect Ontology (Wilmering et al., 2016). We reuse the Music Ontology for the knowledge representation in our time based media annotation platform (see Section 3). A different approach for annotating time-based media on the Semantic Web is discussed in Haslhofer et al. (2014), where the annotation of segments of multimedia Web resources is based on Media Fragments[7]. The more expressive approach of using the timeline ontology for defining segments is taken in Yu et al. (2012) in the context of video annotation for distance learning. A survey of media annotation tools for semantic media is given in Nixon and Troncy (2014).

Overviews of applications for teaching musical instrument performance with the help of computer technology is given in Brandao et al. (1999) and Konecki (2015).

Finally, we shall consider some educational ontologies. Sébastien et al. (2011) discussed the development of a collaborative score annotation platform for e-learning in musical instrument education, then they proposed a Musical Forms and Structure Ontology and a Musical Performance Ontology (Sébastien et al., 2013), reusing concepts of the Music Ontology (Raimond et al., 2007). However, work on the project is incomplete, and the ontologies have not been published.

---

[3]http://motools.sourceforge.net/timeline/timeline.html
[4]http://motools.sourceforge.net/event/event.html
[5]http://purl.org/ontology/chord/
[6]http://purl.org/ontology/temperament/
[7]https://www.w3.org/TR/media-frags/

Cavitt (2003) analyzed error correction in instrumental music rehearsals by segmenting recorded rehearsals into frames where each frame is characterized by a "performance goal" or "target," i.e., the desired outcome of the teacher's instructions (Duke, 1999). The set of labels that were used to characterize the lesson segments were developed from previous research in music instrument performance instruction and included "negative feedback," "off task talking," "pitch accuracy," and "technical facility."

The focus of this work, and the previous work by Buckner (1997) from which the labels were derived, was an analysis of the frequency of different types of teacher feedback, as opposed to the development of new theory to describe the modes and types of teacher-student interaction. The music teacher's ontology presented in **Figure 2** was developed independently of this work, but it contains many of the labels used therein.

Considering less formal ontologies, Hanna (2007) developed a mapping between national standards for music education and Bloom's taxonomy (-Anderson and Sosniak, 1994). For example, *recognize* from the Bloom taxonomy maps to "recognize music vocabulary and note symbols" in the music version. In general, though, there seems to be a lack of ontology creation in the education literature and we present our contribution of two informal and one formal ontology to address this.

In summary of the above, we have discussed methods to create ontologies, methods to compare ontologies and also pre-existing ontologies. There are many methods for creating ontologies, some complimentary and some contradictory. The situation is similar for ontology comparison. There are no existing ontologies that are completely suitable to describe the scenario of music education. In the following section, we will describe how we went about developing three new ontologies.

3. CREATING THE ONTOLOGIES: METHODS AND RESULTS

3.1. Case Study and Technical Platform

We developed the teacher's ontology and the grounded ontology in the context of an EU funded project which was investigating the use of social media technology for enhancing feedback in music education. We carried out a case study wherein 49 undergraduate music students, and their tutors used a bespoke online audio annotation system for ten weeks. The audio annotation platform is shown in **Figure 1**.

The students were taking a course in group music performance, and they used the platform to upload, share and collaboratively annotate recordings of their group practice sessions. The platform itself was developed iteratively with music tutors using a user-centered design process. The students were not required to engage in the annotation activity to pass the course, but the majority did do so, posting over 2,000 comments during the case study.

3.2. Creating the Teacher's Ontology

The music teacher who was the course leader for the case studies was also the researcher who created the music teacher's ontology. The ontology aims to represent a hierarchy of topics on which a music teacher might provide feedback to a music instrument learner. The ontology was an output of a wider study considering the role of feedback in undergraduate musical education (Maria Krivenski, 2012). This study involved the collection of data over two academic years from semi-structured interviews and focus groups with music tutors with differing levels of experience and responsibility. The researcher developed the music teacher's ontology based on this data and her own experience as a professional musician and music teacher.

We fed the ontology into the design of the audio annotation platform, during the requirements gathering phase. The intention was that the platform would aim to enable as many of the categories of feedback found in the ontology as possible. The ontology is presented in **Figure 2**.

3.3. Creating the Grounded Ontology

The researcher who created the music teacher's ontology also created the grounded ontology. She used a grounded theory method to develop the ontology. The annotations made by the students and tutors involved in the case study were the input to the analysis. An annotation on an audio file has a reference to which audio file it pertains, a start point, an end point, a username and a text-based comment. The aim was to create a set of categories that effectively described the set of annotations, starting with an empty set. Thus the set of categories is grounded in the data, not in an external theory. The researcher followed this process in creating the grounded ontology:

1. Select 10% of the 2000 annotations.
2. For each annotation, select applicable categories from the existing set of categories (a set which starts off empty).
3. Create and apply new categories if the annotation has not been sufficiently categorized.
4. Re-apply any newly created categories back to previously categorized annotations where they apply.
5. Continue until all of the initial 10% have been categorized.
6. Validate the category set by iterating over the uncategorized annotations and applying the categories. If category set is insufficient, repeat the process above with another subset of the annotations.

The output of the grounded discourse analysis is presented in **Figure 3** in the form of a formal OWL ontology.

3.4. Creating the Platform Ontology

We developed the platform ontology after the other two ontologies. The aim of the platform ontology is to describe the data and interactions that can be found and carried out in the music annotation platform. We developed the platform ontology using the method described by Sure et al. (2009), which involves problem identification, semi formal ontology description, then an iterated process of formal specification and evaluation. This process is shown in **Figure 4**. In the evaluation phase of the process we developed a set of tasks or facilities which the ontology should support, which are as follows:
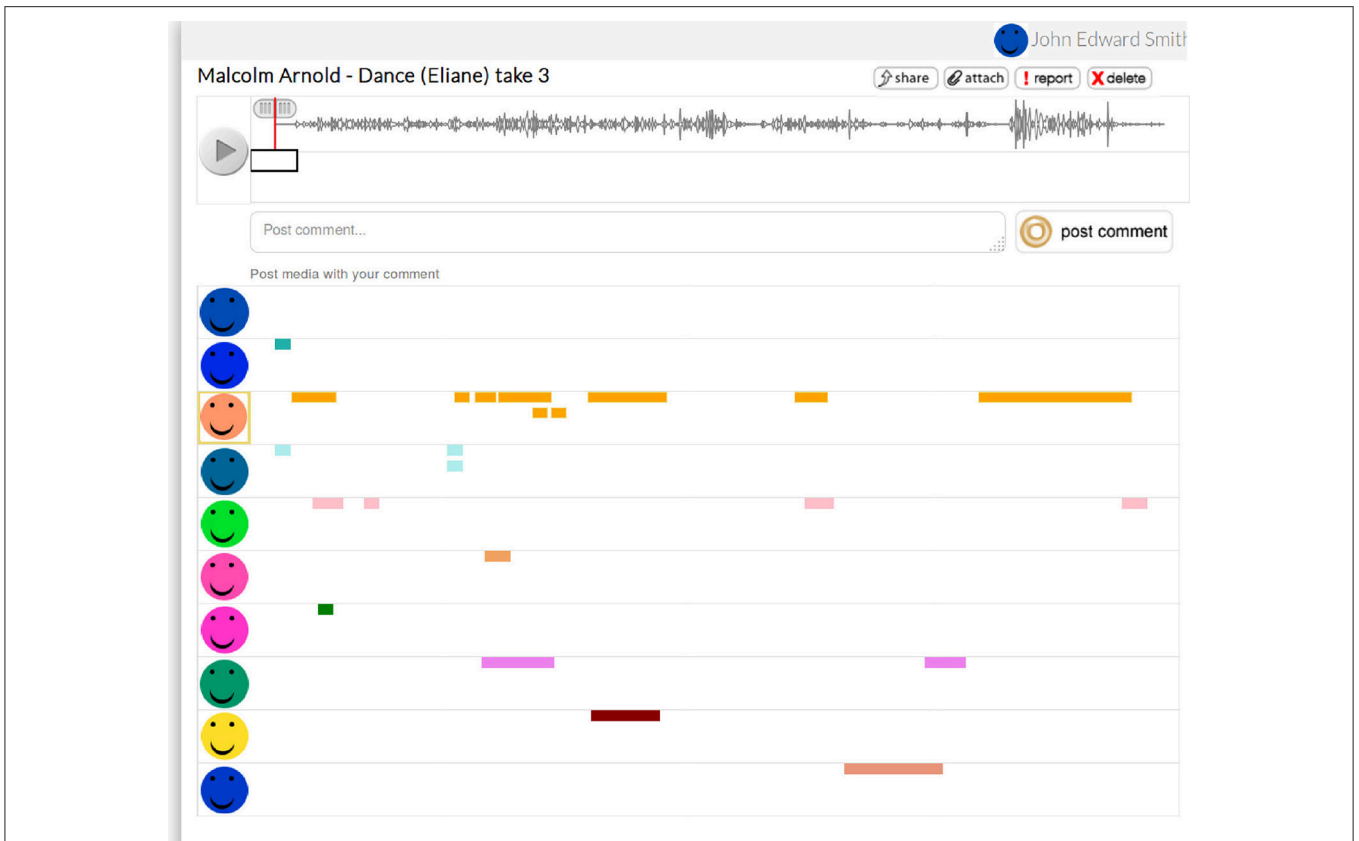
FIGURE 1 | The audio annotation platform. The recording is shown at the top. The colored boxes are annotations, with each color representing a different user.
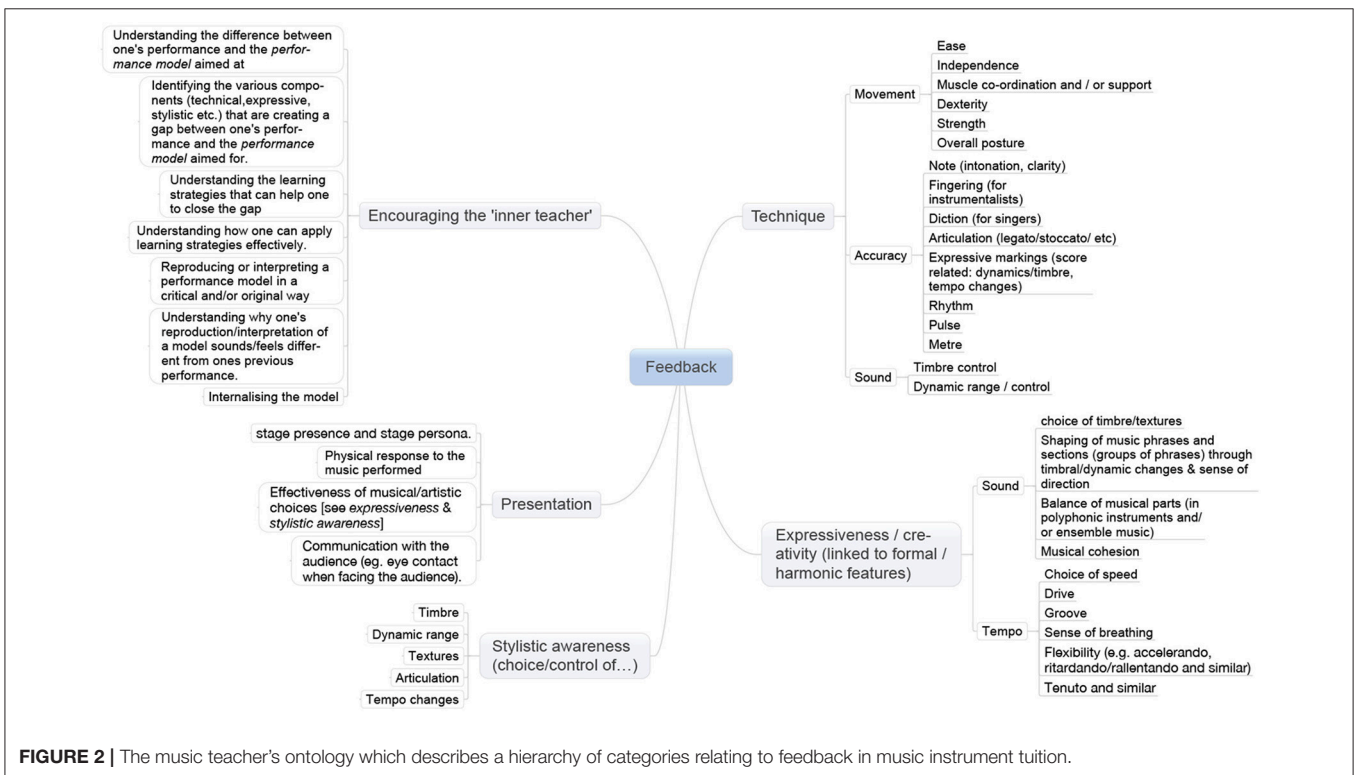

FIGURE 2 | The music teacher's ontology which describes a hierarchy of categories relating to feedback in music instrument tuition.
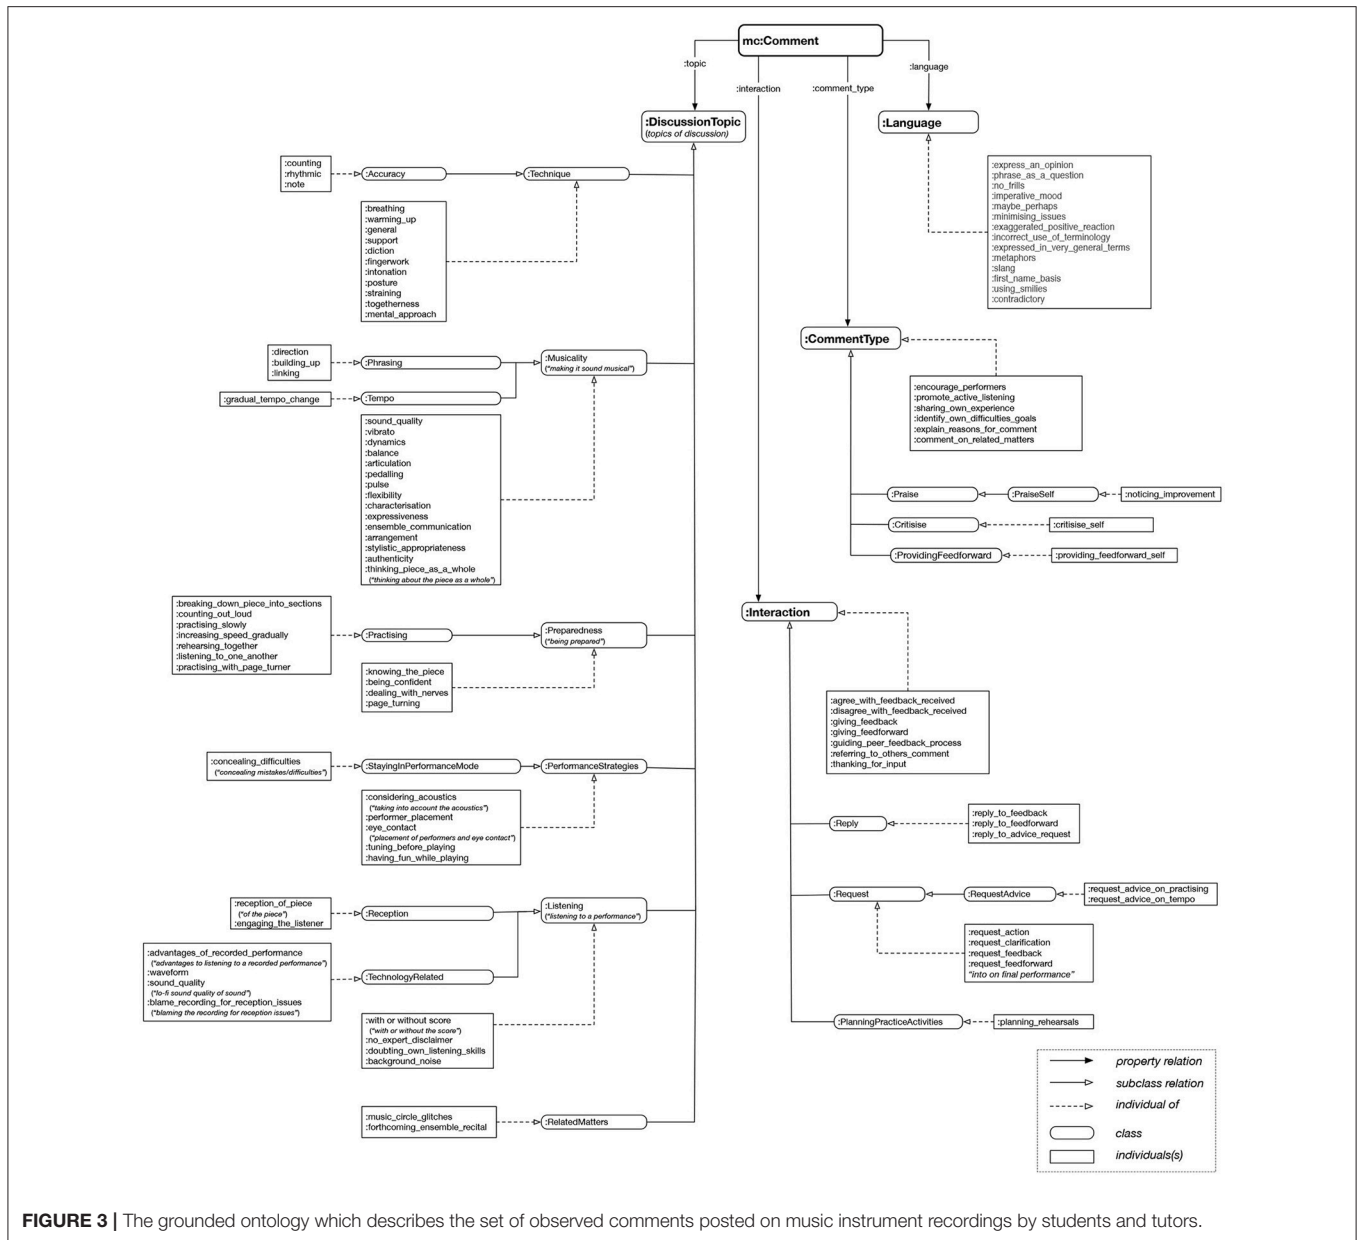
FIGURE 3 | The grounded ontology which describes the set of observed comments posted on music instrument recordings by students and tutors.

- Descriptions of audio material in the system, including provenance and descriptions of its musical content and musicians involved.
- Comments made on the platform annotating certain parts of a performance, directed at specified students.
- Linking comments to the feedback classification encoded in the grounded ontology described in §3.3.
- Associating performances that are influenced by previous feedback to the comments made on the respective earlier performances.

In section 5, we present an evaluation of the platform ontology which establishes that it does support these tasks.

Figure 5 shows the platform ontology's main classes and properties. It reuses several existing ontologies, such as:

- The Music Ontology for the description of music-related data, such as compositions, music performances, and audio recordings.
- The Timeline Ontology for identifying temporal instants and intervals on the audio timeline.
- The Provenance Ontology (PROV-O) (Lebo et al., 2013) for capturing users (agents) and user activity.

Table 1 shows the properties of the Music Ontology and Timeline Ontology that are relevant for the platform ontology, i.e., for the description of recordings of music performances and for the specification of temporal entities (instants and intervals) on their timelines.

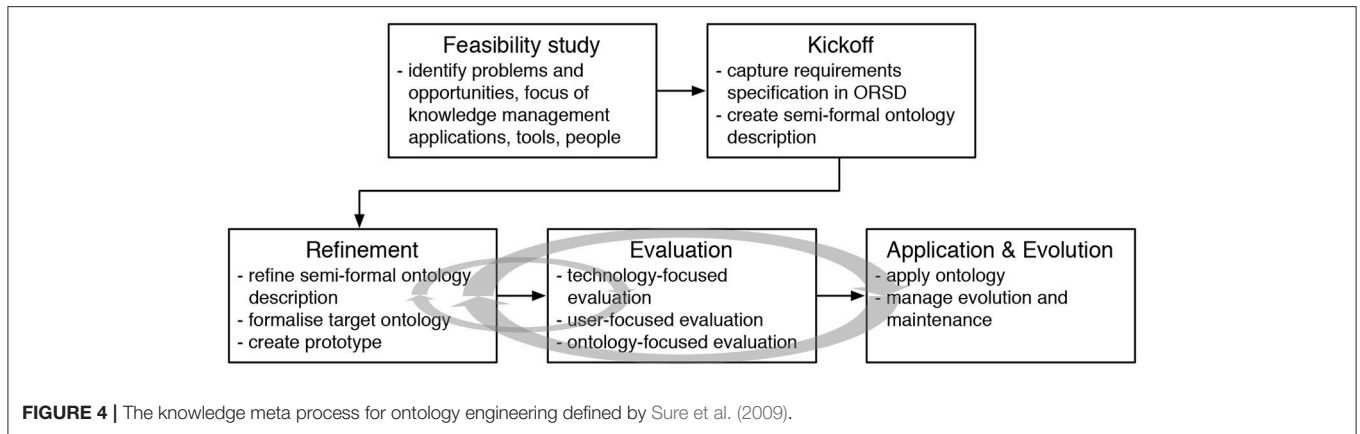The platform ontology defines the class mc:Performance as a subclass of mo:Peformance and prov:Activity.

FIGURE 4 | The knowledge meta process for ontology engineering defined by Sure et al. (2009).
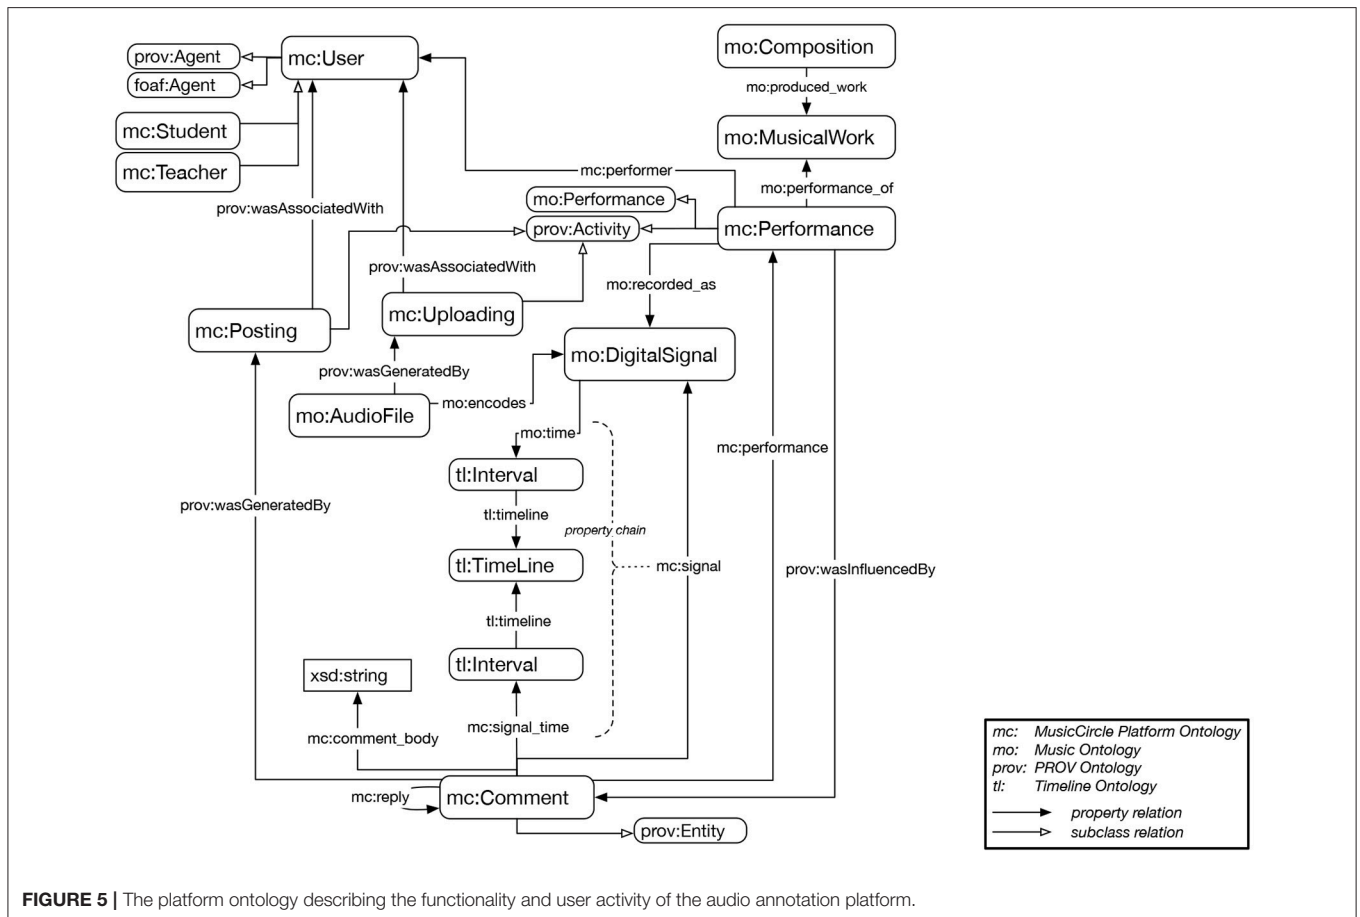

FIGURE 5 | The platform ontology describing the functionality and user activity of the audio annotation platform.

A comment made by a user to give feedback on a music recording is conceptualized as an annotation of uploaded audio material. The ontology defines subclasses of `mc:User` for students and teachers. Performances are linked to performers with `mc:performer`, a property subsumed by both `mo:performer` and `prov:wasAssociatedWith`. A comment (`mc:Comment`) is generated by the user's activity of *posting* the comment defined by the `mc:Posting` class, a subclass of `prov:Activity`. We define concepts to relate a comment (`mc:Comment`) made by a user to a specific interval or instant on the signal's timeline (`mc:signal_timeline`). Moreover, we provide properties to link the comment to the performance (`mc:performance`) and digital signal as a whole (`mc:signal`). The `mc:signal` relation can be inferred by a property chain, linking a comment associated to a temporal event on a given audio timeline to the complete audio signal. The property definition of `mc:signal` is given in Listing 1.

```
mc:signal a owl:ObjectProperty ;
        rdfs:domain mc:Comment ;
        rdfs:range mo:DigitalSignal ;
        owl:propertyChainAxiom ( mo:signal_time
                                 tl:timeline
                                 [ owl:inverseOf tl:timeline ]
                                 [ owl:inverseOf mo:time ]
                               ) .
```

Listing 1 | Property definition of `mc:signal` in the **platform ontology**.

TABLE 1 | Music Ontology properties relevant for the platform ontology.

Property	Domain	Range
mo:recorded_as	mo:Performance	mo:Recording
mo:time	mo:Signal	time:TemporalEntity
tl:timeline	tl:Instant, tl:Interval	tl:TimeLine

The classes `tl:Instant` *and* `tl:Interval` *are subsumed by* `time:TemporalEntity`.

The property chain enables easier querying for comments associated with a given audio signal. Instead of querying via temporal entities on the timeline of the audio signal, comments associated with the signal independently from temporal information can be retrieved by a concise SPARQL query.

The `mc:Comment` class provides the link to the grounded ontology for the classification of feedback (see Section 3.3).

4. ANALYSIS

In this section, we perform an analysis of the music teacher's ontology and the grounded ontology. We compare these two and not the platform ontology as they are designed to describe very similar things, i.e., the topics one might feedback upon when teaching a musical instrument. We identify the similarities and differences between them and highlight the merits and weaknesses of each approach in the process. The platform ontology, which aims to describe activities and data in the online audio annotation platform, is evaluated through use cases in the following section.

4.1. Method for Comparing Ontologies

As specified in an earlier section (2.3), it is possible to compare ontologies at a structural and a semantic level. Here, wo synthesize a method covering structure and semantics, then add another layer of expressiveness, inspired by the work of Raimond et al. (2007). As we shall see, adding expressiveness allows us to identify the type of data and use cases that can be specified with both ontologies. Our method consists of the following elements:

Structural analysis: We carry out an structural analysis of each ontology and compare them based on meta-data features as follows: (i) by identifying the number and types of entities and relationships in the ontologies as suggested in Krogstie (2004), in

order to compare the structural configuration of the ontologies, and (ii) through a measure of information content as defined in Resnik (1995)—which suggests that the higher a concept is in a taxonomy, the more abstract it is, decreasing its degree of informativeness—in order to identify how abstract/specific each ontology is by calculating the depth of detail of each ontology.

Semantic analysis: We carry out a semantic analysis by mapping concepts from both ontologies based on their synonymity as suggested in Rodriguez and Egenhofer (2003) and Bai (2013), in order to identify equivalent entity classes. We also detect concepts unique to each ontology, providing in this way a detailed analysis of the similarities and differences between the them.

Expressiveness analysis: We perform an expressiveness analysis inspired by the work carried out in Raimond et al. (2007) by identifying the types of data that can be described with the ontologies, and subsequently identifying the type of information that can be conveyed through them. In this way, we can compare the comprehensive scope of the ontologies; that is, the kind of scenarios that can and cannot be expressed through their use.

4.2. Results of Comparing Ontologies

Structural analysis: The music teacher ontology is composed of 53 concepts: 11 (%21) classes and 42 (%79) instances or individuals; while the grounded ontology is composed of 128 concepts: 26 (%21) classes and 102 (%79) individuals. This suggests that the grounded ontology is substantially more dense concerning the information that is categorized. A closer inspection of the ontologies reveal that within the grounded ontology, 76 concepts, categorized under the property "DiscussionTopic," overlap significantly with the teacher ontology (they specify topics that are relevant when giving feedback in music education), while the remaining 51 concepts reveal aspects associated with the social nature of the feedback process (e.g., user interaction, types of comments, etc.), which are absent from the teacher music ontology. From the concepts categorized under the property "DiscussionTopic," 14 (%18.4) are classes and 62 (%81.6) represent individuals, which, at first sight, suggests that the grounded ontology covers a larger range of topics than the music teacher ontology (we will analyse this further in the semantic analysis). Furthermore, the grounded ontology establishes four types of high-level properties (topic, interaction, comment_type and language) while the music teacher ontology is defined exclusively following a hyponymy relation between concepts that are based on a *IS-A* relation

TABLE 2 | Structural analysis of the music teacher ontology, the grounded ontologies and a sub-ontology "DiscussionTopic" extracted from the grounded ontology (because of its overlap with the music teacher ontology).

Ontology Feature	Music teacher	Grounded	Discussiontopic (Grounded ontology)
Total concepts	53	128	76
Total classes	11	26	14
Total individuals	42	102	62
Total properties	0	4	0
Avg. Information Content	1.53	1.94	1.72

where concepts share a supertype-subtype association (as is the DiscussionTopic sub-ontology). Thus, the structural composition of the grounded ontology allows for more expressivity than the teacher ontology as it will be evidenced later in this section.

Table 2 shows a summary of the structural composition of the ontologies alongside their average information content. This measure is calculated for a concept c as negative the log likelihood of its probability; i.e., $-log\ p(c)$–where $p(c)$ represents the probability of finding an instance of concept c (measured as the number of concepts subsumed by concept c over the total of concepts in a taxonomy) (Resnik, 1995). In other words, the more concrete a concept, the higher its information content. Although the grounded ontology presents the higher score in information content—which logically follows as this is the larger ontology with the higher number of individuals—the information content of all the ontologies is very similar, which means that their hierarchical taxonomy is similarly distributed (as evidenced by the similar percentages of classes and individuals specify above). Specifically, an average of %80 of the concepts in all ontologies are individuals, which suggest a high level of concreteness in the taxonomies.

Semantic analysis: The underlying semantic difference between the ontologies can be observed immediately from their root concepts: while the teacher ontology starts with the *Feedback* concept, whose focus is on describing the content of feedback, i.e., topics discussed when teaching/learning about music; the grounded ontology starts with the *Comment* concept, which comprises additional features about the social nature of the platform (e.g., type of language used in a piece of feedback, sentiment of its content, etc.). Specifically, as described in Yee-King et al. (2014), two broad types of feedback have been identified in the teacher ontology: (i) feedback connected to desirable traits in a musical performance (described by the concepts under *Technique*, *Expressiveness\Creativity* and *Presentation and Stylistic awareness*), and (ii) "information for guiding tactics and strategies that process the domain-specific information" after Butler and Winne (1995) (under the concept *Encouraging the "inner teacher"*). The grounded ontology, on the other hand, covers a greater range of aspects related to the process of giving feedback within a social machine environment; i.e., it is not focused solely on the content. Firstly, alike the teacher ontology, it provides a categorization of the topics discussed by users of the platform (these are laid out under the *DiscussionTopic* concept as mentioned previously). These

are mostly connected to desirable traits in musical performance (as is the case in the music teacher ontology), with some additional concepts associated with the use of technology, mainly focusing on feedback connected to recorded pieces and aspects related to the functionality of the platform. The ontologies overlap in some classes and instances, in some cases identifiable by name; e.g., concepts in the chain *Technique-Accuracy-Note*, in other cases identifiable by similar meaning and shared sub-concepts; e.g., *Expressiveness/Creativity-Tempo* in the music teacher ontology and *Musicality-Tempo* in the grounded ontology. Secondly, the grounded ontology specifies different social aspects concerning the structure and dynamics of the users of the platform. Specifically, we have identified feedback here as fitting into three general categories: one describing the types of interactions that emerge between users (e.g., comments of encouragement/approval/disagreement, etc.), a second category underlying the collaborative process of creation of feedback (requests/replies, asking_for/offering advice/clarification, etc.), and a third category describing different linguistics resources that are used in order to give feedback (e.g., through figures of speech, by exposing further questions, etc.).

In some aspects, the teacher ontology is more formal and deep than the grounded ontology, highlighting the experience in the domain from the perspective of the expert. For instance, both ontologies address topics related to strategies for self-assessment; however, the teacher ontology does it from the point of view of a *target model*; e.g., "Understanding the difference between one's performance and the performance model aimed at," while the grounded ontology presents more "informal" aspects; e.g., "dealing with nerves," "having_fun_while_playing", etc. This again highlights the nature of the platform as a social machine. Introducing the notion of a target model (a noticeable feature of the teacher ontology) to the platform would be an interesting feature that could be used as a baseline measure for an automated agent to provide feedback.

Expressiveness analysis: Regarding the type of data described, both ontologies lack information about some of the main entities involved in the process of giving feedback in music instrument education, in particular about people (e.g., student, teacher, artist, etc.) and musical devices (e.g., instruments, audio files, etc.). However, the grounded ontology anchors with the platform ontology (through the *Comment* concept), which is itself built on top of the Music Ontology, an ontology able to express such information. Additionally, the teacher ontology lacks transactional categories such as those defined in the grounded ontology that describes diverse types of interactions and exchanges of information carried out in the feedback process. That is (as it has been pointed out before) the teacher ontology is strongly focused on the content of feedback rather than in the process of providing and receiving it.

To conclude we identify that both ontologies can express feedback related to:

● **Traits of musical performance** which describes different topics related to music education; for instance, "work on the *joins between sections* to get them nice and smooth" (phrasing), "try to follow your "*quaver pulse*": can you hear what I

mean?" (counting), "The trumpet could be a bit more *legato*" (articulation), etc.

However, the grounded ontology provides additional levels of expressiveness that can be exploited by a computational agent when interpreting and suggesting feedback. Specifically it can convey information related to:

- The **sentiment** associated with feedback; i.e., whether a piece of feedback is intended as a criticism, an appraisal, a suggestion, etc. For instance, this level of expressiveness can be used to state: "*Sounds great!* In terms of balance, I wonder if the cello could be louder throughout? *It sounds so nice*, especially the deep/low bits, *my ear wants more! :)*" (appraisal).
- The **linguistic** resources being used in a comment. For example, the use of comparative and slang elements can be identified in expressions such as "Pick up the pieces sounds *so much better than* the previous recording. We're *much more together* and *less messy*" (comparison) and "that hi-hat work *is sick!*" (slang).
- The nature of **collaborations** involved in the feedback process. For instance, we can state: "*Could you please* attach a copy of the score? Thanks!" (request).

4.3. Summary of the Analysis

The comparison between the music teacher and the grounded ontologies show that both approaches of creation of ontologies reveal relevant and complementary aspects of a domain. We hypothesized that the systematic nature and amount of data required to perform a grounded theory study made evident conceptual categories that are intrinsic to a domain but may not be evident for an expert when doing the analysis: in our case study this corresponds to the grounded ontology revealing social aspects regarding the feedback process which were not specified in the music teacher ontology. On the other hand, having an expert categorized a domain has numerous benefits; in particular, giving insights about formal aspects of the domain (e.g., the use of a target model as described in Section 4.2) as well as providing a deeper view of some conceptual categories (e.g., *Encouraging the "inner teacher"* in the music teacher ontology).

5. USE CASES FOR EDUCATIONAL ONTOLOGIES

We evaluate the platform ontology using a task-based approach (Porzel and Malaka, 2004), i.e., we test the ontology's performance for use-cases it should be able to support based on the requirements we formulated in Section 3.

5.1. Uploading a Recorded Performance

A user uploads a recorded performance of the Piano Trio by Chopin to the platform. The piece involves three music students, each playing a different instrument (piano, violin and cello). The students are registered as users in the system (`mc:Student`). Reusing concepts of PROV-O, the uploaded audio file is represented as the result of the *uploading* activity of one of the users. The Music Ontology allows us to describe details about the musicians involved in the performance and which

instruments they were playing. The Event ontology provides the mechanism to decompose the performance into sub-events, each conceptualized as an individual performance associated with a performer and an instrument. Using the layered model of the Music Ontology, the musical work (`mo:MusicalWork`) and the audio signal produced by the recording is included in the description. Composers and musical works are linked to external resources, in this case dbpedia[8], providing structured data about the resources. Listing 2 shows RDF for the description of the use case.

Once connected to external data sources, it is possible to retrieve other recordings of the same piece, other combinations of the same instruments and so on. Such capabilities would provide easy access to relevant resources for the students learning the piece.

5.2. Giving Feedback by Commenting on a Performance Segment

The student having uploaded the audio file to the system, the music teacher registers as a user of type `mc:Teacher` on the system and listens to the performance. He or she posts a comment (`mc:Comment`) on the section of the audio material starting at 58s and ending at 120s. The comment addresses the piano performance. In the RDF database this is represented as shown in Listing 3. The property `mc:signal_time` links the comment to an interval on the timeline of the audio signal. The comment is associated with the piano performance with the `mc:performance` property.

Since the performance is linked to the performer, the student can perform queries over the data, for instance, to retrieve all comments associated with his performances. A SPARQL query retrieving the comments linked to performances by `:ex_student_1` is given in Listing 4 (for simplicity, this query is designed to only retrieve the comments linked to sub-events as described in §5.2).

With this functionality, it would be possible to support a range of pedagogical activity. For example, retrieving all comments given by students to evaluate their contribution to a group learning activity or gathering all feedback provided to a student to evaluate the impact of the feedback. Since the comments are attached to audio segments, it would also be possible to carry out automated audio analysis and feature extraction to examine the features indicating performance improvement, for example.

5.3. Linking Comments to the Grounded Ontology

In §3.3 we described the grounded ontology which provides a classification system for feedback in a music education scenario. The platform ontology is linked to the grounded ontology via the `mc:Comment` class. It provides properties to specify the comment type (`mcg:comment_type`), the discussion topic (`mcg:topic`), the type of interaction (`mcg:interaction`), and to describe the language of the comment (`mcg:language`). In Listing 5 we use these properties to link a comment to individuals from the grounded ontology.

---

[8]http://dbpedia.org/

```
:ex_student_1 a mc:Student .
:ex_student_2 a mc:Student .
:ex_student_3 a mc:Student .
:ex_composition mo:composer <http://dbpedia.org/resource/Fr%C3%A9d%C3%A9ric_Chopin> ;
    mo:produced_work :ex_work .
:ex_work a mo:MusicalWork ;
    owl:sameAs <http://dbpedia.org/resource/Piano_Trio_(Chopin)> .
:ex_performance mc:performance_of :ex_work ;
    mo:recorded_as :ex_signal ;
    event:sub_event :piano_performance , :violin_performance, :cello_performance .
:piano_performance a mc:Performance ;
    mo:instrument mit:Piano ;
    mc:performer :ex_student_1 .
:violin_performance a mc:Performance ;
    mo:instrument mit:Violin ;
    mc:performer :ex_student_2 .
:cello_performance a mc:Performance ;
    mo:instrument mit:Cello ;
    mc:performer :ex_student_3 .
:ex_signal mo:time [ a tl:Interval ;
                     tl:timeline :ex_timeline ] .
:ex_audio_file mo:encodes :ex_signal ;
    prov:wasGeneratedBy :ex_uploading .
:ex_uploading a mc:Uploading ;
    prov:wasAssociatedWith :ex_student_1 .
```

Listing 2 | Description of an audio file uploaded to MusicCircle, linked to the musical work, the composer and the performers.

```
:ex_teacher a mc:Teacher .
:ex_comment a mc:Comment ;
    prov:wasGeneratedBy :ex_posting ;
    mc:signal_time [ tl:timeline :ex_timeline ;
                     tl:start "P58S"^^xsd:duration ;
                     tl:end "P120S"^^xsd:duration ] ;
    mc:performance :ex_student_1 ;
    mc:comment_body """You really improved in this part!"""
:ex_posting a mc:Posting ;
    prov:wasAssociatedWith :ex_teacher .
```

Listing 3 | RDF description of a comment annotating a specified audio segment, directed at a specified student.

```
SELECT ?comment
WHERE {
 ?comment mc:performance ?myperformance .
 ?myperformance mc:performer :ex_student_1 .
 ?performance event:sub_event ?myperformance ;
   mc:performance_of ?work .
 ?composition mo:produced_work ?work ;
   mo:composer <http://dbpedia.org/resource/Fr%C3%A9d%C3%A9ric_Chopin> .
}
```

Listing 4 | Example of a comment classified using the four properties of the grounded ontology.

```
:ex_comment mc:topic mcg:rehearsing_together ;
    mc:interaction mcg:giving_feedback ;
    mc:comment_type mcg:noticing_improvement ;
    mc:language mcg:express_an_opinion .
```

Listing 5 | Example of a comment classified using the four properties of the grounded ontology.

5.4. Reacting to Feedback

In education, a main purpose of feedback is for the student to improve in his performance. Thus, it is useful to be able to associate given feedback to subsequent performances of a student. In the platform ontology this can take two forms: (1) a comment by a user is a reply to an existing comment, and (2) a student uploads a new performance acting on the feedback given by another user on a previous performance by the student. For the first case, a comment can be linked to another comment with the `mc:reply` property, for which the `mc:Comment` class is defined both as domain and range. For the second case we reuse the `prov:wasInfluencedBy` property, linking a `mc:Performance` to a `mc:Comment` that is linked to a previous performance, as shown in Listing 6. For further reification of the relationship between comment and performance this model could be extended with a qualified association including the `prov:Plan` concept of PROV-O.

```
:ex_performance_1 a mc:Performance .
:ex_comment_1 mc:performance :ex_performance_1 ;
    mcg:comment_type [ a mcg:Critisise ] ;
    prov:wasGeneratedBy :ex_posting .
:ex_posting a mc:Posting ;
    prov:wasAssociatedWith :ex_teacher .
ex_performance_2 a mc:Performance ;
    mo:recorded_as :ex_signal ;
    prov:wasInfluencedBy :ex_comment .
:ex_audio_file mo:encodes :ex_signal ;
    prov:wasGeneratedBy :ex_uploading .
:ex_uploading a mc:Uploading ;
    prov:wasAssociatedWith :ex_student_1
```

Listing 6 | RDF description.

With this functionality, it is possible to identify which areas of their instrument technique a student needs to improve on. It would also be possible to identify the areas of expertise within a community of music learners, and therefore to identify ideal mentor-mentee relationships.

We have shown how it is possible to use the platform ontology to enable four use cases, including use of other ontologies such as the music ontology and our new, grounded ontology. These use cases reveal a small part of the potential of a fully specified system, and how it might impact on the educational experience of its users.

6. CONCLUSION

This work has come from our belief that feedback is at the heart of learning, and that this is especially so when learning to play music. The ability to give and receive feedback on performance is a critical part of what it is to be a musician, and what it is to learn to get better. Feedback is a deeply human, a deeply creative act. Moreover, even though there are certain limitations of what can be involved in an online context, there are opportunities too. It becomes affordable for students to get feedback on work where the cost of 1-1 tuition is prohibitive as it is in many schools. Feedback can be given in a group context, it is permanent once uploaded, it can specifically relate to a particular passage of the performance, we can set up conversations which start with feedback, we can use the latest audio analyses techniques to provide feedback on pitch, rhythm and timbre. Al of this is what we have been exploring in our work, and to achieve it, we have relied on the development of several ontologies related to the development of a system for supporting feedback in music performance learning.

In the introduction, we enumerated three key aims of this paper. We shall now consider each of these aims in turn to reflect on the contributions we believe we have made:

Descriptions of three different approaches to creating ontologies. We presented an ontology created by a music teacher, which resulted in the categorization of different topics for which a music teacher might provide feedback to a music instrument learner (see Section 3.2). We presented an ontology classifying recurring conceptual categories in our audio annotation platform, through the use of a qualitative grounded theory analysis of a set of 2000 annotations made by students and tutors (see Section 3.3). We presented an ontology describing MusicCircle, developed by an ontologist, which describes both the audio material and interactions between users in the platform (see Section 3.4). To the best of our knowledge, these ontologies provide the first categorization of feedback and feedback interactions, in the context of music education, and in general for the feedback process in learning.

A comparison of two informal ontologies specifying acts of feedback in music education. A thorough comparison, presented in Section 4, was carried out between the music teacher and the grounded ontologies. We analyzed and contrasted the ontologies at different levels. First, at a structural level, we found that although the grounded ontology is significantly bigger in the number of classes, subclasses and individuals. Both approaches yield hierarchical taxonomies that are similarly distributed concerning how concrete their concepts are. At a semantic level, we found that both ontologies specify a number of topics of feedback, some overlapping concepts, and some complementary. The grounded ontology specified additional conceptual categories associated with the social nature of the feedback process, allowing for a more throughout analysis of interactions between students, tutors and peers. At an expressiveness level, we found that both ontologies lack intrinsic elements such as the role of people in the feedback. However, we could easily correct these drawbacks with integration with existing ontologies (e.g., MusicOntology). Moreover, we identified that the grounded

ontology provides additional expressiveness levels that can be beneficial at a computational level, specifically through conceptual categories that specify concepts such as the sentiment of a piece of feedback, the type of linguistic resources used and the type of interactions. An automated agent could exploit this information to generate automatic feedback.

An evaluation of a formal ontology suitable for describing an audio annotation platform, informed by analysis of the two informal ontologies. In Section 5 we followed a task-based approach to test how the platform ontology copes with a set of given use cases derived from scenarios where feedback in the education context takes an important role. We found that the ontology allows for these use cases, and showed their relevance to a pedagogical context.

6.1. Future Work

Artificial Intelligence (AI) makes it possible to implement systems that can automatically reason about ontologies (McNeill and Bundy, 2007). However, with increased intelligence, AI systems should offer increased transparency. The development of techniques to make AI technologies more transparent has become recently important (Baehrens et al., 2010; Bach et al., 2015; Ribeiro et al., 2016; Montavon et al., 2017; Samek et al., 2017). As part of future work, we plan to study models of explanation and justification for the decisions made by automatic agents in the context of music learning. Models such as these, to the best of our knowledge, have not been explored to date. With these models, users and developers can understand what, why and how suggestions are made in social machine environments such as our audio annotation platform.

Lastly, we plan to investigate automatic extraction and displaying of audio features in the annotation platform. By extending our ontologies, it is possible to associate feedback topics with audio features defined in the Audio Feature Vocabulary (Allik et al., 2016). For instance, commenting on rhythm may invoke an onset detector displaying note onsets, while a pitch detector can assist in visualizing issues concerning pitch accuracy in a performance. Moreover, high-level features describing musical structure may further enhance our annotation platform with regards to communication about the nuances of musical performances in the context of music education.

ETHICS STATEMENT

The case study described in this paper involving students at Goldsmiths College was approved by the Goldsmiths Research Ethics and Integrity Sub-Committee, Goldsmiths, University of London, New Cros, London, SE14 6NW, United Kingdom.

AUTHOR CONTRIBUTIONS

MY-K overall paper design and lead postdoc on original music annotation platform research project. MR ontology analysis and comparison. TW formal ontology design and evaluation. MK case study and informal ontology design. MdI oversight on and input to paper and PI on research project.

FUNDING

REFERENCES

Allik, A., Fazekas, G., and Sandler, M. B. (2016). "Ontological representation of audio features," in *Proceedings, Part II The Semantic Web – ISWC 2016: 15th International Semantic Web Conference*, eds P. Groth, E. Simperl, A. Gray, M. Sabou, M. Krötzsch, F. Lecue, F. Flöck, and Y. Gil (Kobe: Springer International Publishing), 3–11.

Anderson, L. W., and Sosniak, L. A. (1994). *Bloom's Taxonomy.* (Chicago, IL: University Chicago Press).

Antoniou, G., and van Harmelen, F. (2009). "Web ontology language: OWL," in *Handbook on Ontologies, International Handbooks on Information Systems*, eds S. Staab and R. Studer (Berlin; Heidelberg: Springer-Verlag), 91–110.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE* 10:e0130140. doi: 10.1371/journal.pone.0130140

Baehrens, D., Schroeter, T., Harmeling, S., Kawanabe, M., Hansen, K., and Müller, K.-R. (2010). How to explain individual classification decisions. *J. Mach. Learn. Res.* 11, 1803–1831.

Bai, F. (2013). *Collaboration Support for the Distributed Development of Ontologies.* Schriften zu Kooperations- und Mediensystemen. (Cologne, MN: Josef Eul Verlag GmbH).

Brandao, M., Wiggins, G., Ão Geraint Wiggins, and Pain, H. (1999). "Computers in music education," in *In Proceedings of the AISB'99 Symposium on Musical Creativity* (Edinburgh).

Brenton, H., Yee-King, M., Grimalt-Reynes, A., Gilles, M., Krivenski, M., and d'Inverno, M. (2014). "A social timeline for exchanging feedback about musical performances," in *Proceedings of the 28th International BCS Human Computer Interaction Conference on HCI 2014 - Sand, Sea and Sky - Holiday HCI*, BCS-HCI '14 (BCS), 281–286.

Bright, M. W., Hurson, A. R., and Pakzad, S. (1994). Automated resolution of semantic heterogeneity in multidatabases. *ACM Trans. Database Syst.* 19, 212–253.

Buckner, J. (1997). *Assessment of Teacher and Student Behavior in Relation to the Accomplishment of Performance Goals in Piano Lessons.* Doctoral Dissertation, The University of Texas at Austin, 1997, Dissertation Abstracts International.

Burck, C. (2005). Comparing qualitative research methodologies for systematic research: the use of grounded theory, discourse analysis and narrative analysis. *J. Fam. Ther.* 27, 237–262. doi: 10.1111/j.1467-6427.2005.00314.x

Butler, D. L., and Winne, P. H. (1995). Feedback and self-regulated learning: a theoretical synthesis. *Rev. Educ. Res.* 65, 245–281. doi: 10.3102/00346543065003245

Cavitt, M. E. (2003). A descriptive analysis of error correction in instrumental music rehearsals. *J. Res. Music Educ.* 51, 218–230. doi: 10.2307/3345375

Charmaz, K. (1996). "The search for meanings-grounded theory," in *Rethinking Methods in Psychology*, eds J. A. Smith, R. Harré, and L. Van Langenhove (Los Angeles, CA: SAGE Publications Ltd), 27–49.

Cho, M., Choi, C., Kim, W., Park, J., Canedo, P. K. A., Abderazek, B., et al. (2007). "Comparing ontologies using entropy," in *2007 International Conference on Convergence Information Technology (ICCIT 2007)*, (Gyeongju) 873–876.

Corbett, D. (2003). "Comparing and merging ontologies: a concept type hierarchy approach," in *Foundations of Intelligent Systems*, eds N. Zhong, Z. W. Raś, S. Tsumoto, and E. Suzuki (Berlin; Heidelberg: Springer), 75–82.

Cox, S., and Little, C., editors (2017). *Time Ontology in OWL.* Available online at: http://www.w3.org/TR/owl-time/

Davis, I., and Newman, R. (2005). *Expression of Core FRBR concepts in RDF.* Available Online at: http://vocab.org/frbr/core

Duke, R. A. (1999). Measures of instructional effectiveness in music research. *Bull. Coun. Res. Music Educ.* 143, 1–48.

Ehrig, M., Haase, P., Hefke, M., and Stojanovic, N. (2005). "Similarity for ontologies - A comprehensive framework," in *Proceedings of the 13th European Conference on Information Systems, Information Systems in a Rapidly Changing Economy, ECIS 2005*, eds D. Bartmann, F. Rajola, J. Kallinikos, D. E. Avison, R. Winter, P. Ein-Dor, J. Becker, F. Bodendorf, and C. Weinhardt. (Regensburg) 1509–1518.

Fankhauser, P., and Neuhold, E. (1993). "Knowledge based integration of heterogeneous databases," in *Interoperable Database Systems (Ds-5), IFIP Transactions A: Computer Science and Technology,* eds D. K. Hsiao, E. J. Neuhold, and R. Sacks-Davis (Amsterdam: Elsevier B.V) 155–175.

Glaser, B. G., and Strauss, A. L. (1967). The discovery of grounded theory. *Int. J. Qualit. Methods* 5, 1–10.

Goldkuhl, G., and Cronholm, S. (2010). Adding theoretical grounding to grounded theory : toward multi-grounded theory. *Int. J. Qualit. Methods* 9, 187–206. doi: 10.1177/160940691000900205

Gómez-Pérez, A., and Benjamins, V. R. (1999). "Overview of knowledge sharing and reuse components: ontologies and problem-solving methods," in *Proceedings of the IJCAI-99 Workshop on Ontologies and Problem-Solving Methods (KRR5)* eds V. Benjamins, B. Chandrasekaran, A. Gómez-Pérez, N. Guarino, and M. Uschold (Stockholm), 1.1–1.15.

Gruber, T. R. (1993). Toward principles for the design of ontologies used for knowledge sharing. *Int. J. Hum. Comp. Stud.* 43, 907–928.

Håkansson, A., Hartung, R., Moradian, E., and Wu, D. (2010). "Comparing ontologies using multi-agent system and knowledge base," in *Knowledge-Based and Intelligent Information and Engineering Systems*, eds R. Setchi, I. Jordanov, R. J. Howlett, and L. C. Jain (Berlin; Heidelberg: Springer), 124–134.

Hanna, W. (2007). The new bloom's taxonomy: Implications for music education. *Arts Educ. Policy Rev.* 108, 7–16. doi: 10.3200/AEPR.108.4.7-16

Haslhofer, B., Sanderson, R., Simon, R., and van de Sompel, H. (2014). Open annotations on multimedia web resources. *Multimed. Tools Applic.* 70, 847–867. doi: 10.1007/s11042-012-1098-9

Hattie, J., and Timperley, H. (2007). The power of feedback. *Rev. Educ. Res.* 77, 81–112. doi: 10.3102/003465430298487

Hovy, E. (2002). *Comparing Sets of Semantic Relations in Ontologies.* (New York, NY: Springer). Start

Jiang, J., and Conrath, D. (1997). "Semantic similarity based on corpus statistics and lexical taxonomy," in *Proceedings of the Intersnational Conference of Computational Linguistics* (Taipei).

Kawamoto, K., Houlihan, C. A., Balas, E. A., and Lobach, D. F. (2005). Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *BMJ* 330:765. doi: 10.1136/bmj.38398.500764.8F

Klein, M., Fensel, D., Kiryakov, A., and Ognyanov, D. (2002). "Ontology versioning and change detection on the web," in *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*, eds A. Gómez-Pérez, and V. R., Benjamins (Berlin;Heidelberg: Springer), 197–212.

Konecki, M. (2015). "Self-paced computer aided learning of music instruments," in *Proceedings of the 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO 2015)*, (Opatija) 809–813.

Krogstie, J. (2004). "Analysing and comparing ontologies with meta-models," in *Information Modeling Methods and Methodologies: Advanced Topics in Database Research: Advanced Topics in Database Research*, eds J. Krogstie, T. Halpin, and K. Siau (Pennsylvania, PA: IGI Global based in Hershey) 1–16.

Lebo, T., Sahoo, S., and McGuinness, D. (2013). *PROV-O: The PROV Ontology. W3C Recommendation.* Available Online at: https://www.w3.org/TR/prov-o/t

Léger, A., Bouillon, Y., Bryan, M., Dieng, R., Ding, Y., Fernandéz-López, M., et al. (2003). *Best Practices and Guidelines.* OntoWeb deliverable 2.2, France Télécom R&D.

Levenshtein, V. I. (1966). *Binary Codes Capable of Correcting Deletions, Insertions, and Reversals.* Technical Report 8. Doklady.

López, M. F., Gómez-Pérez, A., Euzenat, J., Gangemi, A., Kalfoglou, Y., Pisanelli, D. M., et al. (2002). *A Survey on Methodologies for Developing, Maintaining, Integrating, Evaluating and Reengineering Ontologies.* OntoWeb deliverable 1.4, Universidad Politecnia de Madrid.

Maedche, S. (2001). *Comparing Ontologies-Similarity Measures and a Comparison Study.* Internal Report, University of Karlsruhe, Institute AIFB.

Maria Krivenski June 2012 PALATINE project report, (UK: Higher Education Academy York)

McGuinness, D. L., Fikes, R., Rice, J., and Wilder, S. (2000). "The chimaera ontology environment," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on on Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000* , eds H. A. Kautz, and B. W. Porter (Austin, TX: AAAI Press/The MIT Press), 1123–1124.

McGuinness, D. L., and Van Harmelen, F. (2004). Owl web ontology language overview. *W3C Recommend.* 10:2004. Available online at: https://www.w3.org/TR/owl-features

McNeill, F., and Bundy, A. (2007). Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution. *Int. J. Semant. Web Inform. Syst.* 3, 1–35. doi: 10.4018/jswis.2007070101

Montavon, G., Lapuschkin, S., Binder, A., Samek, W., and Müller, K.-R. (2017). Explaining nonlinear classification decisions with deep taylor decomposition. *Patt. Recogn.* 65, 211–222. doi: 10.1016/j.patcog.2016.11.008

Nixon, L., and Troncy, R. (2014). "Survey of semantic media annotation tools for the web: towards new media applications with linked media," in *The Semantic Web: ESWC 2014 Satellite Events: ESWC 2014 Satellite Events, Anissaras, Crete, Greece, May 25-29, 2014, Revised Selected Papers*, eds V. Presutti, E. Blomqvist, R. Troncy, H. Sack, I. Papadakis, and A. Tordai (Anissaras: Springer International Publishing), 100–114.

Noy, N. F. (2004). *Tools for Mapping and Merging Ontologies.* Berlin;Heidelberg: Springer.

Porzel, R., and Malaka, R. (2004). "A task-based approach for ontology evaluation," in *Proceedings of the ECAI Workshop on Ontology Learning and Population* (Valencia).

Raimond, Y., Abdallah, S., Sandler, M., and Giasson, F. (2007). "The music ontology," in *Proceedings of the International Conference on Music Information Retrieval* (Vienna).

Resnik, P. (1995). "Using information content to evaluate semantic similarity in a taxonomy," in *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Vol. 1*, IJCAI'95 (Morgan City, LA: Kaufmann Publishers Inc.) 448–453

Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). ""why should i trust you?": Explaining the predictions of any classifier," in *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 (Paris: ACM), 1135–1144. .

Rodriguez, M. A., and Egenhofer, M. J. (2003). Determining semantic similarity among entity classes from different ontologies. *IEEE Trans Knowledge Data Eng.* 15, 442–456. doi: 10.1109/TKDE.2003.1185844

Samek, W., Wiegand, T., and Müller, K.-R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *ITU J.* 1, 1–10.

Sébastien, V., Sébastien, D., and Conruyt, N. (2011). "Dynamic music lessons on a collaborative score annotation platform," in *The Sixth International Conference on Internet and Web Applications and Services (ICIW 2011)* (St. Maarten), 178–183.

Sébastien, V., Sébastien, D., and Conruyt, N. (2013). "Annotating works for music education: propositions for a musical forms and structures ontology and a musical performance ontology," in *Proceedings of the International Conference on Music Information Retrieval* (Curitiba).

Suárez-Figueroa, M. C., de Cea, G. A., Buil, C., Dellschaft, K., Fernández-López, M., García, A., et al. (2008). *Neon Methodology for Building Contextualized Ontology Networks.* Technical report, NeOn Deliverable D5.4.1.

Sure, Y., Staab, S., and Studer, R. (2009). "Ontology engineering methodology," in *Handbook on Ontologies*, eds S. Staab, and R. Studer (New York, NY: Springer) 135–152.

Uschold, M. (1996). "Building ontologies: towards a unified methodology, in *Proceedings of Expert Systems the th Annual Conference of the British Computer Society Specialist Group on Expert Systems* (Cambridge).

Wilmering, T., Fazekas, G., and Sandler, M. B. (2016). "AUFX-O: novel methods for the representation of audio processing workflows," in *The Semantic Web–ISWC 2016: Proceedings of Part II 15th International Semantic Web Conference on October 17–21, 2016*, eds P. Groth, E. Simperl, A. Gray, M. Sabou, M. Krötzsch, F. Lecue, F. Flöck, and Y. Gil (Kobe: Springer International Publishing), 229–237.

Yee-King, M., Krivenski, M., Brenton, H., Grimalt-Reynes, A., and d'Inverno, M. (2014). "Designing educational social machines for effective feedback," in *Proceedings of the 8th International Conference on e-learning* (Lisbon).

Yu, H. Q., Pedrinaci, C., Dietze, S., and Domingue, J. (2012). Using linked data to annotate and search educational video resources for supporting distance learning. *IEEE Trans. Learn. Technol.* 5, 130–142. doi: 10.1109/TLT. 2012.1

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

APPENDIX

Namespace Bindings

TABLE A1 | Namespace prefixes and corresponding IRIs for RDF and SPARQL listings.

Namespace Prefix	Namespace IRI
=event:=	=http://purl.org/NET/c4dm/event.owl#=
=foaf:=	=http://xmlns.com/foaf/0.1/=
=mc:=	=http://example.com/musiccircle/ontology/=
=mcg:=	=http://example.com/musiccircle/grounded/=
=mit:=	=http://musicontology.com/specification/#term-Instrument=
=mo:=	=http://purl.org/ontology/mo/=
=prov:=	=http://www.w3.org/ns/prov#=
=rdf:=	=http://www.w3.org/1999/02/22-rdf-syntax-ns#=
=rdfs:=	=http://www.w3.org/2000/01/rdf-schema#=
=time:=	=http://www.w3.org/2006/time#=
=tl:=	=http://purl.org/NET/c4dm/timeline.owl#=
=xml:=	=http://www.w3.org/XML/1998/namespace
=xsd:=	=http://www.w3.org/2001/XMLSchema#=