A Model for Searching Musical Scores by Instrumentation

Michel Beigbeder

A Model for Searching Musical Scores by Instrumentation

Michel Beigbeder

École Nationale Supérieure des Mines de Saint-Étienne
michel.beigbeder@emse.fr

Abstract. We propose here a preliminary study on the definition of a search model that allows to look for musical scores that exactly or approximately match a query where the query defines the exact instrumentation wanted by the user. We define two versions of approximate matchings. In the first one, the ranking is done with a crisp matching of the instruments. In the second one we relax this constraint and we use a similarity between instruments. We present a first experiment and envision future works.

1 Introduction

There are many library or bookshop sites on the Web that propose an access to musical scores. Typically the access is provided either by browsing or searching by the composer, the arranger, the publisher or the title of the piece or any conjunctive combination of these fields. This search facility is efficient if the user is aware of some precise musical piece and is able to provide enough information to retrieve the corresponding score and parts.

Our goal is different and we want to give the possibility to search the pieces that a given ensemble is able to play in terms of the instrumentalists that belong to the ensemble. This would be particularly useful for nonprofessional ensembles or for pupil groups in music schools where the composition of the ensemble is not much under control. So the key point in our work is to match the description of an ensemble against the descriptions of the pieces in terms of the number of parts and their instruments.

2 Related Works

Music information retrieval is a large domain because music has many aspects of interest for the searcher. Downie quotes the following facets: Pitch, Temporal, Harmonic, Timbral, Editorial, Textual, Bibliographic[1]. He defends that the instrumentation[1] information is part of the Timbral facet but that it is sometimes considered as part of the Bibliographic facet with an enumeration of the instruments used in a piece. We adopt here the latter point of view. The Downie's survey paper does not quote any work on retrieval by instrumentation.

[1] The author uses the term *orchestration* in place of *instrumentation*.

The paper by Ferrara et al. is focused on genre classification of musical pieces and it defines four facets to describe them: Ensemble, Rhythm, Harmony, and Melody[2]. The Ensemble facet is the one that is related to the instrumentation of a piece, and it is described in terms of performers rather than in terms of parts in the score. In their semantic point of view, they want to transform a list of instruments in more general and compact information. The trivial example is to transform the list *2 Violins, 1 Viola, and 1 Cello* to *String Quartet*. In terms of retrieval they do not address the problem we tackle with, but they are interested in finding pieces that resemble to a given one and they argue that instrumentation is a useful hint for that purpose.

We do not know any publication that address our problem, so we looked at some Web sites to build a tentative state of the art on that matter. We present two commercial sites and a cooperative one that propose scores. Their functionnalities are representative of the best possibilities that we found on any other Web site.

2.1 Free Sites

Regarding instrumentation there is currently only a browsing facility in the site `http://imslp.org/`. This browsing facility is built over a taxonomy where the top level classes are "Keyboard", "Chamber-Instrumental", "Orchestral", "Vocal" and "Featured Instruments". The second level for "Chamber-Instrumental" is mainly composed of the number of instruments (with or without continuo). The second level of the "Orchestral" class distinguishes combinations of i) "orchestra" or "strings", ii) with (or without) soloist, iii) with (or without) continuo. There are also subclasses concerning "toy instruments". It is very difficult to search pieces for a given ensemble. For instance, there are 181 categories of "Chamber-Instrumental" "For 5 players". Moreover in this site it seems that there is not a clean description of the instrumentation, but some "categories" are assigned to each piece just like categories are assigned to articles in Wikipedia.

2.2 Commercial Sites

The site `http://www.di-arezzo.com/` is the site of a bookshop specialized in selling musical scores. Here one can specify that the scores have to contain parts for at least some instruments[2], and the count of these instruments can be specified. For instance, one can search a piece that needs two flutes, two violins, one viola and one violoncello. At the time of writing, this query returns four pieces. The first one needs two clarinets and two pianos too. The second one needs a keyboard in addition. The third one needs two oboes and two more violins and one more violoncello. The last one needs "13 players" whose list can be found on the page of this item. So we can suppose that none of them is a relevant answer for the given ensemble.

[2] More precisely, it is possible to specify at most four pairs of one instrument with a count. So it is not possible to add a trumpet, for instance, to the subsequent example.

If browsing is used, selecting one instrument lists the scores that needs at least one instance of this instrument. This it not very useful, for instance browsing for Flute returns 13212 items mixed with flute solos, many flute and piano pieces, concertos for flute and orchestra, etc.

The search possibility let us infer that the exact instrumentation is stored in the database underlying this site.

On the site `http://www.ewh.dk/` one can specity quite precisely an intrumentation which is developped as a list of instruments with a count. Moreover there are some shortcuts as for instance, "str4tet" (the short for String Quartet) which is developped as the list: 1 "Violin 1", 1 "Violin 2", 1 "Viola", 1 "Cello". Given this list four matchings are proposed:

- "exact match",
- "approximate match", which returns scores with at most one more instrument, or those where one instrument is not used, or those where one intrument is replaced by another one,
- "any combination", which matches all scores whose instrumentation is included in that of the query,
- "any size match", where any piece whose intrumentation is larger matches the query; this is the behavior of the `di-arezzo.com`'s search.

3 Modelization

In terms of mathematical modelling, the first and most simple approach is to model the instrumentation of a musical piece as a function from the set I of instruments to the set $\mathbb{N}$ of positive integers. We will call *multiplicity* this function. In fact the mathematical notion of *multiset* corresponds to that and we will use this vocabulary. So to describe a *string quartet* the multiplicity function maps *Violin* to 2, *Viola* to 1, *Cello* to 1, and any other instrument to 0. The same simple model is applicable to the ensembles: the composition of an ensemble is a multiset of instruments.

We distinguish two kinds of approximate matching. The first one concerns the number of parts and the number of instrumentists where each player has to play a part written for his instrument. The second one relax the latter constraint and will allow instrumentists to play parts written for another instrument. We will present a model for these two matchings in the following sections.

Cardinality Approximate Matching. The matchings used on the site `//www.ewh.dk/` are based on exact matching of instruments and on relationships between multisets that are analogous to those used in the boolean model of document retrieval. To refine the set theoretic point of view used on the site, we can adopt an usual method used in information retrieval to present the result with the most relevant ones in top of the list[3].

[3] *Relevance* as computed by the information retrieval system, which is not necessarily the human relevance.

To model the "any combination" match, the boolean answer is to return each piece p such that $p \subset e$. The larger is p (with $p \subset e$), the closer to the ensemble. So ranking the pieces with decreasing cardinality of p does make sense as the pieces will be ranked with those that employ the maximum of people of the ensemble ranked first. We can notice that the same ranking is obtained if the pieces are ranked with increasing cardinality of $e \setminus p$, and the latter multiset is equal to $e \Delta p$ as $p \subset e$[4].

Symmetrically, for the "any size match", pieces which verify $p \supset e$ should be ranked by increasing cardinality of p, which is also the increasing cardinality of $e \Delta p$ as $e \subset p$.

If both matchings are merged in a single one, ranking by the increasing cardinality of $e \Delta p$ will rank first 0) pieces which matches exactly the ensemble, then 1a) the pieces where one player don't play 1b) with those where one extra musician is needed, then pieces 2a) with two silent players, 2b) or two extra musicians, 2c) or one silent player and one extra musician, etc. Union of cases 0, 1a, 1b, 2c corresponds to the "approximate match" of site `www.ewh.dk`.

Instrument Approximate Matching. The second type of approximation is concerned when a part written for an instrument is played by another one. We introduce operators that can be used in a query, so a user might explicitly state her preferences.

We introduce now the OR operator as in the following example: `flute OR oboe`. There are at least three usages for this operator. The first one is for describing an ensemble where a player is able to play different instruments: for instance someone who can play the flute and the oboe. The second usage is a decision of the user for relaxing his query, probably because he did not find enough answers with a previous query. The last one is when the user knows that in a database two instruments are in fact the same, for instance `flute` and `flauto`. Such an operator would be trivially inserted in the previous modelling. But it also could be seen as a degenerated case of a similarity function in the instrument space, where the user explicitely says that the similarity between the oboe and the flute is one, and it is implied that all other similarities that are not explicitly mentioned in the query should be set to zero.

Finally, our most precise model considers as before a set of instruments I, and defines both ensembles and pieces as multisets of instruments. We now add a structure in the set of instrument with a symmetric similarity function $\mathrm{sim} : I \times I \to [0,1]$ with the condition $(\forall x \in I)(\mathrm{sim}(x,x) = 1)$. This similarity measures to what extent an instrument can replace another one. This function could be either user defined or system defined. We also introduce a pseudo-instrument ϵ which does not belong to I, a kind of *empty* instrument. We also extend the similarity function with $\mathrm{sim}(i, \epsilon) = 0$ for each $i \in I$. This empty instrument will be useful in the mathematical formulation at the end of this section.

[4] The inclusion ($\subset$), the difference between sets ($\setminus$), the symmetric difference (Δ), and the cardinality should be understood in the multiset meaning.

Before giving the full formula for the similarity function between an ensemble e and a piece p, we will first present a formal example. Let us consider:

- an ensemble E with two instruments x and y, each with a multiplicity of 1;
- a similarity function with $\mathrm{sim}(x, a) = x_a$ and $\mathrm{sim}(y, a) = y_a$; $\mathrm{sim}(x, b) = x_b$ and $\mathrm{sim}(y, b) = y_b$;
- a piece P_1 with two instruments x and a;
- a piece P_2 with two instruments y and a;
- a piece P_3 with one instrument x;
- a piece P_4 with two instruments a and a.
- a piece P_5 with two instruments a and b.

The question is how to rank these five pieces for the ensemble E. With P_1, x matches exactly and y matches a with value y_a, so: $score(P_1, E) = 1 + y_a$. With P_1, y matches exactly and x matches a with value x_a, so: $score(P_2, E) = 1 + x_a$. With P_3, only x matches (and we can say that y matches ϵ with a null similarity), so: $score(P_3, E) = 1$. Finally, with P_4, x matches a with value x_a and y matches a with value y_a, so: $score(P_4, E) = x_a + y_a$. Thus in the final ranking, we will have P_1 and then P_2 or the reverse depending on the relative values of x_a and y_a. And then either P_3 or P_4 depending on the relative value of $x_a + y_a$ to 1. To compute the score of piece P_5 we can either affect part a to player x (and thus part b to player y) or the reverse *i.e.* part b to player x (and thus part a to player y). In the former case, the score would be $x_a + y_b$, and in the latter case the score would be $x_b + y_a$, so we can consider that the score of piece P_5 is the best score: $max(x_a + y_b, x_b + y_a)$.

We will now formalize for the general case. We have to consider every pair of one instrument in the ensemble with one instrument in the piece. To formalize that we artificially consider any numbering of the elements of the two multisets p and e. These numberings must have the same length. To do that the smallest multiset is filled with ϵ so as to have the same length as the largest one. After that their common length is $n = \max(|e|, |p|)$. For instance for $E = \{x, x, y\}$ and $P = \{a, b\}$, we have $n = 3$ and we can consider $e_1 = x, e_2 = x, e_3 = y$ and $p_1 = a, p_2 = b, p_3 = \epsilon$, but any other numbering of the elements would give the same result. Given these numberings, we define the score of the piece p for the ensemble e with:

$$score(p, e) = \max_{\sigma \in \mathcal{S}_n} \sum_{i=1}^{n} \mathrm{sim}(e_i, p_{\sigma(i)})$$

where $\mathcal{S}_n$ is the set of all permutations of order n. With the same example we have to consider all the following affectations between parts and instrumentists: (x,a)(x,b)(y,ϵ), (x,a)(x,ϵ)(y,b), (x,b)(x,a)(y,ϵ), (x,b)(x,ϵ)(y,a), (x,ϵ)(x,a)(y,b), (x,ϵ)(x,b)(y,a).

4 Prototype

We have implemented a first version of the model with ranking on the cardinality of $p \Delta e$ as described in Sec. 3, *i.e.* without the generalization of approximation with a similarity function between instruments. For the collection of pieces we used the site http://www.baroquemusic.it/. This site is organized with a browsing by composer. We downloaded the pages of each composer and with simple regular expressions matching the HTML code we were able to extract the name of the pieces and their parts. The site provides 247 pieces. For instance, the very first piece lists the following five parts: "Flauto traverso", "Violino primo", "Violino secondo", "Viola", and "Cembalo". But sometimes the parts were not for a single instrument, for instance some parts (twenty exactly) are for "Violino I e II", and there are many other combination of instruments in the part names. So the main difficulty was to build the instrument sets for the pieces. We used ad'hoc rules to unify many variations of "Violino", "Violini", "Violino primo", "Violini ripieno", "Violons", etc. to a single instrument "Violino"; And, for instance, "Violino I e II" to a multiset with two "Violino".

The data extracted from the Web site were stored in a simple text file. A file containing the ad'hoc rules for converting part names to instrument multisets was manually built. When launching the query process, it first loads the rules and the list of pieces and builds the multisets, this is humanly instantaneous. Query processing considers each piece and compute the cardinality of the symmetric difference between the query and the pieces, and rank the pieces accordingly. With this number of pieces, this is also humanly instantaneous.

5 Future Work

The next step consists in the full implementation of the score formula between a piece and an ensemble for approximate matching. Particularly its complexity in terms of computation with a naive implementation would be in $O(n!)$, which is only usable for very small values of n. We will try a greedy algorithm which will first affect exact matches between an instrument and a part.

We also want to work on a larger database, and we think that imslp.org would be a good choice. Though it perhaps contains too many instrument/part names. An alternative choice would be a commercial site such as di-arezzo.com where the panel of scores is more limited and thus the instrumentations are controlled. We also think that the exact instrumentations are stored in their database and this would eliminate the problem of many names for the same instrument and also the splitting of part names to instrument names.

References

1. Downie, J.S.: Music information retrieval. Annual Review of Information Science and Technology 37, 295–340 (2003)
2. Ferrara, A., Ludovico, L.A., Montanelli, S., Castano, S., Haus, G.: A semantic web ontology for context-based classification and retrieval of music resources. ACM Trans. Multimedia Comput. Commun. Appl. 2, 177–198 (2006)