



Deliverable D2.2

Draft Ontology Specification

Grant agreement nr	688382
Project full title	Audio Commons: An Ecosystem for Creative Reuse of Audio Content
Project acronym	AudioCommons
Project duration	36 Months (February 2016 - January 2019)
Work package	WP2
Due date	31 October 2016 (M7)
Submission date	31 October 2016 (M7)
Report availability	Public (X), Confidential ()
Task leader	QMUL
Authors	Mathieu Barthet, Gyorgy Fazekas, Damir Juric, Johan Pauwels, Mark Sandler and Lasse Vetter
Document status	Final (v1.1)





Table of contents

Table of contents	2
Executive Summary	3
Background	5
1 Introduction	6
1.1 Main objectives and goals	6
1.2 Terminology	6
2 Music Related Ontologies	8
2.1 Sound Ontology	9
2.2 PROV Ontology	10
2.3 MUTO Ontology	12
2.4 Music Ontology	13
2.5 Media Value Chain Ontology	14
2.6 Europeana Sounds sound data model	16
2.7 Creative Commons	17
3 Audio Commons Ontology	18
3.1 Actions	19
3.2 Agents	26
3.3 Web resources	28
3.4 Musical expression	30
3.5 Audio Commons Objects	32
4 Conclusion	34
5 References	35
APPENDIX 1: Audio Commons Ontology excerpt	36
APPENDIX 2: Audio Commons Ontology Schema	37





Executive Summary

This deliverable is part of WP2, a work package that will describe the Audio Commons Ontology and develops the Audio Commons API specification. Building the Audio Commons Ontology requires gathering knowledge from real and anticipated user requirements and already existing ontologies in the music and audio domain, as well as knowledge about the workflows and tasks that people who work in music domain use in their everyday work. The version of the Audio Commons ontology presented in this deliverable is the first draft of the ontology. This version of the ontology should provide the basis for all future iterations. As the project advances, the ontology will be continuously evaluated against new requirements. This means that future versions of the Audio Commons ontology will contain new classes and relationships or whole new taxonomies will be attached to the base classes.

This deliverable will inform the definition of the Audio Commons API specification. The Audio Commons API will define how the different components of the Audio Commons Ecosystem will be technically interconnected (i.e., how production tools will be able to access Audio Commons content and users will be able to communicate for the licensing process). The ontology will be mainly used by the semantic mediator, an application layer that will orchestrate different services or components integrating the Audio Commons Ecosystem, providing the required technology layer for the interconnection of the different components, and providing the required software packages and guidelines to facilitate the incorporation of new actors in the Audio Commons Ecosystem, inform research on rights management and help focus the work on sound and music analysis algorithms and end user prototypes.

In this deliverable we analysed existing ontologies and vocabularies that describe music and audio related domains. Following the guidelines of Semantic Web about the reuse of as much of existing data as possible (ontologies and vocabularies like the Music ontology, Media Value Chain Ontology, Provenance Ontology, Sound ontology, Europeana profile for sound, Modular Unified Tagging Ontology, etc.) we created the Audio Commons ontology as a bridge between these existing ontologies. Since the Audio Commons project strives to connect different tasks carried out in a typical workflow of practitioners working in the audio domain and provide the context for different inputs and outputs of that workflow, modelling the relevant tasks or actions is very important. The Audio Commons ontology deals with different domains: creative, legal and physical domain (see Section 3). Defined action groups in the Audio Commons ontology are replicating that distinction between quite separate domains from the real world. The Audio Commons ecosystem will gather different providers and technology platforms that could possibly offer a wide range of services. In line with the Provenance ontology model, the Audio Commons ontology will be able to represent and capture the complex interaction between the agents, actions and produced entities, that will ultimately lead to creation of web service orchestration based on that knowledge. The Audio Commons ontology will contain metadata that will describe the sound model in a way that metadata coming from different providers can be merged into one unique Audio Commons sound model. Additionally, the Audio Commons ontology will contain metadata that describes and facilitates rights management. This kind of metadata could be used by some Audio Commons services that will create semantic ledgers (i.e., linked data that will describe the lifecycle of Audio Commons objects).





Background

This deliverable is part of the task T2.2 from work package WP2. WP2 will develop the Audio Commons Ontology. Building the Audio Commons Ontology (and ontologies in general) requires a clear definition of scope and extensive knowledge about the domain that the ontology describes. Knowledge about the domain includes the vocabulary of the domain and knowledge about the workflows (processes) that are being carried out by various roles involved in them. Deliverable D2.1 presented the results of the survey that we circulated among professionals and amateurs in the music creation domain. Participants were asked to answer questions about the daily tasks which can give us an insight into the complex network of entities and tasks that are being carried out in the music domain. Those entities and tasks were used as a starting point for the Audio Commons ontology. Another important source of information for the Audio Commons ontology were the already existing music and audio related ontologies. Those ontologies were analysed and connection points were identified. The Audio Commons ontology is developed from these connection points.

Deliverables D2.3 and D2.4 of work package WP2 will provide the Audio Commons API specification, which will define how the different components of the Audio Commons Ecosystem will be technically interconnected (i.e., how production tools will be able to access Audio Commons content and users will be able to communicate for the licensing process) and will enable the orchestration of the different services or components integrating the Audio Commons Ecosystem. A semantic mediator will act as the required technology layer for the interconnection of the different components, and providing the required software packages and guidelines to facilitate the addition of new actors in the Audio Commons Ecosystem. The Audio Commons ontology will play an important role in the task of service orchestration because it will provide necessary data and knowledge about the agents and actions that will be orchestrated.





1 Introduction

1.1 Main objectives and goals

The main objective of this deliverable is the description of the first draft of the Audio Commons ontology. The task was planned around the following objectives:

- Examination of the existing music and audio related ontologies
- Examination of the metadata provided by the survey described in deliverable D2.1
- Connecting the data from the previous two objectives into Audio Commons ontology

This task will inform the definition of the Audio Commons API, inform research on rights management and help focus the work on sound and music analysis algorithms and end user prototypes. The main goal of the ontology is to make the Web service orchestration possible by providing necessary knowledge about entities in that process.

In this deliverable we present the first version of Audio Commons ontology. This means that the ontology will go through a certain number of iterations and evaluations depending of new requirements arising during the implementation stage of the application layer that will use the Audio Commons ontology.

In line with the objectives above, this deliverable features two main chapters. Chapter 2 briefly describes the existing music and audio related ontologies and comments about the potential usefulness of each ontology for the Audio Commons ontology and within the Audio Commons Ecosystem (ACE). Chapter 3 presents the main entities of Audio Commons ontology with modeling examples of some activities that are important for the ACE.

1.2 Terminology

AudioCommons: reference to the EC H2020 funded project AudioCommons, with grant agreement nr 688382.

Audio Commons Initiative: understanding of the AudioCommons project core ideas beyond the lifetime and specific scope of the funded project. The term "Audio Commons Initiative" is used to imply i) our will to continue supporting the Audio Commons Ecosystem and its ideas after the lifetime of the funded project, and ii) our will to engage new stakeholders which are not officially part of the project consortium.

Audio Commons: generic reference to the Audio Commons core ideas, without distinguishing between the concept of the initiative and the actual funded project.





AudioCommons

D2.2 Draft Ontology Specification

Audio Commons Ecosystem (ACE): series of technologies and actors involved in publishing and consuming Audio Commons content.

Audio Commons content (AC): audio content released under Creative Commons licenses and enhanced with meaningful contextual information (e.g., annotations, license information) that enables its publication in the ACE.

Content creator: individual users, industries or other actors that create audio content and publish in the ACE through content providers.

Content provider: services that expose content created by content creators to the ACE.

Content user: individual users, industries or other actors that use the content exposed by content providers and created by content creators in their creative workflows.

Ontology: In the context of computer and information sciences, an ontology defines a set of representational primitives with which to model a domain of knowledge or discourse. The representational primitives are typically classes (or sets), attributes (or properties), and relationships (or relations among class members). The definitions of the representational primitives include information about their meaning and constraints on their logically consistent application. In the context of database systems, ontology can be viewed as a level of abstraction of data models, analogous to hierarchical and relational models, but intended for modelling knowledge about individuals, their attributes, and their relationships to other individuals. Ontologies are typically specified in languages that allow abstraction away from data structures and implementation strategies; in practice, the languages of ontologies are closer in expressive power to first-order logic than languages used to model databases.

Tool developer: individual users, industries or other actors that develop tools for consuming (and also potentially publishing) Audio Commons content.

Embeddable tools: tools for consuming Audio Commons content that can be embedded in existing production workflows of creative industries.





2 Music Related Ontologies

An ontology defines a set of representational primitives with which we can model a domain of knowledge or discourse. Those representational primitives are typically classes (or sets), attributes (or properties), and relationships (or relations among class members). The definitions of the representational primitives include information about their meaning and constraints on their logically consistent application.

One of the main imperatives of the Semantic Web is that ontologies should not be built entirely from scratch and that previously used vocabulary should be reused as much as possible. For the Audio Commons ontology, outside of the knowledge collected from the survey, the source of knowledge comes from already existing ontologies describing music and audio related domains. Figure 1 shows ontologies that have been partially reused to create the Audio Commons ontology. Audio Commons ontology can be seen as a bridge between those ontologies.

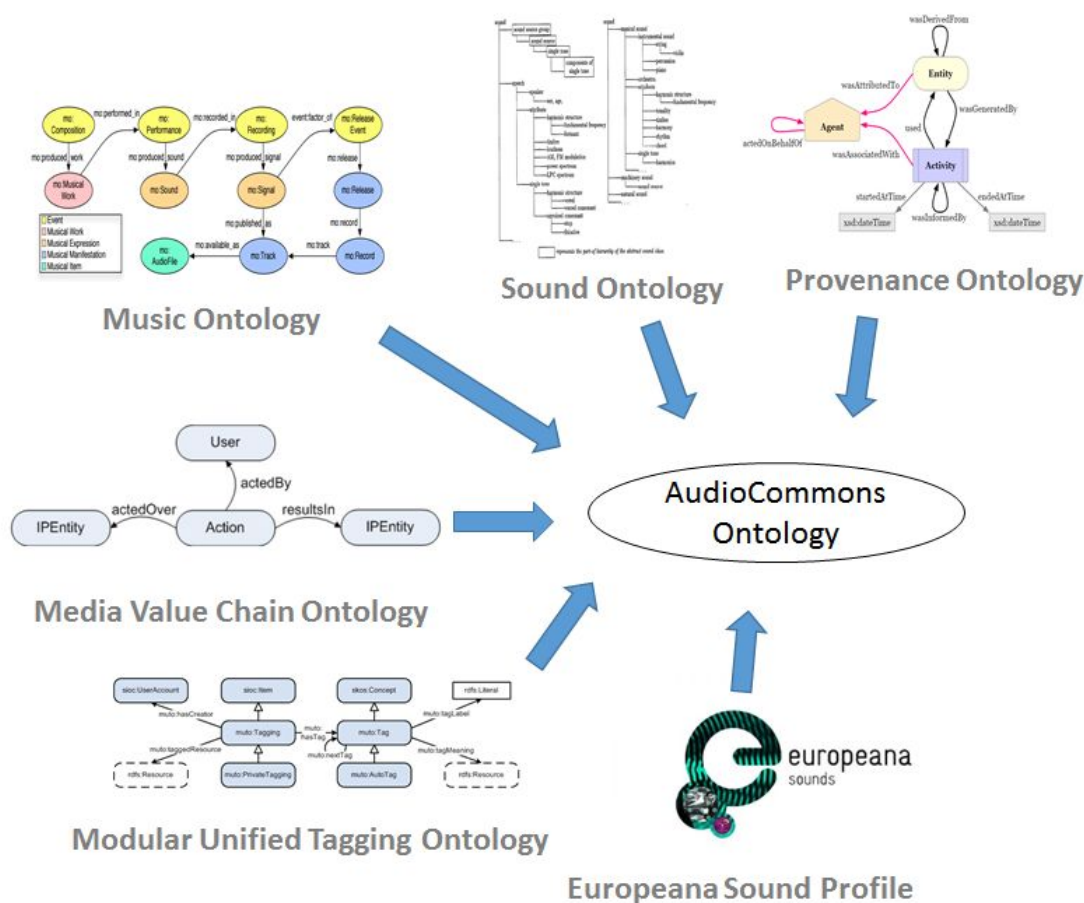


Figure 1. Audio Commons ontology sources





2.1 Sound Ontology

In [SoundOnt] the authors present the Sound ontology that is composed of sound classes, definitions of individual sound attributes, and their relationships (Figure 2). They define the hierarchy by using the following two attributes:

- Part-of hierarchy - a hierarchy based on the inclusion relation in sound and
- Is-a hierarchy - a hierarchy based on the abstraction level in sound

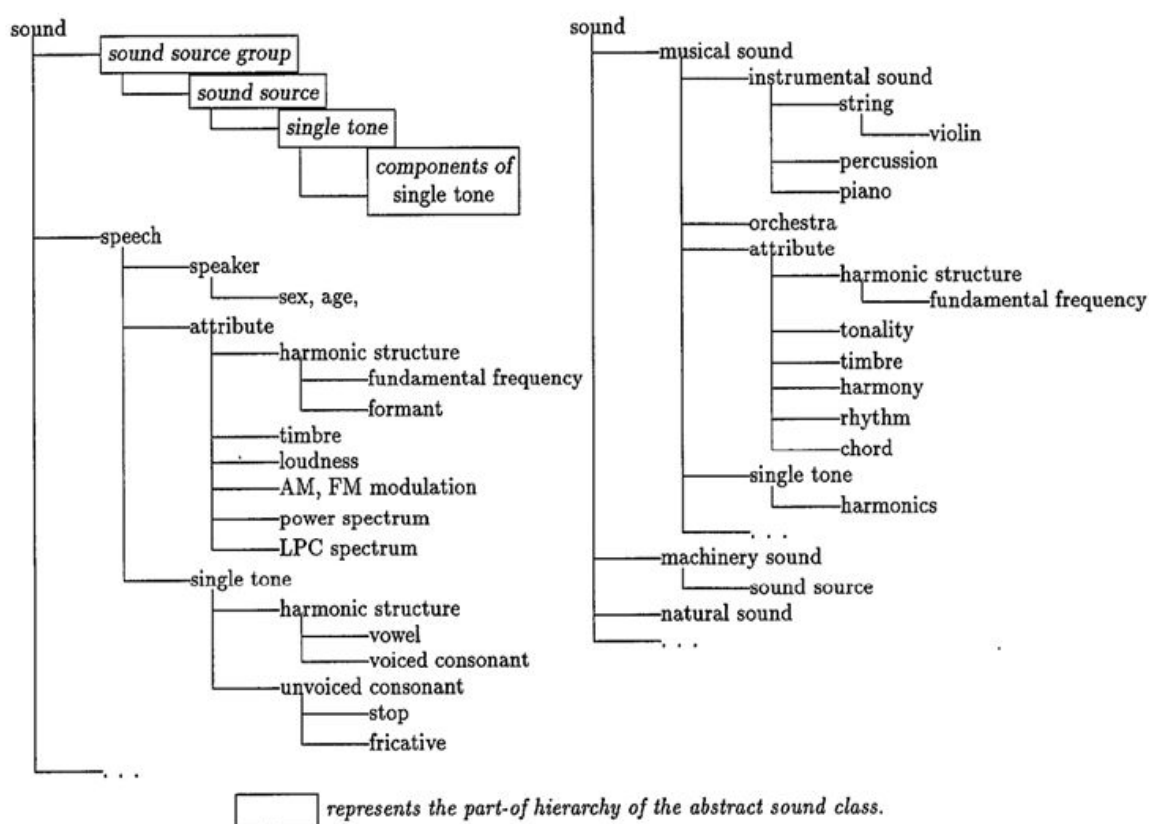


Figure 2. Sound ontology depicted in [SoundOnt]

This ontology is a result of the authors work on a research in the field of computational auditory scene analysis. The goal of the research is the understanding of an arbitrary sound mixture including non-speech sounds and music. The authors' intention was the development of a system that can recognize both speech and music from a mixture of voiced announcements and background music. The Audio Commons ontology will not adapt this ontology directly. This ontology will only serve as a guideline for defining relationships between attributes that are important for Audio Commons (like timbre, harmony, chord, etc.) and a useful starting point as an enumeration of relevant concepts.



2.2 PROV Ontology

PROV ontologies describe *provenance*. Provenance is information about entities, activities, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness. This is important in the context of ACE because knowledge about how content is produced and distributed is useful for search, retrieval and rights management purposes.

The PROV ontology [ProvOnt] defines three main classes:

- A [prov:Entity](#) is a physical, digital, conceptual, or other kind of thing with some fixed aspects; entities may be real or imaginary.
- A [prov:Activity](#) is something that occurs over a period of time and acts upon or with entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating entities.
- A [prov:Agent](#) is something that bears some form of responsibility for an activity taking place, for the existence of an entity, or for another agent's activity.

These three primary classes relate to one another and to themselves using the properties as can be seen in figure 3.

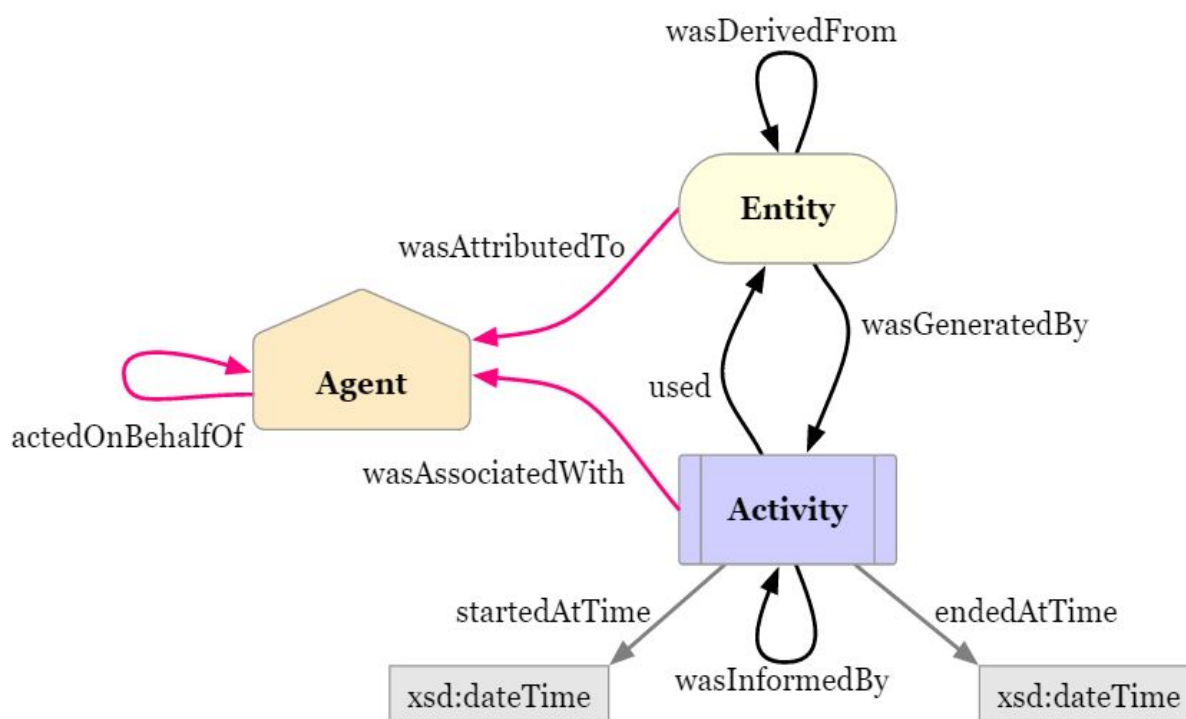


Figure 3. Provenance ontology [ProvOnt]



AudioCommons

D2.2 Draft Ontology Specification

As mentioned above, the PROV ontology describes metadata that will provide information about any kind of activity and the participants involved in that activity. The main classes shown in Figure 3 are very general and the Audio Commons ontology will define a subset of subclasses that are specific to the Audio Commons ecosystem.

The Audio Commons ecosystem is focused on connecting different agents and orchestrating different activities. The list of agents and activities that will be implemented in the Audio Commons ecosystem will grow as Audio Commons develop with time and new services are added. What follows are two examples of the use of the PROV ontology to model two basic services that the ACE will potentially support, searching and uploading:" instead.:

- searching - in line with PROV ontology, search action should be treated as *Activity* (*searchActivity* - this is now an Audio Commons class that is a subclass of PROV Activity class) that is associated (*wasAssociatedWith*) with certain *Agent* (Audio Commons user) at certain date and time (*startedAtTime*). The *query* typed by the user would be a subclass of PROV *Entity* class. The *query* and *searchAction* should be connected with relationships defined in PROV (*used* and *wasGeneratedBy*).
- uploading - similar as previous example, uploading action can be described in PROV ontology through creation of *uploadAction* class and *Agent* class (AudioCommon end user of software) and *Entity* class that will be the thing being uploaded (any Audio Commons object).

Any new service that will be added to the Audio Commons ecosystem should be described similarly to the search and upload services in the example.





2.3 MUTO Ontology

The MUTO ontology (Modular Unified Tagging Ontology) [MUTO] describes the action of tagging resources, which is very important for audio related services. The MUTO ontology describes what is a tag and what is tagging action in a few axioms. Tag axioms include:

- A tag has always exactly one label - otherwise it is not a tag. (Additional labels can be separately defined, e.g. via skos:Concept.)
- Tags with the same label are not necessarily semantically identical. (Each tag has its own identity and property values.)
- A tag can itself be a resource of tagging (tagging of tags).

Tagging axioms include:

- A tagging is always linked to exactly one resource and one user account (the latter can be omitted in case of automatic tagging).
- The number of tags per tagging is theoretically unlimited (though it is practically limited by the constraints of the tagging system).
- A tagging can itself be a resource of another tagging (tagging of taggings).

Figure 4. shows the main concepts of the MUTO ontology.

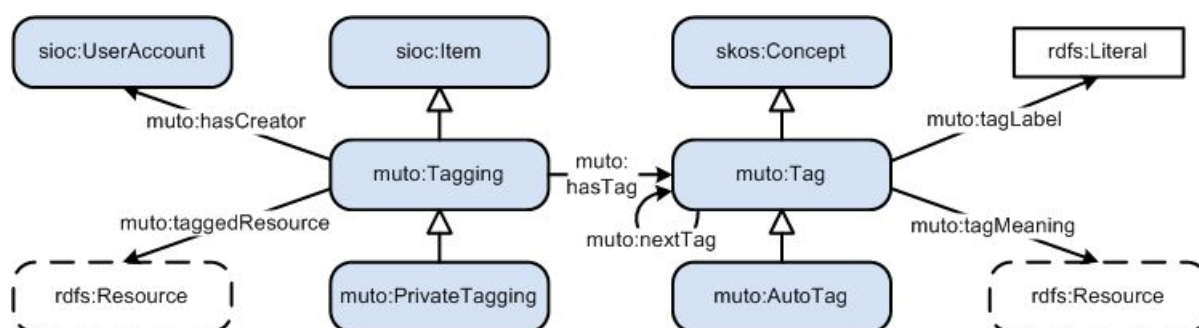


Figure 4. Main concepts of the MUTO core vocabulary [MUTO]

The MUTO ontology defines a special type of activity called *tagging*. While the authors of MUTO ontology never stated that they were inspired by the PROV ontology, the PROV model can be applied to the tagging activity. Unlike activities like *searching* or *uploading* that will require the creation of a special Audio Commons class to represent those activities in Audio Commons as subclasses of PROV *Activity*, the tagging activity does not require creation of special Audio Commons class. The class *muto:Tagging* will take the role as a subclass of PROV *Activity*, while *muto:hasCreator* will point to AudioCommons agent and *sioc:Item* will represent Audio Commons object (and PROV *Entity*). This way the Audio Commons ontology will reuse the MUTO ontology for any activity that involves tagging.



2.4 Music Ontology

The aim of the Music Ontology framework [Raimond] is to provide a comprehensive, yet easy to use and easily extended domain specific knowledge representation for describing music related information. Integration of music related resources (Web services and data repositories) on the Semantic Web, and facilitation of service integration and data communication in distributed music processing environments are among its existing applications. It has certain properties which make it particularly suitable as basis for a general semantic audio information management framework as well as data collection in recording and production. For instance, it relies on, and extends the full FRBR (Functional Requirements for Bibliographic Records [FRBR]) model, and provides an event based conceptualisation of music production workflows as shown in Figure 5.

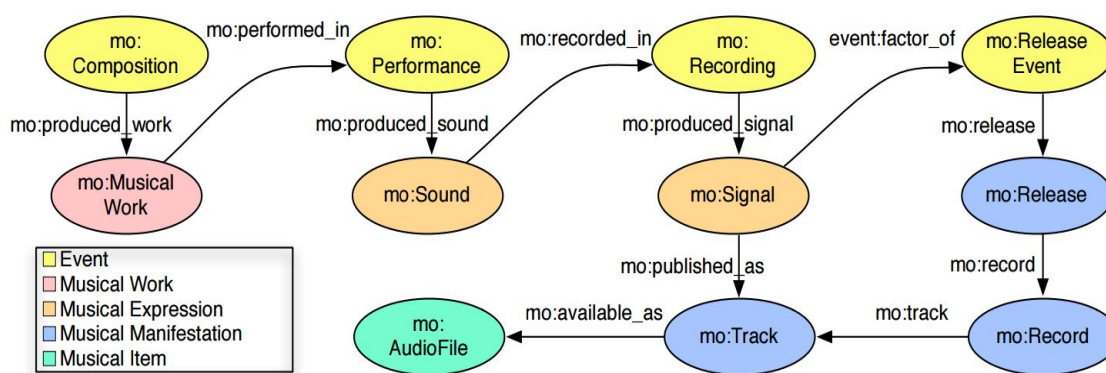


Figure 5. Music Production Workflow Model [Fazekas]

The Music ontology can represent the following things [Fazekas]:

- Editorial metadata: Concepts and relationships involving artists, bands, labels, albums, tracks, audio files or downloads and their identifiers in various databases.
- Music production workflow: The life cycle of musical works from composition through performance, to the produced sounds and recorded signals and their publication.
- Event decomposition: Further details about particular events in the production workflow such as individual performances by different musicians in a recording.
- Content annotation: Audio signals and temporal annotation of their content.

In the broader context (outside the domain music production workflow), the Music ontology lacks the connection with the concepts coming from the legal and intellectual property side of the audio domain, as well as web and web services concepts or a generic sound model that can represent audio content other than music pieces.

2.5 Media Value Chain Ontology

The Media Value Chain Ontology (MVCO) is an ontology for formalizing the representation of the Media Value Chain. The MVCO represents the Intellectual Property (IP) along the Value Chain. Originally, the Value Chain represents a high-level model that describes the process by which businesses receive raw materials, add value to the raw materials through various processes to create a finished product, and then sell that end product to customers. There are different kinds of objects of the Intellectual Property (called IP Entities) and different actions that are performed on them, which define the different roles that users can play regarding these IP Entities. These elements, along the permissions to execute the actions, constitute the essence of the MVCO. The most important are IP Entities, Actions and User roles.

The Media Value Chain ontology is a part of a MPEG-21 multimedia framework which makes this ontology relevant to the audio domain. The author of the Media Value Chain ontology even provides some mappings between this ontology and the Music ontology (Figure 6) and mappings between Creative Commons RDF representation and Media value Chain ontology entities. Some of these mappings are useful and some of them are questionable and open for discussion. In Audio Commons we will include only the mappings that are straightforward and clear.




Figure 6. Action in Media Value Chain ontology [MediaValueChain]

The Media Value Chain ontology models its classes in line with the PROV ontology model. Instead of the *Activity* class as in the PROV ontology, the Media Value Chain ontology defines a class labeled *Action* that takes central part in the model (*Action* class in MVCO is a subclass of *Activity* class in PROV). The Audio Commons ontology will adopt the Media Value Chain ontology and extend it with Audio Commons subset of classes as necessary.

The MVCO defines a certain number of agents in its vocabulary. These agents may trigger certain types of actions. Here, the MVC ontology does not specify very clearly what types of actions can be triggered by a certain Agent. Also the set of actions that are defined in MVCO is limited and Audio Commons needs to extend this set with the actions that are relevant to the ACE (and especially relevant to licensing choices). The best example for this problem is the action type class labeled *MakeAdaptation* which is described as: "The Action of making an Adaptation". The meaning of the word adaptation is very general and it is not finely grained enough to allow us to represent the level of detail we would like to see in the Audio Commons ecosystem. The Audio Commons ontology should include more finely



grained representation of actions that can be done over the Audio Commons object (like tweaking, remixing, etc.).

Figure 7. is a good illustration of how complex the relationships that describe changes of creative works may become (this example is taken from the Functional Requirements for Bibliographic Records (FRBR) ontology).

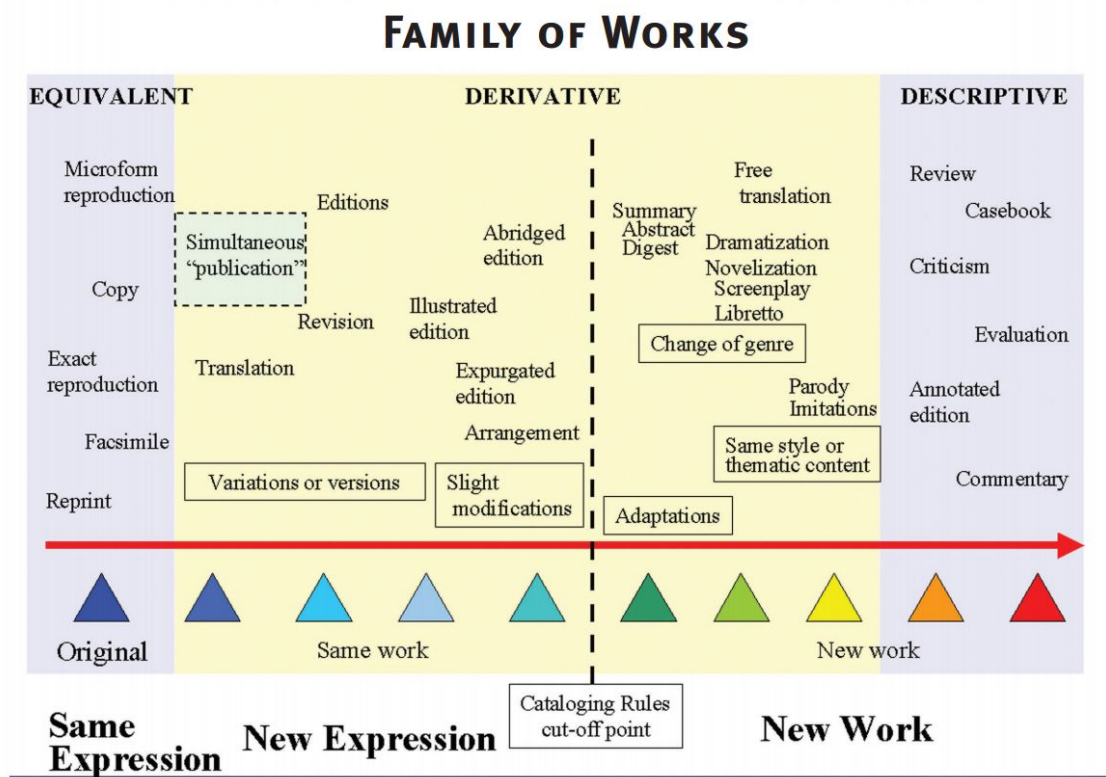


Figure 7. Derivation of creative works [FRBR]

As previously stated, for now the Audio Commons ontology will define licence changing actions defined in the Creative Commons framework.



2.6 Europeana Sounds sound data model

Europeana (<http://www.europeana.eu>) is an initiative co-financed by the European Union to incentive member states to digitise and make Europe’s cultural heritage accessible through a digital platform. The portal provides free access to a wide array of digital content including audio. In the context of Europeana, Europeana Sounds (<http://www.europeanasonsounds.eu>) is a project funded by the European Commission whose aim is to increase the amount of audio content available via Europeana, to improve geographical and thematic coverage by aggregating recordings with widespread popular appeal, to enrich this content and to promote the creative reuse of the content. Europeana Sounds strives to add meaningful contextual knowledge and medium-specific metadata to 2 million items in Europeana’s audio and audio-related collections, developing techniques for cross-media and cross-collection linking.

Work carried out in the context of Europeana Sounds consisted in identifying a set of requirements essential to the description of sound objects from the Europeana collection and extending the Europeana Data Model (EDM) with new properties to better accommodate audio and audio-related objects [Europeana]. Figure 8. shows the modelling of the concept of a Web resource in Europeana, where anything that is part of the Europeana ecosystem will “live” as a Web resource. Similar concepts will be present in the Audio Commons ontology.

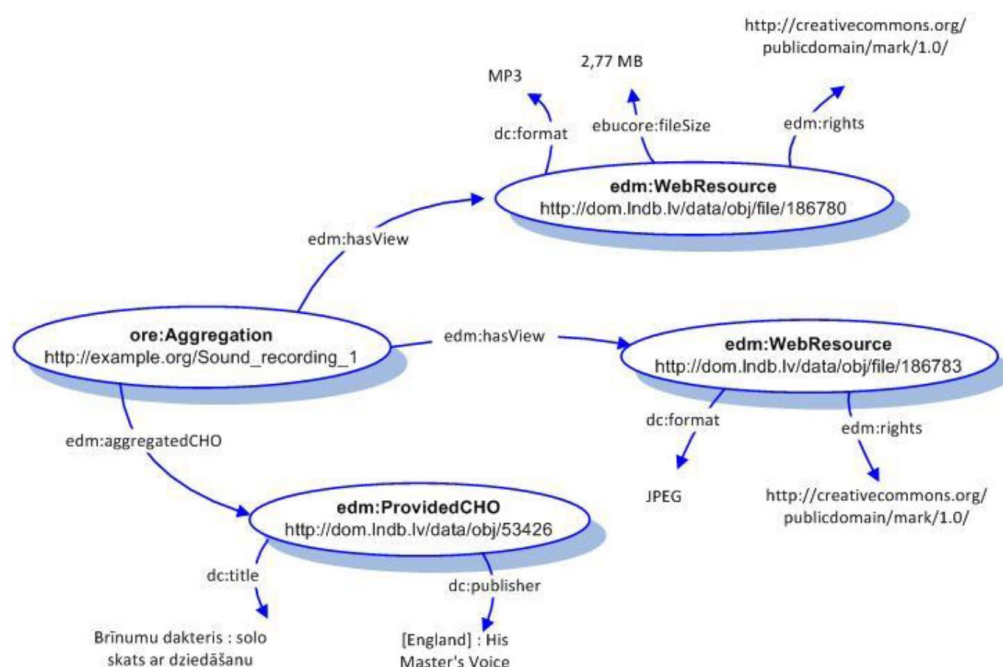


Figure 8. Web Resources in Europeana Sound [Europeana]

2.7 Creative Commons



Creative Commons is a global nonprofit organization that enables sharing and reuse of creativity and knowledge through the provision of free legal tools. Creative Commons licenses provide an easy way to manage the copyright terms that attach automatically to all creative material under copyright. Those licenses allow that material to be shared and reused under terms that are flexible and legally sound. Creative Commons offers a core suite of six copyright licenses (Figure 9). The Audio Commons ontology will map relevant entities of Creative Commons to the concepts in the Audio Commons ecosystem. Information about the Creative Commons licensing framework and its implications in the context of Audio Commons is provided in the publically available deliverables D3.1 and D3.2 [Del].




Figure 9. Core suite of six copyright licences [CreativeCommons]

As previously mentioned when discussing the Media Value Chain ontology (Section 2.5), the Audio Commons ecosystem will include more finely grained representation of actions that can be executed over the Audio Commons objects. The set of these actions should be able to represent the actions that will trigger licence change of an Audio Commons object. Creative Commons describes six types of copyright licences using natural language (Figure 15). Additionally, Creative Commons identifies actions like tweaking, remixing, downloading, etc., which will need to be represented in the Audio Commons ontology.



3 Audio Commons Ontology

This section describes Audio Commons ontology. The ontology (as shown on Figure 1) builds upon existing audio and audio related (or relevant) ontologies. Building an ontology that would encompass the whole audio domain (and all other domains that are connected with the audio domain) in all its complexity would be a very significant task that is beyond the scope of this work. The Audio Commons ontology will, for this reason, be an implementation driven ontology that is evaluated and evolved in use. This means that the Audio Commons ontology will be growing depending on the demand for new services in the Audio Commons ecosystem. Nevertheless, this first version of the ontology will set the path for all future iterations of the Audio Commons ontology because it defines high level concepts and describe how those high level concepts should be interconnected.

The most important concepts in the Audio Commons ontology are the concepts of Action, Agent and Entity. With these three concepts (as defined in PROV ontology) its is possible to describe almost any type of workflow in the ecosystem. These three concepts will have subsets of more specific classes that will describe the context (for example, if service is dealing with licensing information, the Audio Commons ontology will provide a set of actions and agents that are relevant for that context). Entities are object of actions (things we may conduct tasks on). In the Audio Commons ecosystem they will mostly be audio files. The meaning of the word *context* here is quite important because the Audio Commons ontology is not strictly a domain but an interdomain ontology. Audio Commons services will deal with different domains (and different contexts), like the audio domain (sound model), the copyright domain (licensing process and intellectual property) and the physical domain (user profiles, audio files, etc.)

As mentioned before in this deliverable, the Audio Commons ontology is primarily intended to be used as a guideline for a semantic mediator that will orchestrate different services in the ecosystem. The Audio Commons ontology draft mainly defines a data model. The ways in which the ontology will actually be used, will be defined in future deliverables more focused on service orchestration process. The Audio Commons ontology as a data model, describes different groups of metadata and describes how those different groups are connected. In this deliverable we describe the most important elements of the Audio Commons ontology. Actions and agents are described in sections 3.1 and 3.2 respectively. In section 3.3 we introduce the concept of Web resources as prominent entities of the ontology. This section deals mostly with metadata that is used to describe properties of audio files. Audio files themselves are represented as Audio Commons objects, and the most important types of audio files are described in section 3.5. Since presumably a significant number of audio files published in the Audio Commons ecosystem will be music samples or music pieces, we introduce in section 3.4 the concept of musical expression (reused from the Music Ontology) and what kind of properties can be attached to the musical expression class.



3.1 Actions

An important part of the Audio Commons ontology is based on the concept of actions. But what is an action? Both the Provenance ontology and the Media Value Chain ontology provide definitions for such concept:

- The Provenance ontology uses the label *Activity* (*prov:Activity*) and defines the concept as follows: An activity is something that occurs over a period of time and acts upon or with entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating entities. The relationship between Audio Commons *Actions* and the relevant PROV concepts is shown in Figure 11.
- The Media Value Chain ontology defines action in a different context as follows: Actions are the process of doing something over IP Entities. Actions can be applied over the IP Entities themselves or over their representations including both analogue and digital. The result of some Actions may imply the creation of another IP Entity (for example, a *MakeAdaptation* Action generates a new *IPEntity* of the kind *Adaptation*) while others do not as in the case of *Render*. Each action can be exercised over only one kind of *IPEntity*, and it can only be performed by one *Role*. The User who performs an *Action* has to be the rights owner of the *IPEntity*. Further *Permissions* may be required by virtue of the *Creator's* moral rights (Figure 10).




Figure 10. Users, IP Entities and Actions [MVC0]

From the definition of *Activity* in the PROV ontology and the definition of *Action* in the MVC ontology it can be concluded that *Action* in Media Value Chain ontology is just a specialisation of the *prov:Activity* class which provides a very broad and general definition.

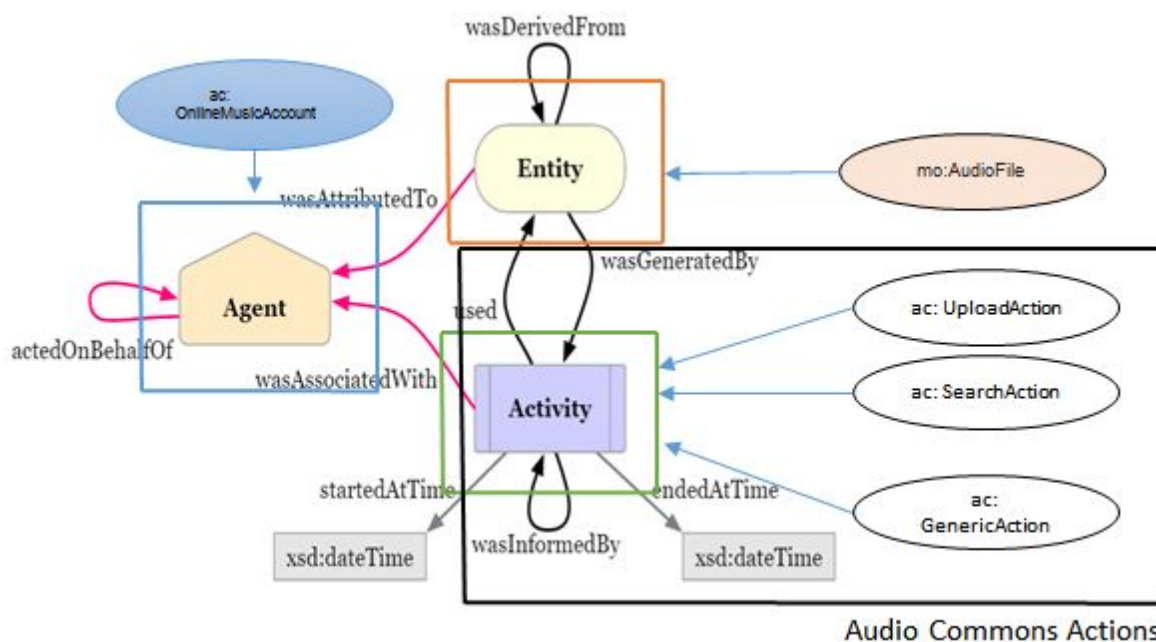


Figure 11. Audio Commons *Actions* in relation to the PROV *Activity* class

Since the Audio Commons project strive to connect different tasks carried out in a typical workflow of people working with audio resources and provide the context for different inputs and outputs of that workflow, recognizing the relevant tasks or actions is very important. The Audio Commons ontology deals with different domains, creative, legal and physical. Actions in the Audio Commons ontology should replicate that distinction between quite separate domains. This is a reason for adopting the grouping of actions in the Audio Commons ontology. Figure 13 is showing the initial representation of actions in the Audio Commons ontology, including the three types of actions that will be recognized in the Audio Commons ecosystem:

- Production actions:** these actions are carried out by roles (agents) working in music composition or audio production process. A large number of these actions were already defined in the Music ontology (where they are classified as events: Composition, Performance, etc.). Some of the actions should be mentioned in Music ontology are not classified as actions or events but just as properties. The Audio Commons ontology will define those actions like *EncodingAction* and *TranscodingAction* (represented on Figure 12).

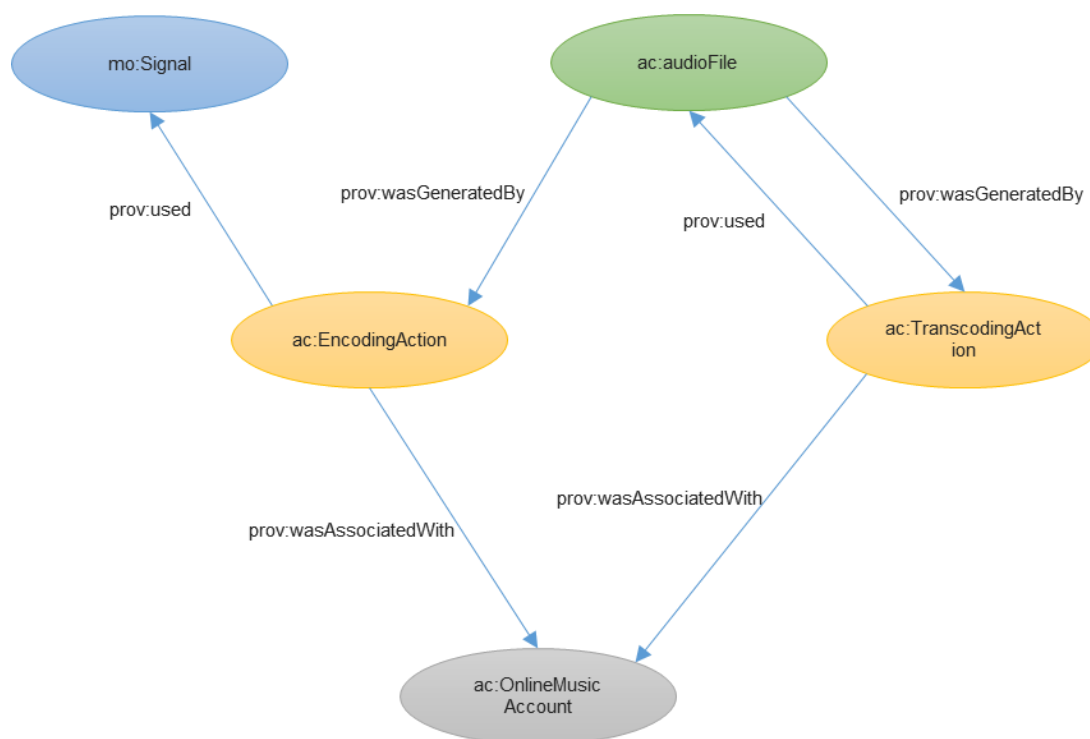


Figure 12. Encoding action and Transcoding actions

- Intellectual property actions (IPActions):** *IPActions* should be distinguished from the production actions because they cover different domains. IP actions describe only those activities that are somehow triggering the change of copyright information (some of activities are shown on Figure 14). The MVC ontology defines a limited set of actions. This set may be extended in the future depending of the requirements arising in the Audio Commons ecosystem. For now, it is agreed that Audio Commons will use the Creative Commons set of metadata to represent copyright information. Figure 15 shows which actions are actually defined in Creative Commons (actions that are relevant to specific level of rights are underlined).

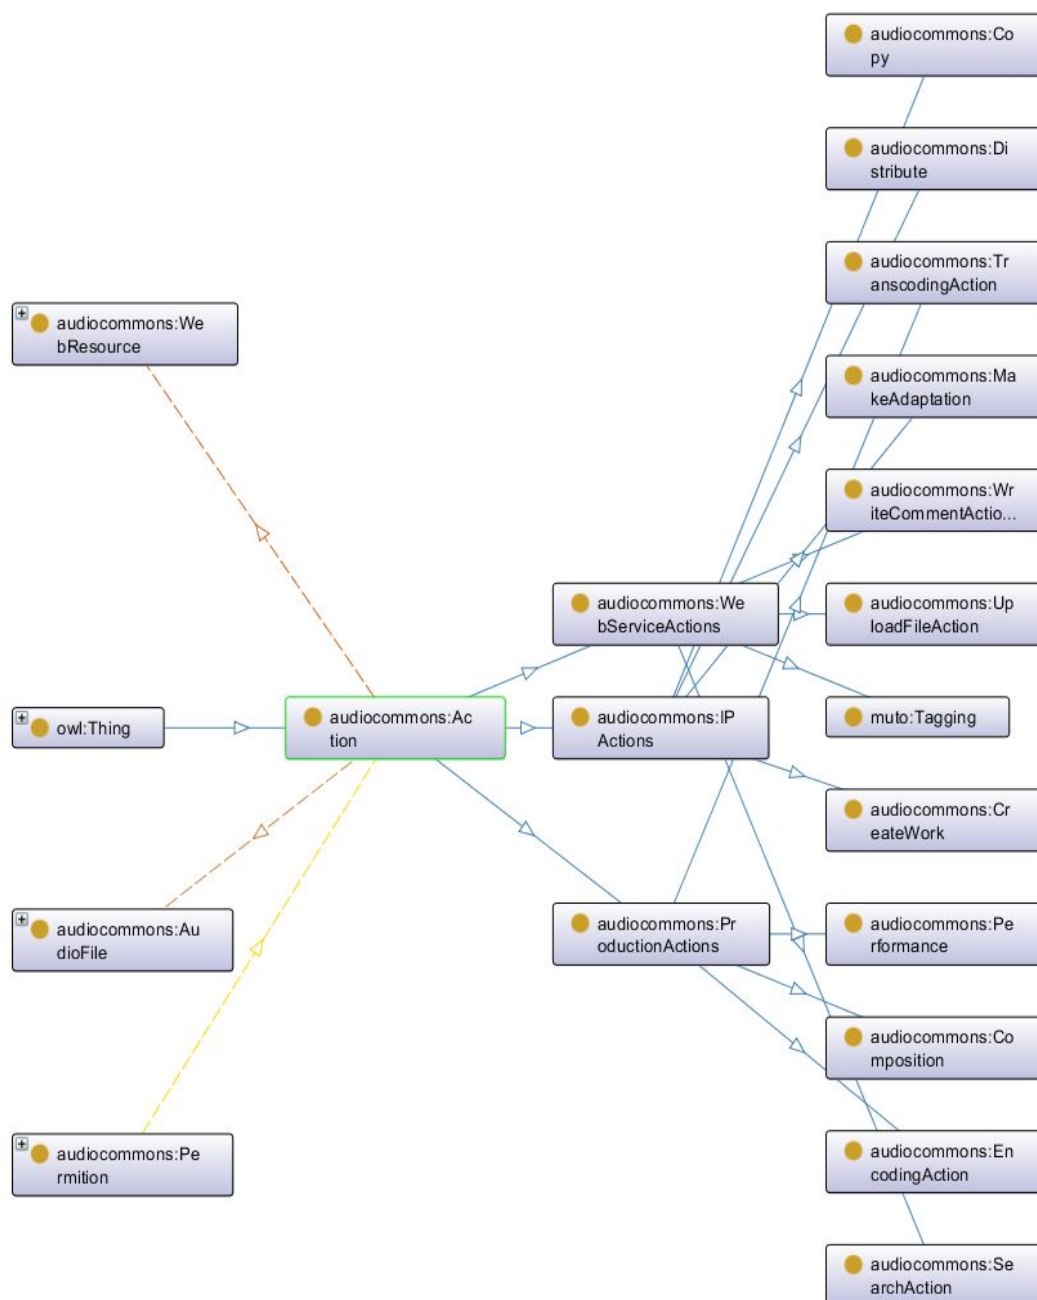


Figure 13. Actions in the Audio Commons ontology



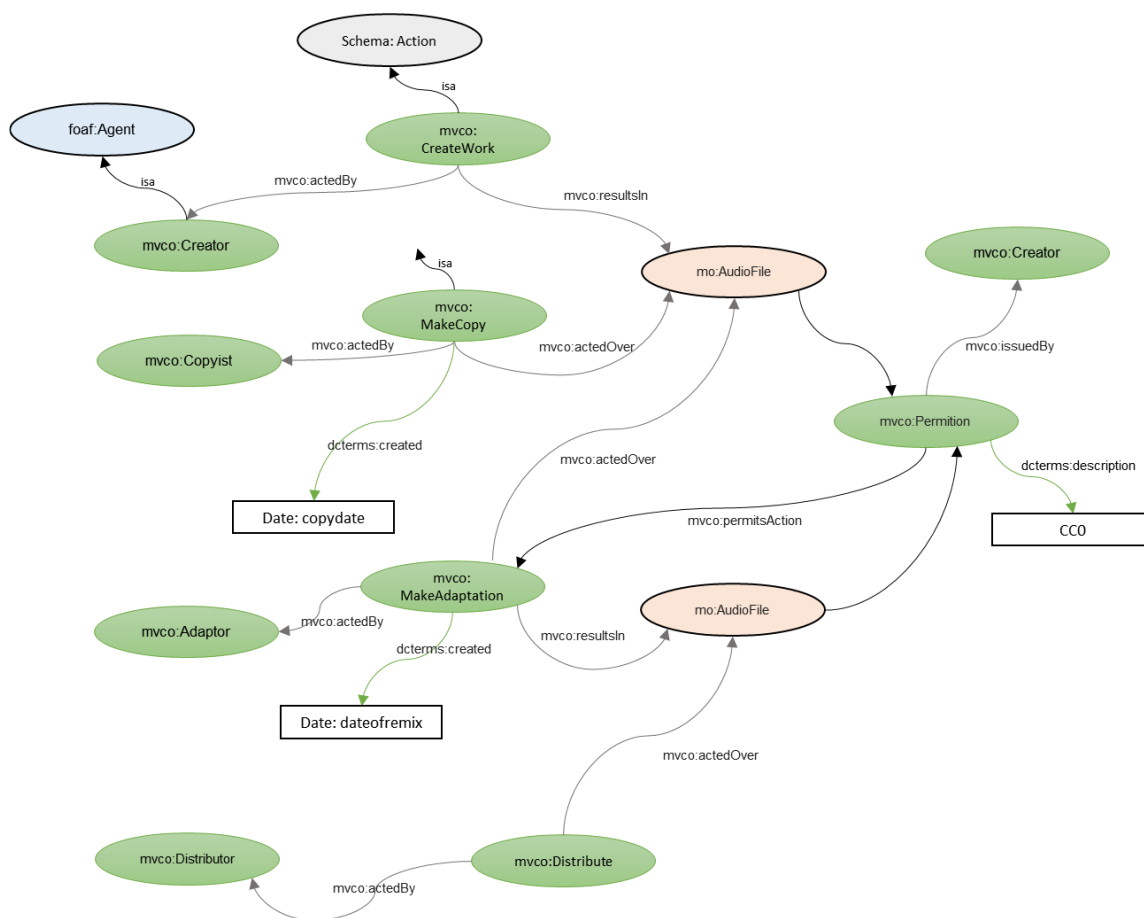


Figure 14. Modelling IP actions





AudioCommons

D2.2 Draft Ontology Specification

[CC0 Public Domain Dedication](#) - By using CC0, you waive all copyright and related rights to a work to the extent possible under the law.

[Attribution](#) - This license lets others distribute, remix, tweak, and build upon your work, even commercially, as long as they credit you for the original creation. This is the most accommodating of licenses offered. Recommended for maximum dissemination and use of licensed materials.

[Attribution-ShareAlike](#) - This license lets others remix, tweak, and build upon your work even for commercial purposes, as long as they credit you and license their new creations under the identical terms. All new works based on yours will carry the same license, so any derivatives will also allow commercial use.

[Attribution-NoDerivs](#) - This license allows for redistribution, commercial and non-commercial, as long as it is passed along unchanged and in whole, with credit to you.

[Attribution-NonCommercial](#) - This license lets others remix, tweak, and build upon your work non-commercially, and although their new works must also acknowledge you and be non-commercial, they don't have to license their derivative works on the same terms.

[Attribution-NonCommercial-ShareAlike](#) - This license lets others remix, tweak, and build upon your work non-commercially, as long as they credit you and license their new creations under identical terms.

[Attribution-NonCommercial-NoDerivs](#) - This license is the most restrictive of our six licenses, only allowing others to download your works and share them with others as long as they credit you, but they can't change them in any way or use them commercially.

Figure 15. Level of rights in Creative Commons

- **Web Service Actions:** Since Audio Commons will mainly manifest itself as a set of Web services it is important to identify the actions that manipulate objects in the Web domain. One of such actions will be tagging action, an action defined in the MUTO ontology. The list of actions describing this domain can be very substantial and it will be included in further specification deliverables of the Audio Commons ontology. A few obvious actions that will be important to identify in the Audio Commons ecosystem include the Search action and the Upload action. An example of using these actions can be seen in Figure 16.



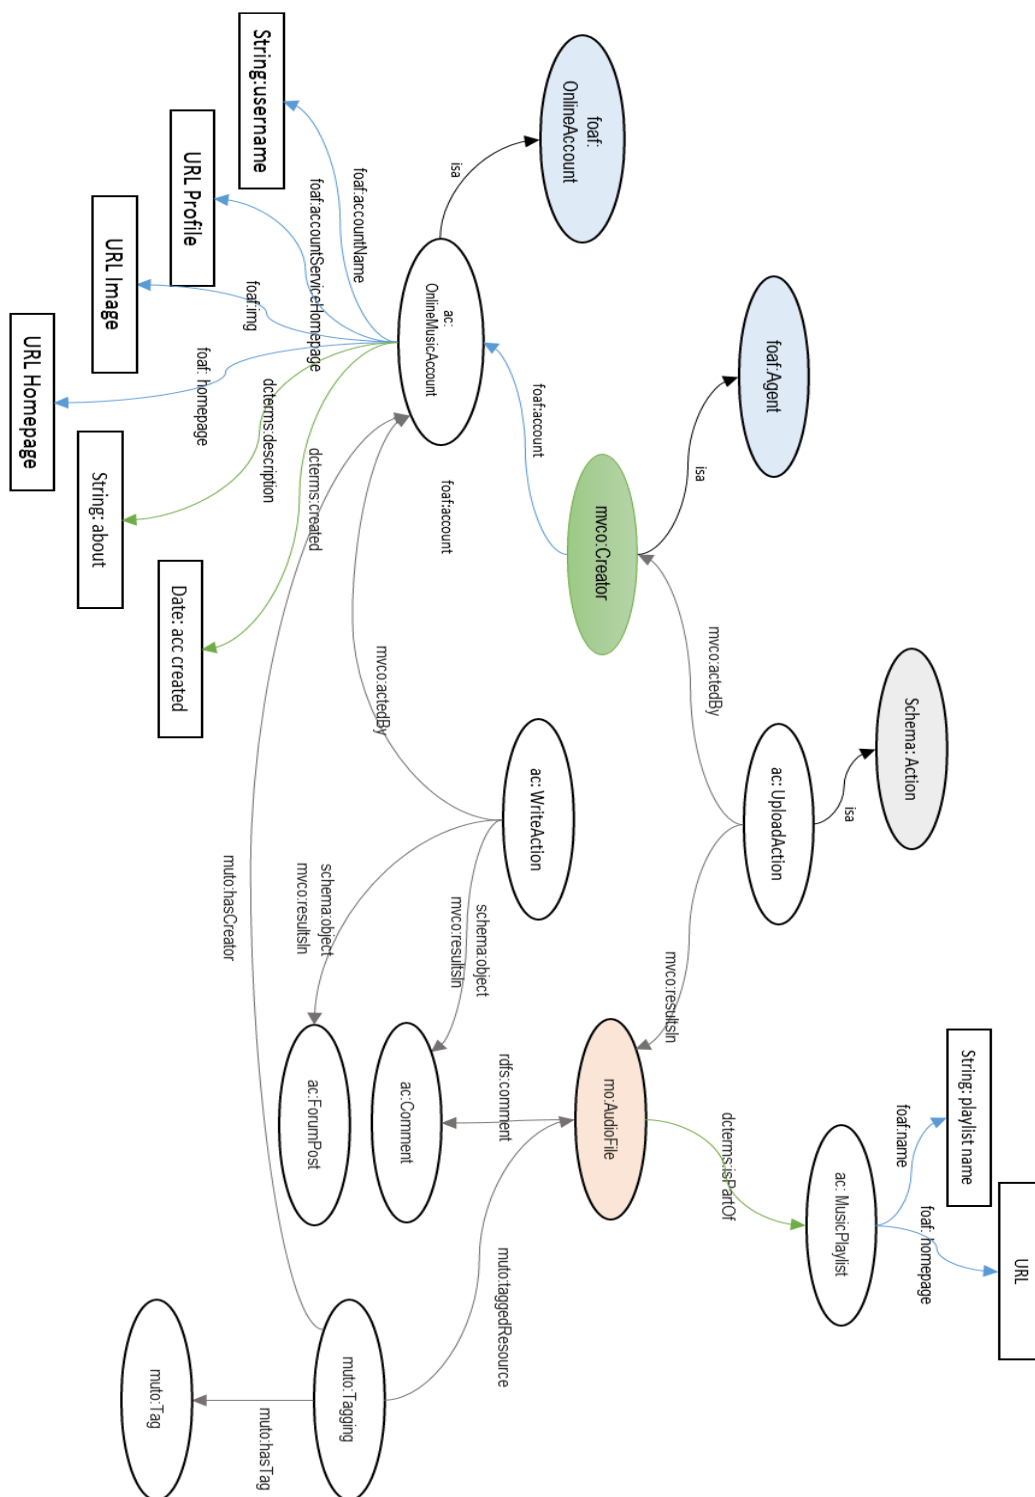


Figure 16. Modelling the Upload and Tagging actions



3.2 Agents

As mentioned before, an agent is something that bears some form of responsibility for an activity taking place, for the existence of an entity, or for another agent's activity. Since the Audio Commons ontology will adopt the PROV ontology representation of agent–activity relationship it is obvious that for each action defined in Audio Commons ontology we need to define who is “in charge” or who is responsible for that action (Figure 17.).

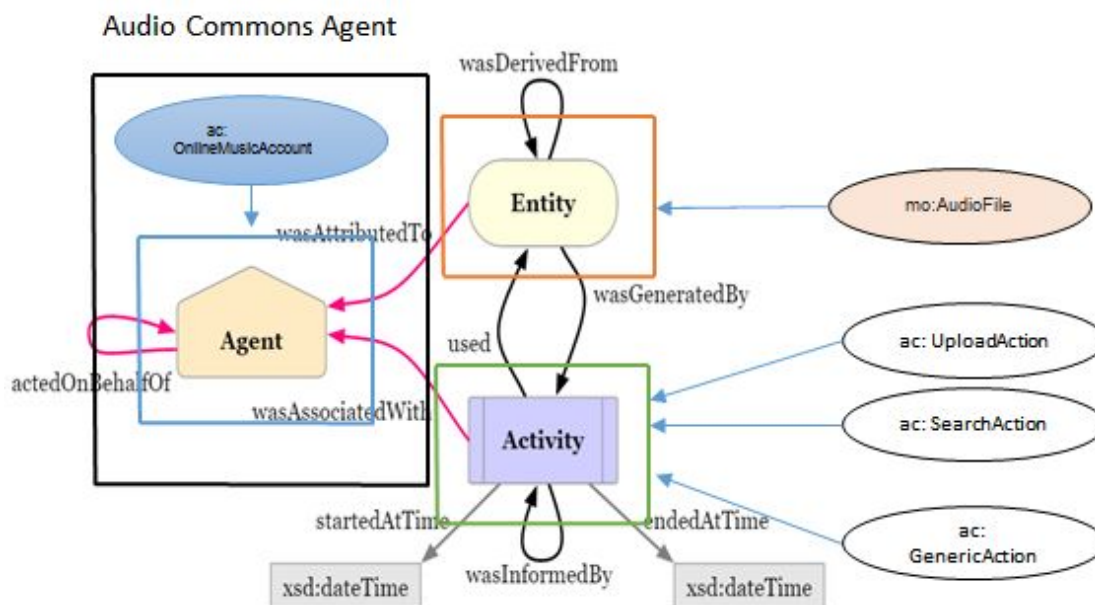


Figure 17. Audio Commons Agent

For example, actions related to copyright issues will have agents defined and labelled depending of the verb that implies the existence of a certain action or task. The action of distribution of an audio file inside the Audio Commons ecosystem should imply the existence of an agent that can be labelled as a Distributor. The class labelled as distributor can then be connected with other higher level classes like *foaf:Agent* or Audio Commons defined class *ac:OnlineMusicAccount*.

The class representing the online music account (*ac:OnlineMusicAccount*):

- can represent an end-user or developer: Implementation issues dictates this distinction (issuing of keys to access the Audio Commons services)
- end-user and developer are defined as agents
- *ac:OnlineMusicAccount* is subclass of *Agent* class and *WebResource* class
- *ac:OnlineMusicAccount* class should define various relationships (Figure 18) that are describing the nature of the account (requirements for these relationships are coming from API vocabulary of Freesound and Jamendo):
 - o temporal information (creation date)
 - o profile information (profile URL, image URL, username, etc.)

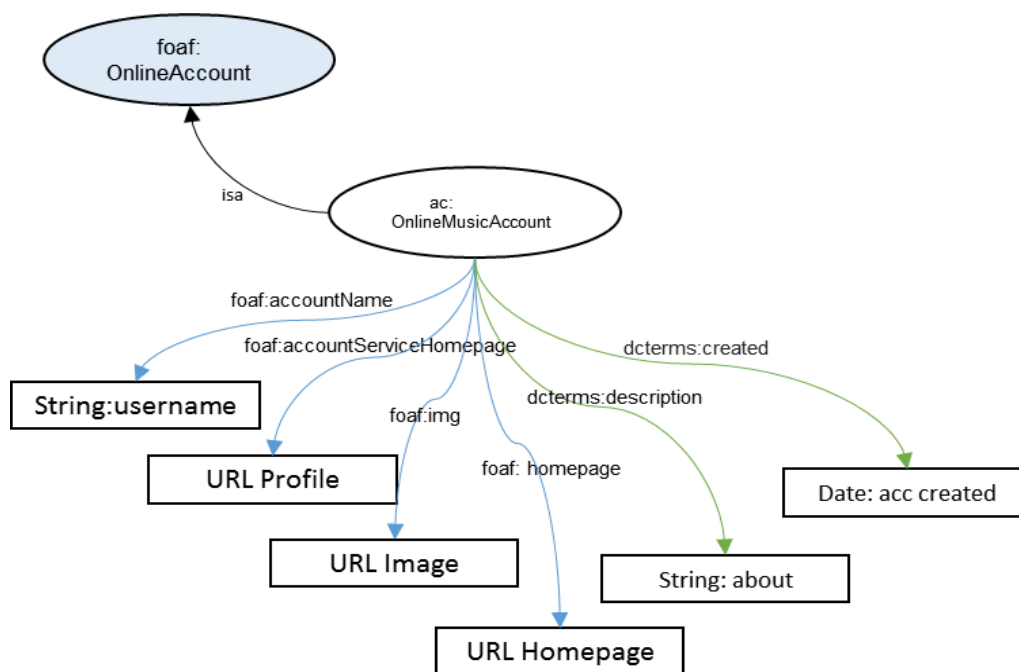


Figure 18. An online music account representation

3.3 Web resources

The Audio Commons ecosystem will provide various services over Web. In line with the PROV ontology model we can say that certain agents will ask for specific actions on specific entities. It is already stated that an agent can be a Web resource if an agent is represented with an online music account (these interconnections are shown of Figure 19.). Entities that are objects of certain Audio Commons related actions will almost certainly be Web resources in the Audio Commons ecosystem.

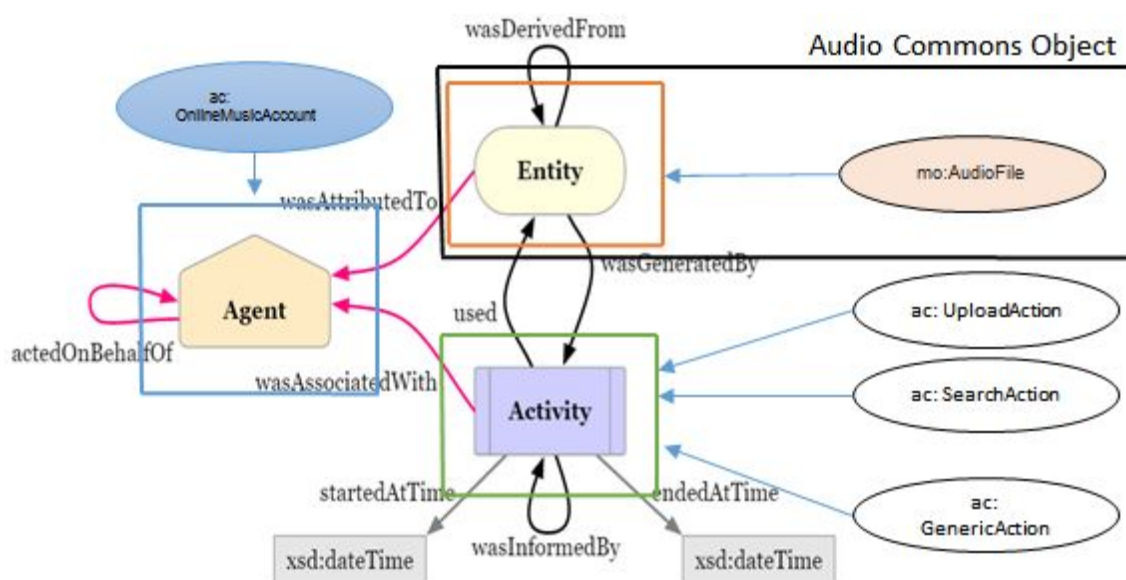


Figure 19. Audio Commons object

The central entity in Audio Commons is the audio file (*mo:AudioFile*) (Figure 20). *AudioFile* is a subclass of Music ontology concept of Medium and it comes to existence as a product of an encoding action. Properties connected with this class are:

- o temporal information (creation date, etc.)
- o profile information (profile URL, image URL, labelname, etc.)
- o technical information (duration, file size, etc.)

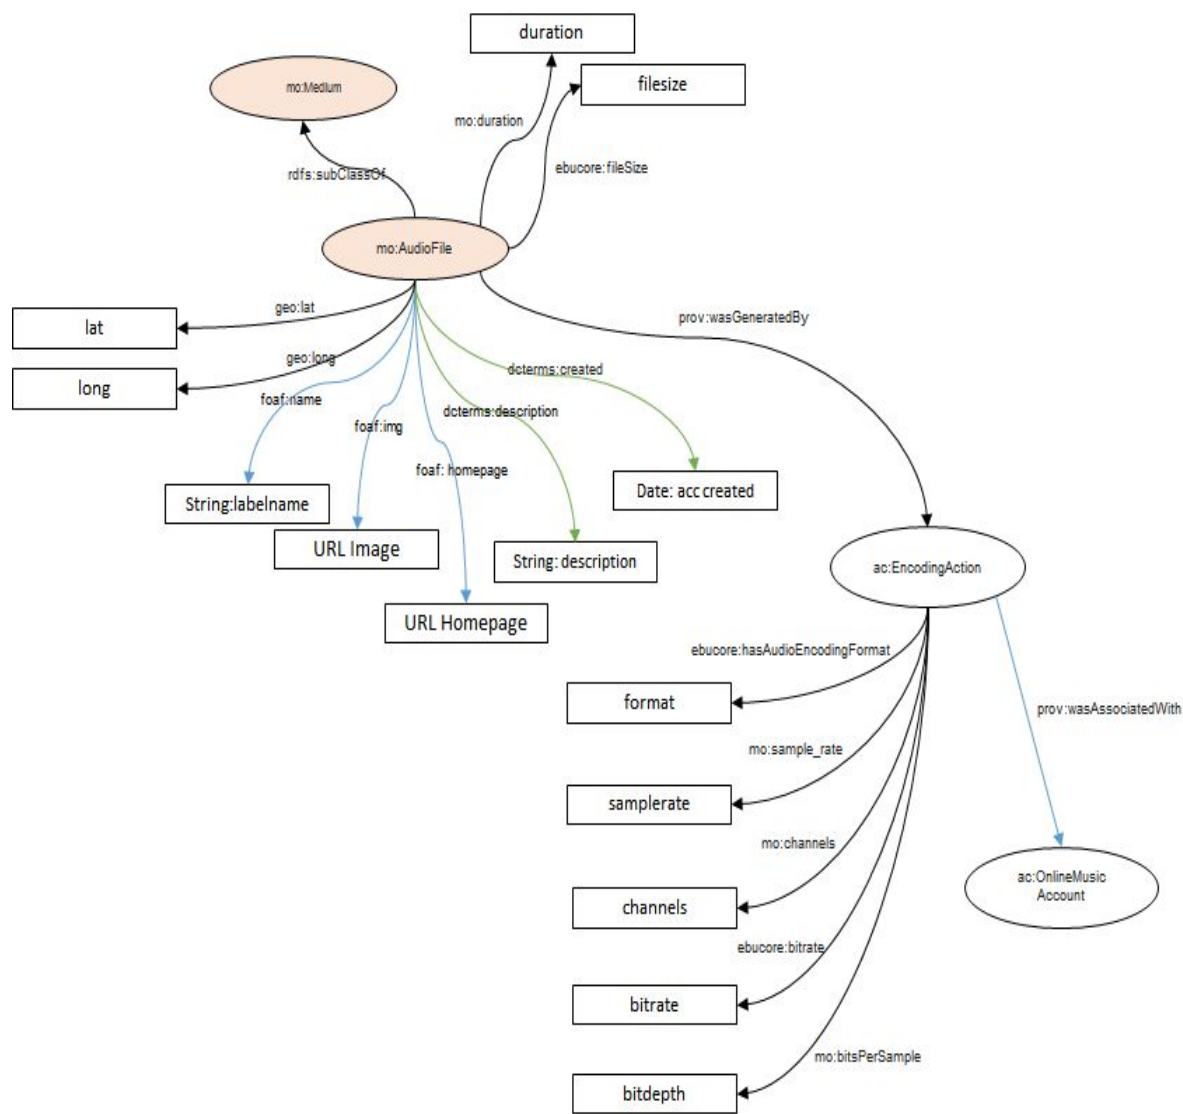


Figure 20. Audio file and encoding features



3.4 Musical expression

The Music ontology defines musical expression as the intellectual or artistic realization of a work in the form of alphanumeric, musical, or choreographic notation, sound, etc., or any combination of such forms. The following example is provided:

- Work #1 Franz Schubert's Trout quintet
 - Expression #1 the composer's score
 - Expression #2 sound issued from the performance by the Amadeus Quartet and Hephzibah Menuhin on piano
 - Expression #3 sound issued from the performance by the Cleveland Quartet and Yo-Yo Ma on the cello

The Music ontology treats the concept of sound as a subclass of the concept of musical expression (together with other subclasses as lyrics, score, signal, etc.). The Music ontology stops at the definition of sound and does not go any further into describing the attributes of the sound or types of the sound. The Audio Commons will define a general taxonomy to classify sounds in a similar fashion as the taxonomy defined in the Sound Ontology [SoundOnt]. As an example, the Sound Ontology distinguishes between different types of sounds:

- natural sound
- machinery sound
- speech
- musical sound:
 - instrumental sound
 - orchestra
 - single tone

For the purposes of the first draft, The Audio Commons ontology will reuse the class `mo:sound` that is a subclass of `mo:MusicalExpression`. The Audio Commons ontology will introduce the property labelled `ac:hasSoundAttribute` that will connect the sound concept with various sound attributes defined in sound ontology under `SoundAttribute`:

- harmonic structure
- chord
- harmony
- rhythm
- timbre
- tonality

The Timbre class will represent the connection point to the timbral terms taxonomy defined in deliverable D5.1 Profiling timbral metadata [Del].

Figure 21. shows a complete diagram of the `MusicalExpression` class as currently represented in the draft ontology.





3.5 Audio Commons Objects

Audio Commons objects define entities that will identify the types of audio files (content types) in the Audio Commons ecosystem. Audio Commons objects in the Audio Commons ecosystem are always audio files. This category is highly debatable and it will require the consensus from all the parties in the ecosystem. Nevertheless, an initial recommendation for the classification of Audio Commons objects is as follows:

- Audio Commons object
 - Musical work
 - Song
 - Musical piece
 - Musical building block
 - Musical sample
 - Instrument sample
 - Sound effect
 - Soundscape
 - Speech
 - edm:Collection

Basically, the Audio Commons ontology will distinguish between music-related objects (Musical work and Musical building block), and non-music objects such as Sound effect(s), Soundscape and Speech recordings. edm:Collection is also initially included as a meta object that aggregates other audio files. This split coincides with the types of data that initial Audio Commons content providers (Freesound and Jamendo) are providing. Figure 22 shows a diagram of Audio Commons objects and related classes. Different types of objects are expected to have different sound attributes and metadata.



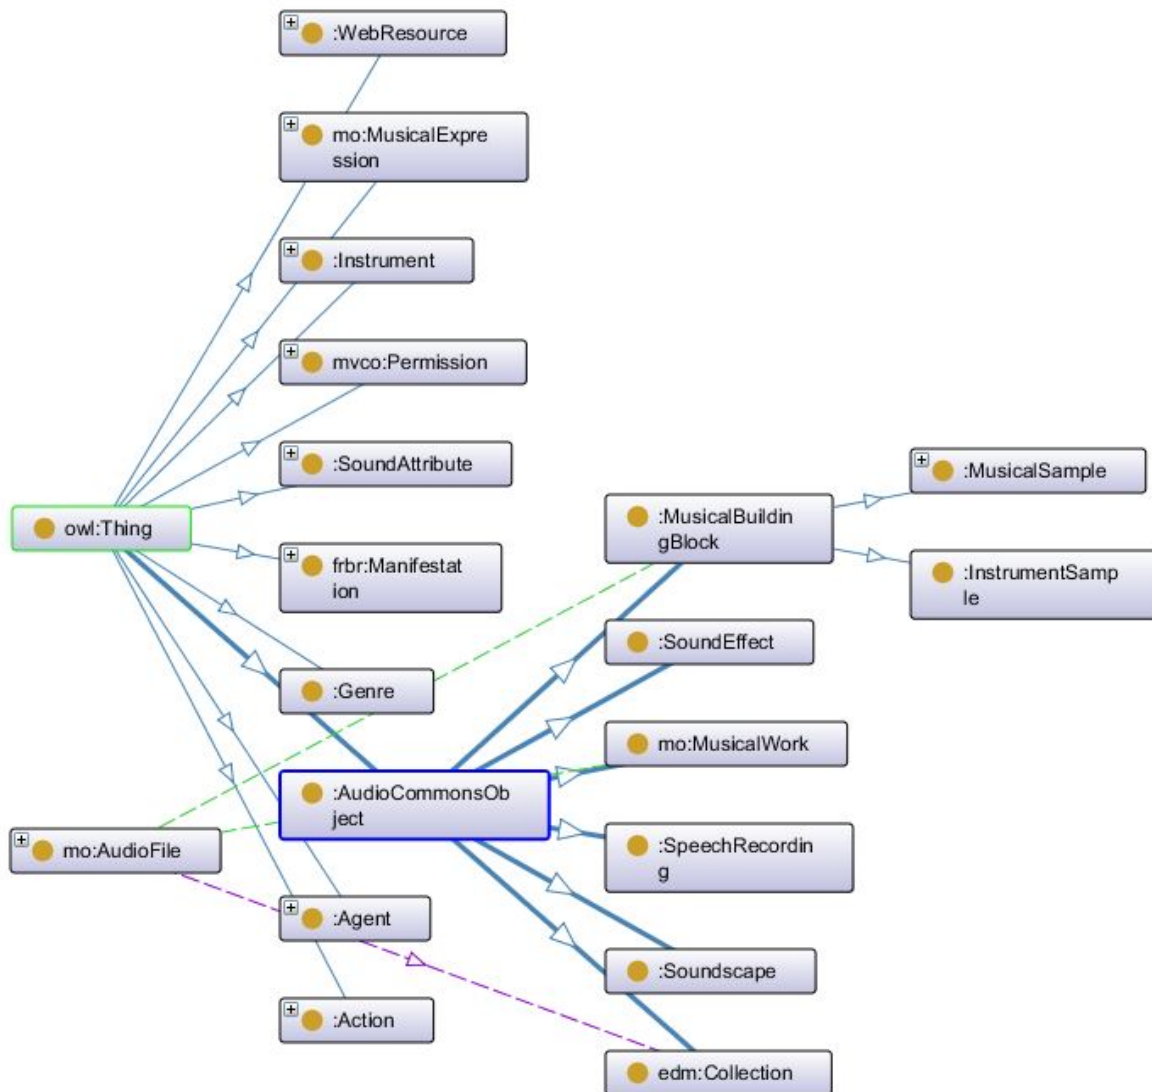


Figure 22. Audio Commons objects in ontology





4 Conclusion

In this deliverable we presented the initial draft of the Audio Commons ontology. The Audio Commons ontology will gather different audio and music related ontologies and create relationships and classes that will bring them together, to cover all sorts of sounds that may exist in the ecosystem. One example is the Music ontology, which provides a comprehensive yet easy to use and easily extended domain specific knowledge representation for describing music related information, that is being merged with terms from the Sound ontology (an ontology that is composed of sound classes, definitions of individual sound attributes, and their relationships).

The Audio Commons ontology focuses on the concept of action (where an action is an activity or something that occurs over a period of time and acts upon or with entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating entities). The concept of an action is then modelled in line with the Provenance ontology (an ontology that provides a set of classes, properties, and restrictions that can be used to represent and interchange provenance information generated in different systems and under different contexts) where certain agents/roles are involved in those certain action that are producing different entities.

Ontologies are shared vocabularies and they should grow and evolve through time, depending of the artefacts and services that will be involved in Audio Commons ecosystem. This deliverable presents the core of the Audio Commons ontology. Subsequent versions of the ontology will contain new classes and relationships or redefine existing ones depending on the project requirements.





5 References

[CreativeCommons] Creative Commons - <https://creativecommons.org/>

[Del] Audio Commons deliverables: <http://www.audiocommons.org/materials/>

[Europeana] Europeana Data Model - <http://pro.europeana.eu/share-your-data/data-guidelines/edm-documentation>

[Fazekas] G. Fazekas. "Knowledge representation issues in audio related metadata model design", 133rd Audio Engineering Society Convention 2012, AES 2012

[FRBR] B. Tillett. "What is FRBR? A conceptual model for the bibliographic universe", The Australian Library Journal Vol. 54 , Iss. 1,2005

[MediaValueChain] Media Value Chain ontology - <http://dmag.ac.upc.edu/ontologies/mvco/>

[MUTO] Modular Unified Tagging Ontology - <http://muto.socialtagging.org/core/v1.html>

[ProvOnt] Provenance ontology - <https://www.w3.org/TR/prov-o/>

[Raimond] Y. Raimond, S. Abdallah, M. Sandler, F. Giasson. [The Music Ontology](#), Proceedings of the International Conference on Music Information Retrieval (ISMIR), 2007

[SoundOnt] T. Nakatani and H.G. Okuno, "Sound ontology for computational auditory scene analysis", AAAI-98, pp. 1004-1010, 1998.





APPENDIX 1: Audio Commons Ontology excerpt

This XML file does not appear to have any style information associated with it. The document tree is shown below.

```

<rdf:RDF
  xmlns="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#"
  xmlns:schema="http://schema.org/"
  xmlns:mo="http://purl.org/ontology/mo/"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:mvco="http://purl.oclc.org/NET/mvco.owl#"
  xmlns:fabio="http://purl.org/spar/fabio/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:frbr="http://purl.org/vocab/frbr/core#"
  xmlns:audiocommons="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#"
  xmlns:edm="http://www.europeana.eu/schemas/edm/"
  xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#"
  xmlns:muto="http://purl.org/muto/core#"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:ebucore="http://www.ebu.ch/metadata/ontologies/ebucore/ebucore#"
  xmlns:dcterms="http://purl.org/dc/terms/"
  xmlns:foaf="http://xmlns.com/foaf/0.1/"
  xml:base="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons">
  <owl:Ontology rdf:about="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons"/>
  <!--
  //
  // Object Properties
  //
  //
  -->
  <!-- http://purl.oclc.org/NET/mvco.owl#actedBy -->
  <owl:ObjectProperty rdf:about="http://purl.oclc.org/NET/mvco.owl#actedBy">
    <rdfs:domain
      rdf:resource="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#IPActions"/>
    <rdfs:range rdf:resource="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#Agent"/>
  </owl:ObjectProperty>
  <!-- http://purl.oclc.org/NET/mvco.owl#actedOver -->
  <owl:ObjectProperty rdf:about="http://purl.oclc.org/NET/mvco.owl#actedOver">
    <rdfs:domain
      rdf:resource="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#IPActions"/>
    <rdfs:range rdf:resource="http://purl.org/ontology/mo/AudioFile"/>
  </owl:ObjectProperty>
  <!-- http://purl.oclc.org/NET/mvco.owl#issuedBy -->
  <owl:ObjectProperty rdf:about="http://purl.oclc.org/NET/mvco.owl#issuedBy">
    <rdfs:domain rdf:resource="http://purl.oclc.org/NET/mvco.owl#Permission"/>
    <rdfs:range rdf:resource="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#Agent"/>
  </owl:ObjectProperty>
  <!-- http://purl.oclc.org/NET/mvco.owl#permitsAction -->
  <owl:ObjectProperty rdf:about="http://purl.oclc.org/NET/mvco.owl#permitsAction">
    <rdfs:domain rdf:resource="http://purl.oclc.org/NET/mvco.owl#Permission"/>
    <rdfs:range
      rdf:resource="http://www.semanticweb.org/sergio/ontologies/2016/8/audioCommons#IPActions"/>

```

The whole ontology can be found at following link:

<https://drive.google.com/open?id=0B58BaC8vR2RLS1dYTVhFTDNnSlk>





APPENDIX 2: Audio Commons Ontology Schema

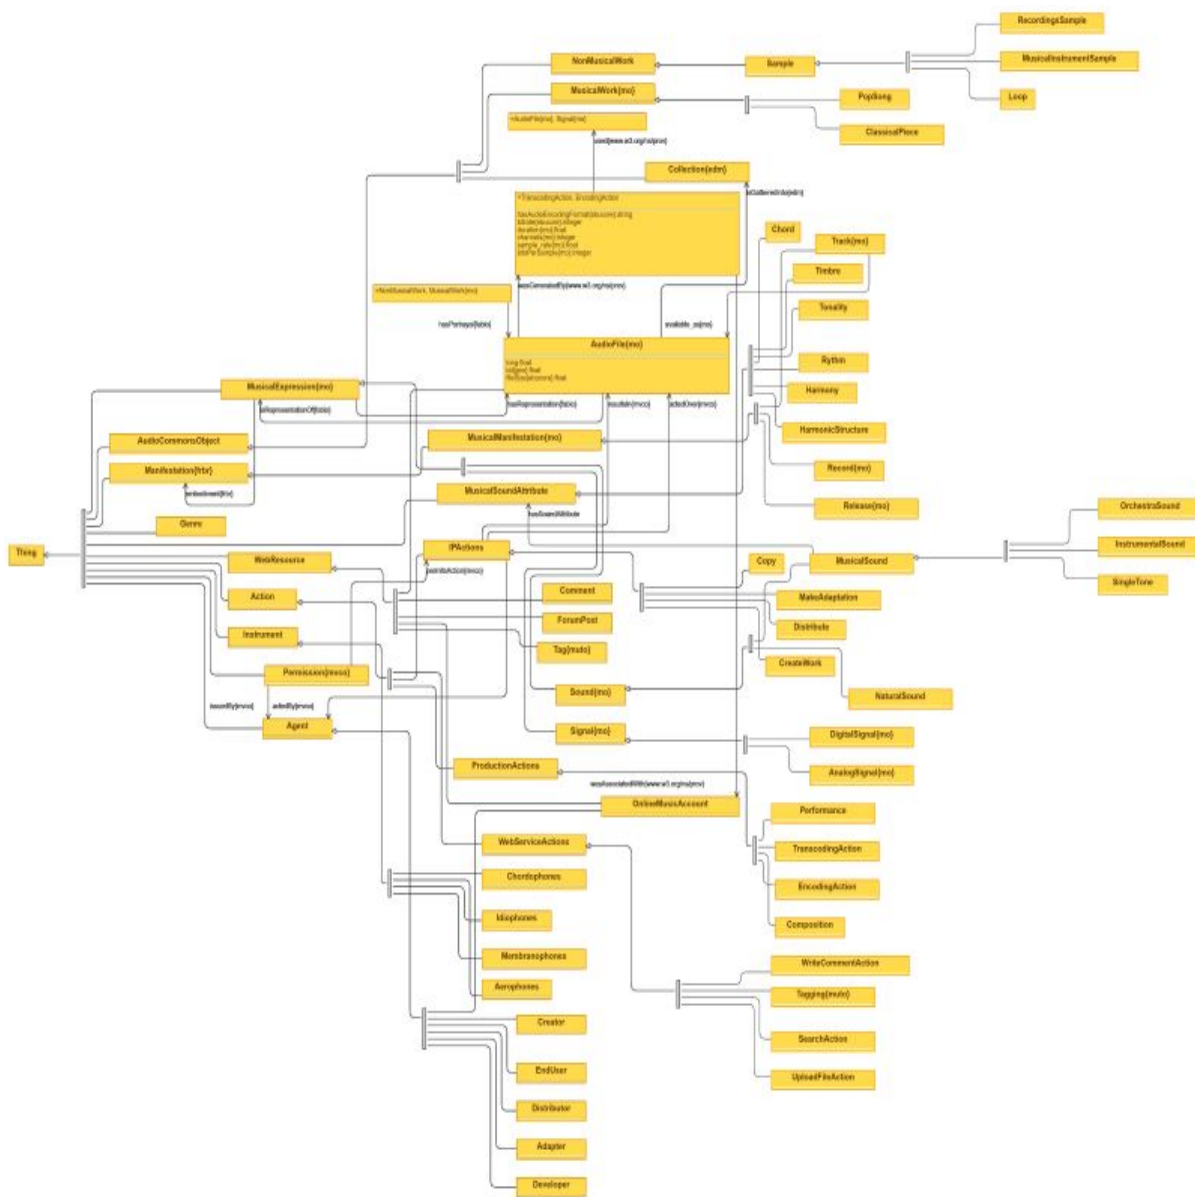


Figure 23. Audio Commons ontology schema
(<https://drive.google.com/open?id=0B58BaC8vR2RLS0pEXy1pT2pOdDg>)

