

Whose future is it?

Philippe Aigrain*

12 October 2004, keynote speech at the ISMIR 2004
Conference, Barcelona

© Philippe Aigrain, 2004. This text and the associated slides can be used under the terms of the Creative Commons Attribution License 2.0 <http://creativecommons.org/licenses/by/2.0/>

Good morning. It is a great honour to deliver this talk to a research community of which I have been a member when we were not sure that it was a domain of its own. It has since then developed in an impressive manner, and today it offers all the signs of a maturing scientific community in a technical field: a yearly conference with growing attendance, some structuring paradigms, progress in modelling its challenges, attraction for bright graduate students, evaluation testbeds and benchmarks. That's exactly when a community generally puts itself on tracks that will lead it to deliver little or much progress for science, little or much usefulness for society. That's when it is worth for such a community to think a while about where it tries to go, to take **[Slide: one step aside]**.

This honour is also a challenge. If I were totally ignorant about information retrieval and content processing of music, it would be much more comfortable. I would share with you some general thoughts about property and commons for information and information technology, about how research funding mechanisms and intellectual property rules influence the targets of research efforts. And from my blissful ignorance, I would leave you with the task of figuring out what it can mean for you and the ISMIR community. But I can't. I know just enough about music information retrieval to feel forced to connect my general thinking with concrete examples in this domain and neighbouring fields. Inevitably, I will say a few stupidities in this process. I hope you won't stop there, at identifying these mistakes, that you will still try to understand the issues I am raising, and only then, decide if they are worth your consideration.

*The author is presently CEO of Sopinspace, Society for Public Information Spaces www.sopinspace.com. His personal page is at <http://www.publicdebate.net/Members/paigrain/>

Little boxes

In 1998, Andrew Odlyzko¹ and myself² had the same outcry of intellectual revolt: [Slide: “Content is not king!"]. We actually meant something different. I meant simply “people are”, while he meant “communication is”. However our message was consistent, and though it was not intended specifically for music, it applies to it. Music is a relationship between people, between gesture and sound, between mind and percept. Music is not contents. Of course, there have been *music boxes* (I mean tools that automatically produce music) for many centuries, and *boxed music* (I mean recorded music on carriers that is sold as a commodity) for a little more than a century. However, we should not let our minds be framed in the *little boxes* that the father of Anthony Seeger (last year’s keynote speaker) has made famous by interpreting Malvina Reynolds’ song [Slide]. Don’t worry if you think my implicit statement is demagogic, there is also a version by Woody Guthrie (the “beatnik version”) where he makes fun of the critics. Anyway, music is not boxed music. Music is so much a relationship that if someone knows to play any of it, he or she can earn some money in the streets of many cities in the world, which is sort of a miracle since at the same time, boxed music is forced upon us in the stores and restaurants of the same cities. Of course this miracle is possible only because and when street musicians don’t pay performance rights.

So what has happened around us in the past 5 years? What has happened to people, the way they access music and use it ? What has happened to people, the way they do research and use it?

The ground beneath our feet

[Slide: Societal peer-to-peer file sharing networks have become a more efficient source of access to published and broadcasted material than libraries and archives]. Raise your hand if you don’t believe me. Of course this is a very general statement. It is absolutely not true for unpublished primary source material. But for material that has been published in any form, if you take a real benchmark, not an hypothetical researcher that would instantly travel to all libraries and archives of the world and have credentials there, but a real person, who suddenly wants access to one performance of “La Ultima Noche”, or a documentary movie that was broadcasted 2 years ago on some channel, the chances he or she will succeed in getting it are much higher on P2P networks. Of course this is only recently developing, and it is under threat. Legal and police threats, but also threats from the private war gangs who inject fakes under the name of P2P warfare. I will come back later to how one can regain the benefits of this wonderful societal mechanism of P2P file sharing, while appeasing some of the pains it is

¹Andrew Odlyzko, Content is not king, First Monday, 6(2), February 2001, revised version on an original 1998 paper.

²Philippe Aigrain, “The Best of Both Worlds: Can Mesia Quality be Combined with Open Contents?”, invited talk at IFTA’s world conference, Vienna, 1998, <http://www.debatpublic.net/Members/paigrain/texts/IFTA.pdf>.

claimed to cause.

[Slide: Remix and hiphop scenes are competing with “official” contemporary music composition as promising users for content-based synchronisation and interaction] I guess Paull Miller’s quote speaks for itself. This is not a new development, but research and creation organisations are slow to react to it. Many people in your community are involved as musicians or technicians in these communities, and some are even deriving inspiration from them, such as the idea of *query by beat boxing*.

[Slide: Creative commons have gone from idea to project to reality in many domains]. Many people in this community are active supporters or would-be supporter of musical commons. However, very few are actively supported by their organisations in building them. Creative commons licenses are moving us out of the false dilemma between respecting music by restricting the relationship that people build can build to it, and making it a commodity that it is freely lootable for the boxed music business.

In WIPO and UNESCO, there is presently a debate about the so-called “protection of traditional knowledge, folklore and genetics resources”. This is supposed to be an answer to the well-grounded complaints by developing countries that powerful businesses are pirating (in the real sense) their traditional culture and natural resources. What WIPO -and at a lesser extent UNESCO- propose to them is to participate into the loot. Enclose them into property, and become little looters of the common goods. Of course the little looters will be locally powerful guys, and they will have to make deals with the big looters for deriving true income, and you can be certain that this income will flow to about anywhere except into the pockets of the poor people who created and cared for these treasures. Creative commons, and fair global redistribution are not an easy path, but at least it is a path in the right direction. The same is true in the developed world. In the era of information, inequality of power just as inequality of wealth is fractal: you find it between countries, between regions in a country, between neighbourhoods in a city, and even between companies and musicians.

[Slide: Software and information patentability are the object of fierce political debates in Europe and globally]. ISMIR is a community of open exchange, where all proceedings are accessible free of charge, and usable freely for some usage (even if it would be better defined in my opinion through a Creative Commons license than by the present clause). But this open exchange is stopped at the border of becoming fact. Technology transfer departments of public research centres are patenting algorithms like hell, and intellectual property divisions of companies are patenting software systems components and application features.

After 20 years, reviews of the impact of the Bayh-Dole act conclude to an insignificant contribution of licensing as a funding source for public labs. However, the same technology transfer departments conclude from their failure that you should help them do more of the same. The funny thing is that while they carefully enclose the public domain of research, innovation disseminates from other sources, even if they

are initially much less sophisticated technically, simply because they have chosen cooperation and sharing instead of enclosures.

Inventory of fixtures

While preparing this talk, I read through the summaries of communications at this conference, and in a number of cases through the proceedings paper version. I produced a taxonomy [Slide], which is slightly more refined than the one the program committee has done to organise the sessions. Which insight can we derived from it?

Don't worry, you do not have to read all this stuff. The table is in the written version of my talk. Let's look at the top-level categories. First, good news, the ISMIR community is doing lots of basic technical work, more than a third of all papers, more if you count the evaluation benchmarks. However, the mass of accumulated know-how is far from being available as components for practical music applications. This is where the phenomena highlighted in my slide on software patents and technology transfer hits stronger. Nothing specific of the music information field: as soon as something seems to have some real potential, the motto seems to be "let's make sure than nobody can become rich with it other than us". Why is this so detrimental? Because then nobody does. Because the idea that in a field like music information any single organisation, even with its industry partners, will succeed in disseminating for usage a global platform is ridiculous. Nobody can have a clue of what people will really do with a full chain of music information processing, retrieval, re-creation, exchange. The last who can know, or more precisely the last who can accept to know are the major companies. You don't believe me, well why do you think that for years Sony and Philips have carefully buried all the nice ideas and technology that we all know exist in their labs? Only when we have a set of freely re-usable and adaptable music information technology modules, that are more or less combinable, and with which people can explore new ways of usage, will we start to see some real usage. I will come back to that later.

At the other end of the technology spectrum, the systems and testbeds end, there is also some very nice work, but it suffers from a related disease: restrictive rights and their implementation by DRMs. Don't mistake the dreams of majors and their trusted technology suppliers for reality. DRMs have no future, though they can wreck ours. Only creative commons and digital rights information are worthwhile, because only them enable the relationship which is music. Don't let yourselves be boxed by boxed music. What's wrong with DRMs? Access and usage control by technology. The user as dangerous enemy. And their complexity. No idea that is that complex, and that interacts with human behaviour ever worked. The Internet would not work if it had been designed with that much determinism, that much a-priori rules about what people can and can't do. And decreasing returns. DRMs and fingerprinting surveillance have a cost per item, and the returns are decreasing. DRMs may be sold under the name of fair reward, but

Topic	Number	Topic	Number
Feature extraction, classification in one feature space, similarity measures, structure extraction	45	Systems	12
Tempo and beats (including tracking)	7	Distributed music processing	1
Timbre incl. drum sounds and gestures	5	Metadata databases and browsers	2
Instrument	2	Digital libraries and Web services MIR systems	3
Drum patterns, rhythm	4	Fingerprinting for DRMs	3
Voice (singing detection), singing language	2	Real-time synchro (w/contextual info, w/accompaniment)	2
Tonality, key	3	Toolboxes (Matlab)	2
Chords	2	Evaluation, testbeds, and benchmarking	9
Motivic and melodic lines	10	Feature extraction and similarity	2
Ornamentation	1	Query-by-humming based systems	2
Other segmentation	2	MIR systems	2
Structure extraction	5	MIR systems with DRMs aspects	3
Sound synthesis meta-language	1	Classification (artists, songs, genres)	7
Optical music recognition	1	Clustering pieces and collections	2
Multi-feature and sequence matching	5	Artists	1
Usage paradigms and their support	23	Genre	4
Sampling to MIDI instrument	1	Empirical studies of usage	6
Query by humming	1	Personal music collections	2
Query by voice percussion	2	Digital libraries MIR systems	2
Annotation	3	P2P MIR systems	1
Notation, pitch spelling	2	Social networks (artists)	1
Transcription, music to score, mapping to patterns	4		
Alignment to score	4		
Production and usage of textual contextual info	3		
Document space representation	3		

Table 1: Topics at ISMIR'20204

all they will do is augment the already huge concentration of reward on bestsellers. I can hear you thinking “as we were the ones deciding”. Yes, you are. Many of your research centres do produce music, all of it, from composition to publishing. Many of your organisations care for public archives, and have accepted for years without enough reaction that the various extensions of copyright duration would stop anything to enter the public domain. Many of your organisations are producing music education software that is used in schools on some predefined music corpora. Is it acceptable that this is often proprietary software (paid by public money), and music corpora with terms of use that forbid any real-world usage? Do you think music learning stops at the walls of the class-room?

In between, we have what I have called “usage paradigms and their support”. This is often the most important item in a technical community: the structuring meeting point. The ISMIR community has gone a long way there, from the dominance of the initial Query by Humming paradigm to things like alignment to score (often a great enabler), and document space representation, that has proved a powerful approach in other fields of information retrieval. There is clearly still a long way to go, many other paradigms to discover, so I can address you with a less contentious statement: just let’s keep our minds open.

Finally, there is a new category: real world usage studies and the connected music classification. This is a very welcome category, at least when work there starts with an open mind. For instance, this is the only domain of MIR where people consider P2P file sharing as an interesting element of reality, and not a compromising devil. I have been surprised however about how many people consider automatic genre classification. It seems to me that “genre” is a concept that music information technology is destroying quickly. It is a concept that was invented for formatting (for instance in radio) and marketing purposes (for instance in record stores), because there was no better way to point people to what they might like. So there is a risk that the target of recognition will evaporate faster than recognition technology will develop.

A selection of music information dreams

I will conclude on a [\[slide\]](#) of dreams. User dreams, I am no longer trying to build these things (in case you wonder, that’s not because I am retired, that’s because I am trying to build other things).

I would like to see a set of free software modules, for various music content processing, indexing & retrieval, and interaction functionality, with a basic framework for combining them into practical chains of usage. I know various people are working on it, both in this community and outside it, so I have at least one realistic dream. One challenge here is that we need a framework for combining them, but this framework must be minimalistic. If the complexity of understanding the framework, of adapting a module to it is too high, we will lose the benefits of wide cooperation, of exploratory usage. Another challenge is that

we need to get rid of any notion that the algorithms or the functionality in these modules could be patented. This may be obtained in the European legal framework, though it is a hard fight.

Considering the limited resources that are available to build large-scale MIR systems and experiment with them, it is necessary to focus on music that be accessed, exchanged and re-used without high transaction costs. Public domain (when there is one), voluntarily created musical commons, and music that can be reused under legal licensing schemes are the only candidates. And the way to start is by constructing music commons. Once again, there are people in your community who work on it, but they are faced with many obstacles. So I dream that there would be more of a community-wide approach at it.

The tragedy of the present repressive approach to P2P file sharing is that it prevents it from maturing by putting it under siege, and then uses this immaturity to further justify warfare. Installing P2P file sharing as a respected social and technical paradigm for archiving and reusing music is a very valuable aim, to which the MIR community can, if it wishes, contribute key elements. This is closely linked to the previous point (musical commons). If you love fingerprinting, use it to control injection instead of controlling usage, for instance control that no fakes are injected. Develop new schemes of connecting contextual metadata with P2P access to raw music (by the way don't forget to look at what the P2P users are already doing in that respect). Stand for attribution, for making available information about usage rights without trying to force or monitor their respect by individual users by technical means. Work with people who construct P2P services with added-value information for media commons, such as the Réseau National d'Echanges that Tariq Krim is trying to create in France.

Finally, the retrieval paradigm however valuable it is may be too narrow. There is more to music information than retrieval, as can be seen from this conference's topics.

[Slide] In case you share any of these dreams, don't believe somebody else will remove the obstacles.

Thank you

After many potentially contentious statements, to install back some serenity, or to finish destroying it, I propose you a very short passage in music before questions, easy listening music that illustrates how the genre category is suffering a bit. Thank you for your attention.