

FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO



Social Software for Music

Cláudio Miguel Teixeira da Costa

Project Report

Master in Informatics and Computation Engineering

Supervisor: Ademar Aguiar (PhD)

July 2009

Social Software for Music

Cláudio Miguel Teixeira da Costa

Project Report

Master in Informatics and Computer Engineering

Approved in public examination by the committee:

Chair: António Lucas Soares (FEUP / PhD)

External Examiner: António Rito Silva (IST / PhD)

Internal Examiner: Ademar Aguiar (FEUP / PhD)

31st July 2009

Abstract

The work presented in this report refers to the project “Social Software for Music”, developed in cooperation with INESC Porto and Palco Principal. The project had a duration of 4 months (between March and June 2009) and has as goal to improve Palco Principal, a social software platform focused in the music industry.

After a study of Palco's platform, as well as of the several other platforms existing, it was possible to understand the needs and opportunities of Palco Principal, in order to level and even excel its competitors. From this study, three main areas of focus were elected: music recommendation, music classification and event management.

A brief analysis on the context of this systems is made, from the definition and applications of social software to the context of the music industry and the Music Information Research field of study.

Several aspects of music recommendation and classification are studied, as well as their applications and associated issues. Also a study of the existing recommendation techniques is developed, along with the metrics used to evaluate their success.

We then try to understand how music classification is made and we explore some novel forms of classification, apart from the traditional genre categorization: acoustic fingerprinting and mood classification.

Some examples of popular technologies of classification and recommendation are provided, as well as some brief information regarding the project's technological aspects.

In order to solve the proposed problem – improve social software for music in Palco Principal – several features are proposed for implementation, being carefully analysed and explained. We also try to understand how these features may help addressing the several problematics and issues of social software recommendation and classification.

Regarding event management, a study about existing solutions is made and a few features are also proposed for implementation, in order to improve the interaction and usability of this important part of the system.

Several prototypes are then presented, in order to specify how the implementation and integration of the several proposed features may be done.

Finally, an analysis of the main conclusions from this work is made and some future and follow-up work is presented.

Resumo

O trabalho apresentado neste relatório tem como âmbito o projecto "Software Social de Música", desenvolvido numa parceria entre o INESC Porto e a empresa Palco Principal. O projecto teve uma duração de 4 meses (entre Março e Junho de 2009) e tem como objectivo melhorar o Palco Principal, uma plataforma do software social focada na indústria da música.

Após um estudo da plataforma Palco, bem como de várias outras plataformas já existentes, foi possível entender as necessidades e oportunidades do Palco Principal, a fim de o conseguir nivelar com os seus concorrentes e até superá-los. Deste estudo, três principais áreas de enfoque foram definidas: recomendação, classificação e gestão de eventos relacionados com música.

É feita uma breve análise acerca do contexto destes sistemas, partindo da definição de aplicações de software social até ao contexto da indústria da música e da área de investigação designada por *Music Information Research*, aonde este trabalho se insere.

Vários aspectos da recomendação e classificação de música são estudados, bem como as suas aplicações e problemáticas. Também um estudo das técnicas de recomendação existentes é desenvolvido, juntamente com um estudo das métricas utilizadas para avaliar o seu sucesso.

De seguida tenta-se compreender como a classificação de música é feita e são exploradas algumas novas formas de classificação, para além da tradicional categorização por género musical: *acoustic fingerprinting* e *mood classification*.

Alguns exemplos de tecnologias de classificação e recomendação mais conhecidos são apresentados, bem como algumas informações relativas aos aspectos tecnológicos do projecto.

A fim de solucionar o problema proposto – melhorar a plataforma de software social para música do Palco Principal – várias funcionalidades são propostas para implementação, sendo cuidadosamente analisadas e comentadas. É feito também um esforço no sentido de entender como estas funcionalidades podem ajudar a enfrentar as diversas problemáticas e questões associadas à recomendação e classificação de software social.

No que diz respeito à gestão de eventos, um estudo sobre as diferentes soluções existentes é feito e algumas funcionalidades são também propostas para implementação, a fim de melhorar a interacção e usabilidade desta importante parte do sistema.

Vários protótipos são então apresentados, a fim de especificar a forma como a implementação e integração de várias funcionalidades propostas será feita.

Finalmente, é feita uma análise das principais conclusões deste trabalho e é apresentado algum trabalho futuro e de aperfeiçoamento.

Agradecimentos

“Without music, life would be a mistake”

Friedrich Nietzsche (1844-1900)

Em primeiro lugar e antes de todos os outros, tenho de agradecer ao Prof. Ademar Aguiar, pela orientação e acompanhamento sempre presente, e pela compreensão e apoio demonstrado.

Tenho também de agradecer a algumas pessoas e entidades importantes na realização deste projecto, nomeadamente ao INESC Porto e equipa da USIC, na pessoa do Eng. António Gaspar, e à equipa do Palco Principal, em particular ao Eng. João Carvalho, Eng. Pedro Trindade e Eng. Daniel Botelho, pelo apoio e confiança depositados em mim.

Um agradecimento especial vai para os dois directores de curso, Prof. Raul Moreira Vidal e Prof. António Augusto Sousa, por conduzirem sempre o curso no melhor interesse dos alunos e por serem, mais do que directores ou professores, mas antes de tudo, um apoio amigo para todos os alunos deste curso e alguém a quem nos podíamos (e podemos) dirigir quando precisávamos de um conselho ou de uma palavra em determinado momento. Uma nota especial vai também para o *nosso* secretariado do *nosso* curso (Mónica, Lina, Marisa e Sandra), que sempre com simpatia, prontidão e carinho nos ajudaram, aqui e ali quando foi necessário.

Presto ainda a devida homenagem a todos aqueles que dentro e fora do curso foram meus colegas de FEUP, de curso, de estudo e de mais alguma coisa, em especial às comunidades GT e VMT pela ajuda prestada e por tornarem os momentos difíceis mais suportáveis.

A muitos mais haveria que agradecer, pelos momentos e ensinamentos dentro e fora das salas desta faculdade. A esses levo-os junto no peso e nas manchas da capa, já um pouco negra, de saudade.

Por fim, mas não em último, mais do que agradecer, cabe-me a obrigação de dedicar não só este trabalho mas todas as vitórias, dificuldades e aventuras à minha família: mãe, pai e irmã. Não porque o deva fazer, mas porque lhes devo isso, e muito mais. Os últimos tempos foram difíceis, mas o futuro só pode ser bom. Juntos, claro.

O Autor,
Cláudio Miguel Teixeira da Costa

Index

Chapter 1 Introduction.....	1
1.1 Context.....	1
1.2 Project description.....	2
1.2.1 Palco Principal.....	3
1.2.2 INESC Porto.....	3
1.3 Motivation and objectives.....	4
1.4 Document overview.....	4
Chapter 2 Problem description.....	6
2.1 Overview.....	6
2.2 Project requirements.....	7
2.3 System description.....	8
2.4 Assumptions and constraints.....	10
2.5 Project planning.....	10
2.6 Summary.....	11
Chapter 3 Review on social software for music.....	12
3.1 Social software.....	13
3.1.1 Brief history.....	13
3.1.2 Definition and characterization.....	13
3.1.3 Social software applications.....	14
3.2 Social software for music.....	16
3.2.1 The new music industry.....	16
3.2.2 Existing platforms.....	19
3.2.3 Comparison and analysis.....	21
3.3 Music Information Research.....	24
3.4 Music recommendation.....	26
3.4.1 Formalization of the recommendation problem.....	27
3.4.2 Applications of recommendation.....	32
3.4.3 Issues with recommendation.....	34
3.4.4 Recommendation techniques.....	39
3.4.5 Evaluation metrics.....	43
3.5 Music classification.....	44
3.5.1 Social tagging.....	46

3.5.2	Annotation games.....	50
3.5.3	Expert classification / Surveys.....	51
3.5.4	Web mining.....	52
3.5.5	Automatic content-based classification.....	54
3.5.6	Mood classification.....	55
3.5.7	Issues with classification.....	56
3.6	Event management and recommendation.....	60
3.7	Recommendation technologies.....	61
3.7.1	Last.fm's Audioscrobbler.....	61
3.7.2	Music Genome Project	62
3.7.3	MyStrands.....	62
3.7.4	The Echonest.....	64
3.7.5	Itunes' Genius.....	64
3.8	Project technologies.....	65
3.8.1	PHP.....	65
3.8.2	Symfony.....	66
3.8.3	MySQL.....	67
3.9	Key issues and lessons learned.....	68
Chapter 4	Improving social music software.....	69
4.1	Improving classification.....	69
4.1.1	Tag clouds.....	70
4.1.2	Multi-tag search	70
4.1.3	Tag suggestion.....	70
4.1.4	Acoustic fingerprinting.....	70
4.1.5	Weighted genre classification.....	70
4.1.6	Mood classification.....	70
4.1.7	Wiki-like content editing.....	70
4.2	Improving recommendation.....	71
4.2.1	Music recommendation.....	71
4.2.2	User affinity and recommendation.....	71
4.2.3	Rating system.....	71
4.2.4	Reputation system.....	71
4.2.5	Graphic recommendation.....	71
4.3	Improving event management.....	72
4.3.1	Geographic Information System (GIS).....	72
4.3.2	Event displaying.....	72
4.3.3	Event suggestion.....	73
4.3.4	Venue profile.....	73
4.3.5	User-added media.....	73
4.4	Conclusions.....	73
4.5	Summary.....	76
Chapter 5	High level prototype.....	77
5.1	Homepage.....	77

5.2 Header and Footer.....	78
5.3 Music	79
5.3.1 Artist overview.....	80
5.3.2 Artist biography.....	82
5.3.3 Artist tags.....	83
5.3.4 Artist recommendations.....	84
5.3.5 Artist music	85
5.3.6 Artist events.....	86
5.3.7 Artist photos.....	87
5.3.8 Artist videos.....	87
5.3.9 Artist blog.....	88
5.4 Community.....	89
5.5 Tags	90
5.5.1 Tag search	90
5.5.2 Tag overview.....	91
5.5.3 Tag music	92
5.5.4 Tag events	93
5.6 Events.....	94
5.7 Venues.....	96
5.8 Top videos.....	97
5.9 Top photos.....	98
Chapter 6 Conclusions and Future Work.....	100
6.1 General Conclusions.....	100
6.2 Future Work.....	102
References.....	104
Annex A: Analysis of the project's areas of focus.....	114
Annex B: Platform analysis.....	117
Annex C: Platforms' dimension – data sources.....	139
Annex D: Patents in MIR research.....	141
Annex E: Proposed features.....	144
Annex F: Prototypes.....	173

List of Figures

Figure 3.1: Online music categories distribution in 2008, by Compete23.....	17
Figure 3.2: Top 20 streaming music sites in 2008, by Compete23.....	18
Figure 3.3: A simplified overview of MIR [Fingerhut 04].....	25
Figure 3.4: The music information plane – MIP ([Lamere 07] and [Celma 08]).....	30
Figure 3.5: Relationship between users and items for recommendation [Celma 08].....	31
Figure 3.6: Schema of how Strands Recommender TM works57.....	63
Figure 5.1: Prototype of Palco 3.0 homepage.....	78
Figure 5.2: Prototype of header and footer.....	79
Figure 5.3: Prototype of music page.....	80
Figure 5.4: Prototype of artist overview page.....	81
Figure 5.5: Prototype of artist biography page.....	82
Figure 5.6: Prototype of artist tags page.....	83
Figure 5.7: prototype of artist recommendations page.....	84
Figure 5.8: Prototype of artist music page.....	85
Figure 5.9: Prototype of artist events page.....	86
Figure 5.10: Prototype of artist photos page.....	87
Figure 5.11: Prototype of artist videos page.....	88
Figure 5.12: Prototype of artist blog page.....	88
Figure 5.13: Prototype of community page.....	89
Figure 5.14: Prototype of tags page.....	90
Figure 5.15: Prototype of tag search page.....	91
Figure 5.16: Prototype of tag overview page.....	92
Figure 5.17: Prototype of tag music page.....	93
Figure 5.18: Prototype of tag events page.....	93
Figure 5.19: Prototype of events page.....	94
Figure 5.20: Prototype of event overview page.....	95
Figure 5.21: Prototype of venues page.....	96
Figure 5.22: Prototype of venue overview page.....	97
Figure 5.23: Prototype of top videos page.....	98
Figure 5.24: Prototype of top photos page.....	99
Figure E.1: Last.fm's tagcloud.....	146
Figure E.2: Del.icio.us's tagcloud.....	147
Figure E.3: del.icio.us multi-tag query.....	148

Figure E.4: Tag suggestions in del.icio.us.....	149
Figure E.5: Musicoverly's mood-pad.....	154
Figure E.6: Wiki-like editing in Last.fm.....	155
Figure E.7: Music recommendation on Pandora.....	158
Figure E.8: Music recommendations on iMeem.....	158
Figure E.9: User compatibility in Last.fm, based on musical proximity.....	159
Figure E.10: User suggestions on Facebook, based on social proximity.....	159
Figure E.11: Comment rating in Youtube, using thumbs up/down.....	161
Figure E.12: Song rating in iMeem, using 5-star rating.....	161
Figure E.13: LivePlasma - recommendations and relations between artists.....	164
Figure E.14: TuneGlue's Audiomap - recommendations and relations between artists.....	164
Figure E.15: Last.fm's location-based event listing	166
Figure E.16: GIS information on iMeem.....	166
Figure E.17: Event displayed in Songkick, with RSVP and ticketing information.....	168
Figure E.18: Event displayed in Last.fm, with RSVP and venue information.....	168
Figure E.19: Event search parameters in bandsintown.....	169
Figure E.20: Songkick's event suggestion, based on event location.....	169
Figure E.21: venue profile in iLike.....	171
Figure E.22: venue profile at SongKick.....	171
Figure E.23: user-added media in Last.fm.....	172

List of Tables

Table 2.1: Prioritization of the areas of focus.....	8
Table 3.1: Applications of social software.....	15
Table 3.2: Demographics of an on-line music buyer (data in %).	16
Table 3.3: Clusters of features for analysis.....	21
Table 3.4: Rank of most complete platforms and their dimension.....	22
Table 3.5: Most popular features.....	23
Table 3.6: Less common features.....	24
Table 3.7: Dewey classification for MIR [Fingerhut 04].....	26
Table 3.8: Comparison of the scale of 4 recommender systems: Netflix, Yahoo!Music, Amazon and Last.fm.....	37
Table 3.9: Strengths and weaknesses of tag-based music annotation approaches [Turnbull 08].	45
Table 3.10: Uses of tags in music annotation	49
Table 3.11: Number of Wikipedia articles written in the 10 most used languages.....	59
Table 4.1: Influence of features in issue solving.....	75
Table 4.2: Feature implementation prioritization.....	76

List of Acronyms

3D	3 dimension
AIR	Audio Information Retrieval/Research
AMG	All Media Guide
API	Application Programming Interface
BPM	Beats Per Minute
C2C	Client-to-client
CAPTCHA	Completely Automated Public Turing test to tell Computers and Humans Apart
CB	Content-Based Filtering
CD	Compact Disc
CF	Collaborative Filtering
CNA	Complex Network Analysis
CPU	Central Processing Unit
CSCW	Computer Supported Cooperative Work
DM	Document-based Mapping
EMD	Electronic Music Distribution
FEUP	Faculdade de Engenharia da Universidade do Porto
GIS	Geographic Information System
GMM	Gaussian Mixture Model
GNU	GNU's Not Unix
HTML	Hypertext Mark-up Language
ICT	Information and Communication Technologies
IFPI	International Federation of the Phonographic Industry
INESC	Instituto Nacional de Engenharia de Sistemas e Computadores
ISMIR	International Symposium on Music Information Retrieval
MAE	Mean Absolute Error
MFCC	Mel-Frequency Cepstral Coefficients
MIP	Music Information Plane
MIR	Music Information Retrieval/Research
MIREX	Music Information Research Evaluation eXchange
MIT	Massachusetts Institute of Technology
MP3	MPEG-1 Audio Layer 3
MVC	Model-View-Controller
ORM	Object-Relational Mapping
P2P	Peer-to-peer

PCM	Pagecount-based Mapping
PHP	PHP: Hypertext Preprocessor
PM	Pattern-based Mapping
PMI	Pointwise Mutual Information
RMSE	Root Mean Square Error
ROC	Receiving Operating Characteristic
RSS	Really Simple Syndication
RSVP	Répondez S'il Vous Plaît
SML	Supervised Multi-class Labelling
SOM	Self Organizing Maps
SQL	Structured Query Language
SVM	Support Vector Machines
TF-IDF	Term Frequency–Inverse Document Frequency
UK	United Kingdom
UMA	Universal Matrix of Associations
URI	Uniform Resource Identifier
US	United States
VOIP	Voice Over Internet Protocol
VSM	Vector Space Model
WMA	Windows Media Audio
WYSIWYG	What You See Is What You Get
XML	eXtensible Markup Language

Chapter 1 Introduction

With the transformations in the music industry over the last years, we have gone from the traditional physical formats of music recordings (vinyl records, audio cassette's or compact discs – CDs) to a new era of digital music, with the music now available on on-line music stores and on the fan's MP3 and MP4 players (some authors refer to the upcoming generation as the “iPod generation”, in a reference to the massive use of this type of portable media player).

We are changing from a time where a few individuals had the power to influence the music preferences of the masses, by controlling the distribution channels and only “giving them what they wanted”. But if before this “they” referred to the record companies, labels and radio hosts, nowadays the “they” is much broader and we are quickly moving to a time where all the music will be available for whoever wants it.

1.1 Context

Ironically, this incredible amount of music may become a problem for the music fan, as he might have some trouble to find what he really wants in the middle of such an overwhelmingly big collection. And although the music industry has changed in size and structure, the tools offered to music consumers to find music are still more or less the same. Music consumers still browse music by its genre or try to find a certain artist, album or song, just as they used to do in the brick-and-mortar music store, with a limited scope of resources. This is where music recommendation and classification become a key part of the music ecosystem.

Classification systems are the busy librarians who help the user easily find the content he needs, by organizing and classifying the information. This is an extremely hard and time-consuming task, for which MIR research is aware and where the community of music fans plays a vital role, with the so-called “wisdom of crowds” or “crowdsourcing”.

Recommender systems serve as the middle-man in the music marketplace, helping a music fan to find its favourite bands and songs among the millions of other artists and songs that are not relevant for him. Social music recommenders already play an important role in online music. Websites such as Last.fm¹ and Pandora² provide recommendations for millions of fans every month. Yet, these music recommendation systems still don't use all the possibilities of the advances techniques that are studied and developed in the Music Information Research community.

With the advent of ever increasing technologies, exploring new communication channels such as the Internet and mobile devices, and a larger and more variate offer of products and services in the music industry, many challenges are posed to its stakeholders (may they be artists, labels, retailers or fans). This brings new and exciting challenges to the music industry, but also a number of opportunities that are ready to be seized and explored. All in the name of music.

1.2 Project description

This project (Social Software for Music) is part of a larger one, entitled Palco 3.0, resulting of a joint cooperation between the company Palco Principal and several research teams from INESC Porto (Instituto de Engenharia de Sistemas e Computadores do Porto), FCUP (Faculdade de Ciências da Universidade do Porto) and FEP (Faculdade de Economia do Porto).

The work presented in this report refers to a period of four months, in the second quarter of a planned three year project, Palco 3.0. It is therefore mostly focused on the study and analysis of the several aspects of social music software.

The main project aims to the general implementation of intelligent and collaborative decision support systems and an improved usability and knowledge management to enhance and maximize the automation and innovation in several critical processes. These processes include the identification and selective search of content and the organization, sharing and intelligent processing of data.

In order to coordinate and optimize the efforts of the different institutions and people involved, several teams were created – one of them being the “social software team”, at INESC Porto, to which this report and project refer to. The other teams focus in areas such as data mining, audio analysis, machine listening and other fields of knowledge, working in close cooperation and interconnected to each other.

1 <http://www.last.fm>

2 <http://www.pandora.com>

1.2.1 Palco Principal

Palco Principal is a Portuguese-based company created in June 2006 with the support of NEOTEC, a programme to support the creation and development of technological-based ideas, coordinated by the national agency for knowledge society (Agência para a Sociedade do Conhecimento, UMIC) and the national agency for innovation (Agência de Inovação, AdI).

The company focuses on the music industry, operating in the context of Web 2.0 and offering a technological and highly innovative form of doing business. Its core business is the commercial exploration and development of the portal, being also able to develop tailored software, as well as providing services in knowledge management and decision-aid support mechanisms.

The Palco Principal team gathers high technological skills and a good knowledge of the music industry, combining a wide experience of management and in the business sector of information technologies.

The company has been developing a portal (www.palcoprincipal.com), using proprietary technology, that already hosts the largest community of Portuguese expression music, by bringing together fans and artists from different Portuguese-speaking countries, such as Brazil, Cape Verde, Angola and Mozambique. Palco Principal has already over 6200 bands and more than 19.000 registered users, with 14.000 songs available for listening and downloading and over 30.000 regular visitors.

One example of its consolidation as a corporate project was the commercial partnership, in September 2007, with the Sonaecom³ group, through its portal Clix⁴. Since 2009, Palco Principal has established a partnership with the largest web portal in Portugal and in the Portuguese communities, Sapo⁵, owned by the Portuguese telecommunications company PT⁶.

The market expansion has been done gradually and the audience increased considerably over the last years. The company entered the Brazilian market in 2007, having established, in the 2nd quarter of 2008, a partnership with one of the most important and largest Brazilian webportals, UOL⁷, with already 13 years of existence and a large presence in the Latin America market. The recent partnership with Sapo brought a wider audience in Portuguese expression communities in Cape Verde, Angola and Mozambique.

1.2.2 INESC Porto

INESC Porto - Institute for Systems and Computer Engineering of Porto - is a private non-profit association, based in Porto, Portugal and recognised as Public Interest Institution, being recently appointed as Associated Laboratory.

3 <http://www.sonaecom.pt>

4 <http://www.clix.pt>

5 <http://www.sapo.pt>

6 <http://www.telecom.pt>

7 <http://www.uol.com.br>

Introduction

Its main activities are scientific research and technological development as well as consulting and advanced training in the areas of Telecommunications and Multimedia, Power Systems, Manufacturing Systems Engineering, Information and Communication Systems and Optoelectronics.

INESC Porto is a research institution working as an interface between the academic world and the worlds of industry, services and Public Administration in the areas of Information Technologies, Telecommunications and Electronics (ITT&E).

The Institute seeks to regulate its activity with criteria such as innovation and internationalisation, mainly through the establishment of a set of strategic partnerships that guarantee its institutional stability and economic sustainability.

1.3 Motivation and objectives

The general goal of the Palco 3.0 project is the achievement of the following results:

- creation of the Palco Principal 3.0 product;
- creation of a generic and reusable technologic platform (PTECH);
- creation of Palco Loja, a prototype to explore the potentialities of PTECH to new areas of application;
- creation of Palco Move!, a prototype to explore the potentialities of PTECH on mobile devices and platforms;
- publications in proceedings of relevant international conferences.

For the project “Social Software for Music”, developed within the “social software” project team, the main objectives are the following:

- analyse and compare other existent social music solutions, as well as their features;
- study and analyse the different aspects of social software, with a focus on its application to the music industry and in aspects such as music classification and recommendation and event management;
- analyse and propose a series of improvements for the platform Palco Principal, in cooperation with the other development teams;
- design an integrable prototype, applying the techniques and approaches studied.

1.4 Document overview

This document is organized in six different chapters. Each chapter has several sections and subsections, identified with a numeration system. As an example, the description of INESC Porto is placed in chapter 1 (Introduction), section 1.2 (Project description) and subsection 1.2.2.

An abstract of this report is provided in the beginning, as well as content, table, figures and acronym indexes. By the end of the document, it is possible to find the bibliographical references, as well as the annexes.

Introduction

This first chapter, introductory, presents some general informations about the project, its context, motivation and objectives. It is also made a brief description of both the entities involved in the project: Palco Principal and INESC Porto;

Chapter 2 exposes the problem that is proposed for resolution, providing an analysis on the *status quo* as well as explaining the project's requirements, planning, assumptions and constraints;

In the third chapter, a review on social music software is made, starting by first describing what social software is and immediately after its applications in the music industry. An important and extensive analysis on the existing social music platforms is also made. The areas involving the study and development of these systems are briefly presented, but a focus is done in three main areas: music recommendation, music classification and event management. In the first two some of its applications and issues are analysed, as well as existing technologies that already apply some of the referred techniques. Finally, some key issues are highlighted, as well as the lessons learned;

A proposal for the resolution of the presented problem – the improvement of social music software – is presented in chapter 4, where a number of features are described and proposed for implementation, which is then specified in chapter 5.

The last chapter presents the general conclusions extracted from the work presented in this report, as well as an overview of the future and follow-up work to be done.

Chapter 2 Problem description

Palco already provides many useful features, however the company would like to improve three essential areas: music classification, recommendation and events. Also an overview and comparison with other similar systems would be useful in order to understand what are the users' needs and which features are most relevant for them.

2.1 Overview

Music classification presents many challenges, namely in what respects to music cataloguing into genres or categories. It is not clear how to identify the borders between music genres or styles, and what generates the success of a reduced group of artists, in relation to the large majority that remains in the shadow of the successful hits.

In order to democratically facilitate the access to music, it is necessary to build a system that is flexible enough to allow the classification of music into different “shelves” and with different perspectives, but at the same time provides accuracy and objectivity when looking for a certain musical item.

However, in an industry that releases every year, only in the United States, between 27.000 to 35.000 new albums⁸, even with accurate and extensive classification methods, it is still necessary to guide music fans in this immensity that is the ocean of music. This is why music recommenders have been garnishing so much media and user attention, becoming a very important player in the whole music deal.

⁸ Source: <http://www.tinfoilmusic.net/copyright-laws-a-news/90-dmca/684-riaa-lies-at-dmca-hearing> Published in 15.05.2003 and retrieved in 23.06.2009

Problem description

Such a system must be able to understand the user and the music herself, promptly providing meaningful recommendations, learning with user's activity and behavioural patterns and acting as a personal musical consultant. This is an enormous challenge, as music perception is a rather open and subjective field.

The final piece of the social music puzzle that requires study and improvement is related to live music concerts or events. These are an inseparable part of the music business and a key bridge between music artists and fans, apart from the music records.

A social network for music should be able to take the most out of the relationships established between fans, artists and their music, being able to bring them all together. This is the proposed challenge.

2.2 Project requirements

For the outlining of the project, a few requirements were established, in order to guide and plan the work to be done. In the first place it was necessary to learn about the business and *status quo* of Palco Principal. This included meeting with the company directors and the multidisciplinary team, reading internal documentation and migrating the internal information to a new collaborative project management tool – Redmine⁹.

With the analysis made to other competing or related solutions (cf. section 3.2.2 - Existing platforms) it was possible to understand how the social music market was behaving and have a general panorama of what the users' needs and preferences were.

The features analysed in the several existing platforms were rated in order to evaluate their usefulness for Palco Principal and to prioritize them. This was done having four factors into account: their innovation (i.e. if they were already common features in other platforms or emerging ones), development effort (an estimation made based on the available time and human resources), their relevance for Palco Principal (this evaluation was made in a meeting with the company's directors, who rated the most relevant areas for their strategy) and the author's preferences. The results of this evaluation¹⁰ are depicted in Table 2.1, with the complete analysis available in Annex A: Analysis of the project's areas of focus.

The areas marked in bold, with an asterisk, are the ones on which to focus in this project: event management, music recommendation and music classification.

Albeit the existence of several other aspects that could deserve a thorough analysis and improvement, these were the areas to focus agreed within the project and the social software team.

⁹ Available at <http://www.redmine.org>

¹⁰ The values are presented using a reversed scale – a lower value means higher priority or effort.

Problem description

priority	area	effort
1,5	events*	1,5
1,6	recommendation system*	1,0
1,8	classification system*	2,0
2,0	reputation system	1,0
2,3	monitoring tools	1,5
2,6	media	2,0
2,8	collaborative tools	2,0
4,2	monetize	1

Table 2.1: Prioritization of the areas of focus

2.3 System description

Palco Principal is able to recommend songs, using association rules. Based on the music users have added to their playlists, the system calculates the probability of a fan to enjoy a certain music based on music he already likes. Then, the related songs are listed, according to their “likeable” probability.

These recommendations are only possible for individual songs and in order to work, it is necessary that a significant number of users add the “recommendable” songs to their playlist. The system will then recommend songs based on other user's preferences: “users that liked song X, also liked song Y”. This leads to a small number of songs that may be “recommendable”, as most of the songs available belong to just a few playlists or no playlists at all. It is not possible to recommend similar or related items, like artists or users. Even though it is possible to find artists within the same music genres, the system is not able to suggest artists that share a similar music profile or classification and fans are only able to influence these recommendations remotely, by adding songs to their playlists.

Music classification is also an area that needs improvement. The artist classification is done primarily by genre: when registering on the website, the artist must choose a main genre from a given list, being able to choose also a number of sub-genres (from the same genre-list provided by the system). However, this genre classification is not ideal, as an artist can be influenced by a number of different musical genres.

In order to solve this, Palco has introduced a tagging system (which was recently updated), where artists, fans, blog posts and news can be tagged freely. These tags may then be used to search for tagged content, however it is not clear the “weight” that different tags have, as they all are listed equally. In the end it is not clear what kind of classification is made, as the user feels a little bit lost in the “tag ocean”.

Problem description

Also, tags are not related to each other, even if they are similar or even synonymous. Searching tagged content is also difficult, as it is not possible to filter the search results (most likely when searching for a tag, the results will be presented in a list include artists, blog posts and news). With this approach, the potential of the “wisdom of crowds” that social tagging provides is not being explored properly.

Rating is only possible for songs, by posting reviews, where the fan can nominate the song for one of several music categories (best production, lyrics, melody, originality or performance) as well as instrumentation categories (best male voice, female voice, guitar, bass, percussion, keyboards or setlist) and rate the song, from 1 to 5. Unfortunately these ratings are not visible for other users and don't have an influence on the artist's ranking in the charts available.

The event area (the general name for music shows or “gigs”) of the website has some visibility, with a list of future concerts in the main page and the possibility of the user accessing a list of the next concerts of their favourite bands in his private area, but there is also a lot of room for improvement. The artist page presents a list of the next gigs where the artist will be performing, allowing the user to view detailed information about each event. However, the event pages are mostly static content, posted by the artist only, with information about the type of event (from a static list), the price, location, date, name and description of the event. It is not possible for users to have any form of interaction with the events (by leaving comments, adding media, reviews or attendance information), view ticketing or location information and even to add an event to the user's personal agenda, it must be done indirectly by flagging the band as favourite. There is also no kind of additional information about venues and no information about related or recommended events.

By comparing the several existing platforms, and the different “philosophies” and features they offer to the general music audience (may they be labels, artists or other music professionals or fans – ranging from the casual to the enthusiast fan), it could be interesting to understand what types of users populate these platforms and how they work as a community. It would be important to learn not only what are the user's profiles and needs, but what features they are eager to use and experiment and how they do it (as in some cases, the use that some users make of the features available is substantially different from the use that was first projected when they were designed).

Another important issue has to do with the social perspective of social music, how the users interact with each other and with the system and how can social software help them adjust to the contemporary aspects of music. These have been changing dramatically over the last years and experiencing a huge influence from the new technological tools and paradigms, both in the technical, social and business perspectives of music.

Problem description

A final note on this description has to be made regarding the website's features and its usability. Although this report does not focus on usability issues, reason why no formal study on this subject was made, it is still relevant to mention that the introduction and/or reformulation of the existing features will only be effective if the system's usability is taken into account.

With the website being used by very different types of users, with distinct knowledge and experience with web tools, and a different vision of the website and its functionalities, it is of key importance to have a good usability so that the users can explore the potentialities of the system in full. This is already a concern of Palco Principal, but from the user's point of view, there are necessary improvements to be made, in order to facilitate and enhance the user experience and the exploitation of the system's features and *nuances*.

2.4 Assumptions and constraints

The realization of this project is based in a few assumptions. One is that the project is developed within a multidisciplinary and complementary team, which are to a certain extent autonomous, but not independent. Another assumption is that the work is done on the latest version of Palco Principal's technology, over which new functionalities will be built. Finally, another assumption has to do with the project's duration: this project is being developed from March to June 2009, but it is only part of the general project which is ongoing until 2011, thus still being in an initial stage.

Regarding the technology to be used, a few constraints were fixed in order to allow the integration with the existing version of Palco Principal. The technologies to be used are PHP¹¹ as the scripting language, together with the Symfony¹² framework and MySQL¹³ database management system. As this project is being developed in INESC Porto, most of the presential work is done from there, with regular visits and meetings at Palco Principal's headquarters.

2.5 Project planning

The project officially started on the 2nd March 2009 and ended on the 29th June 2009.

Its initial planning was intended to be very general and merely indicative of the amount of time available/needed for the development of the project.

By the beginning of the 7th week this plan was revised and detailed, in order to better prioritize the work and manage the effort and time available.

The first three weeks of the project, in March, were spent between the natural integration and kick-off related to the start of the project, as well as the collecting and understanding of the project and system that was going to be used as a base to work on.

11 PHP stands for *PHP: Hypertext Preprocessor*. Available at <http://www.php.net>

12 Available at <http://www.symfony-project.org>

13 Available at <http://www.mysql.com>

Problem description

The following three weeks, until mid-April, were used to discover, study and analyse the several existing social music platforms, in order to understand which aspects were more explored and which others were still appearing as new.

In the last two weeks of April, a bigger insight on Palco Principal's architecture and design was pursued, as well as the study of the technologies in use. With this it was possible to start creating an idea of what it would be the solution for the proposed problem.

Additional research on specific topics related to the solution was made in the first three weeks of May, which were followed by two weeks used to document this study in the report writing.

The month of June was used to finish the project's documentation and the writing of this report, with the last week of the month being used for reviewing and evaluating the work developed.

2.6 Summary

In this chapter an overview of the problem to be addressed in this project is presented. It is explained how the decision of focusing the analysis and implementation of new features related to music classification, music recommendation and event management were made.

A general analysis of the importance of these aspects of social music was presented, as well as an analysis of this aspects in the platform Palco Principal, denoting which aspects are subject to improvement.

Also in this chapter are mentioned the project's requirements, assumptions and constraints. Finally, the project plan is described, in order to provide an idea of the time-frame of the project.

Chapter 3 Review on social software for music

Although only in the last decade social software gained worldwide attention and, subsequently, commercial value, the basis of a concept of a virtual social platform dates back to the 1940's. It has then evolved until the 1990's, when a growing interest on this field led to the creation of the first social software systems.

Naturally this concept expanded to several areas connected to the human and psychosocial relations, of which music (as a mean of self-expression) became one of the most prolific. In this chapter we will understand how social software embraced the several aspects of the music industry and culture, involving fans, artists and labels, as well as perform an analysis on existing social music platforms.

The field of knowledge associated to the research and development of technologies of music information retrieval and research is also briefly presented, followed by a deeper analysis on two of its main aspects: music classification and recommendation. Also a brief overview on event management solutions is made. These three aspects will be the basis for the specification of a solution proposal for the problem described in the previous chapter.

We will have an insight on the specificity of each of these topics, analyse the existing issues and the techniques that have been depicted in the literature. We will also present some examples of their applications. Finally, we will have an overview on some of the most relevant existing technologies and we will drawn some conclusions from this bibliographical review.

3.1 Social software

Social Software is used nowadays for a wide range of activities, from entertainment/personal to enterprise/commercial use. This is much due to the creation and growth of online communities, which has been accelerated over the last years.

3.1.1 Brief history

The basis, however, for the creation of this concept dates back to the 1940's, with Vannevar Bush conceiving a memory and communication device called “memex” [Bush 45], which today might be called of “personal computer”. Here, the author also proposed a number of far-reaching ideas, even decades before other researchers actively worked on them, becoming contemporary computing terms, such as augmentation, groupware and computer-supported collaborative work. The author presents the following definition of memex:

"A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory."

Only until the 1960's was it possible to observe more developments on Bush's ideas, with [Engelbart 62] exploring the concept of augmentation and [Licklider 68] outlining the first methods of computer-aided group collaboration. The first collaborative software, EIES (Electronic Information Exchange System) was implemented in the 1970's (an early version of the system is described in [Turoff 72]) and the first groupware system appeared in the 1980's. Finally, in the late 1980's and early 1990's, Computer-Supported Collaborative (or sometimes Cooperative) Work (CSCW) was developed, with the term being coined by Irene Greif and Paul Cashman in a 1984 workshop. In the 1990's, while the term “groupware” was slowly losing its meaning, a new term was beginning to coming into vogue: *social software*. However, for the first 15 years of its existence, mostly in the 1990s, the term was rarely used outside of very specialized groups, until the beginning of the new millennium.

3.1.2 Definition and characterization

Originally, social software had two distinct foci: i) the technological issues, user acceptance and social effects regarding group collaboration and on-line communication and ii) the use of computational techniques to facilitate the study of society and to test policies before they were employed in real-world situations.

Nowadays, there are several definitions for social software, with Clay Shirky defining it as “*software that supports group interaction*”¹⁴ and Tom Coates defining it as an “*augmentation of human's socializing and networking abilities by software, complete with ways of compensating for the overloads this might engender*”¹⁵.

¹⁴ http://www.lifewithalacrity.com/2004/10/tracing_the_evo.html, published on 13.10.2004, retrieved on 09.06.2009

More recently, [Burg 06] state that “*Social Software refers to those online-based applications and services that facilitate information management, identity management, and relationship management by providing (partial) publics of hypertextual and social networks*”.

Some other authors refer to social software as “social computing”, and although this designation is not consensual, it is worth to explore its meaning. An early definition by [Schuler 94] defines social computing as “*describing any type of computing application in which software serves as an intermediary or a focus for a social relation*”. More recently, [Wang 05] has defined social computing as “*computational facilitation of social studies and human social dynamics as well as the design and use of ICT technologies that consider social context*” and [Charron 06] as “*a social structure in which technology puts power in individuals and communities, not institutions*”.

Many advocates of social software tools defend that they create actual communities, using the term “online communities” to describe the resulting social structures that are created.¹⁵ According to them, some types of software seem to facilitate the development of these communities, in a more egalitarian and meritocratic sense, and also with a bottom-up structure, in which membership is voluntary, reputation is earned within the community and the community's management is done by its own members. This contrasts with other “top-down” communities, where a user's roles are determined by an external authority and circumscribed by rigid mechanisms, such as access rights.

We will then explore some of these communities and different approaches to social software, as well as its applications.

3.1.3 Social software applications

Social software encompasses a number of different software systems, allowing users to interact and share data, that can be divided in two main categories: communication tools and interaction tools. While communication tools are mostly focused in capturing, storing and presenting the items that compose communication, by use of several media formats (text, image, sound or video), interactive tools are focused in facilitating the mechanics of the communication, by providing forms of establishing a virtual conversation. Because of this, communication tools are generally asynchronous in opposition to interactive tools, which allow users to communicate in real time (e.g. video conferencing and VOIP) or near-synchronous (e.g. instant messaging and text chat).

These social software systems have different purposes and applications, as described in [Charron 06] and presented in Table 3.1. To what concerns this project, the most important purpose of social software is its application to the music industry.

15 http://www.plasticbag.org/archives/2002/12/on_the_augmentation_of_human_social_networking_abilities/, published on 19.12.2002, retrieved on 09.06.2009

16 http://shirky.com/writings/community_scale.html, published on 06.04.2002, retrieved on 09.06.2009

Review on social software for music

Application	Description	Examples
Social networks	platforms that allow users to leverage personal connections	facebook, myspace, linkedin, orkut
RSS	stands for <i>Really Simple Syndication</i> ; it is a XML standard that lets users collect and read content feeds	feedburner, bloglines, newsgator, pluck
Open source software	public available software that can be copied or modified without payment	openoffice.org, apache, mySQL
Blogs	on-line diaries of text, photos and/or other media	blogger, wordpress, fotolog.com, MSN spaces
Search engines	services that find web content based on user-specific criterias	google, technorati, yahoo, altavista
User review portals	web portals that allow users to search for peer reviews on a product or service	ripadvisor, cnet.com, reviewcentre
P2P file sharing	stands for <i>Peer to Peer</i> ; sharing media files over a network powered by users who act as both client and server	kazaa, emule, bittorrent, gnutella
C2C eCommerce	stands for <i>Consumer to Consumer</i> ; buying and selling among consumers via Internet	ebay, amazon, craigslist
Comparison shopping sites	sites that allow consumers to compare products or services	froogle, shopzilla, pricegrabber
Podcasts	on-line audio or video that users can download to a device	juice, podshow, iTunes
Wikis & collaborative software	shared publishing software or site that allows users to edit content	wikipedia, basecamp, google docs, dokuwiki, gobby
Social bookmarking	metadata assigned to items like photos or web pages to facilitate searching and sharing	del.icio.us, digg, reddit
Social cataloguing	a web application designed to help users to catalogue things owned or otherwise of interest to them	flixster, LibraryThing, shelfari, CiteULike
Instant messaging & videoconference	a form of real-time communication between two or more people based on typed text	skype, google talk, MSN messenger, AOL messenger
Internet forums & Message boards	an online discussion site, as the modern equivalent of a traditional bulletin board	phpBB, 4chan, something awful, fusionBB, TinyPortal
Social on-line storage	allow users to collaboratively create and store file archives	Gdrive, Wuala, Adrive, FEUUpload
Virtual worlds	services where it is possible to meet and interact with other people in a virtual environment reminiscent of the real world	second life, world of warcraft, sims on-line, webkinz

Table 3.1: Applications of social software

3.2 Social software for music

Powered by the changes in the music industry and with this revolution in the use of social software (often referred to as “social web” and “web 2.0”), a phenomena called “social music” arose, with the launch of Last.fm in 2001 and MySpace Music in 2003. Until then, we were only consumers of music, rather than participants, but this all changed with the integration of widgets in the MySpace profiles, allowing its members to embed a widget with an MP3 of their favourite band on their profile. This was the beginning of the social identity of music.

With this new era of digital music, easily and immediately available, and its massification through an immense audience, social media companies were able to innovate and build communities where music fans could communicate with each other, share their preferences and upload content, creating a bi-directional communication, in opposition to the previous one-way model.

Here, we will make an analysis of the novel situation of the music industry as well as an analysis and comparison of some of the most popular platforms of social music.

3.2.1 The new music industry

According to a 2008 study by the Pew Internet Project¹⁷ (Table 3.2), more and more people are using digital music, with a bigger incidence in the young people age group (18-35 years gap), where digital music already accounts for a little over 50% of the music consumers.

	Ages 18-35	Ages 36-50	Age 51+
All CDs	43	65	77
Most are CDs	26	20	13
Most are digital files	12	7	3
All are digital files	11	3	2
About half-half	4	3	1
<i>Number of cases</i>	148	186	286

Table 3.2: Demographics of an on-line music buyer (data in %)

An article in CNN money¹⁸ presents a study from the tracking firm Nielsen Soundscan referring a 14% album sales decrease to 428.4 million units on the United States market (regarding the 52-week period of 2008), following a 15% drop in 2007. On the other hand, digital sales increased by 27%, to a record number of 1.07 billion units, albeit being a slower growth than in the previous year (2007), where the growth was of 45%. Overall, music sales (including albums, singles, music video and digital tracks) have increased 10.5% to 1.5 billion units, after a 14% growth in 2007 and 19% in 2006.

¹⁷ Pew Internet Project at <http://www.pewinternet.org>; Source of data: <http://blogs.zdnet.com/ITFacts/?p=14531&tag=rbxccnbzd1>, published in 01.06.2008 and retrieved in 02.06.2009

¹⁸ Source: <http://money.cnn.com/2009/01/01/news/companies/music.reut/index.htm>, published in 01.01.2009 and retrieved in 02.06.2009




Figure 3.1: Online music categories distribution in 2008, by Compete²³

In the United Kingdom, the International Federation of the Phonographic Industry (IFPI) presented a report¹⁹ where it is mentioned that in the first half of the 2008 there was an increase of digital music sales of about 45%, with a total of 110 million single tracks being legally downloaded, an increase of 42% compared to the previous year. Still according to this study, the digital album sales now account for 7.7% of the album market.

More data from IFPI²⁰ refer that in Asia, the total sales increased 1% in 2008 (representing about 4.77 billion US dollars) and digital music increased 26.1% (in a total of 1.06 billion US dollars), albeit physical sales decreasing 4.9% (an estimate of 3.6 billion US dollars). In the Latin America market, the sales decreased 4.7% to 518.6 million US dollars, with physical sales declining 10.3% (430.3 million US dollars) and digital music increasing 46.6% (62.6 million US dollars).

The responsible for this change in the music business are the on-line music stores, such as iTunes, that as of June 2008 had already passed the 5 million song downloads, from their 8 million songs catalogue²¹, representing over 70% of the digital downloads market share, way ahead, of the remaining competitors, Wal-Mart, Best Buy, Amazon, and Target²².

It is also important to analyse the impact of on-line streaming radios in the music industry, as it is already one of the most-used forms of accessing to on-line music (view Figure 3.1, based on Compete's²³ data on more than 1,000 music sites, showing the most popular options for listening to music online).

19 Source: <http://www.guardian.co.uk/music/2009/jan/17/music-piracy>, published 17.01.2009; retrieved 02.06.2009

20 Source: http://www.billboard.biz/bbbiz/content_display/industry/e3icbbfd34476c899010789c0bf31184e1, published 22.04.2009; retrieved 02.06.2009

21 Source: <http://www.crn.com/retail/208700670;jsessionid=ZED3HIAVQS5K4QSNLPCXHSCJUNN2JVN>, published 19.06.2008; retrieved 02.06.2009

22 Source: <http://www.usatoday.com/printedition/money/20070517/amazon17.art.htm>, published 17.05.2007; retrieved 02.06.2009

Top 20 Streaming Music Sites
(Ranked by US-Based Unique Visitors, Mar '08) 

Site	Rank	March, 2008 Unique Visitors	% Change vs. March, 2007
imeem.com	1	10,268,904	58%
music.yahoo.com	2	9,658,665	-14%
music.aol.com	3	4,631,468	-30%
music.myspace.com	4	4,511,408	-3%
projectplaylist.com	5	4,186,025	-23%
music.msn.com	6	2,843,152	-12%
HYPEM1500*	7	2,344,620	148%
pandora.com	8	1,972,788	126%
like.com	9	1,420,983	149%
last.fm	10	1,401,093	8%
jango.com	11	665,791	New
live365.com	12	476,013	-14%
qloud.com	13	408,543	164635%
mog.com	14	390,228	432%
reverbnation.com	15	238,443	808%
hypem.com	16	127,328	126%
deezer.com	17	92,789	New
songza.com	18	66,352	New
seeqpod.com	19	65,224	2961%
muxtape.com	20	52,118	New

Figure 3.2: Top 20 streaming music sites in 2008, by Compete²³

According to a study by The NPD Group²⁴, the awareness and usage of Pandora doubled year by year to 18% of the Internet users and the usage of Last.fm increased from 15% in the last quarter of 2007 to 19% in the same period of 2008. In an interview²⁵, MySpace claim to have about 5 million bands who have made their music available on-line, by uploading it to their profiles, more than 100 million playlists and over one billion music streams.

Again referring to the study by Compete²³, in Figure 3.2 we see the 20 most used streaming sites in 2008²⁶, as well as the number of unique visitors they had.

About half of the American teenagers engage with music on social networks, representing an increase of 37% since the previous year; among college-age Internet users, the percentage increased from 30% in 2007 to 41% in 2008.

²³ Source: <http://blog.compete.com/2008/05/12/music-streaming-download-imeem-yahoo/>, published in 12.05.2008, retrieved in 02.06.2009

²⁴ Source: <http://finance.yahoo.com/news/The-NPD-Group-Continued-CD-bw-14661904.html>, published in 17.03.2009 and retrieved in 02.06.2009

²⁵ Source: http://news.cnet.com/8301-13577_3-10194538-36.html, published in 11.03.2009 and retrieved in 02.06.2009

²⁶ Ranked by US-based unique visitors, in March 2008. The Hypem1500 presents aggregate traffic to the more than 1500 music blogs tracked by the Hype Machine, a popular blog aggregator. The Hype Machine (<http://www.hypem.com>) itself is ranked separately from its network at #16.

However, online audio is not the only place where social media has made a big impact on music - mobile platforms and applications have also transformed the listening experience. Some of these platforms have now iPhone and Android applications that make it possible to stream music from almost anywhere, opening up the world's music libraries for anyone, anywhere.

Ringtones are also a major new focus of this industry. And although the purchases of the top 100 mastertone ringers decreased 33%, still 43.8 million units were sold, according to the previously mentioned CNN Money article¹⁸.

Of course all these changes brought also some new challenges and issues to the music industry, being one of the most polemic and important issues the amount of illegal downloads and unauthorized copies that are shared by music fans and internet users.

About this, the British Music Rights²⁷ has presented some results, as a result of a 2008 study made by the University Of Hertfordshire²⁸, where about 63% of the respondents admit to illegally download music from peer-to-peer file-sharing networks, although 80% of these also admit they are interested in a legal file-sharing service. Interestingly, looking at the MP3 collection of the respondents, (with an average of 1.770 tracks), it is possible to observe that 52% of the tracks have been paid for (an average of 928,88 tracks per person).

Still, the IFPI report¹⁹ estimates that a roughly 95% of the music available on-line is being downloaded illegally (only in 2008 the estimates point to about 40 billion files illegally shared), representing a 180 million pounds loss for the record industry in 2008, which might scale up to 1 billion pounds in 2012 if nothing is to be done.

3.2.2 Existing platforms

In order to better understand where we stand in terms of the existing social software (or platforms) for music, we will briefly describe 33 platforms that were under analysis, in the beginning of April 2009. These platforms were chosen due to their dimension and popularity, but also using a more casuistic approach, after a period of search and discovery, in virtue of the features they offered.

The list of platforms is presented here, with a more extended description as well as their URI (Uniform Resource Identifier) being presented in Annex B: Platform analysis.

- **Bebo** – a social networking website founded in 2005
- **Blip.fm** – a music recommendation and streaming site, launched in 2008
- **ContraStream** – an independent music community launched in 2007
- **Deezer** – a music streaming service launched in 2006
- **Dimvision's musicmap** – a flow-chart of similar music
- **Fairtilizer** – an on-line music community, based in Switzerland

²⁷ British Music Rights at <http://www.bmr.org>

²⁸ Published in 16.06.2008 and available at <http://www.ukmusic.org/page/press-release-29>

Review on social software for music

- **Finetune** – a music discovery service launched in 2006
- **FIQL** – a music-centred social network founded in 2004
- **Grooveshark** – an on-line music search engine, music streaming service and music recommendation web software application, launched in 2007
- **iJigg** – a social music site, launched in 2007
- **iLike** – a music recommendation service, launched in 2006
- **iMeem** – an on-line social music network, launched in 2004
- **Jango** – an on-line radio and social music site, launched in 2007
- **Lala** – a music community, launched in 2006
- **Last.fm** – an internet radio, and social music website, founded in 2002
- **LivePlasma** – a music and movie visual mapping website, launched in 2005
- **Maestro.fm** – a digital music social networking service, launched in 2008
- **MOG** – a blog-centric platform about music, launched in 2005
- **Mufin** – an automated music recommendation engine and music platform, launched in 2008
- **Musicoverly** – an internet-based radio station and music map, launched in 2006
- **Myspace Music** – a music community website, launched in 2008
- **MyStrands** – a desktop application and music recommender, created in 2003
- **Pandora** – an internet radio and music recommendation website, founded in 2000
- **Project Playlist** – an on-line music community, launched in 2008
- **PureVolume** – a music community website, launched in 2003
- **Radio.blog.club** – a playlist-based music service, launched in 2007
- **ReverbNation** – an internet radio, and social music website, launched in 2006
- **The Hype Machine** – a social music site, launched in 2008
- **Thesixtyone** – a social music site, launched in 2008
- **Tuneglue's audiomap** – an on-line music mapping tool
- **Turnitup Media** – and on-line store and social music website, launched in 2008
- **Wayango** – a music discovery website, launched in 2008

3.2.3 Comparison and analysis

With the analysis made to the several social music platforms, 47 features were analysed on 32 platforms. These features were grouped in 11 different clusters, as shown in Table 3.3.

classification system	tags, genre, mood, historic, rating Classification may be done by using tags (user defined), genre classification (static list), mood (e.g. happy, calm, agitated), historic (e.g. decade, release year) and ratings;
collaborative/integration tools	wiki, software (standalone application), share/embed, mobile integration, open API In this cluster the tools listed allow the user to work collaboratively or to integrate the platform with other tools;
events	calendar, georeferencing, venues If the platform has support for events, it may also have an event calendar, information about the venues and a georeferencing system;
media	video/on-line TV, artist radio, user playlist, photos The media cluster checks whether the platform has support for videos and/or on-line TV, photos (from users and/or artists) and on-line radio (with the possibility of having user-defined playlists or just artist radios);
profiles	label, artist, albums, user If the platform provides individual profiles of each of these elements (with “albums” meaning a list of the artist's albums, eventually with its track listing);
monetize	on-line store, premium accounts Two possibilities for having commercial profit with the platform: an on-line store (or connection to existing on-line stores, e.g. amazon, iTunes) and premium accounts (paid service, with more benefits for the user);
monitoring tools	user abuse/flag content, spam detection, fraud/illegal content detection Here are several monitoring tools, being “ user abuse/flag content” a general report form and the other two features more specific monitoring/reporting tools (automatic or user defined);
other content	lyrics/tabs, news, gear/equipment Other types of content that can be found;
recommendation system	rankings/tops, events, users, recommended artists/songs, graphic visualization The recommendation systems can range from a simple ranking (top, chart) to an automated recommendation of other users, artists/songs or events and even a graphic visualization of these recommendations/relationships (e.g. graph visualization);
reputation system	A reputation system allows the more influential members to gain a bigger status within the community;
social tools	blog/journal, groups/forums, comments, messages, feeds RSS/atom, profile customization, activity stream, play historic, microblogging, follow/become fan, on-line chat These are generic tools that can be found in social software.

Table 3.3: Clusters of features for analysis

The goal of this analysis was to compare the different platforms and evaluate what was their focus, by identifying which features they offer. The results of the individual platform analysis made may be found in Annex B: Platform analysis.

From the analysis made and comparing the different platforms with each other, it is possible to verify that the more complete platforms are Last.fm and iMeem (both providing over 70% of the analysed features). They offer similar services to users and their communities' size is

not very different (with Imeem taking the lead with about 24 million registered users against 21 million of Last.fm).

Interesting also is to notice that although being a much smaller platform, PalcoPrincipal is ranked 3rd in this analysis, competing directly with other platforms such as iLike and MyStrands.

Please refer to Table 3.4 for more information on this analysis. The “features” column represents the percentage of features the platform offers overall, whereas the “cluster” column presents an average of the clusters the platform is able to reach.

Table 3.4: Rank of most complete platforms and their dimension

(the sources of the data presented may be consulted in Annex C: Platforms' dimension – data sources)

	features	cluster	# reg. users
last.fm	80,9%	86,40%	21.000.000
imeem	74,5%	72,6%	24.000.000
palcoprincipal	59,6%	55,4%	10.000
iLike	59,6%	53,7%	30.000.000
mystrands	59,6%	51,5%	49.507*
wayango	57,4%	45,9%	23.036*
bebo	53,2%	41,9%	40.000.000
thesixtyone	48,9%	49,8%	100.000
deezer	48,9%	43,4%	2.000.000
maestro.fm	46,8%	44,9%	3.000+
MOG	44,7%	47,4%	5.500.000
myspace	44,7%	31,2%	130.000.000
jango	42,6%	42,1%	30.000.000*
Lala.com	42,6%	41,5%	238.519*
purevolume	42,6%	35,6%	1390785

* unique visitors

From this analysis we can also tell that some features are very popular in all the website, whereas others are not so popular or are not offered by most of the platforms. This can be an indication of either an uninteresting feature or a potentially unexplored one.

Among the most popular features we can find – as it would be expected – the on-line radio (with user-defined playlists or simply artist radio). Regarding classification, it is mostly done using a static-defined list of genres, under which the songs and artists are classified. Many of this platforms also offer a recommendation system (some more complex than others), where they can recommend similar artists and/or songs. In almost all of these platforms it is possible to view the user and artist profile, also allowing to “follow” or “fave” them. It is also possible to

comment and share/embed content in most of them as well as buy music, mostly by connecting to an existing on-line music store, such as iTunes or Amazon.

Table 3.5 presents the number of platforms where we can find the listed features (ordered by most popular features).

popular features	
user playlist	24
online store	24
user	23
follow/become fan	23
artist radio	22
recommended artists/songs	22
comments	21

Table 3.5: Most popular features

Taking a look from a different angle now, we find that some features are not so popular or common. Although some platforms start to put a bigger effort on linking events to artists, few of them have information about the venues or a geo-referencing tool, which could help to find the events' location easily.

On the verge of the everlasting “fight” with the music industry and labels, we can now find some platforms that provide labels' profiles, where we can find its artists and licensed content. The user profile customization is not an option in various platforms, as well as the on-line chat or micro-blogging tools.

Regarding classification, it is not very common to find a classification by mood and the graphic visualization of recommendation or links between artists or songs is also very limited. Although spam and fraud detection are two big concerns of social music websites, there are few mechanisms to fight them, apart from the traditional and simple “report form”.

Finally, not many platforms make use of reputation systems for its users and by not allowing the use of collaborative tools (as in wiki software) for content production, the information about artists doesn't go much farther than a small biography, whereas more detailed information (for example regarding the artist's equipment and gear) could be also interesting.

Table 3.6 presents the number of platforms where we can find the listed features (ordered by less popular features).

less common features	
microblogging	6
georeferencing	5
reputation system	5
venues	4
label	4
premium accounts	4
graphic visualization	4
profile customization	4
mood	3
historic	3
wiki	3
spam detection	3
event recommendation	2
online chat	2
fraud/illegal content detection	1
gear/equipment	1

Table 3.6: Less common features

3.3 Music Information Research

Music Information Research (MIR) has a longer history than most people realize, with systems having been developed in the 1960's [Kassler 66]. However, only after the late 1990's has it suffered an accelerated development and gained a wider attention, due to the increasing interest in digital libraries and the potential of Internet, reaching a wide commercial audience of music enthusiasts [Lippincott 02]. Also the technological developments of latest years have exponentially increased the availability of musical contents (nearly anything related to music) in digital format, thus needing a big effort of organization and accessibility [Fingerhut 04].

MIR research encompasses a number of different approaches aimed to the development of tools for music accessing, filtering, classification, recommendation and retrieval [Orio 06], having evolved mainly from three different fields: traditional IR, musicology and music perception [Uitdenboger 00]. In general, traditional IR contributes with theories about matching algorithms as well as evaluation methods based upon recall and precision and more broadly the design of functional retrieval systems; musicology's contribution is in the form of research about the physical structure and characteristics of music, contributing ideas about musical representation, similarity and, item matching; music perception or cognition research contributes with theories about how humans perceive and remember music [Downie 03].

In Figure 3.3, [Fingerhut 04] presents a schema which intends to be an attempt to have a general and simplified overview of MIR.




Figure 3.3: A simplified overview of MIR [Fingerhut 04]

In this schema we can see the actions (red arrows) used to extract and produce the information (in yellow) and the data (external in white, stored in blue). The actors are author, composer, performer, librarian, user and rights owner. They not only produce the data, but also manipulate it to produce information which they search and of which they are the subject of information.

According to [Downie 03], music information consists of seven facets, each of which playing a variety of roles in defining the MIR domain. The facets the author refers to are the pitch, temporal, harmonic, timbral, editorial, textual, and bibliographic facets.

Over the latest years, a growing international research community has been formed, drawing upon multidisciplinary expertise from library science, information science, musicology, music theory, audio engineering, computer science, law, and business [Downie 03]. According to the Dewey classification [Mitchell 03] we can better view the true multidisciplinary of MIR, as presented also in [Fingerhut 04] (Table 3.7).

Obviously, in such a dynamic field of research, patenting discoveries and inventions becomes crucial for the success of a business. As addition information, a series of MIR patents are presented in Annex D: Patents in MIR research.

Further information on MIR research can be found at [Uitdenbogerd 00], [Downie 01] and [Futrelle 02].

Review on social software for music

000 – General • 000 Generalities & computer science • 001 Knowledge • 001.4 Research • 003 Systems • 004 Data processing • 004.5 Storage • 004.6 Interfacing and communications • 005 Computer programming • 005.4 Systems programming and programs • 005.7 Data in computer systems • 005.8 Data security • 006 Special computer methods • 006.3 Artificial intelligence • 006.4 Computer pattern recognition • 006.5 Computer sound synthesis • 020 Library and information sciences • 025 Library operations • 025.3 Bibliographic analysis and control • 025.4 Subject analysis and control • 025.5 Services to users • 025.6 Circulation 100 – Philosophy and Psychology • 150 Psychology • 152 Perception • 152.1 Sensory perception • 152.3 Movements and motor functions, facial expressions, laterality • 152.4 Emotions and feelings • 152.5 Physiological drives • 152.8 Quantitative threshold, discrimination, reaction-time studies • 153 Mental processes and intelligence • 153.7 Perceptual processes	300 Social sciences • 300 Sociology and anthropology • 302.2 Communication • 306 Culture and Institutions • 310 General Statistics • 340 Law • 341.7 Law of international cooperation • 380 Commerce • 384 Communications; Telecommunications 500 – Natural Science and Mathematics • 510 Mathematics • 516 Geometry 600 – General Technology • 620 Engineering & allied operations • 620.2 Sound and related vibrations • 621.3 Electric, electronic, magnetic, communications, computer engineering • 681 Precision instruments & other devices 700 – The Arts • 780 Music • 780.7 Education, performances • 780.92 musicians, composers • 781 General principles & musical forms • 781.1 Aesthetics, appreciation, taste • 781.2 Elements of music • 781.3 Composition • 781.4 Techniques of music • 781.5 Kinds of music • 781.8 Musical forms
--	--

Table 3.7: Dewey classification for MIR [Fingerhut 04]

3.4 Music recommendation

As mentioned before, music recommendation plays a very important role in social music software. Based on music classification and on the users' profile and activity, a recommendation system is able to suggest meaningful content to users. However, this is a complex process, which involves many fields of knowledge and skills.

We will start by formalizing the recommendation problem, analyse some of the applications of recommendation as well as the issues which are associated to it. The most common recommendation techniques will be presented, as well as some evaluation metrics.

3.4.1 Formalization of the recommendation problem

The goal of a recommendation system is to suggest new items which may be interesting to a specific user or to predict the items' utility. This calculation is based on the user's personal opinions (which the user inserts on the system by various means) and previous likings, in the system analysis of the user's behaviour or in the behaviour of other like-minded users.

It is possible to address the recommendation problem from two complementary perspectives: the prediction problem and the estimation problem.

The first refers to the estimation of an item's "likeliness" for a given user, meaning that an item which the user is more likely to enjoy will have a higher rating than other items. This measure can be computed through several techniques, as we will see further on.

According to [Sarwar 01], the prediction problem can be formalised by the following: Let $U = \{u_1, u_2, \dots, u_m\}$ be the set of all users, and let $I = \{i_1, i_2, \dots, i_n\}$ be the set of all possible items that can be recommended. Each user u_i has a list of items I_{u_i} , which the user has expressed his/her interest about (i.e. items which the user has rated himself or a rating which can be implicitly derived by the system, through an analysis of the user's behaviour). Note that $I_{u_i} \subseteq I$ and it is possible for I_{u_i} to be a null-set, $I_{u_i} = \emptyset$. Then, P_{u_a, i_j} is the predicted likeliness of item i_j for the active user u_a , such as $i_j \notin I_{u_a}$. This prediction value is within the same scale as the opinion values provided by u_a .

The recommendation problem has to do with the recommendation of "likeable" items. Once the prediction problem is solved, the recommendation problem comes down to a simple list of the highest N rated items, $I_r \subset I$, of the "likeliness" estimation, i.e., the ones with higher P_{u_a, i_j} value, which is to say the N items the active user will like the most. Note that the recommended list must be on items which are not from other user's interests, i.e., $I_r \cap I_{u_i} = \emptyset$.

We should be aware that both the spaces I of possible items and the user space U can be enormous, which can be very consuming in terms of computational resources.

In most recommender systems, the prediction function is usually represented by a rating. This rating is nothing more than a triple $\langle u, i, r \rangle$ where r is the assigned value (implicit or explicitly) by the user u to a particular item i . Normally this value comes in the form of a real number (e.g. from 0 to 1), a binary variable (e.g. like/dislike) or a value in a discrete range (e.g. from 1 to 5).

As we can see, both users and items play an important role in the recommendation problem. It is then important to describe them and understand how they interact. We will do so in the next two sub-chapters.

User description

[Jennings 07] points out four types of fans (or degrees of interest in music), identified in the UK 2006 Project Phoenix 2 study, based on the analysis of different types of fans, with ages ranging from 16 to 45 years. These four types are:

- **Savants.** For these users, music and life itself blend into one single existence. Everything for them seems to be related to music, thus their musical knowledge is through and extensive. They represent about 7% of the 16-45 age group;
- **Enthusiasts.** For this type of user, music has a very important part in their life, but it is also balanced with several other interests. They tend to have a good knowledge of music and represent 21% of the analysed group;
- **Casuals.** Music has a welcome role in their life, but there are other things that grab their attention and have a greater importance. These are about 32% of the analysed group;
- **Indifferents.** For them, a world without music would still be perfectly liveable. They represent 50% of the 16-45 group and are a predominant type of fans of the whole population.

By identifying these different groups of users, it is easy to understand that they look at music in very different ways, that will influence their musical knowledge (or will to know), their listening habits and preferences and the type of relation they have with music. Of course that in terms of recommendation this represents different profiles that need to be taken into account.

Savant users will probably look for less obvious relations in music items, allowing them to discover new things or establish connections and relationships that they hadn't noticed or thought before. They do not care so much about popular or mainstream recommendations, but prefer more customized searches and suggestions. They are clearly the most difficult fans to appease with music recommendations, as they are very exigent.

Enthusiasts appreciate a balance between popular and less popular recommendations. Since they have a good knowledge of music they seek for unknown and different items, but they also appreciate mainstream recommendation and items.

Both casual and indifferent users (representing a 72% slice of the population, according to [Jennings 07]) do not expect any complex or difficult recommendations, as they are mostly interested in mainstream items. Thus, a recommendation system which is able to detect basic relations and offer some interesting and rather obvious suggestions will fit their musical needs.

Also important for recommendation is the user's individuality (i.e. the user's personality traits, historical and cultural background. For instance some aspects that tend to correlate different users are age, gender, origin, occupation, socio-economic background, musical education, among others as mentioned in [Uitdenbogerd 02]) and integration in the community, as they seem to have a big connection to his musical preferences and profile.

Some work relating to this has been done, namely by [Lesaffre 06] on the context in music perception, where an analysis of the demographic and musical background for different users is made, concluding that music search and retrieval systems should distinguish between different categories of users. On the subjective perception of music similarity [Vignoli 05], [Sotiropoulos 07] and [Sandvold 06] present ideas for music recommendation systems based on the user's perceived similarity, by analysing different features and characteristics. Finally, regarding the involvement of the user in a community, [Kazienko 06], [Baumann 07] and [Firan 07] present different recommendation frameworks that seek to implement some social aspects of music information retrieval, using social filtering and exploring relationships among users.

Finally, an interesting field of research has to do with the user's mood, as a way to influence his preferences at a given point of time. This issue will be further discussed in section 3.5.6.

Item description

The items used in MIR may be of different types, as for example an artist's biography or genre (text), a song's tempo or duration (audio), an album's cover or artist's photo (image) or a song's video clip (video).

For all of them, it is important to identify their uses and characteristics, in order to make a proper use of them in recommendation. This recommendation is done by computing the similarities found between the items and their relation to users.

However, it is possible to observe that these different types of items may be perceived in different levels of abstraction – a tag used to classify an artist is different from his biography, which is different from a song's frequency, although they may all be represented in the form of text. To help us understanding this, [Lamere 07] and [Celma 08] present a model called the Music Information Plane (MIP), as depicted in Figure 3.4.

This model is described in two dimensions: the different types of media which characterize music contents (audio, text and image²⁹) and the level of abstraction in the information extraction process of this content.

The model proposes a layered approach, with different levels of granularity and abstraction:

- **Signal features.** This is the lower level of the model, relating to more basic features of the objects, such as the loudness or the sampling rate of an audio file;
- **Content objects.** A mid-level, regarding semantic features, such as a song's tonality (e.g. key and mode) or its rhythm;
- **Human knowledge.** The higher level of the model, relating to human understanding, which uses reasoning methods and semantic rules in order to retrieve more conceptual results, as for instance, two similar guitar solos over the same key from different songs.

²⁹ Although the author doesn't clearly state it, video plays an important part of music nowadays (e.g. music clips or live concert recordings). It is included in this model under the image category, as video is the conjugation of the other forms of media presented (audio, text and image).




Figure 3.4: The music information plane – MIP ([Lamere 07] and [Celma 08])

Nonetheless it is clear the gap between the upper layer (human knowledge) and the other two layers (content objects and signal features), as pictured in the figure. This is explained by the difficulty to bridge the gap between the techniques and technologies used to describe audio and music content and the concepts that music fans use to relate with music collections – known as “semantic gap”.

Regarding the input media types, in the horizontal axis, they range from audio files, to text (which can be in the form of lyrics, reviews, biographical information, or others) and image (e.g. album covers, music clips, etc.). These allow us to create a knowledge about music, which [Pachet 05] has classified into three categories:

- **Editorial metadata.** This consists of information manually introduced by an editor. It is mainly decided by an expert or group of expert (such as the artist itself or its label), but with the use of social music platforms, it may also be done by the community itself (some social music platforms allow the users to edit information about the artists, by providing a wiki tool).




Figure 3.5: Relationship between users and items for recommendation [Celma 08]

Here we can find simple creation and production information (such as the artist, song title, album and track number), but also biographical or classification information (artist's genre, biography, etc.).

Editorial information is not always objective, as it is sometimes difficult to agree on some content, (e.g. the music genre of an artist), as it is many times directly related to a certain taxonomy in use. In terms of music recommendation, editorial metadata is used mainly for non content-based methods for music recommenders;

- **Cultural metadata.** This is defined as the information implicitly present in huge amounts of data (such as comments and reviews, among others), mainly based on the aggregation of personal opinions. This information is usually gathered from the Internet, via weblogs and other platforms, through several Web-MIR techniques;
- **Acoustic metadata.** The information in this category refers to the content analysis of audio files, which is obtained by computing low-level signal features . This is done by describing music content according to several “music facets” (such as tempo, timbre, and rhythm), by incorporating higher-level semantic descriptors. These can be computed directly from the audio signal combining signal processing, machine learning and musical knowledge. In terms of music recommendation, acoustic metadata is used mainly for content-based methods for music recommenders;

In Figure 3.5, [Celma 08] presents a schema in which we can better understand how users and items interact in the recommendation system.

3.4.2 Applications of recommendation

Having defined the recommendation problem it is important to define what exactly can we recommend - this is commonly known as recommendation use cases. Here it is possible to make a distinction between general applications of a recommender and its specific applications to music recommendation.

We will start by describing the general aspects involving a recommender system, which have been analysed by [Herlocker 04]. There, the following usages of a recommender have been defined:

- **Annotation in context.** This case refers to the intention of filtering through structured discussion postings to decide which ones are worth reading. This would require retaining the order and context of message, and accordingly using predictions to annotate messages in context. With this, users could use the displayed predictions to decide which messages to read (or which links to follow);
- **Find good items.** This use case aims to provide a ranked list of items and an associated prediction of how much the user would like each item. Normally, the user should receive a list with some novel items, unknown to him till then, as well as some familiar items;
- **Find all good items.** The main difference of this use case from the previous is about its coverage. In this case, the false positive rate (the proportion of negative instances that were erroneously reported as being positive) should be lower, thus presenting items with a higher precision;
- **Recommend sequence.** The aim of this case is to provide the user with a sequence of items that will, as a whole, please him. In music recommendation this can be seen in the automatic generation of playlists;
- **Just browsing.** In this case, users find it pleasant to browse around the system, even if they are not willing to purchase any item. This can be regarded a learning activity or simple entertainment. In this cases, the accuracy of the recommender may be less important than the interface, the usability or the information provided;
- **Find credible recommender.** Users do not automatically trust a recommender – many of them “play around” for some time in order to evaluate the recommender's performance and to try to find any signs of bias. A user interacting with a music recommender will probably search for one of his favourite artists and check what the recommender provides him in return, as its similar artists, associated genres or generated playlists;
- **Improve profile.** Here, the users contribute ratings because they believe that by doing this, they are improving their own profile and the quality of recommendation that they will receive;

- **Express self.** For some users, more than the recommendation themselves, it is important to express their opinions. A recommender that has the ability to offer a way to interact with other users allows their self-expression. By doing this, other users can get more information about the items being recommended to them, by means of tags, comments or other tools;
- **Help others.** Although this use case is often related to the previous one, not always does it happen. Here, some users are happy to contribute ratings into the recommender, simply because they believe that the community will benefit from these contributions;
- **Influence others.** This use case is seen as the most negative of all the ones presented, as in some situations a user might want to explicitly influence the community in order to gain attention for himself, thus also influencing the recommender in an “artificial” sense.

[Herlocker 04] also refer that *Annotation in Context* and *Find Good Items* are overwhelmingly the most commonly evaluated tasks in literature, although the other generic cases are also important, even though not explored as much as the first.

Regarding the practical applications of music recommendation, [Celma 08] refers the following:

- **Artist recommendation.** This follows the user-item relationship, where the user will receive recommendations on items (artists), according to his profile and preferences. These recommendations could come as well in the form of related news, upcoming gigs, new releases, among others;
- **Neighbour recommendation.** Here, the recommendation is applied to other users, following a user-user relationship. The recommender will suggest other like-minded users, with similar tastes and preferences, from whom other recommendations can also be mined. The recommender may also compare two users and evaluate their musical compatibility (e.g. Last.fm). By doing this, the social network will also be strengthened, by connecting more users with shared interests;
- **Playlist generation.** This is also an important application, because by generating playlists of songs, this will allow the user to immediately listen and give his opinion on the recommendations he received, feeding back the system, in order to react accordingly. The playlist generation can be done in two forms: creating shuffled, random playlists or creating personalized playlists.
 - **Shuffle, random playlists.** This recommendation is done by using a song or artist as a seed (i.e. the starting point from where the recommendations will follow)
 - **Personalized playlists.** Here, the recommendation is done based on a user's profile and his like-minded neighbours. This is also known as *radio-a-la-carte* and the user can provide personalized feedback by skipping songs or demanding more or less songs of a certain type

Both these types of playlist can be generated by using tracks drawn from the user's collection (a typical use of the shuffle play function) or by using tracks drawn from the celestial jukebox (i.e. tracks from outside the user's collection), where personalized playlists may make more sense, as the random playlist is likely to have no connection whatsoever to the user itself;

We can also include in these applications the use of tag suggestions, which are a form of recommendation to help classification of content. Here, the recommender will suggest the user some tags which have been applied by other users and might be appropriate for the content the user is classifying. This topic will be briefly mentioned in section 3.5.5.

3.4.3 Issues with recommendation

As Music Information Research – and specially recommender systems – are still recent fields of knowledge, which only until recently emerged, together with the social music phenomena, there are still several issues that need to be explored and discussed. Here we will have an overview of some of them.

Feedback Loops

Both in [Lamere 09] and [Lamere 07], the authors explain how the “Harry Potter effect” (also known as “similarity vortex”) can create feedback loops on the recommender. By using collaborative filtering techniques, fans can sometimes be sucked into a set of artist (the vortex) which are so closely linked that the recommendations keep coming back to themselves, with very little discovery taking place.

As an example, in [Lamere 09] the authors present the following recommendation example: “If you enjoy Java RMI you may enjoy Harry Potter and the Sorcerers Stone”. This happens as certain items – such as the Harry Potter series of books – are extremely popular. Because of this they may be associated with unrelated items, as people that buy Harry Potter books also buy any other kind of book, thus creating inaccurate associations.

Popularity Bias / The Long Tail Issue

An effect associated to the “Harry Potter effect” is the one known as the “Long Tail” problem. The Long Tail is composed by a very few popular items (the well-known hits), and the rest (novel, unpopular or unknown items), located in the heavy tail [Anderson 06].

The popular items tend to be recommended much more often than the novel or unknown items (located in the heavy tail) and because of this they also tend to be purchased more frequently. As this happens, they will receive even more recommendations, closing the feedback loop. This is known as the “rich-get-richer” effect or “cumulative advantage” – if one item happens to be slightly more popular than another at just the right point, it will tend to become more popular still³⁰. An interesting discussion about the power of cumulative advantage versus the wisdom of crowds is presented on [Watts 07].

30 Curiously, this phenomena has some similarities with the famous “butterfly effect” of the Chaos Theory, of which the author is a particular enthusiast.

The popular items become more popular at the expense of unpopular items, diminishing the overall diversity of recommendation. Instead of helping users find novel or unknown items, the system does exactly the opposite, by recommending us the same set of popular items, leaving the unpopular or unknown items in obscurity.

According to [Celma 08b], the Long Tail of a catalogue is measured in terms of frequency distribution (e.g. views, downloads, purchases), ranked by item popularity. The model of a long tail curve can be simulated by the following function $F(x)$ [Kilikki 07], which equals to the share of total volume covered by objects up to rank x , as shown in Formula 3.1.

$$F(x) = \frac{\beta}{\left(\frac{N_{50}}{x}\right)^\alpha + 1}$$

Formula 3.1

Here, α is the factor that defines the S-shape of the function, β is the total volume share (and also describes the amount of latent demand), and N_{50} is the number of objects that cover half of the total volume, that is $F(N_{50}) = 50$. With this model it is also possible to divide the curve in three parts (head, mid and the tail), as explained in [Celma 08b].

Early Rater Bias

Also another issue related to the Popularity Bias and the Feedback Loops is the Early Rater Bias [Avery 97]. [Lamere 07] observes that the success of a song depends on the decisions of a few early-arriving individuals. These individuals will classify and influence the social influence (i.e. the relation with the community) of a song, in a way that it may influence its rise to the “head”, or the permanence in the obscurity of the “long tail”. Here, the social influence has a similar importance as the quality itself. This can be seen in the study pictured in [Watts 07].

Cold Start

Although it might seem related to the previous issue, the Cold Start problem of a recommender (also known as the “learning rate curve” or the “bottleneck problem”) must be seen from a different perspective. It happens when a new item (e.g. a new user or artist) enters the system and has not yet been rated and/or has very little content.

Because the system does not have enough information about this new item, the last will have to dedicate a bigger amount of effort on using the system by establishing relations and completing its profile before getting some reward, in the form of useful recommendations. In order to be able to recommend him to other user, when this new item is added to the collection the system should have already enough informations about him. For this to happen, the input of content information can help bridging this gap from existing items to new items, by inferring similarities among them [Schein 02].

Novelty and serendipity

The novelty of items and recommendations is an important aspect of this problem. According to [Celma 08], obvious recommendations are not generally well received by users, as they are probably already familiar to them, thus being seen as rather useless. Nonetheless, these recommendations may be of some use for new users, as they like to receive some recommendations they are already familiar with. Therefore, a balance between novel and obvious recommendations must be reached, in order to provide some credibility to the recommender, on the eyes of the user.

Associated with novelty, we also have the serendipity effect. This refers to the experience of receiving an unexpected and fortuitous item recommendation [McNee 06]. In other words, we can say that serendipity is the good fortune of making unexpected and surprising discoveries, by receiving a recommendation which otherwise would not lead to this discovery. Discoveries which are serendipitous are also novel and relevant for a certain user.

Data sparsity and high dimensionality

It is rather frequent to find in social music software a have a very large number of items and users, an issue commonly known as “high dimensionality”.

Related to this, an important issue arises [Celma 08] – the coverage of the user's interaction with the system items. In other words, when we have a large dataset, the number of items with which we interact (e.g. the number of artists we view or the number of songs we listen to) tends to be a small percentage of it. The coverage of a recommender measures the amount of items in the data set which the system can make recommendations for or form predictions. A low coverage might be less valuable to users, as it limits the space of possible recommendation items. This is an important feature for the “Find all good items use case” application – mentioned in the previous section – and clearly affects the novelty and serendipity factors already mentioned. The coverage is also called data sparsity and is an inherent property of the dataset.

Scale

Associated with the previous issue, there is also the problematic of the scale in music recommenders. The size of the dataset is an important factor in the efficiency of the recommender and although one can be tempted to immediately assume that a larger dataset will provide better or more accurate recommendations, this is not necessarily true.

As an example, in Table 3.8 we can see a comparison of the dimension of 4 different recommenders: Netflix³¹, Yahoo!Music³², Amazon³³ and Last.fm [Lamere 07].

31 <http://www.netflix.com>

32 <http://music.yahoo.com>

33 <http://www.amazon.com>

	Netflix	Amazon	Yahoo!Music	Last.fm
# users	5 million	30 million	25 million	20 million
# items	50 thousand	1 million	600 million mins/month	100 million (tracks)
# ratings	2 million/day	-	7 billion	2 million tags/month

Table 3.8: Comparison of the scale of 4 recommender systems: Netflix, Yahoo!Music, Amazon and Last.fm

Gray sheep

Also related with the issue of data sparsity and high dimensionality is the gray sheep issue [Claypool 99]. This issue relates to users who have atypical tastes, differing from the norm, as a gray sheep in a flock of white-sheep. Because these users don't have many neighbours (similar users), this could lead to poor recommendations.

Temporal effects / inertia

Another issue of recommendation systems has to do with temporal effects [Celma 08] - normally older items will be less relevant than new ones, reason why the timestamp of an item (when it was added to the dataset) is an important factor for the recommendation algorithm. Also the information provided in the user's profile becomes outdated if not regularly updated, which may result in bad recommendations or predictions. This issue can also be regarded as a kind of inertia force that acts on the system: traditional charts have a built-in decay, which new “play” charts can resist to.

Trust and recommender hacking

The fact that a recommendation system works in a dynamic way and that it is influenced by users and their actions, leads to the issue of the trust in the recommender. A trust-aware recommender should be able to determine which users are reliable or trustable and which are not. In [Massa 07], the authors propose an approach designated “web of trust”, in which users express which other users they see as more or less trustable, creating a sort of reputation system among users. This is also used to alleviate the data sparsity issue. An empirical evaluation of this approach shows that by using this trusted information, both the coverage and accuracy of recommendations is improved.

On the other hand, the lack of trust in the recommender may be related to what is called “recommender hacking” under the form of profile injection attacks – these attacks involve the insertion of biased profile into the ratings database with the purpose of changing the recommender's behaviour [Williams 07]. Even a small number of malicious profiles can significantly bias the system.

Transparency

Related to the previous issue of the recommender's trust, transparency (also known as explainability) comes as a way to gain this trust and user loyalty, by providing a detailed explanation of why a certain item was recommended in a specific situation.

Traditionally, commercial music recommenders provide item recommendations based on the approach “people who liked item X, also likes item Y”, which is a rather opaque recommendation that does very little to gain the user's trust [Maillet 08]. Instead, a recommender should be able to explain why a certain recommendation was made, giving a feeling of transparency towards the user, which may lead to an improved search performance. Transparency allows users to meaningfully revise and correct their input in order to make recommendations more efficient [Sinha 02].

According to [Herlocker 00], providing explanations about recommendations to users improves their acceptance of the recommender itself, being the explanations as important as the recommendations themselves. [Tintarev 07] defend that explanations should be personalized to the user and present the following aims of explanations about recommendations:

- **Transparency.** Explain how the system works;
- **Scrutability.** Allow users to tell the system it is wrong;
- **Trust.** Increase users' confidence in the system;
- **Effectiveness.** Help users make good decisions;
- **Persuasiveness.** Convince users to try or buy;
- **Efficiency.** Help users make decisions faster;
- **Satisfaction.** Increase the ease of usability or enjoyment.

Steerability

In a traditional recommendation system, the user can't do much to guide the recommendations he receives according to what he desires [Maillet 08]. These recommendations are usually done using a base item (e.g. an artist, song, genre) as a seed, but they can not decide whether the recommender is giving too much weight to a certain aspect of the recommendation and too little to another. This interactivity with the recommendation can help the user find new and interesting content and also make the whole experience more engaging and enjoyable.

Music representation

Music representation is an intricate issue, as it is not clear how it is done. Music can be seen as the audio representation of a record or as the symbolic representation of a score. Unfortunately this symbolic representation falls short in some important music information and is not user-friendly; on the other hand “full-text” digital audio is way too complex and cannot be stored as a sound file [Lippincott 02].

Because of this, several approaches have been proposed in the literature in order to represent music in such a way that a particular piece can be distinguishable from another and vice-versa, in a similar for to a query. Only when this issue is solved adequately, can the content-based recommendations reach another level of accuracy.

Music queries

According to [Lippincott 02], the issues here may be divided in representation and user creation. The first is directly related to document representation, as content-based MIR system algorithms are designed to match query and documents representation, thus treating them as equals to traditional IR. Regarding user creation, the issue is more about how to create a query, because this depends in a certain extent on who the user of the system will be. Some queries may be done by humming a song and recording the audio input, others are done by inserting information using the keyboard. In both cases the query will need to be translated in order to match the representation of the music stored. This is, of course, a big challenging to what respects the design of matching algorithms.

3.4.4 Recommendation techniques

There are various techniques used to solve the recommendation problem. One widely used approach is collaborative filtering. Here, the system stores the interaction (implicit or explicit) between a user and the item set. The system can provide informed guesses based on the interaction that all the users have provided.

Another approach is to collect information describing the items and then, based on the user preferences, the system is able to predict which items the user will like the most. This approach is generally known as content-based filtering, as it does not rely on other users' ratings but on the description of the items.

Finally, the hybrid approach combines some of the previous approaches. We will present a short overview on each of them.

Collaborative filtering techniques

The collaborative filtering (CF) approach predicts user preferences for items, by learning past user-item relationships. The user gives feedback to the system, so that it can provide informed guesses based on the feedback that other users have provided. It relies on the "wisdom of crowds", by which it is possible to receive the typical recommendation "people who bought X also bought Y".

The core of a CF system is an item-to-item similarity matrix that is used to show how similar or dissimilar items are. Each row represents a user profile, whereas the columns are items.

The first system that implemented the collaborative filtering method was the Tapestry project at Xerox PARC [Goldberg 92]. The project coined the collaborative filtering term. Other early systems are: a music recommender named Ringo [Shardanand 94] [Shardanand 95], and Group Lens, a system for rating USENET articles [Resnick 94].

CF recommenders have a number of advantages. One of them is that they work very well for popular artists. When there is a big listeners group to a set of artists, the overlap is a good indicator of overall preference. Another advantage of CF systems is that they are fairly easy to implement. The mathematical computation involved is pretty straight forward and conceptually it is rather easy to understand.

However, this approach also has a reverse side of the coin. Scaling a CF system to work with millions of artists and billions of tracks for millions of users is an engineering challenge. The volume of data involved is overwhelmingly big and in order to process this amount of information several design decisions must be made adequately. Also another big issue of CF is the long-tail distribution in music recommendation. A CF-based recommender cannot make good recommendations for new and unpopular items, as there is no “wisdom of crowds” to support it. In order to provide collaborative-filtering recommendations, several techniques may be used. They may be generally grouped into item or user based.

Item-based similarity

Item-based collaborative filtering predicts a test user’s interest in a test item based on rating information from similar item profiles. Each item (column vector) is sorted and re-indexed according to its dis-similarity towards the test item in the user-item matrix, and, ratings from more similar items are weighted stronger. Item similarity can be computed applying:

- **Cosine-based distance**, a measure of similarity between two vectors of n dimensions, which is calculated by finding the cosine of the angle between them;
- **Dimensionality reduction techniques**, in order to deal with the sparsity problem, such as Singular Value Decomposition (SVD) or Non-negative matrix factorisation (NMF).

In order to compute text similarity, two techniques are commonly used:

- **Term frequency-inverse document frequency (TF-IDF)**, a traditional text information retrieval technique. Here, the term frequency (TF) is a measure of the frequency (or weight) of a term (or word) within a document, and the inverse document frequency (IDF) is a measure of the term's rarity across the set of documents. This weighting scheme is frequently used with the vector space model and the cosine similarity in order to determine the similarity between two documents;
- **Pointwise mutual information (PMI)**, a measure of association that estimates the semantic similarity of two terms, analysing their co-occurrence frequency. The PMI of a pair of terms x and y quantifies the discrepancy between their joint distribution probability versus their individual distribution probability, assuming independence [Turney 01].

User-based similarity

User-based approaches apply the same idea, but use similarity between users instead of items. Each user profile (row vector) is sorted by its similarity towards the test user's profile. Ratings by more similar users contribute more to predicting the test item rating. The set of similar users can be identified by employing a threshold or selecting top-N. The most common approaches to find similar users are:

- **Pearson correlation**, representing the correlation between two variables that reflects the degree to which the variables are related;
- **Adjusted cosine similarity**, since for item-based similarity, the cosine similarity does not take into account the differences in rating scale between different users. On the other hand, the adjusted cosine similarity makes use of user average rating from each co-rated pair, and copes with the limitation of cosine similarity;
- **Clustering based on stereotypes**, in which a user stereotype represents a certain kind of user who exhibits a set of specific characteristics - an abstraction of a group of similar users. Looking at statistical data gathered from a population of users, these groups can be identified either manually or using automated clustering techniques, and constructed by generalizing the most significant features of the identified groups [Montaner 03].

Content-based filtering techniques

In the content-based (CB) filtering approach, the recommender collects information describing the items and based on the user's preferences, it predicts which items the user could like. This approach does not rely on other user ratings but on the description of the items. The process of characterising the item data set can be automatic (e.g. extracting features by analysing the content), based on manual annotations by the domain experts, or even using the tags from the community of users (e.g. using those tags from the folksonomy that clearly describe the content of the items). The key component of this approach is the similarity function among the items.

Initial CB approaches have its roots in the information retrieval (IR) field. The early systems focused on the text domain, and applied techniques from IR to extract meaningful information from the text. Yet, recently have appeared some solutions that cope with more complex domains, such as music. This has been possible, partly, because the multimedia community emphasised on and improved the feature extraction and machine learning algorithms.

The CB approach presents several advantages. With these techniques, the early-rater bias issue virtually disappears. Because there is no human intervention in the process, all the items are considered (in principle) to be of equal importance and thus, not biased.

Also, when adding a new item into the collection – and computing the similarity among the rest of the items – it can be recommended without being rated by any user.

Although CB offers many possibilities, it still has several limitations. One has to do with the limitation of the features that can be (automatically) extracted from the objects. It is still difficult to extract high-level descriptors with a clear meaning for the user. Music analysis is not ready yet to accurately predict the mood but, on the other hand, it does the job well when dealing with descriptors such as harmony, rhythm and pitch, among others. Thus, an item description is not close enough to the user, but still the automatic description is useful to compute item similarity. Also, the CB methods are limited by the features that are explicitly associated with the items, meaning that subjectivity (or personal opinions) is not taken into account when recommending items to users, unlike CF methods.

Most of the distance metrics deal with numeric attributes, or single feature vectors. Some common distances to compute CB similarity are:

- **Euclidean distance**, the "ordinary" distance between two points that one would measure with a ruler, which can be proven by repeated application of the Pythagorean theorem. By using this formula as distance, the Euclidean space becomes a metric space. It is also called distance as the crow flies or 2-norm distance;
- **Manhattan distance**, in which the usual metric of Euclidean geometry is replaced by a new metric, where the distance between two points is the sum of the (absolute) differences of their coordinates. The Manhattan distance is also known as Manhattan length, taxicab metric, rectilinear distance, city block distance, L_1 distance or L^1 norm;
- **Chebychev distance**, is a metric defined on a vector space where the distance between two vectors is the greatest of their differences along any coordinate dimension. It is named after Pafnuty Chebyshev and is also known as Tchebychev distance or Maximum value distance;
- **Cosine distance for vectors**, as previously defined in CF's item-based similarity;
- **Mahalanobis distance**, is based on correlations between variables by which different patterns can be identified and analysed. It is a useful way of determining similarity of an unknown sample set to a known one. It differs from Euclidean distance in that it takes into account the correlations of the data set and is scale-invariant, i.e. not dependent on the scale of measurements [Mahalanobis 36].

Both Euclidean, Manhattan and Chebychev distance are assuming that the attributes are orthogonal. The Mahalanobis distance is more robust to the dependencies among attributes, as it uses a covariance matrix.

In some cases, items are not modelled with a single feature vector, but using a bag-of-vectors, a time series, or a probability distribution over the feature space. Yet, similarity measures are not always objective. In some domains, similarity is very context-dependent. Actually, the subjective part is a big factor of the measure, and these measures do not take this into account.

Hybrid techniques

The main purpose of a hybrid method is to achieve a better prediction by combining some of the previous stand-alone approaches. Most commonly, collaborative filtering is combined with other techniques. There are different methods to integrate different approaches into a hybrid recommender. [Burke 02] defines the following methods:

- **Weighted.** A hybrid method that combines the output of separate approaches using, for instance, a linear combination of the scores of each recommendation technique;
- **Switching.** The system uses some criterion to switch between recommendation techniques. One possible solution is that the system uses a technique, and if the results are not confident enough, it switches to another technique to improve the recommendation process;
- **Mixed.** In this approach, the recommender does not combine but expand the description of the data sets by taking into account the users' ratings and the description of the items. The new prediction function has to cope with both types of descriptions;
- **Cascade.** The cascade involves a step by step process. In this case, a recommendation technique is applied first, producing a coarse ranking of items. Then, a second technique refines or re-rank results obtained in the first step..

A hybrid method can alleviate some of the drawbacks that suffer a single technique.

3.4.5 Evaluation metrics

In order to measure the quality of the items recommended to a given user and to measure how good the music collection is, evaluation metrics are used. According to [Celma 08], the evaluation metrics (or strategies) may be divided into three different groups, according to their focus:

- **System-centric.** Here, the metrics are used to evaluate the system's capability to provide accurate predictions based on the actual values which users have already assigned. This has been used particularly in collaborative filtering techniques with explicit feedback, such as rating;
- **Network-centric.** This type of metrics is used for measuring the topology of the item's network similarity. It uses metrics from complex network analysis (CNA);
- **User-centric.** The evaluation metrics focus on the user's perceived quality and usefulness of the recommendations.

Regarding system-centric metrics, one can find several approaches in the literature:

- **Predictive accuracy.** This is used to measure the deviation between the predicted ratings and the actual user rating. The result is the average over the deviations. They are particularly important for evaluating tasks in which the predicting rating will be displayed to the user such as Annotation in Context. Some commonly used techniques are the Mean Absolute Error and the Root Mean Square Error;

- **Decision-based methods.** They measure the frequency with which a recommender system makes correct or incorrect decisions about whether an item is good or not, i.e. the selection of high-quality items from the whole list (the top-N recommendations for a given user). These metrics are also known as Classification Accuracy Metrics and some of the techniques used are the Mean Average Precision, Recall (also known as sensitivity, true positive rate - TPR, or hit-rate), F-measure, and Receiver Operating Characteristic – ROC. The main drawback of decision-based metrics is that do not take into account the ranking of the recommended items. Thus, item at top-1 has the same relevance as an item at top-20. To avoid this problem, one can use rank-based metrics;
- **Rank based metrics.** These metrics measure the ranking quality of the recommendation list, by combining the hit and its position in the list. Ideally, the top items should be considered more relevant than the items in the bottom of the recommendation list. Unlike classification metrics, ranking metrics are more adequate to evaluate algorithms that will be used to present ranked recommendation lists to the user, in domains where the user's preferences in recommendations are non-binary. They might be overly sensitive for domains where the user just wants an item that is “good enough” (binary preferences), as the user won't be concerned about the ordering of items beyond the binary classification (good / not good). The rank-based metrics are also known as prediction-rating correlation and some examples of techniques used here are the Spearman's ρ and Kendall's τ [Herlocker 04].

3.5 Music classification

Content classification has gained importance over the last few years, being a component of the on-line publishing process, allowing editors to extend the reference labels, associations and search keywords by which any type of content can be found.

In fact, the use of keywords pre-dates the internet, being brought to the first websites as a way for publishers to help users find relevant content, until in 2003 the social bookmarking website delicious³⁴ started to use tags for users to classify their bookmarks (so that they could easily find them when accessing later on) and for browsing through the bookmarks of other users. Also Flickr³⁵, an image hosting website, allowed its users to use tags to classify their photos, making them easily and highly searchable.

Following the success of these websites, the concept of tag became very popular and other social websites also decided to use it, exploring even more possibilities with it. This new form of keyword classification allowed publishers to provide reference classification terms for nearly any digital content that was published on-line (e.g. news, blogs, articles, photos, presentations, videos), regardless of the content-management system used.

³⁴ <http://delicious.com>

³⁵ <http://www.flickr.com>

Nowadays, tagging is a rather big phenomena/habit of Internet users, as according to [Weinberger 07], about 28% of them have tagged or categorized online content, and still 7% does it everyday.

However, classification of content is no longer an activity done by humans alone. One of the biggest challenges in this area is to create an effective system of automatically classifying content, using computer-aided techniques.

According to [Turnbull 08], there are five different approaches to collect music annotation: harvesting social tags, deploying annotation games, conducting an expert survey, mining web documents and autotagging audio content.

The author presents a series of strengths and weaknesses for each of this approaches, presented in Table 3.9.

Approach	Strengths	Weaknesses
Social Tags	collective wisdom of crowds unlimited vocabulary provides social context	create & maintain popular social website ad-hoc annotation behaviour, weak labelling sparse/missing in long-tail
Annotation Games	collective wisdom of crowds entertaining incentives produce high-quality annotations fast paced for rapid data collection	“gaming” the system difficult to create viral gaming experience listening to short-clips, rather than entire songs
Expert Survey	custom-tailored vocabulary high-quality annotations strong labelling	small, predetermined vocabulary human-labour intensive time consuming approach lacks scalability
Web Documents	large, publicly-available corpus of relevant documents no direct human involvement provides social context	noisy annotations due to text-mining sparse/missing in long-tail weak labelling
Autotagging	not affected by cold-start problem no direct human involvement strong labelling	computationally intensive limited by training data based solely on audio content

Table 3.9: Strengths and weaknesses of tag-based music annotation approaches [Turnbull 08]

We will then start by having an overview over each of these different classification approaches, explaining what “tags” and “social tagging” is, as well as trying to understand what makes users tag content and what are these tags used for. We will also have an insight on some issues related to classification and finally we will give some examples of classification systems.

3.5.1 Social tagging

The term “social tagging” (sometimes called a folksonomy) represents the aggregation of individual sets of tags, a user-created bottom-up categorical structure development with an emergent thesaurus. Social tagging is done in an open environment, shared with other users, by the person consuming the information. Let us first start by understanding what a tag is.

What are tags?

Tags are short and free text labels (or keywords) applied to content, with no structure or vocabulary limits, which are typically applied by the generator (or the consumer) of the item being tagged [Lamere 08b]. Some other characteristics of tags include the fact that they are instructive, democratic, dynamic, collaborative, empowering, user-oriented, cheap and distributed. Unlike a taxonomy, it has a weak labelling, an unstructured and free-for-all vocabulary. Wikipedia³⁶ presents the following definition for tag:

“A tag is a non-hierarchical keyword or term assigned to a piece of information (such as an internet bookmark, digital image, or computer file). This kind of metadata helps describe an item and allows it to be found again by browsing or searching. Tags are chosen informally and personally by the item's creator or by its viewer, depending on the system. On a website in which many users tag many items, this collection of tags becomes a folksonomy.”

The value in this external tagging comes from people using their own vocabulary and adding explicit meaning, which may come from inferred understanding of the information/object. People are not so much categorizing, as providing a means to connect items (placing hooks) to provide their meaning in their own understanding. Thus, one urges for the answer to the following question: why do people tag?

Why do people tag?

Several motivations for tagging have been suggested by [Marlow 06], [Ames 07] and [Lamere 08b], whereas the last defines the motivation as follows:

- **Memory and context.** Items are tagged in order to assist the user in personal retrieval of an item or a group of items. As an example, a user can use different tags in a song as a way of better describing it;
- **Task organization.** In this situation, tags are used as a way to assist in the organization of discovery tasks. As an example, a user can tag an artist as a reminder to look better for him later on or tag a group of artists with the same tag, creating a kind of set to group them all together;

36 [http://en.wikipedia.org/wiki/Tag_\(metadata\)](http://en.wikipedia.org/wiki/Tag_(metadata)), retrieved on 04.06.2009

- **Social signalling.** This is when users tag items as a way to express their tastes to others. This happens when, for instance, an artist is tagged as “seen live”, indicating (possibly) which artists a given user has seen at a live concert;
- **Social contribution.** Items are tagged in order to create or insert knowledge, for the benefit of the community. An example of this is when a user classifies a certain artist with a given musical genre;
- **Play and competition.** In this situation, tags are used as a result of the participation in classification games, either explicitly (e.g. Major Miner, The Listen Game or Tag-a-tune) or implicitly (e.g. “tag the saddest song you know”)
- **Opinion expression.** Here, users express their personal opinion using tags. For example, an user can express his opinion by tagging songs as “music I enjoy”

Also, what makes people using tags, is its simplicity, compared to the traditional categorization, being a channel for narrative and social interaction and allowing one to draw a circle around the things s/he cares about. And although tags are personal, they can benefit everyone, in a broad sense of community.

What are tags used for?

Tags can be used to give insights to user behaviour and language usage, group items based on tags or taggers, find social groups with shared interests, generate user profiles from tagging behaviour and provide information on how items change over-time. [Lamere 08b] enumerates the following applications of tags:

1. **Search and discovery.** One of the most common uses of tags is the search and discovery of content, by mean of tag clouds or simple recommendations of similarly tagged content;
2. **Directed search.** With this technique, the search space is successively narrowed down as the user adds tags to the query, making successive filters and reducing the number of results. This technique is very popular in some social websites, such as delicious;
3. **Tag similarity and clustering.** In order to help the discovery and presentation of meaningful tags, some systems use a tag clustering system [Hassan-Montero 06]. These tag clusters can help the user in either applying tags (as in a suggestive tagging system) or in identifying alternatives that may be useful in the user’s search. Similarity is typically determined by taking into account tag co-occurrence. For this, tags that are frequently applied to the same items are considered to be more similar than tags applied to disjoint sets of items.

Some similarity measures such as Jaccard, Overlap, Dice and Cosine Distance are used to compute such similarities, as explained in [Begelman 06];

4. **Latent semantic analysis.** One way of addressing the vocabulary problem³⁷ [Fumas 87] encountered with tags is to use a technique called latent semantic analysis (LSA), as explained in [Deerwester 90]. The underlying idea is that the aggregate of all the word contexts in which a given word does and does not appear provides a set of mutual constraints that largely determines the similarity of meaning of words and sets of words to each other, defining a semantic space. By using statistical techniques, tag noise can be reduced or removed, helping to mitigate some issues related to tags, such as synonymy and polysemy [Levy 07];
5. **Tag overlap.** By using a folksonomy derived from social tags, it is possible to classify content in various ways, allowing a bigger accuracy and a better representation of the fuzziness at the boundaries of the classification. For example, when classifying an artist, it may not be very clear the musical genre he belongs to. In this case, a traditional genre hierarchy would not be accurate and limit the artist's description, but using tags it is possible to classify the artist as “87% rock” overlapping with “60% pop”;
6. **Genre hierarchy.** Social tags may be used to automatically create a genre hierarchy, by using tag frequencies and tag similarity. And although genre classification and hierarchization does not capture all the nuances of the folksonomic space, it is still a familiar technique for users to browse the genre space;
7. **Faceted navigation.** This is an alternative technique to explore the tag space, where the tags are organized into related groups (facets) [Morville 05]. With a faceted browser, a user is able to browse by picking any of the facets, which applied to music may be some of its aspects such as genre, mood, theme, instrumentation or others;
8. **Item similarity.** By calculating tag similarity (using techniques such as Pearson's correlation, Jaccard distance or simple co-occurrence), it is possible to infer item similarity, relating items that share common tags. However, some frequently used tags tend to dominate these metrics, resulting on biased results.

Regarding music, tags can be used for multiple purposes, such as building playlists (by tagging tracks), categorizing a profile and the global music catalogue, getting recommendations based on tags or summarizing a profile. Babs_05, a Last.fm user³⁸, points out a number of uses for tags in music annotation, as presented in Table 3.10.

Even though tags can represent such distinct information, it is mostly used to classify audio content according to its genre, mood and instrumentation (accounting for 77% of the tags), as it is mentioned in [Lamere 08b], regarding the distribution of the types of tags for the 500 most frequently applied tags at Last.fm.

³⁷ The vocabulary problem refers that the information relation systems rely on users typing the right query words, but often this is not the case, as users use wrong words or misspelling errors

³⁸ http://www.last.fm/group/Track-Tag+Bitches/forum/24779/_/319957/1#f4621887

Use	Examples
personal meaning	“seen live”, “totally awesome”, “I own it”
playlists with unusual names	“wherein i compile a comprehensive list of songs about zombies”
regional/locale	“russian”, “manchester”, “balkan music”
decade	“60s”, “eighties”
group-specific	“trio”, “female singer”
descriptive	“gorgeous”, “comedy”
tempo	“energetic”, “slow”
mood	“sad”, “happy”, “energetic”
instruments	“guitar”, “piano”
films	“bollywood”, “anime”
music to...	“music to fall asleep to”, “music to read to”, “music to study”
album labels	“virgin”, “sony music”
other/misc	“composers”, “Coldplay”

Table 3.10: Uses of tags in music annotation

Tag clouds

Tags are often displayed in the form of a tag cloud (or word cloud), a visualization of the word frequency in a given text (or page) as a weighted list. Usually, tags have different sizes and/or colours, as these characteristics represents the number of times that the tag has been applied to a single item (the more important a tag is, the bigger and louder it appears). Typical tag clouds have between 30 and 150 tags. The use of tag clouds was firstly made popular by Flickr⁵⁰ and Technorati³⁹, among others, and has since then evolved with slight variations.

Albeit the existence of variety of ways to implement tag clouds (with some methods being more popular than other), most clouds are primarily sorted alphabetically. What really distinguishes tagclouds, though, is their meaning rather than appearance.

There are three main types of tag cloud applications in social software. In the first type, size represents the number of times that tag has been applied to a single item. This is useful as a means of displaying metadata about an item that has been democratically 'voted' on and where precise results are not desired. A good example of this is Last.fm, which uses this method as a means of displaying the genre with which an artist or track has been tagged.

In the second, more commonly used type, size represents the number of items to which a tag has been applied, as a presentation of each tag's popularity. Examples of this type of tag cloud are used on the image-hosting service Flickr and the blog aggregator Technorati.

In the third type, tags are used as a categorization method for content items. Tags are represented in a cloud where larger tags represent the quantity of content items in that category.

³⁹ <http://technorati.com>

Tag clustering

Albeit being a great visualization paradigm, by using tagclouds it becomes a bit difficult to find more than one or two popular tags to click on and all their related tags, which tend to dominate the whole cloud. Tags are not grouped, there is too much information, so that many related tags are found scattered on the tag cloud. One way to solve this issue is to show several semantically more cohesive tag clouds, using tag clustering.

The key in building an effective exploration space seems to be able to group and show related items and to explain how the items are related. This may be achieved by using traditional clustering techniques, as the ones found in the literature:

- **Jaccard index**, a statistic used for comparing the similarity and diversity of sample sets. The Jaccard coefficient measures similarity between sample sets, and is defined as the size of the intersection divided by the size of the union of the sample sets. It is also known as the Jaccard similarity coefficient;
- **Dice's coefficient**, a similarity measure related to the Jaccard index, that gives twice the weight to agreements;
- **Matching distance**, a metric on the space of size functions. It is the observation that the information contained in a size function can be combinatorially stored in a formal series of lines and points of the plane, called respectively cornerlines and cornerpoints;
- **Overlap similarity coefficient**, a measure to compute overlap between two sets. If a set X is a subset of Y or the converse, then the overlap similarity coefficient is equal to 1 (maximum). It is also known as overlap distance;

3.5.2 Annotation games

Annotation games (also known as “games with a purpose”) are a technique used to annotate content, generating data that will be used for machine learning. They provide an entertaining and appealing interface for users that, by playing the games, produce high-quality semantic annotations. von Ahn refers to this technique as “human computation”, as these games harness the power of many human collaborators to solve problems that otherwise would be hard to compute or explore, even in areas such as Artificial Intelligence.

Music annotation games were firstly presented at the 2007 ISMIR conference⁴⁰: ListenGame [Turnbull 07], MajorMiner [Mandel 07] and Tag-a-Tune [Law 07]. These games were developed following the work of von Ahn et. al of creating a suite of games for collecting semantic information about images [von Ahn 04], [von Ahn 06a] and [von Ahn 06b].

⁴⁰ ISMIR 2007 - 8th International Conference on Music Information Retrieval, September 23rd-27th 2007, in Vienna, Austria (<http://ismir2007.ismir.net>)

Listen Game [Turnbull 07] is a multi-player, web-based game that provides real-time feedback for players who actively labelling music with words. The purpose of the game is to find associations between audio content and words, by asking the players to choose the best and worst tags for describing a song. If a large group of players agree on a tag, the song becomes strongly associated with it and over time, the game collects a semantic word association for a large set of songs. Users are presented with songs and a list of tags, consisting of preselected music-relevant words, related to genre, instrumentation or emotional content and words that the users suggest while they are playing. Like in a music survey, this has the advantage of using a structured vocabulary of tags, which may be considered a strong labelling approach, as it also collects information reflecting negative semantic associations between tags and songs.

In MajorMiner [Mandel 07], the tags entered by the player are compared against the database of previously collected tags in an off-line manner, by describing 10 second clips of songs and scoring points whenever their descriptions match those of other players. The main goal of this game is to encourage users to describe music thoroughly, in an original, yet relevant manner, with the rules being designed to stimulate objective and specific judgements. Another goal of the authors is to avoid cheating or any forms of manipulate the scoring system and the data collected.

Finally, Tag-a-Tune is a two-player audio-based game, where the players listen to a song and are asked to enter tags freely, until they both enter the same tag. Players are paired randomly and anonymously from a pool of available players and are given three minutes to agree on the descriptions of as many sounds as possible. There is no communication between them, although being notified of the partner's activity. The system also displays a word which specifies the kind of description the game is looking for, from a set of eight categories: Object, Place, Action, Colour, Mood, Film Genre, Opposite (describe what the sound is not), and Freebie (unrestricted description).

This game is different from the previous two, as in Tag-a-Tune players are not limited to music clips, but to a variety of audio samples, essential for producing usable sound CAPTCHA's⁴¹. Also, players in this game must collaborate to label a given tune, unlike the other games where users were playing off-line against a database or on-line against multiple players. Finally, Tag-a-Tune allows a more discriminative audio search, as it collects comparative information about sounds.

3.5.3 Expert classification / Surveys

The classification by experts or through the realization of surveys seeks to obtain objective and accurate information about music content. Unfortunately, this is done at a very big cost of time and resources. Still, some companies – such as Pandora, All Media Guide⁴² (AMG) and Moodlogic⁴³ – have invested big amounts of money and human resources to annotate their

41 CAPTCHA stands for “Completely Automated Public Turing test to tell Computers and Humans Apart” and is a type of challenge-response test used to ensure that the respondent is a “human” and not a machine

42 <http://www.allmusic.com>

43 acquired by All Media Guide in May 23rd 2006

music databases, but because of this huge investment, these companies are not likely to share their data.

The allmusic website (previously known as All Music Guide) was created in 1995 by AMG and encompasses hundreds of music critics who review and rate albums and songs, describe related products and write artist biographies. Its database is licensed and used by some music retailers (both on-line and brick-and-mortar stores), consisting of:

- **Basic data.** e.g. names, genres, credits, copyright info, product numbers
- **Descriptive content.** e.g. styles, tones, moods, themes, nationalities
- **Relational content.** e.g. similar artists and albums, influences
- **Editorial content.** e.g. biographies, reviews, rankings

As of 2007⁴⁴, Allmusic claimed to have the largest digital archive of music, with information on over 8.5 million songs and 900 thousand individual albums. They also claimed to have around 6 million streaming sound files and 500 thousand cover images.

The Pandora website (owned by Pandora Media, Inc.) offers an automated music recommendation system and internet radio, created by the Music Genome Project in January 2000. Its founders strived to capture the essence of music at the fundamental level, by analysing over 400 attributes to describe songs and a patented mathematical algorithm⁴⁵ to organize them.

A song is represented by a vector containing about 150 “genes”. Each one of these “genes” corresponds to an attribute of the music. Different genres of music have different number of genes to describe them, normally ranging between 150 and 500. Given the vector of one or more songs, a list of similar songs is then constructed by computing their distance. To create a song's genome, a musician analyses the song for about 20 to 30 minutes, with about 10% of the songs being analysed by more than one technician to ensure better reliability.

3.5.4 Web mining

Web mining techniques, used in web-MIR research, focus on the analysis and exploitation of cultural information from the web, normally associated with artist biographies, album reviews, and song reviews. This can be done using several techniques, in order to gather music-related information: querying a standard search engine [Knees 08], monitoring MP3 blogs [Celma 06] or crawling music sites [Whitman 04]. The origins of the work with web mining may be found in [Whitman 02], [Baumann 03] and [Whitman 03].

Most of the web mining systems developed use a set of documents associated with a song and convert them into a single document vector (for example using the “term frequency–inverse document frequency”, or TF-IDF, weight representation).

This representation using a vector space (known as Vector Space Model approach, or VSM) is useful for tasks like calculating music similarity [Whitman 02] and indexing content

⁴⁴ <http://www.metrotimes.com/editorial/story.asp?id=10087>, published on 24.01.2007 and retrieved on 06.06.2009

⁴⁵ US Patent number 7003515, issued on Feb 21, 2006. Inventors: William T. Glaser, Timothy B. Westergren, Jeffrey P. Stearns, Jonathan M. Kraft. Assignees: Pandora Media, Inc.

for a text-based music retrieval system [Knees 07]. An alternative to this approach is using a document-centred technique known as “rank-based relevance scoring” (or RRS) [Knees 08], which attempts to retrieve music pieces relevant to arbitrary natural language queries, improving the quality of the resulting document rankings.

In VSM, after obtaining a term weight vector for each track in the music collection, natural language queries to the system are processed by adding the constraint music to the query and sending it to the search engine. For instance, when querying a search engine for artist-related info we can use the artist's name and, in order to limit the size of the results dataset, add some keywords such as “music” and “review”, retaining only the terms that relate to its music and associating weights to these terms. From the top-ranked web pages, a query vector is constructed in the feature space. This query vector can then be compared to the music pieces in the collection by calculating cosine distances. Based on the distances, a relevance ranking is obtained.

On the other hand, RRS proposes to directly use the Web content that has been retrieved in the data acquisition step, instead of using the search engine, as in VSM. This is done by creating an off-line index of all pages, allowing us to apply an alternative relevance ranking method (as all the indexed documents are at least relevant to one of the music pieces in the archive) and exploit their relations. As we know for which content the documents have been retrieved, it is possible to simply create a relevant set of music pieces to the query, by gathering all the pieces associated with at least one of the returned documents. It is also possible to exploit the ranking information of the returned documents to introduce a very simple – yet effective – relevance scoring function.

In order to describe and categorize artists into categories using web data, [Geleijnse 06] propose three different techniques, based on co-occurrence analysis: Pagecount-based mapping (PCM), Pattern-based mapping (PM) and Document-based mapping (DM).

PCM uses the total number of hits retrieved by the search engine; this can be done, for instance, in order to classify an artist according to its genre, analysing the total number of hits of the query combining the artist name and the genre we want to classify him into. The authors also provide a normalised approach of this technique, in order to avoid the problem of some terms appearing more often than others.

Using Pattern-based Mapping (PM), the occurrences of combination of terms in phrases on the web are analysed. For doing this, some English phrase patterns are used, which may be composed manually or automatically, using a train set. As an example, the search pattern could be “(artist) sounds like (genre)” or “(artist) is one of the best (genre) artists”.

With Document-based mapping (DM), the content of the top-K pages returned by the search engine are analysed, with the algorithm downloading the most relevant pages, according to the query executed, and counting the related terms found in the first k pages.

3.5.5 Automatic content-based classification

Automatic content-based classification is a method that, unlike the previous ones, does not require music content to be annotated by humans. By automatically analysing audio content, these methods try to classify music, mostly focusing in genre, emotion, rhythm and instrumentation. In order to achieve this, it is necessary to analyse audio files and to create a profile for each of them. This is commonly known as an acoustic fingerprint.

Acoustic fingerprinting

According to the ISO-IEC TR 21000-11's definition of fingerprinting [ISO/IEC 04]:

"Fingerprinting is the term used for a type of pattern-recognition techniques when applied to identifying content and associating information with content, albeit without modifying the content. It works by extracting characteristics of a piece of audio content and storing them in a database. When the technology is presented with an unidentified piece of audio content, characteristics of that piece are calculated and matched against those stored in the database."

Some of the common characteristics analysed included in audio fingerprints include average zero crossing rate, estimated tempo (in BPM), average spectrum, spectral flatness, prominent tones across a set of bands, and bandwidth.

A robust acoustic fingerprint algorithm must take into account the perceptual characteristics of the audio. If two files sound alike to the human ear, their acoustic fingerprints should be equal or very similar, even if their binary representations are quite different (Therefore, acoustic fingerprints are not really fingerprints in the strict sense of the term --- which must be sensitive to any small changes in the data).

Most Audio compression techniques (e.g. MP3, WMA, Vorbis) will make radical changes to the binary encoding of an audio file without affecting the way it sounds. A robust acoustic fingerprint will allow a recording to be identified after it has gone through such compression, even if the audio quality has been reduced significantly. On the other hand, a good acoustic fingerprint algorithm must be able to identify a particular master recording among all the productions of an artist or group.

The audio features extraction can be done using several techniques, such as electronic music distribution (EMD), mel-frequency cepstral coefficients (MFCC), log spectrogram, autocorrelation or multidimensional scaling.

Applications

Several authors have been exploring this area, using different approaches, with encouraging developments over recent years. One of the first works in this area was done in [Tzanetakis 02], where the authors describe some algorithms for automatic genre categorization of audio signals. Here they also propose a set of features for representing texture and instrumentation and define a set of features for representing rhythmic structure and strength. The performance of those feature sets is evaluated by training statistical pattern recognition classifiers, using real world audio collections.

In [Eck 07], the authors use a discriminative approach by using a set of boosted classifiers, mapping audio features onto social tags collected from the web. The resulting tags provide information about music that has not been tagged (or has few tags), helping to avoid the cold-start problem and smoothing the tag space from which similarities and recommendations are extracted.

Related to the previous work, [Eck 07b] investigate the automatic generation of tags with properties similar to those generated by social taggers and propose a method for predicting social tags using audio features and supervised learning. This is done by using a machine learning algorithm (that uses the meta-learning algorithm AdaBoost [Freund 96]) that takes as input acoustic features and predicts social tags mined from the web.

Another approach is used in [Sordo 07], where a content-based measure of music similarity is used to propagate tags from annotated songs to similar songs that have not been annotated. This non-parametric approach is also applied to mood labels, with interesting and encouraging results.

Finally, in [Turnbull 08b], a supervised multi-class labelling (SML) model is used to automatically annotate songs with a diverse set of tags, based on audio content analysis. This SML model is parametrized by a Gaussian Mixture Model (GMM) distribution over an audio feature space for each tag in the vocabulary, thus evaluating their likelihood. The result of this analysis is a vector of probabilities that, after being normalized, may be interpreted as the parameters of a multinomial distribution over the tag vocabulary, representing a compact and interpretable index for a song, where the large parameter values correspond to the most likely tags.

3.5.6 Mood classification

Different forms of classifying music have appeared over the latest years, as an alternative to the tradition genre classification. One of the most popular and, perhaps, most interesting, is the mood classification. This classification takes advantage of the emotional bond that music has with its consumers, being able to fit and influence a person's state of mind.

In order to create mood models, some ontologies have been created. Initially [Hevner 36] and [Farnsworth 58] ontologies proposed the creation of eight and ten mood clusters, respectively. Nowadays, ontologies have expanded, reaching many more clusters, as in [AMG 09], a large edited music information database, with a total of 179 distinct (but not unrelated) mood labels.

An alternative approach to the modelling of human emotions views affect as a combination of orthogonal continuous sub-features. The most popular such representation of musical affect is Thayer's two-dimensional valence-arousal space [Thayer 89], which itself is derived from Russell's general model of human emotions (pleasant-unpleasant vs. arousal-sleep) [Russell 80]. Thayer's model decomposes emotion in music according to two principal dimensions: valence (positive vs. negative; e.g., happy vs. sad) and arousal (high- vs. low-energy; e.g., energetic vs. calm).

According to this model, music can be broadly categorized into one of four quadrants: high valence and arousal (joy, exuberance), high valence and low arousal (contentment), low valence and high arousal (anger), and low valence and arousal (depression). But the model also views the axes as continuous features, allowing for an unlimited combination of overall moods.

A few existing techniques for automatically classifying audio content into mood categories exist in the literature: Li and Ogihara [Li 03] used acoustic features related to timbre, rhythm, and pitch to train Support Vector Machines (SVMs) to classify music into 13 mood categories derived from Farnsworth's emotion groupings, achieving an accuracy of about 45%, with 50% of the database used for training and testing, respectively. Later, in [Li 06], they extended this work, considering 3 bipolar adjective pairs: cheerful/depressing, relaxing/exciting and comforting/disturbing. Here, the authors reported an accuracy of around 60%.

The same 13 clusters as in [Li 03] were used in [Wieczorkowska 06], where the authors modified the k Nearest Neighbours algorithm in order to handle multi label data directly. However, they found that the predictive performance was low.

Lu, Liu, and Zhang [Lu 06] pursued mood detection and tracking (following dynamic mood changes during a song) using a variety of acoustic features related to intensity, timbre, and rhythm. Their classifier used Gaussian Mixture Models (GMMs) for Thayer's four principal mood quadrants in the valence-arousal representation. Their system achieved an accuracy of about 85% when trained on 75% of the clips and tested on the remaining 25%.

The same emotion classes were used in [Yang 06]. The authors experimented with two fuzzy classifiers, using the 15 features proposed in [Schubert 99]. They also experimented with a feature selection method, which improved the overall accuracy (around 78%), but they do not mention which features were selected.

In 2007, the Music Information Research Evaluation eXchange (MIREX) first included a "beta" task on audio music mood classification with 8 systems submitted. The audio clips used for this task were assigned to one of 5 mood clusters, aggregated from AMG mood labels (adjectives), and 600 30-second hand-labelled clips distributed equally among the 5 mood clusters were used in the evaluations. All participants performed reasonably well (far higher than chance) with the highest performing system – MARSYS [Tzanetakis 07] – achieving correct classification slightly over 60% of the time.

3.5.7 Issues with classification

As already pointed before, classification – and in particular, the use of tags – have several issues associated. [Turnbull 08], based on [Lamere 08b], divides these issues into two categories: scalability issues, and quality issues.

Issues regarding the *scalability* of classification systems have three major impacts: their financial costs, human involvement and computational resources. If we are to have music manually annotated by a group of experts (as in the case of Pandora and All Music), the financial costs might be enormous, as in the first a music needs an average of 45 minutes do be categorized, with a classification of up to 500 parameters, requiring a big human involvement.

According to [Lamere 07], the manual annotation of one million songs by 100 people would take up to 3 years and cost about 10 million US dollars. This clearly poses a scalability problem, as it would take about 5 years to analyse 1 year of new releases. Nonetheless, this issue might be tackled by using automatic annotation, using autotagging techniques, for instance, which could reduce the annotation cost up to one thousand US dollars per 1 million songs and only 8 hours with 100 CPU⁴⁶s. However the automatic annotation of audio content still poses many challenges and deals with huge amount of data, thus representing an increase of the computational resources needed.

Regarding *quality* issues, they relate to several aspects: cold start, strong and weak labelling, vocabulary structure and size, annotation accuracy, the source and focus of tags and their polysemy, synonymy and noise.

Cold Start

One of the biggest issues regarding classification refers to the cold-start problem, in which a song that is not annotated cannot be retrieved. This sparsity of tags for new and obscure artists is related to the popularity bias problem, in which popular items are abundantly described by tags (strongly labelled), while less popular items might not have enough tags (both in quantity and diversity) to have meaningful and stable descriptions. If the tag information is used for retrieval, the less popular items will probably be retrieved less frequently, degenerating in a retrieval starvation effect and if these tags are using in music recommendation, they end up by biasing recommendations [Sarmiento 09].

This issue has been explored by many authors, including [Schein 02], [Lamere 07], [Turnbull 08] and [Gunawardana 08].

Strong vs. Weak Labelling

As mentioned in the previous paragraph, when a song has been explicitly labelled with a number of tags (or not labelled if the tags are not relevant) we can say that the song has a strong labelling. In contrast, a song with little or none tags, or with irrelevant ones, indicates a weak labelling. This might be problematic if we intend to design a system with high recall or to collect a training dataset for supervised learning, using discriminative identifiers, as in [Eck 07] or [Whitman 04].

Vocabulary structure & size

The structure, size and extensibility of the tag vocabulary is also an important aspect of music classification. The ideal vocabulary has a large and diverse set of semantic tags, describing and characterizing music meaningfully. This tag vocabulary may be fixed or extensible (as it will be fixed if the user must choose from a given list of tags and extensible if the user is free to enter any tag he wants) and structured or unstructured (if the classification obeys to some kind of structured organization).

⁴⁶ CPU - Central Processing Unit

As an example, a survey may be regarded as a fixed and structured tag vocabulary, as the set of tags available and their grouping into coherent categories is predetermined by experts using domain knowledge. On the other hand, a social tagging community determines an extensible and unstructured vocabulary, as users may suggest any tag to describe music, with this tags not being organized in any form.

But if having a large vocabulary may allow a better classification, this high dimensionality of the datasets may also become a drawback. To illustrate this, we can refer to an example presented in [Knees 04], where a typical web crawl found over 200.000 terms for a set of 200 well-known artists. [Levy 07] affirm that such information may be noisy, as an automated retrieval system does not guarantee that all the web pages retrieved are relevant (specially if the keywords used have other meanings) and even on the relevant pages, much of its content will be useless anyway, as it does not refer directly to the information that is being looked for.

Annotation accuracy

Evaluating the accuracy of the annotation of a song may be extremely difficult and prone to different opinions. There is no ideal ground truth about music classification (mostly talking about the music genre) and not rarely fans don't agree with a certain tag being applied to a given song. This is also known as the subjectivity problem, as presented in [McKay 06].

The use of social networks and social tagging has raised this question even further: is it better to have a small group of trained musicologists annotating music instead of a rather large group of untrained music enthusiasts? Up to what extent is the wisdom of experts better than the wisdom of crowds? Do users need to be educated in respect to tagging? And by doing so and imposing rules, won't it cause users to feel demotivated and stop tagging content?

According to [Guy 06], this issue comes down to the use we give to folksonomies: we want to serve two masters at once: the personal and the collective collection and it is hard to find the balance in the middle of the two worlds. As in Amy Gahran's comment on an article⁴⁷: "A folksonomy merges, diverges, and evolves much the way language does, through usage and interaction".

Sources of tags

With the prominent usage of the English language on the Internet and the vast amount of content available in this language, we tend to lighten the importance of other languages on social on-line communities. However the usage of other languages has a big impact on the way people communicate and the classification of content is directly affected by this.

The use of web-mining techniques, social tagging and others mainly focus on this one-language world, but it is easy to understand that much of the knowledge available (and specifically knowledge related to music) is available in other languages.

⁴⁷ Amy Gahran's comments on "Technorati tags: Good idea, terrible implementation", available at http://www.intuitive.com/blog/technorati_tags_good_idea_terrible_implementation.html, published on 23.03.2005 and retrieved on 08.06.2009

This is specially important when we talk about, for instance, artists relatively new or unknown, whose information on the web is not available in English.

One example of this is the social encyclopedia Wikipedia⁴⁸, which has, as of 08.06.2009, 2.904.000+ articles in English language and over 5.399.000 articles written in other languages⁴⁹, as shown in Table 3.11.

Language	# of articles available	Language	# of articles available
English	2.904.000	Italian	574.000
German	914.000	Dutch	540.000
French	812.000	Portuguese	483.000
Polish	609.000	Spanish	480.000
Mandarin (Chinese)	591.000	Russian	396.000

Table 3.11: Number of Wikipedia articles written in the 10 most used languages

Also, when talking about the sources of tags, we must remember that the sources may be varied, not only in terms of language but also of who provides the information, and how accurate (or even true) the information is. The use of on-line encyclopedias or the extraction of information from news providers or official webpages must also be regarded with some care.

Focus

Another issue comes from the focus of annotation, where most tags are used to annotate artists rather than individual songs. In [Lamere 08b], the author mentions a statistic where Last.fm claims to have about 150 million tracks in its catalogue and 40 million tags, yielding an average of 0.26 tags per track. According to this, a typical track is tagged 100 times less often than a typical artist at Last.fm.

This poses a problem, as an artists may have songs with very distinct styles and characteristics, but since the songs are not properly classified, they end up by falling in the general classification of the artist, making it difficult to retrieve semantically relevant songs from eclectic artists.

Polysemy, synonymy and noise

Although the advent of social tagging has many advantages, their unstructured and free-form nature, as well as the use of ad-hoc techniques when annotating music also raises some issues. One of them has to do with the polysemy of some artist's names, as in the music artists Kiss, Bush and Porn [Schedl 05]. It may become problematic to find information regarding these artists, as the result of the searchable dataset is extremely noisy, with results that are not related to what we are trying to find.

⁴⁸ <http://www.wikipedia.com>

⁴⁹ This does not represent unique articles, as in many cases one article is translated into several languages. However, the author believes this is already a good reference in relation to the problem of the source of tags

Noisy results may also result from tags which have very little apparent relation to music or are related to personal opinions or classifications. Some examples from Last.fm are “Asdf”, “Random”, “Lazy eye”, “d”, “stuff a donut would like”, “woot” and “Sh-t that my sister listened to on my pc grrrrr”. These tags are not typically applied very often and as such are easy to filter out.

Finally, some tags are the result of misspellings, spelling variants and synonyms. These variants can dilute the dataset of tags, contributing to the sparseness of the tag space. However, these variants can also improve the “findability” of items. As an example, we can consider the multiple variants of the tags “hip hop”, “hip-hop”, “hiphop” and “rap”.

3.6 Event management and recommendation

Based on the analysis made in chapter 3.2.2, it is possible to conclude that the management of music events is still a novel focus of interest in social music software, with only a few platforms (such as Last.fm, iMeem, iLike, PalcoPrincipal, MySpace, Bebo and ReverbNation) effectively integrating this aspect of social music.

In some sites like PalcoPrincipal, Bebo and MySpace it is only possible to view a listing of the artist's shows and add selected shows to the user's personal calendar, with little information added about the event itself. However other sites such as Last.fm, ReverbNation and iMeem provide additional detailed information, such as maps and information about the show's venue, as well as ticketing information.

A growing number of other websites is specializing in live concerts (or “gigs”), associated to ticket sales, as in gigsandtours.com⁵⁰, ticketmaster.com⁵¹, livenation.com⁵², however none of these websites allow the social dynamics of an on-line community like the previously mentioned. Two exceptions however, are worth mentioning: [songkick](http://songkick.com)⁵³ and [bandsintown](http://www.bandsintown.com)⁵⁴.

Songkick was launched in October 2007, providing a social network of on-line live music resources for fans. It aggregates ticketing information from over 29 global ticket vendors, allowing users to add concerts and is able to index other concert listings, such as the ones from Last.fm. As of 09.06.2009⁵⁵, the company unveiled its new social concert database, with about one million concerts and complete tour histories for thousands of artists, dating as far back as a Bob Dylan's performance in 1960 to a series of concerts to be performed by Michael Jackson in 2010. The site encourages the documentation of live music events by empowering users to create their own personal “gigographies”, by flagging concerts they've attended and sharing their experiences with photos, videos, setlists and other concert ephemera.

50 <http://www.gigsandtours.com>

51 <http://ticketmaster.com>

52 <http://livenation.com>

53 <http://songkick.com>

54 <http://www.bandsintown.com>

55 <http://www.foxbusiness.com/story/songkick-launches-social-concert-database/>, published and retrieved on 09.06.2009

Bandsintown was launched in July 2007. It is an online community where users can discover scheduled live music concerts in their city (or in any city for that matter), as well as find tickets for the show and access media information, using a very simple and straightforward user interface. Here, the user can customize his search by distance, ticket price and date of the event, allowing him to receive email reminders and to link his bandsintown profile with his iTunes or Last.fm favourites and recommendations.

From the analysed platforms, Last.fm seems to be the most complete one, regarding event management and recommendation, presenting similar features as the ones in Songkick and Bandsintown; it is possible to comment/review the event, add information about attendance (commonly known as “RSVP⁵⁶”) and view other attendants, share event's media (e.g. photos, videos, setlists) and see event's recommendations.

There are three ways to provide these recommendations: based on the venue (by presenting other shows at the same venue), based on the artist (suggesting other events where the artist – or similar artists – will perform) and based on the user (by using the user's profile information and his preferences). All the 3 platforms mentioned before are able to provide event recommendations by using the venue or artist information, but only Last.fm is able to suggest events based on the user's profile, using its comprehensive information about the user, which allows the system to provide customized recommendations about events the user might be interested in attending.

Another interesting feature worth mentioning is the possibilities to track an artists and/or venues, in both Songkick, Bandsintown and ReverbNation, which might be very useful for users that are willing to follow the activity of both an artist or venue.

As far as the author is aware, there are no scientific publications or studies undertaken regarding the recommendation or management of music events. However, the author believes that it is possible to apply some of the principals of the recommendation of music to events.

3.7 Recommendation technologies

Here we will present some of the most notable recommendation technologies applied to music recommendation systems.

3.7.1 Last.fm's Audioscrobbler

Last.fm's “Audioscrobbler” desktop software, after having been installed on a user's PC, starts tracking the user's music listening behaviour. The listening data is sent to the Last.fm website, where a profile of the user's musical tastes is created. Listening statistics for each song, album, artist, and genre are aggregated and made available online.

⁵⁶ RSVP is an acronym for “Répondez s'il vous plaît”, a French phrase that translates to “please respond”, being used to reply to an invitation request, whether confirming attendance, declining or indicating an intention of attending (“maybe”)

In this way, Last.fm has created the world's largest open music database, comprising over 80 million songs, accessible by everyone and the largest ontological classification (and the largest open database) of music in the world, by aggregating the musical tastes of its 20 million users and then data-mining it for similar musical genres, artists, and songs.

The user-generated data also lays the foundation for personalization, collaborative filtering, and ontological classifications. Users can:

- view timelines and statistics about their own listening behaviour;
- view the most popular tracks for each of their favourite artists, and most popular artists for their favourite genre;
- receive music recommendations based on the song library already existing on their PC;
- discover similar tracks/artists for every track/artist in the Last.fm database.

3.7.2 Music Genome Project

The Music Genome Project, the technology behind the Pandora recommendation system, was created in January 2000 by Will Glaser, Jon Kraft and Tim Westergren to "*capture the essence of music at the fundamental level*" using over 400 attributes to describe songs and a complex mathematical algorithm to organize them.

A given song is represented by a vector (a list of attributes) containing approximately 150 "genes" (analogous to trait-determining genes for organisms in the field of genetics). Each gene corresponds to a characteristic of the music, for example, gender of lead vocalist, level of distortion on the electric guitar, type of background vocals, etc. Rock and pop songs have 150 genes, rap songs have 350, and jazz songs have approximately 400. Other genres of music, such as world and classical music, have 300–500 genes. The system depends on a sufficient number of genes to render useful results. Each gene is assigned a number between 1 and 5, in half-integer increments.

Given the vector of one or more songs, a list of other similar songs is constructed using a distance function. To create a song's genome, it is analysed by a musician in a process that takes 20 to 30 minutes per song. Ten percent of songs are analysed by more than one technician to ensure conformity with the standards, i.e., reliability.

An implementation of the Music Genome Project seems to be covered by United States Patent No. 7,003,515. This patent shows William T. Glaser, Timothy B. Westergren, Jeffrey P. Stearns, and Jonathan M. Kraft as the inventors of this technology. This patent is also owned by Pandora Media, Inc..

3.7.3 MyStrands

Strands technology is based on a complementary set of statistical analysis, artificial intelligence and complex filtering algorithms, collectively referred to as the Strands Social Recommender.




Figure 3.6: Schema of how Strands Recommender™ works⁵⁷

By exploiting human-generated links between items, implicit and explicit user behaviour and analysis of certain intrinsic characteristics, the algorithms used in Strands generate associations that result in a compelling representation of how society considers and uses the products, content, and information it creates and consumes.

Most of the basic principles used to construct networks of associations are conceptually simple, but the diversity of content types and the sheer amount of data collected often makes the networks themselves huge, requiring intelligent software algorithms to efficiently collect, parse, mine, and analyse large data streams.

In order to address this, Strands⁵⁷ uses the Universal Matrix of Associations (UMA) architectural scheme (Figure 3.6) to represent these networks logically and visually. By capturing user behaviour, aggregating information into discernible patterns and identifying matches with individual taste and preference profiles consistent with the UMA scheme, the patent-pending Strands Social Recommender engine provides great flexibility, power, and scalability.

Several key features distinguish Strands social recommendation technologies: The recommendation system inherently is content and platform agnostic, content awareness is captured in the models for the relationships between items which may be provided explicitly or automatically learned over time.

⁵⁷ http://recommender.strands.com/public/tech_architecture.sbs

Real-time recommendations are available for new users with no training and increasingly personalized as the system accumulates knowledge about the user. Similarly, new items inserted into the models are immediately recommendable based on approximate relationships with previously existing items and then increasingly recommended based on their relationships with other items learned from user behaviour. Finally, Strands' recommenders offer fast response times and robust scalability.

3.7.4 The Echonest

Echonest's hybrid recommender incorporates audio analysis, blog content and traditional collaborative filtering techniques for a seemingly promising mix. However, its results have yet to be fully available in a significant consumer offering.

The audio analysis tool takes an MP3 as input and generates an XML file which describes the track's structure and musical content (including rhythm, pitch, and timbre). All information is accurate at the sub-millisecond level. It uses proprietary machine listening techniques to simulate how people perceive music, unlike most automatic acoustic feature extraction tools that only consider the raw audio signal. All musical dimensions are represented globally and locally, such as key, time signature, loudness, beat, sections, harmony, etc. The analysis can even distinguish between sounds. Finally, rendering the complete analysis of an audio file only takes a few seconds of CPU time. The final output represents only a fraction of the original track size.

Regarding its recommendation mechanism, three types of recommendation are provided:

- **Artist-To-Artist.** Given an artist, it is possible to get recommendations for up to 15 similar artists in the user's catalogue;
- **Bulk or Playlist.** Given a playlist or user favourites artists, it is possible to get a nuanced list of recommended artists;
- **Similar Listener.** Given a listeners listening history or favourite artists, it is possible to get a nuanced list of other people with similar musical tastes.

3.7.5 Itunes' Genius

Genius is a recommender system in the iTunes media player computer software that uses collaborative filtering to generate music recommendations. It was first introduced in iTunes 8 and uses the integrated rating system built into iTunes. This allows for better recommendations because of a larger data set.

Because iTunes users often rate the songs in their library, Apple gets a lot of explicit information about a song's popularity. Users also regularly transmit information about how often they played and skipped a song to Apple's central servers.

Besides this, however, one can only speculate about what Apple looks at to give its recommendations.

It probably evaluates playlists and the similarities between different users' libraries, but one can only assume that Apple uses a mashup of various recommendation techniques to come up with its own suggestions. The results are generally quite good, though often either very predictable or completely random.

3.8 Project technologies

In the development of this project, several technologies are used. Although they are not going to be all presented, a small overview of the main technologies used in Palco Principal will be made: the PHP scripting language, the Symfony framework and the MySQL database management system.

3.8.1 PHP

PHP (standing for “PHP: Hypertext Preprocessor”) was designed in 1995 by Rasmus Lerdorf and is now produced by The PHP Group. It is free software released under the PHP License (which is incompatible with the GNU General Public License due to restrictions on the usage of the term *PHP*), being installed on more than 20 million websites and 1 million web servers⁵⁸, as of April 2007.

According to The PHP Group's website⁵⁹, “PHP is a widely-used general-purpose scripting language that is especially suited for Web development and can be embedded into HTML“. They continue by saying that “instead of lots of commands to output HTML (as seen in C or Perl), PHP pages contain HTML with embedded code that does «something» (...). The PHP code is enclosed in special start and end processing instructions `<?php` and `?>` that allow you to jump into and out of «PHP mode.»

What distinguishes PHP from something like client-side JavaScript is that the code is executed on the server, generating HTML which is then sent to the client. The client would receive the results of running that script, but would not know what the underlying code was. You can even configure your web server to process all your HTML files with PHP, and then there's really no way that users can tell what you have up your sleeve. The best things in using PHP are that it is extremely simple for a newcomer, but offers many advanced features for a professional programmer.”

According to [Lerdorf 02], PHP can be used in three primary ways:

- **Server-side scripting** - the most traditional and main target field for PHP. Three things are needed to make this work: the PHP parser (CGI or server module), a web server and a web browser. It is necessary to run the web server, with a connected PHP installation. Then it is possible to access the PHP program output with a web browser, viewing the PHP page through the server.

⁵⁸ according to <http://www.php.net/usage.php>

⁵⁹ <http://www.php.net>

- **Command line scripting** – it is possible to make a PHP script to run from the command line, much like Perl, awk, or the Unix shell. Command-line scripts may be used for system administration tasks, such as backup and log parsing.
- **Writing desktop applications** - PHP is probably not the very best language to create a desktop application with a graphical user interface, but with a deep knowledge of PHP and by using some advanced PHP features on the client-side applications, it is possible to use PHP-GTK to write such programs. PHP-GTK is an extension to PHP, not available in the main distribution. It is also possible to write cross-platform applications this way.

PHP runs on all major operating systems, including Microsoft Windows, Mac OSX, Unix variants including Linux, FreeBSD and Solaris. It can be used with all leading web servers, including Apache, Microsoft IIS and the Netscape/iPlanet servers. It also has a wide-range support for databases, supporting all the most known ones (including MySQL, PostgreSQL, Oracle, Sybase and ODBC-compliant databases).

PHP also provides a library of PHP code to perform common tasks, such as database abstraction, error handling and so on, with the PHP Extension and Application Repository (PEAR⁶⁰), a framework and distribution system for reusable PHP components.

3.8.2 Symfony

As stated at Symfony Project's website⁶¹, “*Symfony is a web application framework for PHP projects. It aims to speed up the creation and maintenance of web applications, and to replace the repetitive coding tasks by power, control and pleasure. (...) Symfony is aimed at building robust applications in an enterprise context. This means that you have full control over the configuration: from the directory structure to the foreign libraries, almost everything can be customized. To match your enterprise's development guidelines, Symfony is bundled with additional tools helping you to test, debug and document your project*”.

It is also important to mention that Symfony follows the model-view-controller (MVC) paradigm of software engineering. According to wikipedia⁶², “*in MVC, the model represents the information (the data) of the application; the view corresponds to elements of the user interface such as text, checkbox items, and so forth; and the controller manages the communication of data and the business rules used to manipulate the data to and from the model*”.

It is also mentioned that the “successful use of the pattern isolates business logic from user interface considerations, resulting in an application where it is easier to modify either the visual appearance of the application or the underlying business rules without affecting the other. “

Since Symfony was entirely written in PHP 5, it also follows the language's new object-oriented model (this has been a major change from PHP 4 to PHP 5).

⁶⁰ available at <http://pear.php.net>

⁶¹ <http://www.symfony-project.org>

⁶² <http://en.wikipedia.org/wiki/Model-view-controller>

In order to access the project's database in an object-oriented way, an interface translating the object logic to the relational logic is used - this interface is called an object-relational mapping, or ORM. Symfony integrates Propel⁶³ – another open source project – seamlessly, as its default ORM, although other ORM tools are also available (e.g. Doctrine⁶⁴).

As stated in The Definitive Guide to Symfony⁶⁵, *“one benefit of an object/relational abstraction layer is that it prevents you from using a syntax that is specific to a given database. It automatically translates calls to the model objects to SQL queries optimized for the current database. This means that switching to another database system in the middle of a project is easy”*.

According to its creators, Symfony has been thoroughly tested in several real-world projects, being actually in use for various high-demand e-business websites.

Finally, Symfony runs on Unix and Microsoft Windows platforms and is compatible with most of the available databases engines, including MySQL, PostgreSQL, Oracle, and Microsoft SQL Server. It has been released under the MIT license on October 18, 2005 by project founder Fabien Potencier and is currently sponsored by Sensio Labs, a French Web Agency known for its innovative views on web development.

Some of its features, according to The Symfony Project's Webpage², are:

- Easy to install and configure on most platforms (and guaranteed to work on standard (Unix and Windows platforms)
- Database engine-independent
- Simple to use, in most cases, but still flexible enough to adapt to complex cases
- Based on the premise of convention over configuration--the developer needs to configure only the unconventional
- Compliant with most web best practices and design patterns
- Enterprise-ready--adaptable to existing information technology (IT) policies and architectures, and stable enough for long-term projects
- Very readable code, with phpDocumentor comments, for easy maintenance
- Easy to extend, allowing for integration with other vendor libraries

3.8.3 MySQL

MySQL is a relational database management system (RDBMS), with over 6 million installations. The program runs as a server providing multi-user access to a number of databases. The project's source code is available under terms of the GNU General Public License, as well as under a variety of proprietary agreements. MySQL is owned and sponsored by a single for-profit firm, the Swedish company MySQL AB, a subsidiary of Sun Microsystems, which holds the copyright to most of the codebase.

63 <http://propel.phpdb.org>

64 <http://www.doctrine-project.org>

65 http://www.symfony-project.org/book/1_2

3.9 Key issues and lessons learned

From the study and analysis made, there are several key issues to point out. One has to do with the shift in music industry over recent years, with the “digitalization” of music and the way people consume and interact with it. This raises some legal issues which affect social music websites.

There is a clear semantic gap between people and music, in a sense that it is difficult to interpret and understand how people perceive music and which factors determine their likings. This is much due to the different types of music consumers and the different involvement that people have with music. It is important to understand the audience to which we are targeting our system, and Palco has a clear and specified audience.

Other key issues have to do with technological challenges. Machine learning techniques have been evolving over the last years, but they are still not accurate enough in order to allow a proper analysis of audio documents. Acoustic fingerprinting and mood analysis are two interesting techniques for classifying music, but there is still work to be done in this field.

Music classification and recommendation have several inherent issues, of which the most notable are the cold start problem, the long tail effect, the high dimensionality and sparsity of data, the trust and hacking of recommenders and the polysemy and synonymy of tags.

Various valuable lessons were also learned from this review. It was possible to learn that social software has suffered an evolution along time, until nowadays, in the several different forms as we know it, with different applications and philosophies.

Another useful information is regarding the social and digital music consumers: although they have distinct profiles and perspectives, we are talking mostly of an audience belonging to a young age gap. It is very important to adapt the system to the user, and note force the opposite.

A good recommendation is based in three important aspects: classification (tags), users and items. In order to provide good recommendations, several techniques must be combined and evaluated. Likewise, there are several techniques to classify content, each of them with advantages and disadvantages, for which it is not possible to say that there is one better than all the others.

There are several proprietary and commercial recommendation systems, with patented technology. Because of this, the competition to have the best recommender is fierce.

Classification is not only useful in the perspective of the system or system administrators, but also for users it is important. It is a form of users organizing their contents and relate to other users, both in a personal and social dimensions.

Events are fastly gaining importance in social music websites. They are an important component of the music experience and a privileged form of establishing contact between music fans and artists.

The development of geographic information service technologies and the increasing use of smart phones and mobile devices is opening great opportunities and perspectives for the future.

Chapter 4 Improving social music software

In this chapter, several features will be presented, with the objective of fulfilling the main goal of this project: improve Palco Principal's music software.

Based on the system description and main areas of focus of the project presented in chapter 2, and also on the several considerations in literature mentioned in chapter 3, the features proposed in this chapter aim to address many of the existing issues in social music software, by proposing the addition of new functionalities to Palco Principal and some changes to already existing ones. Their implementation will be discussed in chapter 5.

A more complete description of the features may be found in Annex E: Proposed features.

In the end of this chapter, some conclusions are presented, together with an evaluation of the impact that these features may have in the new version of the system.

4.1 Improving classification

Classification is undoubtedly one of the most important areas in social music, since other important areas such as recommendation and navigation directly depend on it.

From the analysis made to Palco Principal, it was possible to learn that classification is an area with a lot of room for improvement in the system, and definitively an area to focus on, by adding new features and restructuring existing ones. For doing this, seven new features are presented: Tag clouds, Multi-tag search, Tag suggestion, Acoustic fingerprinting, Weighted genre classification, Mood classification and Wiki-like content editing.

These features are focused mainly in the content visualization and automatic classification techniques.

4.1.1 Tag clouds

A tag cloud is used to help both artists and fans have a simple and easy-to-read display of how items are classified. This way, the user is able to find which tags are more relevant or compare the popularity of a given tag within the cloud (its lexical relations), and browse through the tagged content.

4.1.2 Multi-tag search

Using a multi-tag search tool enables fans to narrow the set of search results for a specific item, by combining multiple tags.

4.1.3 Tag suggestion

A tag suggestion mechanism is meant to help music fans classifying content, by presenting suggestions of meaningful tags which he can apply to several items, such as artists, songs, blog posts, news, reviews, events and venues.

4.1.4 Acoustic fingerprinting

The objective of acoustic fingerprinting is to provide Palco 3.0 a way to match and associate certain patterns of acoustic data, by extracting and analysing some of its characteristics and matching them to existing patterns, thus creating a “fingerprint” that identifies a song and allows its classification into several categories.

4.1.5 Weighted genre classification

Using a weighted genre classification system allows Palco 3.0 and fans to classify songs and artists into predefined genre categories, having into account the acoustic fingerprinting of the song and the user's classification, by assigning weights to different genres the song or artist might fit in.

4.1.6 Mood classification

This novel type of classification allows the system and fans to classify and discover music, according to their mood or emotional perception. This is an alternative to music genre classification, which has similar goals: classify songs and artists into a mood profile, using automated classification methods and user-based classification.

4.1.7 Wiki-like content editing

A wiki-like content editing tool allows fans and artists to easily improve and increase the amount of information available, using a social, wiki-like approach, making use of the knowledge of the community.

4.2 Improving recommendation

Along classification, recommendation plays a vital role in social music systems. In order to improve Palco's recommendation system, five features will be presented: Music recommendation, User affinity and recommendation, Rating system, Reputation system and Graphic recommendation.

With these features, it is expected to be able to provide more accurate, better and personal recommendations for users, allowing them to use them to navigate and discover new music.

4.2.1 Music recommendation

A music recommendation engine is expected to provide recommendations (artists or songs), according to the user's preferences or, by comparison with other items, recommending similar ones.

4.2.2 User affinity and recommendation

The goal of user affinity and recommendation is to determine similarity and compatibility between users, according to their preferences and listening habits, in order to provide music and user recommendations (possible friends).

4.2.3 Rating system

The objective of this is to improve the current rating and reviewing system, in order to provide a more intuitive and meaningful mechanism for users to rate and review musical content, using the ratings to contribute for the music recommendations.

4.2.4 Reputation system

The goal of creating a reputation system is to achieve a bigger and more fair involvement of fans in the community, allowing them to earn a reputation according to their contribution and activity.

4.2.5 Graphic recommendation

A graphic recommender is intended to provide recommendations using an attractive and interactive graphic display, allowing a better and easier navigation.

While doing this, the listener is able to listen to his recommended songs. This also promotes recommendation steerability, by allowing the user to approve or disapprove the recommendations he receives.

4.3 Improving event management

The focus on music events has been growing over the last years, with the social music platforms understanding that live concerts are an indissociable aspect of the whole music sphere. Because of this and in order to proceed to large improvements in the event management system of Palco Principal, five features are hereby proposed: Geographic Information System (GIS), Event displaying, Event suggestion and Venue profileUser-added media.

4.3.1 Geographic Information System (GIS)

The integration of Geographic Information System (GIS) technologies in Palco 3.0 would allow a bigger focus on the user and provide him more accurate information, based on its location. It would have an impact mostly in two areas: recommendation and event management.

In event management, GIS technology would allow to geographically compute with precision the location of events, displaying maps, distances and directions from the user's (or for that matter, any) location to the event itself. It would also be useful in order to locate venues and search for events, within a range of proximity.

Regarding recommendation, other users whose location is close enough (with the measure of granularity being possibly changed) could be recommended, as well as events of bands (recommended or not for the user) in nearby locations. Also, when providing information about popular artists and/or songs, these could be provided specifically to the user's country of location, in order to better characterize the musical trends of a specific community, as they tend to vary from country to country (e.g. top artists in Portugal or top songs in Angola).

4.3.2 Event displaying

Although Palco Principal provides information about musical events (also known as shows or gigs), it is very limited and mostly static. With this feature it is intended to add more information to the events, such as attendance and ticketing information.

Regarding attendance, users should be able to provide their RSVP⁶⁶ status. On the event, it should be possible to see which users would attend the event. About ticketing information, apart from the ticket price, it would be useful to have information and/or links of how and where to buy tickets.

Also, both the artist, fan and venue should be able to add photos and videos, related to the event, such as posters, tickets, *in loco* photos and videos and others.

⁶⁶ RSVP stands for "Répondez s'il vous plaît", a French phrase that can be translated to "please respond". It is standard practice to reply to an RSVP request whether confirming attendance or declining.

4.3.3 Event suggestion

The system may suggest events, which the user might find interesting to attend to. This can be done by suggesting nearby events, based on the user's location and distance to the events, suggest other events from the same artist in other locations, or from similar artists. Both these recommendations may be done having into consideration the user's music recommendations.

4.3.4 Venue profile

By creating a new type of profile, focused on the venues, these are included in the dynamics of the network, since venues are also an important player. An event needs an artist to perform, fans to attend and a venue to be hosted, so it seems quite logic that, as it happens with fans and artists, venues also have a profile page.

The venue profile should contain basic information about it (such as the website, opening hours, photos, description), but also information about upcoming and past events. Using the GIS technologies, it is possible to easily locate the venue and find directions, allowing the creation of city maps with all its venues.

4.3.5 User-added media

This feature refers to the possibility of the users adding photos and videos related to events, may they be promotional or documentary. This would allow not only a better coverage of the event (together with event reviews and comments), but also the creation of photo and video collections of fans and venues.

4.4 Conclusions

With the proposal of implementation of these features, it is intended to improve Palco Principal, providing better and easier classification tools, a more comprehensive recommendation mechanism, with useful tools and also improve the event management.

In order to evaluate the impact that each of these features may have in the process of the system's improvement, we recall the issues presented and discussed in sections 3.4.3 and 3.5.7 (both in chapter 3) and we try to match the proposed features to the issues, whose impact is more significant. Hopefully, the implementation of these features will substantially help to avoid some of the problems found in music classification and recommendation.

Improving social music software

Let us first recall the discussed issues. The issues related to recommendation are displayed in big caps (A to >N) and the ones related to classification with small caps (from o to t):

- | | |
|--|---------------------------------|
| A. Feedback Loops | K. Transparency |
| B. Popularity Bias / The Long Tail issue | L. Steerability |
| C. Early Rater Bias | M. Music representation |
| D. Cold Start | N. Music queries |
| E. Novelty and serendipity | o. Strong vs. Weak Labelling |
| F. Data sparsity and high dimensionality | p. Vocabulary structure & size |
| G. Scale | q. Annotation accuracy |
| H. Gray sheep | r. Sources of tags |
| I. Temporal effects / inertia | s. Focus |
| J. Trust and recommender hacking | t. Polysemy, synonymy and noise |

The features are divided in classification features (1) and recommendation features (2):

- | | |
|-----------------------------------|-------------------------------|
| 1.1 Tag clouds | 1.7 Wiki-like content editing |
| 1.2 Multi-tag search | 2.1 Music recommendation |
| 1.3 Tag suggestion | 2.2 User affinity |
| 1.4 Acoustic fingerprinting | 2.3 Rating system |
| 1.5 Weighted genre classification | 2.4 Reputation system |
| 1.6 Mood classification | 2.5 Graphic recommendation |

In Table 4.1 the proposed features (with exception to the event features) and issues are summarized, with the bullets establishing the relation between the features and the issues they address.

Improving social music software

	<i>1.1</i>	<i>1.2</i>	<i>1.3</i>	<i>1.4</i>	<i>1.5</i>	<i>1.6</i>	<i>1.7</i>	<i>2.1</i>	<i>2.2</i>	<i>2.3</i>	<i>2.4</i>	<i>2.5</i>
<i>A</i>								•				•
<i>B</i>								•		•		•
<i>C</i>			•	•	•	•					•	
<i>D</i>			•	•	•	•						
<i>E</i>					•	•		•		•		•
<i>F</i>	•	•										•
<i>G</i>	•	•										•
<i>H</i>		•			•	•		•	•		•	•
<i>I</i>		•			•	•		•		•		•
<i>J</i>				•	•	•		•	•	•	•	
<i>K</i>				•				•		•		•
<i>L</i>		•			•	•		•				•
<i>M</i>				•	•	•						
<i>N</i>		•			•	•						
<i>o</i>			•	•	•	•						
<i>p</i>	•				•	•						
<i>q</i>			•	•	•	•	•					
<i>r</i>				•	•	•					•	
<i>s</i>	•			•	•	•						
<i>t</i>	•		•									

Table 4.1: Influence of features in issue solving

Based on this evaluation, it is important to assign priorities, in order to determine which features are more or less critical for the proposed solution.

This prioritization was made separating the features into three different priorities: low, medium or high priority. This was done with the concern of keeping highly related features within the same group of features, in order to not compromise its implementation or success.

The results of the prioritization are displayed in Table 4.2.

Feature	High	Medium	Low
1.1 Tag clouds		●	
1.2 Multi-tag search		●	
1.3 Tag suggestion		●	
1.4 Acoustic fingerprinting	●		
1.5 Weighted genre classification	●		
1.6 Mood classification	●		
1.7 Wiki-like content editing			●
2.1 Music recommendation	●		
2.2 User affinity and recommendation			●
2.3 Rating system	●		
2.4 Reputation system			●
2.5 Graphic recommendation		●	
3.1 Geographic Information System (GIS)		●	
3.2 Event displaying	●		
3.3 Event suggestion		●	
3.4 Venue profile	●		
3.5 User-added media			●

Table 4.2: Feature implementation prioritization

4.5 Summary

In this chapter, several features were proposed for implementation, as an attempt to address the projects main goal – improve the social music software of Palco Principal – having also in consideration the various existing issues in the fields of music classification and recommendation.

From the 17 proposed features, 7 are directly related to music classification, 5 are connected to music recommendation and 5 more are proposed for improving the event management in the platform. Most of the features proposed add tools and mechanism to the system, but some other proposals go in the direction of improving and expanding already existing features in Palco Principal.

Based on this list of features, as well as in the list of issues, a correspondence between the several features and the issues they attempt to address is made in the chapter conclusions, as well as a division of the features, according to their implementation priority, which can be used as a guidance for the incremental development of this set of features.

Chapter 5 High level prototype

In this chapter a series of prototypes is presented, in order to foresee and specify how the implementation of the features described in the previous chapter should process. These prototypes are intended to be schematic and were designed for the single purpose of visualization, thus not having any real implementation.

There are five major areas (or pages) on the prototype: the homepage, music, community, tags and events. Links for these pages are always displayed in the page header. There are also three important components that deserve attention: videos, photos and venues. All these eight areas, just mentioned, are displayed in individual pages. Each of these pages has several tabs, displaying sub-pages.

Since many of the sub-pages are similar (e.g. the event or artist photo sub-page has a similar layout, with only the content provided being different – for the first photos related with a certain event and for the second photos related to a specific artist), only the prototypes for all the sub-pages of the artist page will be presented. As for the rest of the pages, similar sub-pages will be omitted in these representations.

All the figures displayed may be viewed with higher resolution in Appendix F: Prototypes.

5.1 Homepage

The homepage is the main page of the system, from where the users begin exploring it. It is depicted in Figure 5.1.

As in the current version of Palco, there is an area on the top of the page for placing an artist in the spotlight, in order to highlight the artist from the others. There is also an area for placing advertisement (present in nearly every page, under the name “Ads”) and minor boxes that may be used for either purpose a little bit below.

Also in the top of the page it is possible to see a small list of upcoming events, as well as the latest news, reviews or blog posts (choosing from the appropriate buttons) and genre-based radios. These features may also be found in the current version of Palco.

High level prototype




Figure 5.1: Prototype of Palco 3.0 homepage

One improvement, however, is the inclusion of a mood-based radio, on which the user may define what is his state of mind (mood) and receive suggestions of music that correspond to this feeling. This is done by using the mood classification feature described in section 4.1.6.

Another improvement in this prototype is the display of videos and photos in the homepage, in the bottom part of it. This is an important component of the system, since photos and videos generally attract the user's attention.

Finally, just by the side of the video list is a list of the last songs listened in Palco. This is a way of promoting songs, by listing them by latest activity. This feature is also provided in the homepage of the current version of Palco.

5.2 Header and Footer

All the webpages in this prototype are generally divided in three columns (as seen already in the homepage), whereas the leftmost and the centre columns might be merged, if necessary.

All the pages, with no exceptions, have a static header and footer, as depicted in Figure 5.2. These will be used to guide the user while navigating through the website, as it already happens in the current version of Palco.

High level prototype



Figure 5.2: Prototype of header and footer

The header displays the logo of Palco Principal and links for the five main areas, as already mentioned: the homepage, music, community, tags and events. There is also a search field, where the user can directly make any query to the system, in order to find one of the items displayed in the dropbox: music (artists or songs), events, videos, photos, tags, news, reviews or blog posts. Also links to the other communities in Palco (Brazil, Angola, Mozambique and Cape Verde) may be displayed here, but for the sake of simplicity they have not been represented in the prototype.

At the footer is displayed information, as in the current version of Palco, related to the following aspects: help (for providing user assistance), about us (displaying information about the website and the company), genres (where it is possible to directly access a few popular genres), forums (linking to the existing forums of Palco), highlights (displaying the top news, events, etc.), bands/artists (displaying a few popular artists), users (highlighting a few popular users) and “do more” (allowing users to access other functionalities such as Palco desktop software, among others).

5.3 Music

In the music page (Figure 5.3), the user is able to view the highest ranked songs and artists, as already happens in Palco (for artists only). This ranking is automatically calculated by the system and rewards the most active and popular artists/songs.

The major improvement in this prototype has to do with the possibility of filtering the results per country, by pressing the button “in your country”.




Figure 5.3: Prototype of music page

With this option, only the items based in the user's country will be displayed, whereas using the other option shows the top-ranked items, among all the music in Palco. This is done by using the feature described in section 4.3.1 - Geographic Information System (GIS), which enables the system to automatically identify the user's location and display results based on that information.

Also the user is now able to filter the results according to the calendar week that they refer to, as the trends vary along time. For doing this, the user may choose the week directly from the calendar displayed.

Finally, the user is also able to listen to songs from the “popular songs” list, using the incorporated player. He is also able to see information about the song in playback and rate it, by using the star-rating system (introduced in section 4.2.3), also shown in Figure 5.3.

5.3.1 Artist overview

When selecting an artist, the user is redirected to the artist overview page (Figure 5.4).

Here the user is presented with a summary (or overview) of all the contents related to the artist. In the heading of this subpage (maintained in all the subpages of the artist profile; not to be confused with the main header) is clearly displayed the artist's name, as well as its mood (this classification is generated with base on the featureMood classification, described in section 4.1.6), location (city and country's flag), as well as a small view of the tag cloud associated to him (implementing the feature presented in 4.1.1).

High level prototype




Figure 5.4: Prototype of artist overview page

Also in this header it is possible to become a fan of the artist, using the button “become fan”). This will enable the user to receive updated information about the artist's activity. Another button just next to it – “Tag!” – allows the user to add or modify the tags he assigned to this artist, redirecting the user to the “Tags” subpage. Just below the tag cloud, it is possible to rate the artist, using the star-voting system already mentioned (in the example, the user voted the artist with a 5-star rating, with the average rating being 4.4).

Still in the top of the page, benefiting of a good visibility, the main photo of the artist (chose by him) is shown just below the name and above its description. This may be expanded pressing the “more...” button just next to it, redirecting the user to the “biography” subpage.

Next to this it is possible to find a list with the songs performed by this artist and a small player, which is used to playback the selected songs from the list, as already happens in Palco. It is also possible to “favourite” the songs by pressing the star button in front of the song's title and add it to the user's playlist by pressing the “+” button next to it.

High level prototype

In the centre of the page it is possible to find some statistics regarding the artist (these are calculated by the system and may be found in the current version of Palco), a list of the artist's albums and his upcoming events. The last two may be expanded by pressing the respective “more...” buttons, redirecting the user to the “music” and “events” subpages, respectively.

Further below is displayed a list of recommended artists, that are somehow related to the artist displayed. This is an implementation of the feature presented in section 4.2.1: Music recommendation. The explanations of the recommendations are displayed when the user passes the mouse over the artist icon or name.

Below the recommendations is a list of the latest blog entries, reviews, photos and videos related to the artist. All of them may be expanded by pressing the appropriate “more...” button, redirecting the user to the respective subpages.

On the bottom-right box is displayed a list with the artist's fans and in the bottom of the page are displayed the latest comments made to the artist's profile, where the user may also insert his own comment.

5.3.2 Artist biography

Below the profile's header it is possible to view or edit the artist's biography or description (Figure 5.5). This is done by implementing a wiki-like editor, as presented in section 4.1.7. The user may also view the historic of editions or participate in the page's discussion, using the buttons “history” and “discussion”.




Figure 5.5: Prototype of artist biography page

High level prototype

On the rightmost column, some quick facts about the artist are presented (e.g. year of formation, home town, location, webpage, etc.) and the band members. These may be added by manually inserting their identification or by inserting the member's username in Palco. This way the member's profile will be connected to the artist's profile.

5.3.3 Artist tags

In this feature, the user is able to classify the artist, using tag, genre and mood classification techniques. This is illustrated in Figure 5.6.

On the top of the page is presented the artist's tag cloud (an implementation of the feature presented in section 4.1.1). The most popular tags are highlighted and are a result of the global tagging applied to this artist. The user may add or modify his tags by using the box on the right side. In the example, the most popular tags for the artist are “portuguese”, “pop” and “rock”; the user's tags are “pop”, “portuguese” and “lisboa”.

Also some suggestions are presented, with the most appropriate tags from his collection of tags (i.e. from all the tags the user has ever used, the system selects the tags which are somehow more related to this artist). This is an implementation of the Tag suggestion feature explained in section 4.1.3. In the example, the suggested tags are “portuguese” and “pop”.

The user may view the general genre classification of the artist at the centre-right box. He can also give his own opinion regarding which musical genres the artist belongs to, assigning different weights to the several genres. This feature – Weighted genre classification – has been described in section 4.1.5.




Figure 5.6: Prototype of artist tags page

High level prototype

It allows users not only to identify the genre categories to which an artist belongs to, but also to distinguish between them. In this example the user classified the artist in three genres: rock (70%), pop (50%) and metal (20%). However, the general classification (not changeable by the user) differs a little bit: rock (70%), pop (56%), metal (17%) and funk (62%).

Finally, in the bottom of the page, it is possible to classify the artist according to its mood. This is done by choosing from the predefined set of mood. In the example, the users has chosen the “circular”, “harsh” and “stylish” moods. The general evaluation of the artist's mood, based on its acoustic fingerprint (explained in section 4.1.4) and on users' classification is then displayed in the bottom-right graphic, as presented in section 4.1.6: Mood classification.

5.3.4 Artist recommendations

The artist recommendations page (Figure 5.7) is the general implementation of the idea described in section 4.2.1: Music recommendation.

On the top of the page, a list of similar songs is presented to the user. He is able to play them, using the playback device on the page. While playing the song, the user is able to see an explanation of why the song has been recommended to him. This explanation is displayed on the panel on the top-right of the page.




Figure 5.7: prototype of artist recommendations page

High level prototype

Also other artists are recommended to the user. When passing the mouse over the artist, it is possible to view the explanation for that suggestion in the box below the artist list. On the right side of this box, a list with other users who share this and other artists with the user are presented – an implementation of the feature User affinity and recommendation, from section 4.2.2.

Finally, on the bottom of the page, it is possible to see a graphical display of the recommendations provided, allowing the user to interactively browse through the recommendations. This is an application of the feature Graphic recommendation, from section 4.2.5.

5.3.5 Artist music

In this page it is possible to listen to the artists' songs and view detailed information about each of them. The songs may be filtered by the album they belong to. Informations about the album and the song are displayed on the column at the right side.

Also a list of the top listeners is displayed in the middle of the page, with the artist's fans displayed in the right-side column.

Finally, in the bottom of the page it is possible to view and edit the lyrics (using the Wiki-like content editing feature, from section 4.1.7) of the song in playback, as shown in Figure 5.8.




Figure 5.8: Prototype of artist music page

5.3.6 Artist events

In this page it is possible to visualize the artist's events, upcoming or past. The user is able to choose a date using the button on the top of the page and view the events pinpointed in the calendar, on the right side of the page. By selecting the event, the user will be redirected to the event's page.

There is also a button, just above the calendar, on the right side of the page, entitled “come to my town!”. By pressing this button the user manifests its interest in attending a live concert of the band in his city, or another city nearby. This information might be useful for the artists, helping them to better plan their shows.

On the mid-right side of the page there's a small map of the locations of the events displayed, using the Geographic Information System (GIS) feature explained in section 4.3.1.

At the bottom of the page it is possible to see a list of suggested events, an implementation of the Event suggestion feature presented in section 4.3.3.




Figure 5.9: Prototype of artist events page

5.3.7 Artist photos

In the artist photos page (Figure 5.10) are displayed the artist's photos. On the top of the page is displayed a list with thumbnails of the photos. When a photo is selected, its information is displayed on the box on the right of the page and the photo is visualized in the designated area.




Figure 5.10: Prototype of artist photos page

In the bottom of the page it is possible to leave comments on the photo and a tag cloud for the photo is also displayed. The user is also able to add or update his tags for the photo.

5.3.8 Artist videos

The artist videos page (Figure 5.11) is in all similar to the previous page, with the difference that it applies to videos instead of photos of the artist.

High level prototype




Figure 5.11: Prototype of artist videos page

5.3.9 Artist blog

Both the artist's news, blog and comments have the same structure (Figure 5.12).




Figure 5.12: Prototype of artist blog page

High level prototype

They are displayed as a list and the user is able to filter the posts by date (using the option on the top of the page) or use the search option to find a specific item.

The user is also able to tag the post, as well as view its tag cloud.

5.4 Community

Using the community button, on the main header of Palco, the user is redirected to the page illustrated in Figure 5.13.

Here, the user is able to see the most popular users and filter them by country. The most popular artists are determined using a Reputation system, as explained in section 4.2.4. They may also be sorted using another criteria, such as the number of photos uploaded, the number of songs in the playlist or others.

Using the User affinity and recommendation feature proposed in section 4.2.2, the system is able to recommend other similar users and by means of a Geographic Information System (GIS) (depicted in section 4.3.1), determine other nearby users. These users are pinpointed in the small map in the page.




Figure 5.13: Prototype of community page

5.5 Tags

When choosing the “tags” option in the main header the user is redirected to this page – Figure 5.14.

Here the user has an overview of how the tags applied in the whole system are distributed, by means of a tag cloud (explained in section 4.1.1). Also, the user can see his most popular tags (i.e. the tags he used the most), on his personal tag cloud, on the right side of the page. This way, the user is presented with both dimensions of classification – personal and social tagging.

The user can also navigate through tags by using the search button or selecting one of the tags displayed in the clouds, being redirected to the tag overview page.




Figure 5.14: Prototype of tags page

5.5.1 Tag search

Another improvement in this prototype has to do with the possibility of refining the user's search by using multiple tags (Figure 5.15). This is an implementation of the feature described in section 4.1.2.

By using a combination of tags, the user is able to narrow down his search, receiving recommendations of artists, songs, albums, photos, videos and events who share the tags entered.

In order to do this, the user must insert consecutive tags in the top of the page. These tags may be removed using the appropriate button, allowing a flexible way for an user to find the content he is looking for.

High level prototype




Figure 5.15: Prototype of tag search page

5.5.2 Tag overview

When the user selects a tag, he is redirected to this page (Figure 5.16).

In the top of this page, a photo of one of the popular artists classified with this tag is displayed. Also, a list of songs related to this tag are displayed, with the user being able to play the songs using the incorporated playback device.

It is possible to add or edit the description for the tag, using the Wiki-like content editing feature described in section 4.1.7. It is also possible to see which artists, albums, blog entries, reviews, photos, videos and upcoming events are classified with the selected tag, as well as some users that have applied them.

Using the recommendation techniques presented before, the user also received a list with related tags, suggested by the system. These tags are presented as a result of the implementation of the Tag suggestion feature, explained in section 4.1.3.

Finally, users are also allowed to leave comments in the tag page.

Because the subpages photos, videos, news, blog, comments and description are similar to the ones presented in the music page, they will not be illustrated here.

High level prototype




Figure 5.16: Prototype of tag overview page

5.5.3 Tag music

In this page the user is able to see which artists and songs are associated with the specific tag(s) specified.

The songs classified with the selected tag are displayed in a list, and may be played using the playback device. The song details are displayed on the info box, on the right side of the page.

It is also possible to see which artists have been classified with that tag and some popular users.

The prototype for this page may be illustrated in Figure 5.17.

High level prototype




Figure 5.17: Prototype of tag music page

5.5.4 Tag events

In Figure 5.18 is displayed the prototype for the tag events page. In this page, the events classified with the specified tag(s) are listed and pinpointed in a small map.




Figure 5.18: Prototype of tag events page

5.6 Events

The Events page is accessed when the user chooses the “events” button from the main header. In this page (Figure 5.19), the user is able to see a list of events, which may be filtered with several tools.

The user is able to choose all the events or only the events based on one location. In the example, the user is located in Porto (Portugal). It is also possible to filter the results by specifying a distance range, from the defined location. In the example, the events displayed are in a maximum distance range of 15 km from Porto. These features are implemented using the Geographic Information System (GIS) feature, referred in section 4.3.1.

The same applies to the ticket price, with the user defining the maximum ticket price he is willing to pay for the event. All the events whose ticket is more expensive than that limited are not listed.

Finally, it is also possible to specify a data range for the events the user is searching. All the listed are also pinpointed in a small map, on the right side of the page.




Figure 5.19: Prototype of events page

High level prototype

In this event overview page (Figure 5.20) the user is able to see all the information related to a given event.

In the page header is displayed the event name and its location. On the top of the page a photo of the artist (or one of the artists) of the event is displayed and on the right side the event details are presented.

The user is able to see some of the photos and videos of the event, as well as other recommended events, from the same artists or similar ones. These events are displayed in a small map, on the right of the suggestions.

Another improvement here is the possibility of the user providing his availability. This is done by choosing one of the attendance options: “yes”, “no” or “maybe”. The user is also able to see which other users are attending the event.

Finally, users may leave comments about the event, in the bottom of the page.

Because the subpages recommendation, photos, videos and comments are similar to the ones presented at the music page, they will not be illustrated here.




Figure 5.20: Prototype of event overview page

5.7 Venues

The venues page (Figure 5.21) displays the most popular venues.

The list of venues may be filtered by their location or within a certain range, and according to their rating.

When selecting a venue, the user is redirected to the venue overview subpage.




Figure 5.21: Prototype of venues page

The venue overview page (Figure 5.22) provides detailed information about a venue.

On the top of the page, the name and location of the venue are displayed, as well as a photo. On the right it is also possible to see the details and contacts of the venue.

Just below the main photo is displayed the venue's description, with its location being displayed in the map on the right. This map may be enlarged, by choosing the appropriate button and also it is possible to find directions to the venue.

A list of upcoming events is also displayed, along with some related photos and videos. Users may also leave and read other users' comments, rate the venue and see which other venues are nearby, in the box at the left-right side of the page.

High level prototype




Figure 5.22: Prototype of venue overview page

5.8 Top videos

Users may also view the most popular videos at this page (Figure 5.23). The videos' thumbnails are listed at the top of the page, with their details being displayed on the right side. When the user selects a video it is displayed in the designated area.

Users may also tag the video and see the associated tag cloud.

Finally, users may leave comments on the selected video.

High level prototype




Figure 5.23: Prototype of top videos page

5.9 Top photos

The popular photos page (Figure 5.24) is in all similar to the previous one.

High level prototype




Figure 5.24: Prototype of top photos page

Chapter 6 Conclusions and Future Work

6.1 General Conclusions

Palco Principal and its context

The world of music industry has changed dramatically over the last years, with the digital revolution and the widespread of internet. The traditional physical record has been decreasing in sales, in contrast to the new digital distribution formats, catapulted by the massification of the use of MP3 audio-files and MP3-supported portable music players.

This phenomena of shifting music into the “digital world” has also opened new opportunities for the creation of on-line music communities and contents. The end of the 1990's saw the appearance of the first portals, where music started to be democratically available for its fans, with its most prominent case being the one of MySpace. Soon other portals came to seize these new opportunities and between 2003 and 2006 an “eruption” of music portals and social music websites came to the Internet.

From the analysis made in this report to 32 platforms, we concluded that music communities have already achieved a large, international dimension, which reflects on the business itself. The 5 largest communities are MySpace, Bebo, iLike, iMeem and Last.fm, all with over 20 million users each. Comparing in terms of dimension, Palco Principal remains a very small platform (with about 10 thousand users), however its investment in portuguese-speaking markets may be changing this situation soon.

In terms of the number of features provided to its users, Palco Principal directly competes in the 5 most complete features with iLike and myStrands. The three cover about 60% of the analysed features, only behind Last.fm and iMeem, both with over 70% of coverage.

Palco Principal, critical success factors and opportunities

Some of the critical success factors of Palco Principal are related to its ability to reach independent artists and projects, allowing them to easily share their music, providing them with useful social and monitoring tools, such as an on-line radio where they can upload their songs as well as rankings and statistics, which are used to monitor the performance of them. Also the possibility to create a blog and easily leave comments to other artists or users, makes Palco a great place for socialization and interaction. However, some areas still need improvement, out of which music recommendation, music classification and event management point out.

From the analysis made to the several platforms, it was possible to learn which features were more and less common. Most of the platforms provide its users with on-line radios, customized playlists and recommendations, possibilities to access on-line music stores and the possibility to follow and interact with their favourite artists or users.

On the other hand, some features are not so popular yet, as only a few platforms provide them. These are the cases of microblogging and on-line chatting tools, reputation systems, mood classification, geographic localization, event recommendation and fraud or abuse detection.

In both the situations, valuable informations may be extracted for Palco Principal, in order to detect and explore business opportunities. In order to keep up with the remaining companies, Palco Principal should give a bigger focus to music recommendation and user-based playlists, as well as improve its music classification and event management systems. The implementation of a reputation system could bring large benefits for the community and the integration of Palco with mobile platforms and content providers could result in a strategic opportunity as well.

The technological challenges

Music recommendation is one of the most important aspects of social music software, as it guides music fans through the immensity of music, however, it is also a complex area, which involves many knowledge competences in fields such as data mining, machine learning, music information retrieval, acoustic fingerprinting, affinity systems, adaptive web, etc.

Because of the immensity of songs and data that can be extracted from music, as well as the subjective nature of its context, with different people having different perceptions and attitudes towards music, music recommendation poses many technological challenges. Analysing audio and extracting acoustic features is still an area with room for improvement and the semantic gap that exists between people and music is still a big barrier.

There are several open issues in music recommendation, such as the long tail issue, the cold start problem, the difficulty of handling data sparsity and its high dimensionality, the need of trust and transparency on music recommenders. All these have been studied and explored in the Music Information Research community, but there are still many questions left unanswered.

Nonetheless, the recent development on computational performance and development and improvement of techniques and algorithms leave good indications and perspectives towards the future, and the constant growth of social websites will allow more and better solutions for users.

Improving social software for music

In order to improve Palco Principal, placing it at the level of all the major social music platforms, three main areas of focus must be improved: music classification, music recommendation and event management.

The first two areas help both music fans and artists to share and promote music, reaching a growing audience in a much closer and customized form. Each user is different, thus needs to be treated particularly, according to its preferences. With the use of the social tools available, it is much more easier for artists and fans to “have a voice” in the community, and gaining a respectable position or reputation helps the community to have a greater commitment, as well as a better performance, in general.

Regarding event management, it has grown in importance, as live concerts are the closest connection that most music fans have to their favourite artists. Also for artists (specially newcoming ones, as in the case of Palco Principal) these *gigs* or *shows* are important both to promote themselves, but also to create a steady group of followers. Bringing this bond into the social network is then, of great importance.

The proposed solution is an attempt to effectively improve the system, having into account the context in which it is involved, as well as the problems or issues that they address. Its goal is to improve the user experience and involvement and make the social music platform more than just a door to the ocean of music, but a guide to all the different islands and continents that it encompasses.

In order to achieve this solution, some aspects were left for future work, which will necessarily involve all the teams of the project, as well as the company itself, whose strategic and structural decisions will dictate the path to be followed in order to reach the vision that was initially targeted.

6.2 Future Work

As explained before, this report refers to the second quarter of a three year long project – Palco 3.0. For this reason, there is naturally much work to be done, but hopefully the work developed up until now will be a useful instrument for the upcoming challenges.

In order to properly develop the classification and recommendation systems, a more thorough work has to be done in the field of audio and acoustic data analysis. The analysis and classification of Palco music collection's audio spectral data will provide an important contribute in the classification and categorization of artists and their songs.

The work of the data mining team, using machine learning techniques and algorithms to compute similarities between items and users will be the basis for the creation of a recommendation engine, that should be capable of assemble and process the information provided by the audio analysis and the user activity.

Conclusions and Future Work

Both these areas will allow the social software team to build upon the existing tools and features of Palco, improving them and adding new ones, in order to provide its users with more means of interaction and tools to share their music. These improvements must also be done in a logic of improving the community, creating a close bond between users (music fans, artists, venues) and the system, in order to create a proper dynamic that can be expanded to other markets and audiences.

Regarding the area of event management, much can be done, namely to provide a better integration of this tool. By using geo-positional technologies, a user could receive up-to-date and customized information about events (as well as ticketing information, etc.), based on its location. Here, the use of mobile platforms and devices may be a key landmark.

Finally, some work regarding usability is also necessary, in order to make the website more user-friendly and to properly present the new or improved features just mentioned. At least as important as having the features is that the users can actually find and use them, for which is important that some work and time is invested in this area.

References

- [Ames 07] Ames, M. and Naaman, M. ; "Why we tag: motivations for annotation in mobile and online media". In *CHI '07:Proceedings of the SIGCHI Conference on Human Factorsin Computing Systems*, pp. 971-980. New York: ACMPress, 2007.
- [AMG 09] "The All Music Guide" [online]. Available at <http://www.allmusic.com/cg/amg.dll?p=amg&sql=75> (retrieved on 12.06.2009) , 2009.
- [Anderson 06] Anderson, C.; "The Long Tail. Why the future of business is selling less of more". Hyperion Verlag, 2006.
- [Avery 97] Avery, C. and Zeckhauser, R.; "Recommender systems for evaluating computer messages". In *Communications of the ACM*, 40(3), pp. 88–89. , 1997.
- [Baumann 03] Baumann, S. and Hummel, O.; "Using cultural metadata for artist recommendations". In *Proc. Wedelmusic*, 2003.
- [Baumann 07] Baumann, S., Jung, B., Bassoli, A., and Wisniowski, M.; "Bluetuna: let your neighbour know what music you like". In *CHI '07 extended abstracts on Human factors in computing systems*, pp. 1941–1946, New York, NY, USA. ACM, 2007.
- [Begelman 06] Begelman, G., Keller, P. & Smadja, F.; "Automated tag clustering: improving search and exploration in thetag space". 2006.
- [Burg 06] Burg, Thomas N.; Schmidt, Jan (Hrsg.) S; "BlogTalks reloaded : social software - research & cases". In *Norderstedt: Books On Demand*, pp. 31-49. 2006.
- [Burke 02] Burke, R. ; "Hybrid recommender systems: Survey and experiments". In *User Modeling and User-Adapted Interaction*, 12(4), pp. 331–370, 2002.
- [Bush 45] Bush, V. ; "As we may think". In *Perspectives on the Computer Revolution*, Z. W. Pylyshyn and L. J. Bannon, Eds. Ablex Publishing Corp., 1989, pp. 49-61. , 1945.

Conclusions and Future Work

- [Celma 06] Celma, O., Cano, P. and Herrera, P.; "Search sounds: An audio crawler focused on weblogs". In *ISMIR 2006*, 2006.
- [Celma 08] Celma, Oscar ; "Music Recommendation and Discovery in the Long Tail". 2008.
- [Celma 08b] Celma, O.; Cano, P.; "From hits to niches? or how popular artists can bias music recommendation and discovery". In *Workshop on Large-Scale Recommender Systems*, 2008.
- [Charron 06] Charron, C., Favier, J. and Li, C.; "Social Computing: How Networks Erode Institutional Power, and What to Do about It". In *Forrester Customer Report 2006*, 2006.
- [Claypool 99] Claypool, M., Gokhale, A., Miranda, T., and Murnikov, P.; "Combining content-based and collaborative fillters in an online newspaper". In *Proceedings of ACM SIGIR Workshop on Recommender Systems*, 1999.
- [Deerwester 90] Deerwester, S.C., Dumais, S.T., Landauer, T.K., Furnas, G.W. & Harshman, R.A. ; "Indexing by latent semantic analysis". In *Journal of the American Society of Information Science*, 41(6), pp. 391-407, 1990.
- [Downie 01] Downie, J. Stephen; "Music Information Retrieval Annotated Bibliography Website Project, Phase I". In *2nd International Symposium on Music Information Retrieval (2001)*; available at <http://ismir2001.indiana.edu/posters/downie.pdf> , 2001.
- [Downie 03] Downie, J. Stephen; "Music information retrieval (Chapter 7)". In *Annual Review of Information Science and Technology 37*, ed. Blaise Cronin, 295-340. Medford, NJ: Information Today. Available from http://music-ir.org/downie_mir_arist37.pdf, 2003.
- [Eck 07] Eck, D., Lamere, P., Bertin-Mahieux, T. and Green, S.; "Automatic generation of social tags for music recommendation". In *Proc. NIPS Conf., 2007*, 2007.
- [Eck 07b] Eck, D., Bertin-Mahieux, T. and Lamere, P.; "Autotagging music using supervised machine learning". In *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)*, 2007.
- [Engelbart 62] Engelbart, D. C. and Friedewald, M.; "Augmenting Human Intellect a Conceptual Framework: A Conceptual Framework". Bootstrap Alliance, 1997, 1962.
- [Farnsworth 58] Farnsworth, P. R. ; "The social psychology of music". The Dryden Press, 1958.
- [Fingerhut 04] Fingerhut, Michael; "Music Information Retrieval, or how to search for (and maybe find) music and do away with incipits". In *IAML-IASA 2004 Congress. Oslo, August 2004* , 2004.

Conclusions and Future Work

- [Firan 07] Firan, C. S., Nejdl, W., and Paiu, R.; "The benefit of using tag-based profiles". In *Proceedings of the 2007 Latin American Web Conference (LA-WEB)*, pp. 32–41, Washington, DC, USA. IEEE Computer Society., 2007.
- [Freund 96] Freund, Y. and Shapire, R.E.; "Experiments with a new boosting algorithm". In *Machine Learning: Proceedings of the Thirteenth International Conference*, pp. 148–156, 1996.
- [Fumas 87] Fumas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T.; "The vocabulary problem in human-system communication". In *Communications of the ACM*, 30, pp. 964-971. 1987.
- [Futrelle 02] Futrelle, Joe ; "Interdisciplinary communities and research issues in music information retrieval". In *Music Information Retrieval," in Proc. Int. Conf. on Music Information Retrieval (ISMIR 2002)*, 2002.
- [Geleijnse 06] Geleijnse, G. and Korst, J.; "Web-based artist categorization". In *Proceedings of the 7th International Conference on Music Information Retrieval*, pp. 266 – 271, 2006.
- [Goldberg 92] Goldberg, D., D, N., Oki, B. M., and Terry, D. ; "Collaborative filtering to weave and information tapestry". In *Communications of the ACM*, 35(12), pp. 61–70. , 1992.
- [Gunawardana 08] Gunawardana, A. and Meek, C; "Tied boltzmann machines for cold start recommendations". In *Proceedings of the ACM Conference on Recommender Systems 2008*, pp. 19-26. , 2008.
- [Guy 06] Guy, M. and Tonkin, E.; "Folksonomies: Tidying up Tags?". In *D-Lib Magazine, January 2006, Volume 12 Number 1 ISSN 1082-9873*, 2006.
- [Hassan-Montero 06] Hassan-Montero, Y. & Herrero-Solana, V.; "Improving tag-clouds as visual information retrieval interfaces". In *Proceedings of the First International Conference on Multidisciplinary Information Sciences and Technologies(InSciT 2006), Mérida, Spain., 2006*.
- [Herlocker 00] Herlocker, J. L., Konstan, J. A., and Riedl, J. ; "Explaining collaborative filtering recommendations". In *Proceedings of the 2000 ACM conference on Computer supported cooperative work, pages 241–250, New York, NY, USA. ACM., 2000*.
- [Herlocker 04] Herlocker, J. L., Konstan, J. A., Terveen, L. G., and Riedl, J. T.; "Evaluating collaborative filtering recommender systems". In *ACM Trans. Inf. Syst.*, 22(1), pp. 5–53. , 2004.
- [Hevner 36] Hevner, K. ; "Experimental studies of the elements of expression in music". In *American Journal of Psychology*, no. 48, pp. 246–268. , 1936.

Conclusions and Future Work

- [ISO/IEC 04] ISO/IEC JTC 1 Joint Technical Committee; "Information technology. Multimedia framework (MPEG-21). Part 11: Evaluation Tools for Persistent Association Technologies". 2004.
- [Jennings 07] Jennings, D; "Net, Blogs and Rock 'n' Roll: How Digital Discovery Works and What it Means for Consumers". Nicolas Brealey Publishing, 2007.
- [Kassler 66] Kassler, M; "Toward Musical Information Retrieval". In *Perspectives of New Music* 4(2): 59-67, 1966.
- [Kazienko 06] Kazienko, P. and Musial, K.; "Recommendation framework for online social networks". In *Advances in Web Intelligence and Data Mining, volume 23 of Studies in Computational Intelligence, pages 111–120. Springer*, 2006.
- [Kilkki 07] Kilkki, K.; "A practical model for analyzing long tails". In *First Monday*, 12(5), 2007.
- [Knees 04] Knees, P., Pampalk, E. and Widmer, G.; "Artist classification with web-based data". In *Proc. ISMIR 2004*, 2004.
- [Knees 06] Knees, P., Schedl, M., Pohle, T., and Widmer, G.; "An innovative three-dimensional user interface for exploring music collections enriched with Meta-Information from the Web". In *Proceedings of the 14th Annual ACM international Conference on Multimedia (Santa Barbara, CA, USA, October 23 - 27, 2006). MULTIMEDIA '06*, pp. 17-24. ACM, New York, NY, 2006.
- [Knees 07] Knees, P., Pohle, T., Schedl, and Widmer, G.; "A music search engine built upon audio-based and web-based similarity measures". In *ACM SIGIR 2007*, 2007.
- [Knees 08] Knees, P., Pohle, T., Schedl, M., Schnitzer, D. and Seyerlehner, K.; "A document-centered approach to a natural language music search engine". In *ECIR, 2008*, 2008.
- [Lamere 07] Lamere, P.; Celma, O.; "Music recommendation tutorial". In *n International Symposium on Music Information Retrieval, 2007*, 2007.
- [Lamere 07b] Lamere, P. and Eck, D.; "Using 3D Visualizations to Explore and Discover Music". In *Proceedings of ISMIR 2007 Eighth International Conference on Music Information Retrieval, September 2007*, 2007.
- [Lamere 08b] Lamere, Paul; "Social Tagging and Music Information Retrieval". In *Journal of New Music Research*, pp. 101-114, 2008.
- [Lamere 09] Lamere, Paul; Volodkin, Anthony; "Help! My iPod thinks I'm emo". In *SXSW interactive keynote presentation, 17 March 2009*.

Conclusions and Future Work

- [Law 07] Law, E., von Ahn, L. and Dannenberg, R.; "Tagatune: a game for music and sound annotation". In *International Symposium on Music Information Retrieval, 2007*, 2007.
- [Lerdorf 02] Lerdorf, Rasmus; Tatroe, Kevin; "Programming PHP". In Nathan Torkington and Paula Ferguson , O'Reilly Media, Inc., 2002.
- [Lesaffre 06] Lesaffre, M., Leman, M. and Martens, J.-P.; "A user-oriented approach to music information retrieval". In *In Content-Based Retrieval*, 2006.
- [Levy 07] Levy, M. and Sandler, M.; "A semantic space for music derived from social tags". In *Proc.ISMIR Conf.*, pp. 411–416. , 2007.
- [Li 03] Li, T. ; "Detecting emotion in music". In *Proc. 4th International Conference on Music Information Retrieval, Baltimore, Maryland*, 2003.
- [Li 06] Li, T. and Ogihara, M.; "Toward intelligent music information retrieval". In *IEEE Transactions on Multimedia*, 8(3), pp. 564–574. , 2006.
- [Licklider 68] Licklider, J. C. R. and Taylor, Robert W. ; "The Computer as a Communication Device". In *Journal of Science and Technology*, vol. 76, pp. 21-31. , 1968.
- [Lippincott 02] Lippincott, Aura; "Issues in content-based music information retrieval". In *Journal of Information Science*, 28; 137, 2002.
- [Lu 06] Lu, L., Liu, D. and Zhang, H.-J. ; "Automatic mood detection and tracking of music audio signals". In *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 14, no. 1, pp. 5–18. , 2006.
- [Mahalanobis 36] Mahalanobis, P. C.; "On the generalised distance in statistics". In *Proceedings of the National Institute of Sciences of India 2 (1)*, pp. 49–55. , 1936.
- [Maillet 08] Maillet, Francois and Lamere, Paul; "Creating Transparent, Steerable Recommendations". In *Late-breaking Proceedings of the 8th International Conference on Music Information Retrieval Philadelphia, USA: (2008)*, 2008.
- [Mandel 07] Mandel, M. and Ellis, D.; "A web-based game for collecting music metadata". In *International Symposium on Music Information Retrieval, 2007*, 2007.
- [Marlow 06] Marlow, C., Naaman, M., Boyd, D. & Davis, M.; "Proceedings of the Seventeenth Conference on Hypertext and Hypermedia". In *HIPERTEXT'06*, pp. 31-40. New York: ACM Press, 2006.
- [Massa 07] Massa, P. and Avesani, P; "Trust-aware recommender systems". In *Proceedings of the 2007 ACM Conference on Recommender Systems*

Conclusions and Future Work

(Minneapolis, MN, USA, October 19 - 20, 2007). *RecSys '07*. ACM, New York, NY, 17-24., 2007.

- [McKay 06] McKay, C. and Fujinaga, I; "Musical genre classification: Is it worth pursuing and how can it be improved?". In *ISMIR 2006*, 2006.
- [McNee 06] McNee, S. M., Riedl, J., and Konstan, J. A; "Being accurate is not enough: how accuracy metrics have hurt recommender systems". In *CHI '06 Extended Abstracts on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006)*. *CHI '06*. ACM, New York, NY, 1097-1101, 2006.
- [Mitchell 03] Mitchell, Joan S. and Dewey, Melvil; "Summaries Ddc 22: Dewey Decimal Classification". Published by OCLC, 2003.
- [Montaner 03] Montaner, M., Lopez, B., and de la Rosa, J. L.; "A taxonomy of recommender agents on the internet". In *Artificial Intelligence Review*, 19, pp. 285–330. , 2003.
- [Morville 05] Morville, P.; "Ambient Findability". O'Reilly, 2005.
- [Orio 06] Orio, Nicola; "Music Retrieval: A Tutorial and Review". Now Publishers Inc, 2006.
- [Pachet 05] Pachet, F.; "Knowledge Management and Musical Metadata". In *Idea Group*, 2005.
- [Pampalk 02a] Pampalk, E., Rauber, A. and Merkl, D. ; "Content-based Organization and Visualization of Music Archives". In *Proceedings of the ACM Multimedia 2002*, pp. 570-579. , 2002.
- [Pampalk 02b] Pampalk, E., Rauber, A. and Merkl, D.; "Smoothed Data Histograms for Cluster Visualization with Self-Organizing Maps". In *Proceedings of the International Conference on Artificial Neural Networks (ICANN'02)*, 2002.
- [Resnick 94] Resnick, P., Iacovou, N., Suchak, M., Bergstorm, P., and Riedl, J. ; "Grouplens: An open architecture for collaborative filtering of netnews". In *Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work*, pp. 175–186. ACM, ACM Press, 1994.
- [Russell 80] Russell, J. A.; "A complex model of affect". In *J. Personality Social Psychology*, vol. 39, pp. 1161–1178. , 1980.
- [Sandvold 06] Sandvold, V., Aussenac, T., Celma, O., and Herrera, P.; "Good vibrations: Music discovery through personal musical concepts". In *Proceedings of 7th International Conference on Music Information Retrieval, Victoria, Canada*, 2006.
- [Sarmiento 09] Sarmiento L., Gouyon F., Oliveira E.; "Music Artist Tag Propagation with Wikipedia abstracts". In *Workshop on Information Retrieval over*

Conclusions and Future Work

Social Networks, European Conference on Information Retrieval, Toulouse, 2009.

- [Sarmiento 09b] Sarmiento, L.; Gouyon, F.; Costa, B. and Oliveira, E. ; "Visualizing Networks of Music Artists with RAMA". pp. 232-237. INSTICC Press, 2009.
- [Sarwar 01] Sarwar, B., Karypis, G., Konstan, J., and Reidl, J. ; "Item-based collaborative filtering recommendation algorithms". In *n WWW'01: Proceedings of 10th International Conference on World Wide Web*, pp. 285-295. , 2001.
- [Schedl 05] Schedl, M., Knees, P., and Widmer, G.; "A web-based approach to assessing artist similarity using co-occurrences". In *Proceedings of 4th International Workshop on Content-Based Multimedia Indexing (CBMI'05)*, 2005.
- [Schein 02] A. Schein, A. Popescul, L. Ungar, and D. Pennock; "Methods and metrics for cold-start recommendations". In *Proceedings of the 25th Annual International ACM SIGIR Conference*, pages 253-260, 2002.
- [Schubert 99] Schubert, E. ; "Measurement and Time Series Analysis of Emotion in Music". In *PhD thesis, University of New South Wales*, 1999.
- [Schuler 94] Schuler, D.; "Social Computing". In *Comm. ACM*, vol. 37, no. 1, pp. 28-29. , 1994.
- [Shardanand 94] Shardanand, U.; "Social information filtering for music recommendation". In *Master's thesis, Massachussets Institute of Technology*, 1994.
- [Shardanand 95] Shardanand, U. and Maes, P.; "Social information filtering: Algorithms for automating "word of mouth"". In *Proceedings of CHI'95*, 1995.
- [Sinha 02] Sinha, R. and Swearingen, K.; "The role of transparency in recommender systems". In *CHI '02 extended abstracts on Human factors in computing systems*, pages 830–831, New York, NY, USA. ACM. , 2002.
- [Sordo 07] Sordo, M., Lauier, C. and Celma, O.; "Annotating music collections: How content-based similarity helps to propagate labels". In *ISMIR 2007*, 2007.
- [Sotiropoulos 07] Sotiropoulos, D. N., Lampropoulos, A. S., and Tsihrintzis, G. A.; "Evaluation of modeling music similarity perception via feature subset selection". In *User Modeling, volume 4511 of Lecture Notes in Computer Science*, pages 288–297. Springer., 2007.
- [Surowiecki 04] Surowiecki, J.; "The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations". Doubleday, 2004.

Conclusions and Future Work

- [Thayer 89] Thayer, R. E.; "The Biopsychology of Mood and Arousal". Oxford, U.K.: Oxford Univ. Press, 1989.
- [Tintarev 07] Tintarev, N. and Masthoff, J.; "Effective explanations of recommendations: user-centered design". In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 153–156, Minneapolis, MN, USA. ACM, 2007.
- [Turnbull 07] Turnbull, D., Liu, R., Barrington, L. and Lanckriet, G.; "Using games to collect semantic information about music". In *International Symposium on Music Information Retrieval, 2007*, 2007.
- [Turnbull 08] D. Turnbull, L. Barrington, and G. Lanckriet; "Five approaches to collecting tags for music". In *International Symposium on Music Information Retrieval, 2008. In Proceedings of the 9th International Conference on Music Information Retrieval*, pages 225–230, Philadelphia, USA., 2008.
- [Turnbull 08] D. Turnbull, L. Barrington, and G. Lanckriet; "Five approaches to collecting tags for music". In *International Symposium on Music Information Retrieval, 2008. In Proceedings of the 9th International Conference on Music Information Retrieval*, Philadelphia, USA., pp. 225–230. , 2008.
- [Turnbull 08b] Turnbull, D., Barrington, L., Torres, D. and Lanckriet, G.; "Semantic annotation and retrieval of music and sound effects". In *IEEE TASLP*, 16(2), 2008.
- [Turney 01] Turney, P.D.; "Mining the Web for synonyms: PMI-IR versus LSA on TOEFL". In *Proceedings of the Twelfth European Conference on Machine Learning* , pp. 491-502. , 2001.
- [Turoff 72] Turoff, M.; "Delphi Conferencing: Computer Based Conferencing with Anonymity". In *Journal of Technological Forecasting and Social Change*, (3:2), pp. 159-204. , 1972.
- [Tzanetakis 02] Tzanetakis, G. and Cook, P.; "Musical genre classification of audio signals". In *Speech and Audio Processing, IEEE Transactions on* , vol.10, no.5, pp. 293-302. , 2002.
- [Tzanetakis 07] Tzanetakis, G. ; "Marsyas submissions to MIREX 2007". In *MIREX 2007*, 2007.
- [Uitdenbogerd 00] Uitdenbogerd, A.L.,Chattarajam, A. and Zobel, J; "Music IR: past, present and future". In *International Symposium on Music Information Retrieval, Plymouth, MA, 23–25 October 2000*, 2000.

Conclusions and Future Work

- [Uitdenbogerd 02] Uitdenbogerd, A. and van Schnydel, R.; "A review of factors affecting music recommender success". In *Proceedings of 3rd International Conference on Music Information Retrieval, Paris, France.*, 2002.
- [Vignoli 05] Vignoli, F. and Pauws, S.; "A music retrieval system based on user driven similarity and its evaluation". In *n Proceedings of the 6th International Conference on Music Information Retrieval*, pages 272–279, London, UK, 2005.
- [von Ahn 04] von Ahn, L. and Dabbish, L.; "Labeling images with a computer game". In *ACM CHI '04, 2004*, 2004.
- [von Ahn 06a] von Ahn, L., Ginosar, S., Kedia, M., Liu, R. and Blum, M.; "Improving accessibility of the web with a computer game". In *ACM CHI Notes, 2006*, 2006.
- [von Ahn 06b] von Ahn, L., Liu, R. and Blum, M.; "Peekaboom: A game for locating objects in images". In *ACM CHI '06, 2006*, 2006.
- [Wang 05] Wang, F.; "Social Computing: A Digital and Dynamical Integration of Science, Technology, and Human and Social Studies". In *China Basic Science*, vol. 7, no. 5, pp. 5–12. , 2005.
- [Watts 07] Watts, Duncan J.; "Is Justin Timberlake a Product of Cumulative Advantage?". In *New York Times (April 15, 2007)*, 2007.
- [Weinberger 07] Weinberger, D.; "How tagging changes peoples relationship to information and each other". Pew Internet & American Life Project, 2007.
- [Whitman 02] Whitman, B. and Lawrence, S.; "Inferring descriptions and similarity for music from community metadata". In *Proceedings of International Computer Music Conference, Goteborg, Sweden*, 2002.
- [Whitman 03] Whitman, B.; "Semantic rank reduction of music audio". In *Proceedings of the 2003 Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pp. 135–138. , 2003.
- [Whitman 04] Whitman, B. and Ellis, D.; "Automatic record reviews". In *ISMIR 2004*, pp. 470–477. , 2004.
- [Wieczorkowska 06] Wieczorkowska, A., Synak, P. and Ras, Z.W. ; "Multi-label classification of emotions in music". In *Proceedings of the 2006 International Conference on Intelligent Information Processing and Web Mining (IIPWM'06)*, pp. 307–315. , 2006.
- [Williams 07] Williams, Chad; Mobasher, Bamshad; Burke, Robin; "Defending recommender systems: detection of profile injection attacks". In *Service Oriented Computing and Applications*, Vol. 1, No. 3. (30 November 2007), pp. 157-170., 2007.

Conclusions and Future Work

- [Yang 06] Yang, Y.-H., Liu, C.-C. and Chen, H.-H. ; "Music emotion classification: A fuzzy approach". In *Proceedings of ACM Multimedia 2006 (MM'06)*, pp. 81–84. , 2006.

Annex A: Analysis of the project's areas of focus

Here is presented the individual evaluation of the analysed features. This analysis is made in order to obtain a prioritization of each feature, in order to determine the project's areas of focus.

Four factors are analysed in this evaluation: the effort to implement them (an estimation made by the author), the author's preference, their relevance for Palco Principal (a classification made by Palco's directors) and their innovation (based on the analysis made in Annex B: Platform analysis – the more common the feature, the less innovative).

The last column in the table refers to the feature's priority. This value is calculated as a weighted average of the previous four factors, with 50% of the weight attributed to the implementation effort and the remaining 50% for the other three factors, distributed as follows: preference 40%, relevance 40% and innovation 20%.

Regarding the code colour, the pallet of colours ranges from green (good) to red (bad). The individual evaluation is then presented in the following table:

	effort	preference	relevance	innovation	priority
classification system	2,0	2,0	1,5	2,0	1,8
tags	2,0	2,0	1,0	2,0	1,6
mood	2,0	2,0	2,0	1,5	1,9
historic (e.g. decade)	2,0	3,0	3,0	1,5	2,7
rating	2,0	2,0	1,0	3,0	1,8
collaborative tools	2,0	3,0	3,0	2,0	2,8
wiki	2,5	3,0	2,0	3,0	2,4
software (desktop)	1,5	3,0	3,0	2,0	3,1
open API	1,5	3,0	3,0	2,0	3,1
events	1,5	1,0	1,0	2,0	1,5
calendar	1,5	2,0	1,0	2,0	1,9
georeferencing	1,0	1,0	1,0	1,0	1,5
venues	2,0	3,0	1,0	2,0	2,0

Analysis of the project's areas of focus

media	2,0	3,0	2,0	3,0	2,6
video/online TV	1,5	3,0	3,0	3,0	3,3
radio online	2,5	3,0	1,5	3,0	2,2
photos	2,0	3,0	2,0	3,0	2,6
monetize	1,0	3,0	5,0	2,5	4,2
online store	1,0	3,0	5,0	2,5	4,2
premium accounts	1,5	3,0	5,0	2,0	3,9
monitoring tools	1,5	3,0	1,0	2,0	2,3
user abuse/flag content	2,0	3,0	1,0	3,0	2,2
spam	1,5	3,0	1,0	2,0	2,3
fraud/illegal content	1,0	3,0	1,0	1,5	2,4
recommendation system	1,0	1,0	1,0	1,5	1,6
rankings/tops	2,0	2,0	1,0	3,0	1,8
events	1,0	1,0	1,0	1,5	1,6
users	1,0	1,0	1,0	1,5	1,6
recommended artists/songs	1,0	1,0	1,0	1,5	1,6
graphic visualization	1,0	2,0	2,0	1,0	2,3
reputation system - rankings	1,0	2,0	1,0	1,5	2,0
other content					
lyrics/tabs	2,0	3,0	1,0	2,0	2,0
publications/magazines	2,0	3,0	5,0	1,5	3,5
dynamic/automatic webzine	1,5	2,0	1,0	1,5	1,8
news	2,5	3,0	3,0	3,0	2,8
gear/equipment	2,5	3,0	3,0	2,0	2,6
social tools					
blog/journal	3,0	3,0	5,0	3,0	3,3
groups/forums	2,5	3,0	5,0	3,0	3,6
comments	3,0	3,0	5,0	3,0	3,3
messages	2,5	3,0	5,0	3,0	3,6
feeds RSS/atom	3,0	3,0	5,0	3,0	3,3
profile customization	2,5	3,0	2,0	2,5	2,3
online chat	1,0	3,0	5,0	1,5	4,0
activity stream	2,0	3,0	1,0	2,0	2,0

Grouping the features by their main area and calculating the average of their priorities, we obtain the final prioritization list, with the main project's areas of focus highlighted in bold, with an asterisk.

priority	area	effort
1,5	events*	1,5
1,6	recommendation system*	1,0
1,8	classification system*	2,0
2,0	reputation system	1,0
2,3	monitoring tools	1,5
2,6	media	2,0
2,8	collaborative tools	2,0
4,2	monetize	1

Analysis of the project's areas of focus

The list of features, ordered by their computed priority, is the following:

priority	feature	effort	priority	feature	effort
1,5	georeferencing*	1	2,4	wiki	2,5
1,6	tags*	2,0	2,4	fraud/illegal content	1,0
1,6	recommend events*	1,0	2,6	gear/equipment	2,5
1,6	recommend users*	1,0	2,6	photos	2,0
1,6	recommend artists/songs*	1,0	2,7	historic (e.g. decade)	2,0
1,8	dynamic/automatic webzine *	1,5	2,8	news	2,5
1,8	rating*	2,0	3,1	software (desktop)	1,5
1,8	rankings/tops*	2,0	3,1	open API	1,5
1,9	calendar	1,5	3,3	video/online TV	1,5
1,9	mood	2,0	3,3	blog/journal	3,0
2,0	venues	2,0	3,3	comments	3,0
2,0	reputation system - rankings	1,0	3,3	feeds RSS/atom	3,0
2,0	lyrics/tabs	2,0	3,5	publications/magazines	2,0
2,0	activity stream	2,0	3,6	groups/forums	2,5
2,2	radio online	2,5	3,6	messages	2,5
2,2	user abuse/flag content	2,0	3,9	premium accounts	1,5
2,3	spam	1,5	4,0	online chat	1,0
2,3	profile customization	2,5	4,2	online store	1
2,3	graphic visualization	1,0			

Annex B: Platform analysis

The results of the platform analysis, made during April 2009 are shown in the following tables. These platforms were chosen due to their dimension and popularity, but also using a more casuistic approach, after a period of search and discovery, due to the features they offered.

Here is the list of all the platforms analysed, as well as their URI (Uniform Resource Identifier). A more complete description of each of these platforms may be found in the following pages.

bebo	http://www.bebo.com	liveplasma.com	http://www.liveplasma.com
blip.fm	http://blip.fm	maestro.fm	http://www.maestro.fm
contraststream	http://contraststream.com	MOG	http://mog.com
deezer	http://www.deezer.com	mufin	http://www.mufin.com
dimvision's musicmap	http://www.dimvision.com/musicmap	musicoverly	http://www.musicoverly.com
fairtilizer	http://fairtilizer.com	mvspace	http://music.mvspace.com
finetune	http://www.finetune.com	mvstrands	http://www.mvstrands.com
fiql	http://www.fiql.com	palcoorincipal	http://palcoorincipal.sapo.pt
grooves shark	http://listen.grooves shark.com	pandora	http://www.pandora.com
hvpemachine	http://hvpem.com	plavlist.com	http://www.plavlist.com
iJegg	http://www.iiigg.com	purevolume	http://www.purevolume.com
iLike	http://www.ilike.com	radio.blog.club	http://www.radioblogclub.com
imeem	http://www.imeem.com	thesixtvone	http://www.thesixtvone.com
iangoo	http://www.iango.com	tune&ue's audiomap	http://audiomap.tune&ue.net
Lala.com	http://www.lala.com	turnitupmedia	http://www.turnitupmedia.com
last.fm	http://www.last.fm	wavango	http://www.wavango.com

After the platform description, the results of the analysis made are presented. In the first column of each table we can find the characteristics under analysis. Then, each column represents one platform (identified in the first cell). The last column, in the last table, has a count of the number of platforms where a given feature was found, allowing us to analyse its popularity. The platforms are displayed by descendant order of the number of features available.

Bebo

Bebo (an acronym for "Blog early, blog often"), is a social networking website, founded in January 2005. Bebo Music (originally as Bebo bands) was launched on July 11, 2006. This website allows bands or solo artists to create a profile showcasing their music. The band profiles include a band member list which is similar to the friends list on a normal profile, a list of fans (fans who have added the band as though it were a friend), an area for tour dates, a blog, and a list of songs which have been uploaded for playing through Bebo's media player, or which may be added to other user's playlists. These songs can be grouped by the band into albums, along with cover-art. All band members can edit the content of the band profile.

Blip.fm

Blip.fm, launched in March 2008, is a music recommendation and streaming site. Users post short messages on what music they are listening to at the moment, the site then turns them into streaming links using Imeem. Blip is essentially "Twitter for Music", where users post short messages on what music they are listening to at the moment. It allows the user to suggest music and share his thoughts about it within a network of contacts. No file uploads are necessary. All the user has to do is to search for the song he has in mind and Blip will grab it from Imeem. The users' followers ("fans") can hear full versions of the songs he posts using Imeem's player widget.

ContraStream

ContraStream was launched in September 2007. It allows users to tap into the wisdom of crowds, through community filtered independent ("indie") music. The website helps people find the best new indie music quickly, this way having a reliable source to find music through community recommended albums.

Artists upload indie music and others vote on it in a similar system as the one used in Digg⁶⁷ to push the good items to the top of the site. It is at least somewhat similar to iJigg, which also lets users vote on music using such a system. ContraStream will leverage the user-generated voting data to let users search/browse popular music. Each artist and album also gets its own dedicated page on the site. In an effort to "keep the music indie," users are encouraged to flag music that is "too mainstream."

Deezer

Deezer is a free music streaming service, launched in August 2006. Its concept is simple: to give consumers a full and free access to all their favourite songs. Deezer pays artists and their producers through a revenue share based on their advertising revenues, being the first music site based in France to sign an agreement with a publisher rights' organisation to reimburse artists through advertising revenue.

67 <http://www.digg.com>

Platform analysis

There are many features of the current site which includes the ability to create blind tests, and send them to friends on the site. Users can also create a blog page where they can present themselves, copy or share playlists with other members or add other members to their friend list. Also, a mail feature is available to members, allowing them to talk with other members and a music news feature, also available to non-members, which displays recent music news. An Internet radio feature is available where songs belonging to a certain radio are played randomly.

Dimvision's musicmap

Using Amazon.com music information, musicmap creates a living similar-music flow-chart. By typing the name of an artist, a bubble will appear on the empty page and by clicking more information from amazon will be provided. When the expand button is clicked, more little bubbles pop out from this bubble, each representing a similar album. Over time, and many clicks on expand, one should have a moving, criss-crossing page of bubbles connected to other bubbles and artist names.

Fairtilizer

Fairtilizer is an online music community, based in Geneva (Switzerland), aspiring to define a new generation of music media. It uses a simple, clean interface that offers good navigation.

Tracks submitted to Fairtilizer appear in the “Upcoming Section”, where the tracks that get the most support make it to the Charting Section. Based on a member’s status (member, junior member or senior member) the tracks appear in a buffer zone where the community members and an Editorial Board validate the quality and the ownership of the track. Members can see the tracks they have uploaded on their “My Tracks” page and each track is labelled with its status as it traverses the Buffer Zone, Upcoming Section or Charting section.

Standard tools for any Web 2.0 offering include: an embed feature, my playlist, my tracks, ratings, my votes, my profile and very good searchability either worldwide or localized. Sharing via email, embed and RSS are options and tracks may also be downloaded to iTunes player.

Finetune

Finetune is a music discovery service that launched in late 2006 and is focused on letting its users build shareable playlists of popular music. The user is able to browse and search through a fairly large directory of popular artists and pick out full versions of their songs based on 30-second previews. There is also a simple music discovery system that automatically builds user-specific playlists based on the user's artist picks. In addition to its Web-based player, Finetune also has an application for Adobe Air that both Mac and PC users can run on their machines. It lets the user listen to both his Finetune and iTunes playlists, as well as search by name for artists he enjoys and get a custom radio station of their work.

FIQL

FIQL is a social network centred on music playlists, based in New York and founded in early 2004. It allows users to upload, categorize and share playlists with others. FIQL then uses these playlists to grab music videos from third party video sites. They also provide an embeddable player. FIQL member are able to create a blog with photos and videos, comments, tags, and the ability to save playlists, as well as music recs, and contests. It is possible to import playlists from iTunes or anywhere else, and FIQL will find the videos to match.

Grooveshark

Grooveshark is an online music search engine, music streaming service and music recommendation web software application, launched in early 2007 in private beta. It allows users to search for and stream music, for free, that can be played immediately or added to a playlist.

Grooveshark also allows artists and record labels to promote themselves through the application with both advertisement and paying for listens from other users. All data and analytics relating to an artist (or group of artists) are freely available to promoters, such as the number of times their content has been listened to, along with demographic and geographic information about the fans.

One of its most notable feature is the recommendation system called "Autoplay", which finds similar songs to those in a user's playlist and queues them for playback. Similar to Pandora's "thumbs up" and "thumbs down" feedback mechanism, users of Grooveshark can tell the recommendation system whether a particular recommendation was good or not by clicking a "smile" or "frown" icon. When users are satisfied with the current list of songs in their queue, they are able to save the songs as a traditional playlist for later retrieval.

iJigg

iJigg is a social music site that lets users vote on submitted singles using Digg-like voting. It was launched in January 2007. The site's home page shows what popular or recent singles the iJigg community is "jiggling". Users can rate songs, embed songs on MySpace or find similar music fans. Users on iJigg can discover new music, rate music they like using a method similar to the one in digg, download free MP3s and share their music with friends. Bands or artists can post their music tracks, easily sharing their music and attracting new fans.

iLike

iLike (also available as an iTunes plug-in and a Facebook app), launched in October 2006, displays artist info, allows user to stream or download songs, and gives fans the opportunity to connect with their favourite musicians. It uses a wisdom of the crowds philosophy, tracking what users are listening and recommending songs and people with similar tastes. iLike is partnered with ticketmaster, a ticket-selling company, and allows users to directly buy tickets from iLike.

iMeem

iMeem is an online community, launched in late 2004, where artists, fans and friends can promote their content, share their tastes, and discover new blogs, photos, music and video. Users are able to tag, comment, rate and share any content, embed their media on other websites and stay-up-to-date with their personal network of music fans and get in-depth stats for their music contents.

The service has both a social network structure as well as a content browsing/filtering structure similar to that of Flickr and YouTube. It works on an advertising based business model, and is therefore free to use.

iTunes' Genius

Genius is a recommender system in the iTunes media player computer software that uses collaborative filtering to generate music recommendations. It was first introduced in iTunes 8, using the integrated rating system built into iTunes. This allows for better recommendations because of a larger data set.

When clicking on the arrows to view the music information on the iTunes Store, this software provides links on the side to other related music found in the Store. The user must be signed in to the iTunes store to use Genius. Once genius is active in iTunes, the data can be put into current generations of iPod Classic, iPod Nano, iPod Touch or the iPhone. Genius uses data from other genius users' musical tastes.

Jango

Jango, launched in 2007, is an online radio free social music site, whose company's office is based in New York (USA). Jango allows users to create their own custom radio stations, buy songs, look up lyrics, rate and share the songs they like.

Users can create channels that will increase the probability of listening to the songs they like. They can also tune in to any other station. The fan's player will show other people listening to the same artist, allowing the fan to "tune in" to any song from other users, message or share stations with each other.

Lala

Lala.com was launched in June 2006 as a membership service that facilitated CD swaps. It has since then evolved into a "cloud music store", allowing users to share their Lala collection to anyone, anywhere on the web. As in a web-to-PC synchronization (or cloud computing, like in Google Docs), Lala.com matches the music information on one's computer to its online catalogue of licensed music. Whatever a user has at his personal computer or on any playlist that he created can be added to his Lala.com collection. One of the biggest advantages of this website is also its easiness of use and intuitive interface.

Last.fm

Last.fm is a UK-based internet radio and music community website, founded in 2002. It is one of the world's largest social music platforms with over 15 million active users based in more than 232 countries. Using a unique music recommendation system known as "Audioscrobbler", Last.fm builds a detailed profile of each user's musical taste.

It allows users to recommend songs to other people as well as find songs that might interest from other people's recommendations as well. The site lets users tag music, write about it, and most importantly listen to the music online. Last.fm also provides a desktop software (scrobbler), where the user is able to create and manage playlists, view related artists and songs, favourite, share, and tag tracks. Last.fm also has a mobile app, available for both iPhone and Android.

LivePlasma

LivePlasma, launched in 2005, is a website that maps and displays music and movie search results with linkages and groupings, making a good use of Amazon.com's API. After the search term is submitted, it is immediately surrounded by other artists; the closer they are, the more similar they are in style to the target. The user can search, map, discover new movies or artists then save and share their maps. The site is still in beta, but the visual maps are easy to read and aesthetically pleasing.

Maestro.fm

Maestro.fm is a digital music social networking service launched in March 2008 and run by Maestro Music Inc. It allows users to remotely access their entire music library from any browser, interact with their music and their friends' music – viewing related content such as artist bios, reviews, lyrics, photos and videos – and purchase MP3s of the songs. Users can also scan their music collection, sync their iTunes playlists on-line, stream playlists to MySpace or Facebook and send them to friends using their desktop application "Maestro Connector".

MOG

MOG, launched in June 2005, is a blog-centric platform for people to publish thoughts about music, as well as a place to share songs, music library, videos and, using a desktop application – MOG-O-MATIC – that runs in the background and automatically detects songs no matter which music player is used. MOG has integrated Rhapsody's Streaming Music Collection, making it a discovery-oriented interface for streaming music. MOG is able to get instant recommendations and personalized content at the click of a button that starts the MOG-O-MATIC, the key to the MOG-experience. In April 2009, MOG redesigned their platform, presenting an improved user interface and making it easier to use.

Mufin

Mufin, located in Berlin (Germany), was founded as a spin-off of the Fraunhofer Institute and is now a subsidiary of MAGIX AG, having been launched in 2008. It works as an automated music recommendation engine. Its technology is based on digital signal processing (DSP) analysis and fingerprints, which identifies, introduces and showcases music from all over the world to people looking for something new and specific to their taste.

After searching the site's 3.5 million song database for one of their favourite songs, users can press "Similar Tracks" to view a short list of songs that have been determined by the site's powerful algorithm to contain similar characteristics. The algorithm analyses 40 characteristics of each song, including tempo, sound density, and variety of other factors. Users can search for songs both from Mufin's homepage and through an embeddable MySpace widget (with Facebook on the way).

Musicoverly

Musicoverly, whose team comprises the creator of LivePlasma, launched in November 2006, is an internet based radio station with the goal of discover new music for its users, by using a visual approach which offers users a fun and unique way to listen and interact with their custom playlists. Users can choose what music will be played not only by genre and time, but also by mood and temp. In one matrix, users can choose how energetic, calm, dark and positive the music shall be, and in another how danceable and what tempo it shall have. Then Musicoverly populates a musical road map with songs that match the user's interests, allowing him to add favourites, block songs, and narrow in on specific songs to repopulate the map. Also it allows users to purchase songs from eBay, Amazon, and iTunes.

Myspace Music

Myspace, launched in 1999, is more than just a place to upload photos, blog, or communicate. It has offered many musicians another way to reach their fans. Musicians can offer fans the opportunity to stream or downloads their music. Musicians and fans also get the chance to interact through comments, messages and bulletins. Many bands got their start on Myspace, one of the best examples of that may be Panic at the Disco. As a result of this, MySpace created a new website – MySpace Music – in order to clearly focus on the music industry. This website was launched in September 2008.

MyStrands

MyStrands (which started off as MusicStrands in 2003) is a downloaded desktop application that works with the user's music player to build recommendations based on what the user is listening to. It is able to recommend similar songs, albums, and artists, in a similar fashion to last.fm's scrobbler. By tying music on mobile phones, MyStrands is a bit ahead of others and helps tie all of a user's music collections together.

Pandora

Pandora, a company founded in January 2000, launched in July 2007 its own customizable internet radio. It is an automated music recommendation and internet radio service created by The Music Genome Project. Users enter a song or artist that they enjoy, and the service responds by playing selections that are musically similar. Pandora also offers users the ability to purchase the song they are listening to through Amazon.com or iTunes. It is built using OpenLaszlo and provides a clean and enriching experience to the user. It is possible to create stations around artists, songs or albums and to provide feedback (thumbs up or thumbs down) on the songs Pandora chooses. The site lets users create a personalized music listening experience through this customized Internet radio channel.

Project Playlist

Project Playlist, launched in November 2008, is an information location tool devoted entirely to the world of music. Their purpose is to help users find and enjoy music legally throughout the web and they, also adding a social community twist. With Project Playlist users can easily create playlists, share them with friends, and browse playlists of others. By this, Playlist.com can also be called a social networking site for music lovers.

PureVolume

PureVolume, launched in 2003, is a website for bands and musical artists to promote and advertise themselves. Each artist/band has its own miniature website which usually includes news, photos, upcoming shows, background/biographical information, contact information, and music for free listening and downloads. It is somehow similar to Myspace music pages, but differs in the fact that the entire website is set up for artists to promote their bands. It charts the most popular artists and songs so users can know what's been popular that day and maybe hear about the next big act.

Radio.blog.club

radio.blog.club is a music service that builds playlists based on an a specific artist or song, launched in 2007 and based in Paris (France). Radio.blog.club users search for songs, play it in the web based player and create embeddable playlists. Radio.blog.club has a Facebook applet and is accessible through the Opera browser in the Nintendo Wii. A store is also available in collaboration with Amazon. As of June 2009, it appears that radio.blog.club is working on a new version of the software, to be launched soon.

ReverbNation

ReverbNation, launched in October 2006, is a company that is responding to the changing music business paradigm. It provides a central site for musicians, producers, and venues to collaborate and communicate and is set up to handle a band's needs, offering email, gig finder, songs via IM, widgets, an online store, and free stats. Bands willing to add on ReverbNation's services will see increased storage for songs, digital distribution, and press kits.

ReverbNation also allows artists to sell their music. They can both decide that music belongs to the people and give it away for free or electronically distribute it on iTunes, Amazon mp3, Rhapsody, Napster, and e-music, for a small fee.

The Hype Machine

This social music site, launched in 2008, aggregates blog posts from various mp3 and music blogs, and post links to the blogs where the music tracks were mentioned on its main page. Users can preview the music on the Hype Machine site or purchase the track using the link provided. The site also features its own Internet radio called, Hype Radio where the songs that it aggregates are streamed.

Thesixtyone

Thesixtyone is a social music company launched in 2008. As in other social music websites, thesixtyone uses a rating system similar to the one of Digg. A Listener who signs up on the website is given a daily number of “hearts” he can use to vote on songs he likes. If the songs he “hearts” become popular, the user earns reputation points, levels up, and unlocks community abilities. Listeners can also earn achievement badges for various activities, like “hearting” songs in under-explored genres or reaching a playcount milestone. All tracks are fully streamable and can be organized in tag-based playlists.

Besides the rating system, every song can be commented or shared, and the fan can see who else liked the song. There are listening and leaderboards, showing the members with the highest reputations and their playlists of hearted songs. Artists can be subscribed to so that the user can be alerted when they upload new songs to the site. Songs can be shared by email, using a link or embedding it on another page.

Tuneglue's Audiomap

TuneGlue's audiomap is an online interactive music relationship explorer, using a force-directed graph for visualization. It doesn't offer streaming music or any other sort of media – it is just a visual map based on a rotational model for finding related artists. The system delves into the Last.fm database and explores relationships among musical artists. It also fetches data from Amazon.com, so that a user can proceed and buy select items. After a user types in his or her query, the hit is displayed as a node. Clicking it brings up a set of options, among which "Releases" shows an artist's albums and "Expand" displays other nodes for related artists.

Turnitup Media

TurnItUp Media, launched in 2008, is both an online music store and a social networking site where it is possible to purchase or earn legal music downloads by participating in short ad spots from their sponsors. Users can preview the hottest new songs, read about artists and entertainers, keep up on the Billboard charts, or just hang with their friends in a music-inspired setting. TurnItUp Media also acts as a forum for new artists, both unsigned and Indie, allowing them to showcase their music to the masses.

Wayango

Wayango, launched in November 2008, is a music website enabling music discovery, social connections and buying music. With Wayango, users are able to connect each other, use the Music Finder tool to find new music you they might like, read or submit music reviews, rate and track their favourite artists, albums and tracks, create playlists and blogs, upload photos and lyrics and do a music compatibility test, in order to compare the user's music taste with the one of his friends and other music lovers.

Platform analysis

	last.fm	imeem	mystrands	iLike	palcoprincipal	wayango	bebo	deezer	thesixtyone
classification system									
tags	•		•		•	•	•		•
genre	•	•	•		•	•	•	•	•
mood									•
historic	•		•						
rating	•	•	•	•	•	•			•
collaborative/integration									
tools									
wiki	•					•			
software (desktop)	•		•	•	•				
share/embed	•	•	•	•	•	•	•	•	•
mobile integration	•	•	•				•	•	
open API	•	•	•	•	•		•	•	
events									
calendar	•	•		•	•	•	•		•
georeferencing	•	•	•	•		?			
venues	•	•		•		?			
media									
video/online TV	•	•	•	•	•		•	•	•
artist radio	•	•	•	•		•	•	•	•
user playlist	•	•	•	•	•	•	•	•	•
photos	•	•	•		•	•	•	•	•
profiles									
label	•								
artist	•	•	•	•	•	•	•	•	•
albums	•	•	•	•	•	•	•	•	•
user	•	•	•	•	•	•	•	•	•
monetize									
online store	•	•	•	•		•	•	•	
premium accounts	•	•							
monitoring tools									
user abuse/flag content	•	•	•	•	•		•	•	•
spam detection		•	?				•		
fraud/illegal content detection	•		?						
other content									
lyrics/tabs					•	•			
news	•	•		•	•			•	
gear/equipment					•				
recommendation system									
rankings/tops	•	•	•	•	•		•	•	•
events	•			•					
users	•	•	•	•		•	•		•
recommended artists/songs	•	•	•	•	•	•	•	•	•
graphic visualization									
reputation system									
social tools									
blog/journal	•	•	•	•	•	•	•		•
groups/forums	•	•	•		•	•	•	•	•
comments	•	•	•	•	•	•	•	•	•
messages	•	•	•	•	•	•	•	•	
feeds RSS/atom	•	•	•	•	•	•		•	
profile customization		•				•	•		
activity stream	•	•		•		•	•		
play historic	•	•	•	•	•	•		•	•
microblogging		•				•	•		
follow/become fan		•	•	•	•	•	•	•	•
online chat							•		
features									
average features	38 80,9%	35 74,5%	28 59,6%	28 59,6%	28 59,6%	27 57,4%	25 53,2%	23 48,9%	23 48,9%
clusters									
average clusters	9,51 86,4%	7,99 72,6%	5,67 51,5%	5,90 53,7%	6,09 55,4%	5,05 45,9%	4,61 41,9%	4,77 43,4%	5,48 49,80%

Platform analysis

	maestro.fm	myspace	MOG	Lala.com*	jango	purevolume	finetune	fiql
classification system								
tags			?				•	
genre	•	•	•	•	•	•	•	
mood		•	?					
historic			?					
rating			•	•	•			
collaborative/integration tools								
wiki	•		?					
software (desktop)	•		?	•			•	
share/embed	•		?	•	•		•	•
mobile integration		•	?				•	
open API		•	?					
events								
calendar			?		•			
georeferencing			?			•		
venues			?		•			
media								
video/online TV	•	•	•		•	•		•
artist radio	•	•	•	•	•	•	•	•
user playlist	•	•	?	•	•	•		•
photos	•	•	•		•	•		
profiles								
label				•		•		
artist	•	•	•	•	•	•	•	•
albums	•		•	•		•	•	
user	•	•	•	•	•	•	•	•
monetize								
online store	•		•	•	•	•	•	•
premium accounts								
monitoring tools								
user abuse/flag content		•	?					
spam detection			?					
fraud/illegal content detection			?					
other content								
lyrics/tabs	•		•		•	•		•
news	•		•					
gear/equipment			?					
recommendation system								
rankings/tops	•		•	•		•		•
events								
users			•		•		•	
recommended artists/songs	•		•	•	•		•	
graphic visualization								
reputation system								
social tools								
blog/journal		•	•			•		•
groups/forums		•	?			•		•
comments	•	•	•	•	•	•	•	•
messages	•	•	•	•	•	•		
feeds RSS/atom	•	•	•	•				•
profile customization		•	?					
activity stream	•		?	•		•		
play historic	•		•	•	•			
microblogging		•	?					
follow/become fan	•	•	•	•	•	•	•	•
online chat		•						
features	22	21	21	20	20	20	14	14
average features	46,8%	44,7%	44,7%	42,6%	42,6%	42,6%	29,8%	29,8%
clusters	4,93	3,43	5,21	4,56	4,63	3,92	3,35	2,85
average clusters	44,9%	31,2%	47,4%	41,5%	42,1%	35,6%	30,5%	25,9%

Platform analysis

	contraststream	fairtilizer	mufin	hypemachine	musicoverly	playlist.com	iJigg
classification system							
tags		•					•
genre	•	•			•		•
mood					•		
historic					•		
rating	•			•	•		
collaborative/integration tools							
wiki							
software (desktop)			•		•		
share/embed		•				•	•
mobile integration					•		
open API							
events							
calendar	•			•			
georeferencing				•			
venues							
media							
video/online TV	•		•	•			
artist radio		•	•			•	
user playlist		•	•	•	•	•	•
photos						•	
profiles							
label		•					
artist			•				
albums			•				
user		•	•	•		•	•
monetize							
online store	•		•	•	•	•	
premium accounts		•			•		
monitoring tools							
user abuse/flag content	•		•				
spam detection	•						
fraud/illegal content detection							
other content							
lyrics/tabs							
news				•			
gear/equipment							
recommendation system							
rankings/tops	•	•	•				
events							
users							
recommended artists/songs	•		•		•		
graphic visualization					•		
reputation system							
social tools							
blog/journal				•		•	
groups/forums	•						
comments	•	•				•	•
messages	•					•	
feeds RSS/atom	•	•		•			
profile customization							
activity stream							•
play historic			•	•	•	•	•
microblogging							•
follow/become fan		•		•		•	•
online chat							
features	13	12	12	12	12	11	10
average features	27,7%	25,5%	25,5%	25,5%	25,5%	23,4%	21,3%
clusters	2,74	2,39	3,30	2,37	3,30	2,15	1,37
average clusters	24,9%	21,7%	30,0%	21,6%	30,0%	19,6%	12,50%

Platform analysis

	grooveshark	blip.fm	turnitupmedia**	radio.blog.club	tuneglue audiomap
classification system					
tags			•		
genre					
mood					
historic					
rating					
collaborative/integration					
tools					
wiki				•	
software (desktop)				•	
share/embed	•	•		•	
mobile integration					
open API		•			
events					
calendar					
georeferencing					
venues					
media					
video/online TV					
artist radio	•		•		
user playlist	•	•	•	•	
photos					
profiles					
label					
artist			•		•
albums			•		•
user		•	?		
monetize					
online store		•		•	•
premium accounts					
monitoring tools					
user abuse/flag content	•		?	•	
spam detection			?		
fraud/illegal content detection			?		
other content					
lyrics/tabs				•	
news					
gear/equipment					
recommendation system					
rankings/tops	•		•		
events					
users	•	•	?		
recommended artists/songs	•		•		•
graphic visualization					•
reputation system					
social tools					
blog/journal			?		
groups/forums			?		
comments			?		
messages			?		
feeds RSS/atom	•	•	?		
profile customization			?		
activity stream			?		
play historic	•	•	?		
microblogging		•	?		
follow/become fan	•	•	•		
online chat			?		
features	10	10	8	6	5
average features	21,3%	21,3%	17,0%	12,8%	10,6%
clusters	2,00	1,69	1,60	1,82	1,40
average clusters	18,2%	15,4%	14,5%	16,5%	12,7%

Platform analysis

	dimvision musicmap	liveplasma.com	pandora**	TOTAL
classification system				
tags			?	9
genre			?	20
mood			?	3
historic			?	3
rating		•	?	14
collaborative/integration tools				
wiki			?	3
software (desktop)			?	10
share/embed			?	20
mobile integration			?	8
open API			?	9
events				
calendar			?	7
georeferencing			?	5
venues			?	4
media				
video/online TV			?	17
artist radio			?	22
user playlist			?	24
photos			?	14
profiles				
label			?	4
artist	•		?	20
albums	•		?	18
user			?	23
monetize				
online store	•	•	?	24
premium accounts			?	4
monitoring tools				
user abuse/flag content			?	13
spam detection			?	3
fraud/illegal content detection			?	1
other content				
lyrics/tabs			?	8
news			?	8
gear/equipment			?	1
recommendation system				
rankings/tops			?	18
events			?	2
users			?	12
recommended artists/songs	•	•	?	22
graphic visualization	•	•	?	4
reputation system				
			?	5
social tools				
blog/journal			?	14
groups/forums			?	12
comments			?	21
messages			?	16
feeds RSS/atom			?	17
profile customization			?	4
activity stream			?	9
play historic			?	19
microblogging			?	6
follow/become fan			?	23
online chat			?	2
features	5	4	0	
average features	10,6%	8,5%	0,0%	
clusters	1,40	1,10	0,00	
average clusters	12,7%	10,0%	0,0%	

Platform analysis

The last 4 rows in each table contain the following information, for each column:

- **features** - the number of features available

```
for each cell {  
    if cell has "tick" {  
        features = features + 1  
    }  
}
```

- **average features** – the percentage of features available

```
avg_features = features / tot_features
```

- **clusters** – the number of clusters (groups of features) available

```
for each cluster{  
    features_clust = 0  
    for each cell {  
        if cell has "tick" {  
            features_clust = features_clust + 1  
        }  
    }  
    clusters = clusters + (features_clust / tot_features_clust)  
}
```

- **average clusters** – the percentage of clusters available

```
avg_clusters = clusters / tot_clusters
```

Annex C: Platforms' dimension – data sources

Here are the sources from where the data regarding the information with the platforms' dimension provided in Table 3.4 (page 22) was extracted.

	Reference	published in	retrieved on
last.fm	http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9021918	30-05-2007	06-05-2009
imeem	http://news.cnet.com/8301-13577_3-9902273-36.html	24-03-2008	06-05-2009
palcoprincipal	Palco Principal	2009	06-05-2009
iLike	http://blog.last.fm/2009/03/24/lastfm-radio-announcement	24-03-2009	06-05-2009
mystrands	http://siteanalytics.compete.com/mystrands.com/	2009	06-05-2009
wayango	http://siteanalytics.compete.com/wayango.com/	2009	06-05-2009
bebo	http://www.theaustralian.news.com.au/story/0,25197,22858110-7582,00.html	03-12-2007	06-05-2009
thesixtyone	http://www.techcrunch.com/2009/03/09/thesixtyone-is-building-a-digg-for-indie-music/	09-03-2009	06-05-2009
deezer	http://www.deezer.com/legal/about-us.php	2009	06-05-2009
maestro.fm	http://profile.maestro.fm/	2009	06-05-2009
MOG	http://mog.com/about	2009	06-05-2009
myspace	http://asterisk.tmcnet.com/topics/open-source/articles/54804-myspace-co-founder-step-down-as-ceo.htm	23-04-2009	06-05-2009
jango	http://www.jango.com/aboutus	2009	06-05-2009
Lala.com	http://siteanalytics.compete.com/lala.com/?metric=uv	2009	06-05-2009
purevolume	http://www.purevolume.com/_subpages/fan_browse.php?min_age=16&max_age=99	2009	06-05-2009

Annex D: Patents in MIR research

Here is a list with some of the most relevant patents in Music Information Research.

1. Music Recommendation system and method

Filing Date: 13/08/2004 – **Publication:** 17/02/2005

Authors (MusicIP, formerly Predixis): Hicken, Wendell T. (La Verne, CA, US); Holm, Frode (Santa Barbara, CA, US); Clune, James Edmond III (Glendora, CA, US); Campbell, Marc Elroy (Monrovia, CA, US)

Keywords: audio fingerprint, song similarity, playlist generation

United States Patent 20050038819

2. Method and system for music recommendation

Filing Date: 03/10/2003 – **Publication:** 10/06/2004

Authors (Polyphonic HMI): Alcalde, Vicenc Gaitan (Castella del Valles, ES); Ullod, Carlos Maria Lopez (Zaragoza, ES); Bonet, Antonio Trias (Sant Cugat del Valles, ES); Llopis, Antonio Trias (Sant Cugat del Valles, ES); Marcos, Jesus Sanz (Barcelona, ES); Ysern, Daniel Caldentey (Barcelona, ES); Arkwright, Dominic (Barcelona, ES)

Keywords: song similarity (FFT, chunks, avg. values), vector similarity, user's taste vector, relevance feedback

United States Patent 20040107821

3. Sharing music essence in a recommendation system

Filing Date: 23/05/2006 – **Publication:** 23/11/2006

Author (MusicIP): Hicken, Wendell T. (La Verne, CA, US)

Keywords: playlist characterization, playlist sharing, fill-in the gap, modify playlist

United States Patent 20060265349

4. Introducing new content items in a community-based recommendation system

Filing Date: 27/10/2003 - **Publication:** 20/04/2006

Authors (Philips): Bodlaender, Maarten Peter (Eindhoven, NL); Hollemans, Gerrit (Eindhoven, NL); Vignoli, Fabio (Eindhoven, NL)

Keywords: user community, comparing user profiles, generating a recommended user set for the user

United States Patent 20060085818

5. Client-based generation of music playlists from a server-provided subset of music similarity vectors

Filing Date: 27/01/2005 - **Publication:** 25/05/2006

Authors (Microsoft): Platt, John (Redmond, WA, US); Renshaw, Erin (Kirkland, WA, US)

Keywords: music similarity, hybrid graph, MDS embedding, euclidean space

United States Patent 20060112082

6. Method for content recommendation

Filing Date: 04/02/2008 – **Publication:** unknown

Authors (Sony Deutschland GmbH): Jana Eggink, Thomas Kemp, Wilhelm Hagg, Taufan Zimmer, Tomasz Feduszcak

United States Patent: 20090013002

7. System and method for the characterization, selection and recommendation of digital music and media content

Filing Date: 07/07/2008 – **Publication:** unknown

Authors (Rockbury Media International, C.V.): Avet Manukyan, Vartan Sarkissian

United States Patent: 20090013004

8. MyStrands

<http://labs.mystrands.com/patents.html>

16 pending patents (from 2004-2006)

Examples

“Personal music recommendation mapping applet overview”

“Sharing tags between individual user media libraries”

“User to user recommender”

“Freeing space for new media items on a mobile media playback device based on inferred user taste”

“Building and sharing a composite playlist from collective group tastes on multiple media playback devices”

Annex E: Proposed features

When describing the features, their audience is indicated. These are the feature's main target, although it might affect others. For the sake of simplicity, four entities are included in the audience: music fans and music artists (jointly referred to as “users”, as they will be the gross audience), venues (where events take place) and Palco 3.0 (referring to the system, in general).

E.1. Improving classification

E.1.1. Tag clouds

Benefits

- Provides a quick and straightforward overview of the tag collection
- Enables the display of the interrelations between tags
- Improves usability and browsability
- Visually appealing for users

Related features

Weighted genre classification, Mood classification and Tag suggestion

Overview

As mentioned in chapter 3.5.1 - Social tagging, there are three main types of tag clouds, according to their meaning (item-based, global or categorization clouds), and a number of visual displaying implementations.

In this case, the first two types of tag clouds may be applied. The item-based tag clouds apply to a specific item (an artist, song, album and even to news and blog posts), indicating which tags better classifies the item.

Proposed features

Using a global tag cloud, the user has an overview of Palco's content and its most relevant tags, with the tag weights being aggregated over the popularity of the tags applied to all items and by all users.

Also the user is able to see the tags he applied more frequently. This way, the user is presented not only with the social aspect of tagging, but also with a more personal perspective of the tags he commonly uses.

Regarding the visual display of the tag clouds, two aspects are important: the tag font (its size, colour, background colour and font style – bold, italic, underlined) and its position within the cloud.

The most common (and adequate, for this matter) display approach is to display tags with different font sizes, in order to depict their popularity (i.e. the bigger the font size, the most occurrences it has) and font styles, with italic used for the less popular and bold for the most popular tags. The usage of different background or font colours is discouraged, as it creates visual noise and distraction for the user.

As for the tag positioning within the cloud, the most common (and also appropriate in this case) approach is to sort the tags alphabetically. This way, finding a tag both by alphabet and by popularity is possible and quite straightforward.

An alternative approach to the tag positioning is to group tags by similarity, using clustering techniques, such as the ones described in section 3.5.1, in chapter 3.

This way, related tags would be close to each other, preventing the scattering of information in the tag cloud that happens when tags are sorted alphabetically or randomly.

This could be useful as the most popular groups of related tags would dominate the cloud, providing the user with a viewing tool that would not only highlight the most important tags, but also the relations among them, what is commonly known as lexical relations. However, because this is a more complex representation and requires a significant number of tags, as well a big displaying area, in order to have an effective impact, it should remain as an alternative to the traditional, alphabetic sorting display.

Usage

The item-based tag clouds are displayed in two situations: on the header of the artist profile and under the classification tab, also in the artist profile. In the later, two different item-based tag clouds are presented: one displaying all the tags used to classify the artist and another one to display the mood tags applied to the artist (a subset of the artist's tag collection).

The global tag cloud is displayed in two situations as well: in the main page of Palco Principal and on the tag navigation screen. In the later, two global tag clouds are also displayed: one for the global collection of tags (with all tags, applied to all items, by all users).

Implementation details

The tag size s_i may be obtained using a simple rule:

$$s_i = \left\lceil \frac{f_{\max} \cdot (t_i - t_{\min})}{t_{\max} - t_{\min}} \right\rceil$$

with s_i being the display font size, and $f_{\max}$ the maximum font size. t_i represents the current count of items, $t_{\min}$ the minimum and $t_{\max}$ the maximum count of items.

Some tag pre-processing is also necessary, in order to normalize tags, improve their quality and reduce noise. This pre-processing includes normalizing the tags' case (normally converting all the tags to low-case format) and removing punctuation.

Another practice for tag normalisation includes removing common language words (e.g. “the”, “of”, “a”), but since most of the tags applied to music content are related to its classification (into genres or moods), this is not necessary.

Popular examples

- flickr - <http://www.flickr.com/photos/tags>
- technorati - <http://technorati.com/tag>
- Last.fm - <http://www.last.fm/charts/toptags>
- del.icio.us - <http://delicious.com/tag>




Figure E.1: Last.fm's tagcloud

Proposed features

.net 2008 3d advertising ajax and animation api apple architecture art **article** articles artist **audio**
blog blogging blogs **book** books browser **business** car cms code collaboration comics community
computer converter cooking cool **CSS** culture data database **design** desktop development diy
documentation download downloads drupal ebooks economics education electronics email
entertainment environment fashion fic film finance firefox **flash** flex flickr food forum **free** freeware fun
funny gallery game **games** geek **google** government graphics green guide hardware health history
home hosting house **howto** html humor icons illustration images imported information inspiration
interactive interesting internet iphone japan java **javascript** jobs jquery kids language learning library
linux **list** lists literature **mac** magazine management maps marketing math media microsoft mobile
money movie movies mp3 **music** network networking **news** online **opensource** osx **people** phone
photo **photography** photos photoshop php plugin podcast **politics** portfolio privacy productivity
programming psychology python radio rails realestate recipe recipes **reference** religion research
resources reviews rss ruby rubyonrails school **science** search security seo shop shopping social
socialnetworking **software** statistics streaming teaching tech **technology** tips todo tool tools toread
travel **tutorial** tutorials tv twitter typography ubuntu usability **video** videos vim **visualization** web
web2.0 webdesign webdev **wiki** **wikipedia** windows wishlist wordpress work writing youtube

Figure E.2: Del.icio.us's tagcloud

E.1.2. Multi-tag search

Benefits

- Help finding specific items
- Discover new items, by combining different tags and narrowing the results
- Providing a way of steering recommendations, by filtering user recommendations according to the tags used

Related features

Tag clouds

Overview

Tags allow users to freely classify content, as mentioned before. And although this type of classification helps guiding the user to find a certain item, the use of a single tag when searching for content does not allow the use of more specific queries.

As in a search engine, the use of several keywords (tags) helps the user specifying what he is searching for, may it be for finding a specific item or discovering new content.

For example, when searching for an artist whose name one can not recall, the use of several tags used to classify the artist reduce substantially the number of possible results. On the other hand, assuming that the user is a fan of electronic music and wants to discover new artists, he can narrow down his search by specifying that he would like to find artists who are tagged as “electronic” and also “minimalist”, as well as “british” and with a “female” vocalist.

Usage

The multi-tag search feature can be used directly in the search bar displayed in the page header (always visible) or in the main page of tags, by dragging and dropping the tags displayed in the tag cloud into the search box.

Implementation details

The tags can be added to the search query either manually (typing the text) or by clicking in existing tags.

To build the tag query, tags must be separated by spaces, with multi-word tags being written between quotation marks. Also tags can be removed from the query by using a hyphen immediately before the desired tag. For example, a query containing

-jazzy electronic “female vocalist” from -england brazil

would return all the results containing the tags “electronic”, “female vocalist”, “from” and “brazil”, which did not contain the tags “jazzy” and “england”.

Popular examples

- Last.fm's - <http://playground.last.fm/multitag>
- del.icio.us - <http://delicious.com>



Figure E.3: del.icio.us multi-tag query

E.1.3. Tag suggestion

Benefits

- Prevent the increase of the tag collection's dimensionality and the data sparsity of tags
- Reduce the tag spam (or tag pollution), caused by the vocabulary problem (typing errors, synonymous or polysemous tags)
- Facilitate, stimulate and improve the user-based classification

Related features

Acoustic fingerprinting, Weighted genre classification and Mood classification

Overview

With a tag suggestion mechanism, users can receive suggestions of tags that could apply to the item s/he is tagging. The suggestions can be made based on community or personal tags.

The system suggests the most popular tags applied by other users, as well as user-owned tags (by selecting the most popular tags in the community that also exist in the user's tag

Proposed features

collection). This is particularly useful, as the user is “guided” to reuse existing tags, avoiding the necessity of creating new ones.

Autocompletion may be also used to auxiliate and complement the tag suggesting feature. It involves the software predicting a word that the user wants to type in without the user actually typing it in completely.

Usage

The tag suggesting feature is used when the user decides to tag content, either by pressing the “tag this” button on the page header (always visible) or when viewing the classification page of one of the following items: artists, songs, news, blog posts, reviews, events or venues.

Implementation details

Tag suggesting can be achieved using several algorithms, already introduced in the state of the art chapter. Also, autocompletion may be a useful tool, in order to avoid typing errors and auxiliate the user using tags.


Popular examples

del.icio.us - <http://delicious.com>

tagzania - <http://www.tagzania.com>

flickr's zonetag - <http://zonetag.research.yahoo.com>

Tagyu (no longer available) - <http://www.programmableweb.com/api/tagyu>



The screenshot shows the 'del.icio.us Save a Bookmark' interface. At the top, it says 'Signed in as palco30'. The form has fields for 'URL' (http://www.cnet.com/), 'TITLE' (Product reviews and prices, software downloads, and tech news - CNET), 'NOTES' (1000 chars), and 'TAGS'. Below the 'TAGS' field, there is a help text: 'Separate Tags with spaces. E.g., hotels bargains newyork (not new york)'. Underneath, there are two sections: 'Recommended Tags: click to add from your existing tags' with the tag 'blog' listed, and 'Popular Tags: click to add from popular tags on Delicious' with a list of tags: 'technology', 'reviews', 'news', 'tech', 'computer', 'gadgets', 'technews', and 'computers'.

Figure E.4: Tag suggestions in del.icio.us

E.1.4. Acoustic fingerprinting

Benefits

- Automatically classifying content (autotagging)
- Addressing several issues such as the early rater bias, the cold start and weak labelling problem and the feedback loops caused by popular artists
- Perform an uniform analysis of the musical collection
- Create an identifier for each song

Related features

Weighted genre classification, Mood classification and Music recommendation

Overview

By using the content-based and machine learning filtering techniques mentioned in the section 3.4.4 (chapter 3), the music collection existing in Palco Principal can be analysed and classified, according to its acoustic characteristics.

Acoustical fingerprinting analyses the actual content of an audio file, as opposed to a written description of what is contained within it (the so-called metadata). This means that the filtering tool “listens” to the sound recording, and creates a small “fingerprint” that can then be compared against any other file. This fingerprint may be regarded as the song's DNA⁶⁸, which is able to identify it and distinguish it from other songs.

Usage

Acoustic fingerprinting can be used both in classification and recommendation.

The analysis and fingerprinting of audio files allows their classification, providing valuable information about instrumentation (e.g. electric guitar, drums, female voice, male voice, etc.), mood characteristics (e.g. rhythm, tempo, pitch, etc.) and genre (by comparing the fingerprints with predefined genre fingerprint patterns).

These characteristics are then translated into tags, which are automatically applied to the song and are used to create a genre and mood profile (please refer to the weighted genre classification and mood classification features in sections 4.1.5 and 4.1.6 for more detail).

When providing music recommendations, they can be made by computing fingerprint similarity, which may lead to suggesting songs (or artists) that “sound” very much alike. Further information on this may be found in the music recommendation feature, in section 4.2.1.

Implementation details

The development of this feature is mostly done by the audio retrieval team, reason why no further details will be presented in this section.

⁶⁸ Deoxyribonucleic acid (DNA) is a nucleic acid that contains the genetic instructions used in the development and functioning of living organisms. It is often regarded as a set of blueprints or a code, that uniquely identifies the organism it refers to.

Popular examples

- All Media Guide/Macrovision's LASSO is a commercial service that uses acoustic fingerprinting, and other techniques, to recognize music (US Patent 7,277,766) - http://www.macrovision.com/products/ce_manufacturers/metadata_media_recognition/lasso.htm
- Last.fm have their own method of acoustic fingerprinting via the Fingerprinter application - <http://blog.last.fm/2007/08/29/audio-fingerprinting-for-clean-metadata>
- MusicBrainz is a free and open content project for a music database that uses MusicIP's Open Fingerprint Architecture for fingerprinting and the MusicDNS service for identifying audio files - <http://musicbrainz.org/doc/PUID>

E.1.5. Weighted genre classification

Benefits

- Enhance artists or song browsing, together with music recommendation techniques
- Address the cold start problem and weak labelling of some artists, by automatically classifying songs into musical genres

Related features

Music recommendation and Acoustic fingerprinting

Overview

Musical genres are categorical descriptions used to structure and create some sort of hierarchical division for the enormous amount of music available. Classifying music into genres is not a consensual approach, but it is still widely used, both in brick-and-mortar stores and in digital music websites.

However, music classification is not trivial due to its subjectivity, as it is difficult to “fit” an artist or song into one specific music genre. This situation is also valid not only for “human” but also for machine-automated classification.

A particular musical genre has certain acoustic and statistical properties related to its instrumentation and rhythmic structure (its audio fingerprint). Using content-based filtering techniques it is possible to identify some of these properties in a song (or group of songs) and find similarities with some predefined genre patterns.

As in “human” classification, automated methods are not likely to result in a perfect match with a specific music genre, as it is unlikely that a song will fit exactly the predefined pattern, and the extraction and classification techniques are themselves subject to some inaccuracy.

The proposed solution for this problem is to classify a song or artist (with the genre classification of the artist being a mixed classification of its songs) into different weighted genres, in the form of ratings or percentages.

Usage

The genre classification is displayed in the artist's and song's profile. As an example, a song can be classified as being 60% rock, 40% pop and 10% jazz, thus giving an idea of the different music styles the song shares similarities with.

Classification (mainly by genre) is fundamental in music recommendation, and this particular form of genre classification, using weights assumes an even bigger importance as it allows more accurate recommendations. Also, genres are important for users, as they are able to listen to the genre-customized radios.

Implementation details

The weighted classification of a song or artist into musical genres has into account both its acoustic fingerprint and the classification done by the users, according to the various music genres that it resembles the most.

The classification using acoustic fingerprinting has already been explained (cf. section 3.5.5, in chapter 3). As for the user classification, it is done by mean of the user-defined tags, but also by manually assigning weights to the predefined genre categories. This way, a user that considers a song as fitting the “pop” and “rock” genres, steering his classification to determine which genres are more or less dominant.

Popular examples

None

E.1.6. Mood classification

Benefits

- Enhance artists or song browsing, together with music recommendation techniques
- Address the cold start problem and weak labelling of some artists, by automatically classifying songs into mood profiles
- Help bridging the semantic gap between users' perception and music items

Related features

Acoustic fingerprinting and Music recommendation

Overview

By introducing this feature, and using similar techniques to the ones of automated music classification, it is possible to classify music according to the user's “mood”. Using low-level audio analysis, a song can be classified as slow/fast, happy/sad, enthusiastic/calm, etc.

As in genre classification, it is not trivial to classify a song's mood due to its subjectivity. In this case, classifying a song to fit a certain mood is even more difficult, as the same song may cause very different reactions in different people. For example, a heavy-metal song, with a low

Proposed features

tempo may be considered calm for a heavy-metal enthusiast, but may still be regarded as very energetic by an enthusiast of classical music.

Again, the approach of classifying a song into different mood categories by using tags and steerable ratings seems to be the most appropriate, however it has some different *nuances* in relation to genre classification.

In order to facilitate the content-based classification, as well as the user interaction, the proposed model to follow is the one defined in Thayer's two-dimensional valence-arousal space [Thayer 89]. This model decomposes emotion in music according to two principal dimensions:

- valence: positive vs. negative (e.g., happy vs. sad) and
- arousal: high- vs. low-energy (e.g., energetic vs. calm).

According to this model, music can be broadly categorized into one of four quadrants: high valence and arousal (joy, exuberance), high valence and low arousal (contentment), low valence and high arousal (anger), and low valence and arousal (depression).

Usage

Mood classification is. As in genre classification, displayed in the artist's and song's profile. This classification is also important in music recommendation, as it allows more accurate recommendations, according to the fan's mood.

By choosing his current mood, the fan is able to receive music recommendation that fit this state of mind, with customized playlists.

For example, a song classified as 20% sad – 80% happy and 70% energetic – 30% calm would be placed in the high valence and arousal quadrant, representing a joyful or exuberant mood.

The user could then choose what kind of “mood” he is into, for example by choosing music that is happy and enthusiastic, and the system would provide him a playlist of music with these characteristics.

Implementation details

As with genres, mood classification has into account both the acoustic fingerprint of a song and the classification done by the users. By analysing a number of acoustic features the song can be automatically classified in one of the quadrants of the proposed 4-quadrant model, with the user being able to steer this classification, by means of tags (displayed in the form of text or small icons – commonly known as “smileys”) or changing the weighting of each of the dimensions.

The proposed model may be mapped into a 2-axis graph. Each axis represents one dimension (valence and arousal), allowing an unlimited combination of overall moods.

Popular examples

- All Music moods - <http://www.allmusic.com/cg/amg.dll?p=amg&sql=75>
- Musicoverly - <http://www.musicoverly.com>

Proposed features



Figure E.5: Musiccovery's mood-pad

E.1.7. Wiki-like content editing

Benefits

- Improve the amount and quality of the descriptive information available about artists, songs, lyrics, venues, events or tags
- Strengthen the community and the feeling of being part of it

Related features

None.

Overview

Using a wiki tool to maintain descriptive information of several items on the system takes advantage of the so-called *wisdom of crowds* or *crowdsourcing*. A seminal contribution to this topic is made by James Surowiecki in [Surowiecki 04], where the author explores a simple idea that has profound implications: large groups of people are *smarter* than an elite few, no matter how brilliant—better at solving problems, fostering innovation, coming to wise decisions, even predicting the future.

By allowing the community to contribute with information, instead of the current approach, where only one person is allowed to edit information, a contribution for the knowledge and strengthening of the community is made.

Usage

This approach would allow users to add or update information about a given artist, song, event, venue or tag.

The editor would be able to not only add descriptive information (e.g. the biography of an artist or the lyrics of a song), but also factual information (when and where the band was formed, which instruments they use, etc.).

Proposed features

Implementation details

There are several wiki implementations, which could be adapted to fit Palco's needs. In order to preserve a certain “order”, editions would only be permitted to registered users and, if requested or necessary, edition could be blocked (for example, if an artist didn't want anyone else to change his biography or its song's lyrics).

A basic wiki implementation keeps track of the changes made and who made the changes (the editing history), allowing to easily revert to previous versions of the content. A special syntax is used in the text edition, but also some tools for WYSIWYG⁶⁹ editing (e.g. tinymce or FCKeditor) can be introduced, in order to facilitate the task.

Popular examples

- Maestro.fm - <http://www.maestro.fm>
- Wayango - <http://www.wayango.com>
- Last.fm - <http://www.last.fm>




Figure E.6: Wiki-like editing in Last.fm

⁶⁹ WYSIWIG – “What You See Is What You Get”, a system in which content displayed during editing appears very similar to the final output

E.2. Improving recommendation

E.2.1. Music recommendation

Benefits

- Improve user's trust in the system, by providing more meaningful and transparent recommendations
- Provide recommendation steerability by approving or disapproving recommendations
- Help users discover related or novel music (artists or songs)
- Enhance navigation and browsability through the content

Related features

Tag clouds, Tag suggestion, Acoustic fingerprinting, Weighted genre classification, Mood classification, Mood classification, Rating system and Graphic recommendation

Overview

Recalling the recommendation techniques presented in section 3.4.4 (chapter 3), music recommendation can be done using content-based filtering (CB), collaborative filtering (CF) or hybrid techniques.

For CB recommendations, the acoustic fingerprinting feature is essential, since with this analysis it is possible to recommend music based on low-level characteristics, such as tempo, rhythm, timbre, pitch, beats per minute (BPM), etc.

Regarding CF recommendations, they are made based on the items ratings (naturally, well-rated items should be recommended in the first place), peer endorsed preferences (recommendation based on mind-like users preferences), socially endorsed preferences (based on a user's friends preferences) and on general community listening habits (having into consideration aspects such as the number of plays of a song or artist, for instance).

Also, an important part of recommendations is their meaning. By providing recommendation explanations, the system makes them more transparent for users, allowing them to understand why a certain item is being recommended to them and helping to gain their trust.

Information on whether the recommendation was useful or not is also greatly important. This way, users are able to train, steering its recommendations, making them more meaningful and accurate. This can be achieved by adjusting the weight of tags within a tag cloud or list, in order to obtain a new profile of recommendation, in a similar way to multi-tag searching, only with the difference of attributing weights to the various tags. This issue and its application to music is well documented and developed in the Music Explaura project [Maillet 08].

Usage

Recommendations are presented in both the artist and songs profile pages. They may be presented in the form of a list or by filtering which types of recommendation the user wants.

For each recommendation, it is desirable to provide an explanations on *why* (and not *how*) the recommendation was made, but since these might not interest some users, it is recommended to show the explanations only if the user wants to, by choosing that option or displaying the explanation when the mouse is over the artist or song.

Implementation details

In order to implement the recommendation system, it is necessary to use the data mining techniques presented in section 3.4.4 (with some practical examples in section 3.7, both in chapter 3). Since the development of this feature is done in cooperation with the data mining team, no further details on the implementation will be presented. However, some examples of desired outcomes of the recommendations, in the form of explanations, are left for discussion and improvement.

The recommendation explanations may come in various forms. Using CB techniques, recommendations can be made based audio or spectral similarities. They can be made based on the genre classification, as in [1] or based on the mood classification as in [2].

When applying CF techniques, the typical recommendation is the “shopping cart” recommendation, where we receive suggestions based on other user's behaviour, as shown in [3] and [4]. Also, the particular use of tags (and their grouping in tag clouds) is useful for recommendation. We can recommend music based on the number of shared tags, as in [5], or based on the similarity of their tag clouds, as in [6].

[1] X sounds like Y

[2] If you are in the mood for X, maybe you are also in the mood for Y

[3] Users that listened to X also listened to Y

[4] Your friends have been listening to X, maybe you will like it too!

[5] X and Y share the following tags: *tag1*, ..., *tagN* and *M* other tags

[6] If you like artists tagged as *tag1*, ... and *tagN*, you'll like *artist1*, ... and *artistM*

Popular examples

Recommendations with explanations:

- Thisnext - <http://www.thisnext.com> (gadget recommendations)
- Clerck Dogs - <http://www.clerkdogs.com> (movie recommendations)
- Pandora - <http://www.pandora.com> (music recommendations)

Recommendation without explanations

- iLike - <http://www.ilike.com> (music recommendations)
- Musicoverly - <http://musicoverly.com> (visual music recommendations)
- Amazon - <http://amazon.com> (shopping recommendations)
- Flixster - <http://www.flixster.com> (movie recommendations)

Proposed features

- iMeem - <http://www.meem.com> (music recommendations)



Figure E.7: Music recommendation on Pandora



Figure E.8: Music recommendations on iMeem

E.2.2. User affinity and recommendation

Benefits

- Improve the user's social network, by determining mind-like users and recommending them as possible friends
- Help providing peer endorsed recommendations, using the preferences of similar users
- Measure the musical compatibility between users
- Strengthen the network, by creating new (sub) communities of mind-like users

Related features

Music recommendation, Reputation system and Geographic Information System (GIS)

Proposed features

Overview

User affinity (also referred to as user similarity or compatibility) is applied to determine mind-like users, by analysing their profiles' activity. It may be related to their musical or social proximity.

In the first case, users with a high musical compatibility tend to listen to the same artists or share the same genre or mood tastes. This way, mind-like users share a large part of their musical collection, and the items which are not shared may be recommended to the other user, as it might be of his interest.

The users' social compatibility is determined by their social relations. Users that have a large number of common friends are also likely to know each other or to be socially close. This could lead to a “friend-of-a-friend” recommendation.

Also, using GIS techniques, users could be recommended based on their physical proximity. This way, they would be able to meet in person and/or learn to which events they would be attending and eventually meet.

The combination of these two factors, associated to the characteristic dynamics of a social network could improve not only the recommendations provided to users, but also strengthen the network, in a sense of users discovering their “musical soul mates”. This could lead to the creation of sub-groups or sub-communities within the whole “cloud” of users, which could improve the personal relations between users, one very important factor in social networks.

Usage

This feature is meant to be applied in the recommendation engine, in order to provide peer endorsed recommendations. These recommendations may come in the form of artists or friends.

The user compatibility is displayed when visiting another user's profile, and the user recommendations are made in the user's homepage or friend page.

Implementation details

As in the previous feature, to implement the user affinity and recommendation feature, it is necessary to use the collaborative filtering and data mining techniques presented in section 3.4.4 (chapter 3). Since the development of this feature is done in cooperation with the data mining team, no further details on the implementation will be presented.

Popular examples

- Last.fm - <http://www.last.fm>
- Facebook - <http://www.facebook.com>
-



Figure E.9: User compatibility in Last.fm, based on musical proximity



Figure E.10: User suggestions on Facebook, based on social proximity

E.2.3. Rating system

Benefits

- Possibility to rank songs, according to user-based ratings
- Influence the recommendation engine, in order to recommend positively rated items
- Bigger user interaction
- Adding knowledge and “teaching” the system

Related features

Reputation system and Music recommendation

Overview

One of the biggest advantages of social software is the ability to give a voice to the user, and that is also made by providing the users with tools to express their opinions.

The existent rating system in Palco Principal, is a bit limited, in the sense that it is not very intuitive and does not allow much user interaction (users are not able to see other users' ratings, can not change their ratings, etc.).

By facilitating and promoting the usage of the rating system, it would be easier to assess the overall quality of items, through social contribution (the so-called “wisdom of crowds”) and improve overall recommendation quality, by having the item's ratings into account.

Usage

The rating system can be used in a number of different items, such as artists, songs, news, blog posts, photos, videos, events, comments, reviews and lyrics.

Then, it may be applied to generate recommendations as well as influencing the music rankings or charts.

Implementation details

There are several approaches to the implementation of rating systems, based on the format (star-rating, thumbs up/down, numerical rating, etc.) and on the scale (e.g. 3-scale rating, as good, bad, indifferent; 5-scale rating, as bad, not that good, reasonable, good, very good; 10 scale rating; percentage rating).

For the case of music and media rating (i.e. when rating artists, songs, photos, videos, events and lyrics), the most appropriate implementation is the 5-star rating system, as it is very intuitive and is widely spread.

Regarding the user-added content, the most appropriate rating mechanism is the thumbs up/down, which could be applied to comments, reviews, news and blog posts.

Popular examples

- Rate Your Music - <http://rateyourmusic.com>
- Youtube - <http://www.youtube.com>

Proposed features

- iMeem - <http://www.imeem.com>

MangeyYeti (18 hours ago)

[Reply](#) | [Spam](#)

+3  

i have no idea why but the music the slow talking and the black and film seriously freaks me out

Figure E.11: Comment rating in Youtube, using thumbs up/down

E.2.4. Reputation system

Benefits

- Involve the user and increase his participation and usage of the community, as well as the tools available
- Motivate users and strengthen the network
- Positively influence the classification, rating and recommendation of content

Related features

Music recommendation and Rating system

Overview

Reputation systems may be seen as a reward and merit system. As a reward system, it should reward good and punish bad behaviour. Based on the user's merit, the system can determine that users with high reputation may be granted special privileges, whereas users with low or unestablished reputation may have limited privileges.

The reputation may then be used to influence how strong the opinion of the user is, both when classifying and rating content, as in recommendations, where a user with a high reputation may be regarded as a music expert, providing better recommendations.

Also related to the concept of reputation system is the concept of virtual money. Some social networks allow their users to earn or buy “credits” (or “virtual coins”), in order to access limited content or acquire some kind of benefit (e.g. placing his profile in the website's spotlight for a short period of time, garnishing him more visibility and attention). This could be an interesting means of capitalising Palco Principal, and generate additional revenues.

The use of the reputation system could be associated as a means to earn virtual money, stimulating the usage of the website and the transactions of these “coins” within the system. One interesting application of this concept has been done in popular social networks as hi5 and Netlog, and more recently through the Open Social's Virtual Currency project⁷⁰.

Usage

The user's reputation (also known as the user's karma) may be represented by a point (or credit) system, representing his importance within the community. According to his actions, the user may earn or lose points, achieving a certain level of experience (or karma), which translates in a reputation inside the system. Some examples on how to earn these points are presented:

⁷⁰ Available at <http://code.google.com/p/opensocial-virtual-currency/>

Proposed features

- earn points for tagging content (e.g. earn 10 points for every 100 tags applied)
- earn points when successfully tagging untagged content
- earn points by flagging inappropriate content (and lose in the case of being flagged for abuse)
- earn points when uploading content (and earning more points if the content is rated as having good quality or becomes popular)
- earn points for continuous and intensive use of the website (e.g. pages viewed, number of days active, amount of publicity clicks, number of comments and reviews written)
- earn points for contributing to biography updates (using the wiki editing)

These points could then be used to access limited features or promote himself. Here are some examples:

- Being able to upload more content
- Place his profile or songs in the spotlight (home page)
- Buying on-line gifts or profile gadgets
- Download limited content, such as songs

Implementation details

A reputation system uses collaborative filtering algorithms to determine user ratings, given a collection of opinions and actions that users have with each other. Once again, since specific data mining techniques are necessary, with its development being done in cooperation with the data mining team, no further details on the implementation of this feature will be presented.

Popular examples

- Netlog's credit system - <http://en.netlog.com/go/helpdesk/view=question&faqid=343>
- Laconi.ca's user karma - <http://laconi.ca/trac/ticket/506>
- Drupal's user karma - http://drupal.org/project/user_karma

E.2.5. Graphic recommendation

Benefits

- Offer a more intuitive and appealing mechanism of recommendation and listening to music
- Enable recommendation steerability
- View the relations between several artists and songs

Related features

Music recommendation

Overview

Having the music recommendation feature as base, the idea of the graphic recommendation is to provide a more interactive and visual tool for users to receive and adjust the recommendations they receive from the system.

Having an item (artist or song) as a seed, the user can decide to expand its recommendations from there (being provided with similar items). This action is repeated over and over again, ending up by creating a true network of recommendations.

Another important aspect of this features is the possibility to view relationships between artists and songs, which is much more difficult to achieve using text or playlist formats.

Usage

The graphic recommendations may be accessed through the user's recommendation page, by choosing the “graphic visualization” option. Since this feature requires a large screen space, it should be kept as an alternate visualization to the user's recommendations.

Implementation details

The most common approach for graphic recommendations is undoubtedly the usage of undirected graphs (e.g. RAMA, by [Sarmiento 09b]). Each node (or vertex) represents an item (an artist or a song), with the edges representing their relations.

Another existing approach (proposed by [Pampalk 02a] and [Pampalk 02b], further explored in [Knees 06]) is the creation of Islands of Music, where the results obtained, using Self-Organizing Maps (SOM), are presented as the mountains and hills of an island.

Finally, in [Lamere 07b] some 3D visualizations for music are also explored: the album cloud, album grid and album spiral.

As the development of a visual interface for recommendations is out of the scope of this project, no further details on this feature will be provided, since some important design decisions must be made together with the whole project team.

Proposed features

Popular examples

- Musicoverly - <http://musicoverly.com>
- Dimvision's MusicMap - <http://www.dimvision.com/musicmap>
- RAMA - <http://pattie.fe.up.pt/RAMA>
- TuneGlue's Audiomap - <http://audiomap.tuneglue.net>
- LivePlasma - <http://www.liveplasma.com>



Figure E.13: LivePlasma - recommendations and relations between artists




Figure E.14: TuneGlue's Audiomap - recommendations and relations between artists

E.3. Improving event management

E.3.1. Geographic Information System (GIS)

Audience

fans, Palco 3.0, venues, artists

Goal

Determine the user's geographical position (coordinates) and its proximity to other users, events and venues.

Benefits

- Provide recommendations based on the user's geographical proximity
- Provide a local focus to both artist and venues
- Possibility to explore the mobile communications market, by using the GIS capabilities of smart phones, in order to integrate with Palco 3.0

Related features

Music recommendation, Event displaying and Event suggestion

Overview

The integration of Geographic Information System (GIS) technologies in Palco 3.0 would allow a bigger focus on the user and provide him more accurate information, based on its location. It would have an impact mostly in two areas: recommendation and event management.

In event management, GIS technology would allow to geographically compute with precision the location of events, displaying maps, distances and directions from the user's (or for that matter, any) location to the event itself. It would also be useful in order to locate venues and search for events, within a range of proximity.

Regarding recommendation, other users whose location is close enough (with the measure of granularity being possibly changed) could be recommended, as well as events of bands (recommended or not for the user) in nearby locations. Also, when providing information about popular artists and/or songs, these could be provided specifically to the user's country of location, in order to better characterize the musical trends of a specific community, as they tend to vary from country to country (e.g. top artists in Portugal or top songs in Angola).

Implementation details

There are several GIS technologies available, with API's that allow developers to integrate them with nearly any platform. Some of the most popular technologies are Google

Proposed features

Maps⁷¹ and Yahoo Maps⁷². However, as a result of the technological partnership established with SAPO.pt, the usage of its GIS service – Sapo Mapas⁷³ – is highly recommended.

Popular examples

- Last.fm – <http://www.last.fm>
- Imeem – <http://www.imeem.com/discover>




Figure E.15: Last.fm's location-based event listing




Figure E.16: GIS information on iMeem

⁷¹ <http://code.google.com/apis/maps/>

⁷² <http://developer.yahoo.com/maps/>

⁷³ http://developers.sapo.pt/projects/developers/wiki/Services/Maps_EN

E.3.2. Event displaying

Audience

venues, fans, artists

Goal

Improve the existing event data and its displaying, providing information about event attendance and ticketing.

Benefits

- Provide more detailed information about the event
- Provide ticketing information
- Allow the user to enter information about his attendance

Related features

Venue profile and User-added media

Overview

Although Palco Principal provides information about musical events (also known as *shows* or *gigs*), it is very limited and mostly static. With this feature it is intended to add more information to the events, such as attendance and ticketing information.

Regarding attendance, users should be able to provide their RSVP⁷⁴ status. On the event, it should be possible to see which users would attend the event. About ticketing information, apart from the ticket price, it would be useful to have information and/or links of how and where to buy tickets.

Also, both the artist, fan and venue should be able to add photos and videos, related to the event, such as posters, tickets, *in loco* photos and videos and others.

Usage

Events may be displayed in several places. In Palco Principal homepage, it is possible to see upcoming events around the user's location (using the GIS feature). On the artist's profile page, it is possible to see the artist's upcoming and past *gigs*. Similarly, in the venue profile it is possible to see upcoming and past events held there. Also, in the user's profile page, it is possible to see which shows the user is going or has gone to.

Implementation details

Since Palco Principal already contemplates the existence of events, the implementation should mostly focus on the event's profile redesign, as well as making it more interactive for the

⁷⁴ RSVP stands for "Répondez s'il vous plaît", a French phrase that can be translated to "please respond". It is standard practice to reply to an RSVP request whether confirming attendance or declining.

Proposed features

user, with web links for the venue's profile, box-office or ticket-selling websites, event's guest list, among others.

Popular examples

- Songkick – <http://www.songkick.com>
- Last.fm – <http://www.last.fm>



Figure E.17: Event displayed in Songkick, with RSVP and ticketing information



Figure E.18: Event displayed in Last.fm, with RSVP and venue information

E.3.3. Event suggestion

Audience

fans

Goal

Suggest events that might be of the fan's interest.

Benefits

- Discover new events

Related features

Geographic Information System (GIS) and Music recommendation

Proposed features

Overview

The system may suggest events, which the user might find interesting to attend to. This can be done by suggesting nearby events, based on the user's location and distance to the events, suggest other events from the same artist in other locations, or from similar artists. Both these recommendations may be done having into consideration the user's music recommendations.

Usage

Event suggestions may be provided to the user at the homepage, but also when viewing an artist or venue's profile or when searching for a certain tag, or tag combination.


Implementation details

As in music recommendation, event suggestion is made by computing similarities between users and artists. However, by using GIS technologies, the recommendations may also be done based on the distance. In general, the GIS software API provides methods to calculate the distance from the user's location to another point.

However, as it is outside of the project's scope, more specific implementation details will not be provided in this report.

Popular examples

- Songkick - <http://www.songkick.com>
- Bandsintown – <http://www.bandsintown.com>



The image shows a search interface for Bandsintown. At the top, there is a placeholder for a user profile picture and the text "current location: Porto, Portugal". Below this, there is a "Date Range" section with a dropdown menu set to "today" and a date selector set to "Jul 16". Underneath the date range is a horizontal slider. Below the slider are two checkboxes: "Tonight Only" and "This Weekend". Further down, there is a "Distance From City:" section with a value of "75 miles" and a horizontal slider. At the bottom, there is a "Max Ticket Price:" section with a value of "\$25 or less" and a horizontal slider.

Figure E.19: Event search parameters in bandsintown



Figure E.20: Songkick's event suggestion, based on event location

E.3.4. Venue profile

Audience

venues, fans, artists

Goal

Provide detailed information about venues.

Benefits

- Invite the venue owners to be part of the network
- Facilitate event promotion
- Bridge the gap between upcoming artists and venues
- Create communities of fans and regular customers around venues

Related features

Event displaying, Event suggestion and Geographic Information System (GIS)

Overview

By creating a new type of profile, focused on the venues, these are included in the dynamics of the network, since venues are also an important player. An event needs an artist to perform, fans to attend and a venue to be hosted, so it seems quite logic that, as it happens with fans and artists, venues also have a profile page.

The venue profile should contain basic information about it (such as the website, opening hours, photos, description), but also information about upcoming and past events. Using the GIS technologies, it is possible to easily locate the venue and find directions, allowing the creation of city maps with all its venues.

Usage

The event profile page works similarly to the artist and fan page, being used to find more detailed information about the venue.

Implementation details

The implementation of the venue profile would require changes in the system database and in the business logic. However, its implementation and design is similar to the profile pages of artists and fans.

Popular examples

- Last.fm – <http://www.last.fm>
- Jango – <http://www.jango.com>
- iLike – <http://www.ilike.com>
- Songkick – <http://www.songkick.com>

Proposed features




Figure E.21: venue profile in iLike




Figure E.22: venue profile at SongKick

E.3.5. User-added media

Audience

fans, venues, artists

Goal

Allow users to add photos and videos to events.

Benefits

- Provide new media content to the system
- Photo and video sharing

Related features

Event displaying and Venue profile

Proposed features

Overview

This feature refers to the possibility of the users adding photos and videos related to events, may they be promotional or documentary. This would allow not only a better coverage of the event (together with event reviews and comments), but also the creation of photo and video collections of fans and venues.

Usage

The photos and videos are displayed both in the profile of the event that they refer to, but also in the uploader's profile and in the venue profile.

Implementation details

Because it is outside of the scope of this project, no implementation details will be presented here.

Popular examples



Figure E.23: user-added media in Last.fm

Annex F: Prototypes

In this annex the prototypes described in chapter 5 are presented with a higher resolution.


In order to facilitate the discovery and correspondence of the figures presented in the text, the following table summarizes the list of prototypes included in this annex:

Figure	Annex	Description
5.1	F.1	Home
5.2	F.2	Header and Footer
5.3	F.3	Music
5.4	F.3.1	Artist overview
5.5	F.3.2	Artist biography
5.6	F.3.3	Artist tags
5.7	F.3.4	Artist recommendations
5.8	F.3.5	Artist music
5.9	F.3.6	Artist events
5.10	F.3.7	Artist photos
5.11	F.3.8	Artist videos
5.12	F.3.9	Artist blog

Figure	Annex	Description
5.13	F.4	Community
5.14	F.5	Tags
5.15	F.5.1	Tag search
5.16	F.5.2	Tag overview
5.17	F.5.3	Tag music
5.18	F.5.4	Tag events
5.19	F.6	Events
5.20	F.6.1	Event overview
5.21	F.7	Venues
5.22	F.7.1	Venue overview
5.23	F.8	Top videos
5.24	F.9	Top photos

E.1. Home


High resolution version of Figure 5.1.



created with Balsamiq Mockups - www.balsamiq.com

E.2. Header and Footer


High resolution version of Figure 5.2.



created with Balsamiq Mockups - www.balsamiq.com

E.3. Music


High resolution version of Figure 5.3.



created with Balsamiq Mockups - www.balsamiq.com


E.3.1. Artist overview

High resolution version of Figure 5.4.



E.3.2. Artist biography


High resolution version of Figure 5.5.



created with Balsamiq Mockups - www.balsamiq.com

E.3.3. Artist tags


High resolution version of Figure 5.6.



created with Balsamiq Mockups - www.balsamiq.com


E.3.4. Artist recommendations

High resolution version of Figure 5.7.




E.3.5. Artist music

High resolution version of Figure 5.8.



E.3.6. Artist events


High resolution version of Figure 5.9.



created with Balsamiq Mockups - www.balsamiq.com


E.3.7. Artist photos

High resolution version of Figure 5.10.




E.3.8. Artist videos

High resolution version of Figure 5.11.



E.3.9. Artist blog


High resolution version of Figure 5.12.



created with Balsamiq Mockups - www.balsamiq.com

E.4. Community


High resolution version of Figure 5.13.



created with Balsamiq Mockups - www.balsamiq.com

E.5. Tags


High resolution version of Figure 5.14.



created with Balsamiq Mockups - www.balsamiq.com

E.5.1. Tag search


High resolution version of Figure 5.15.



created with Balsamiq Mockups - www.balsamiq.com

E.5.2. Tag overview


High resolution version of Figure 5.16.



created with Balsamiq Mockups - www.balsamiq.com

E.5.3. Tag music


High resolution version of Figure 5.17.



created with Balsamiq Mockups - www.balsamiq.com

E.5.4. Tag events


High resolution version of Figure 5.18.



created with Balsamiq Mockups - www.balsamiq.com

E.6. Events


High resolution version of Figure 5.19.



created with Balsamiq Mockups - www.balsamiq.com

E.6.1. Event overview


High resolution version of Figure 5.20.



created with Balsamiq Mockups - www.balsamiq.com

E.7. Venues


High resolution version of Figure 5.21.



created with Balsamiq Mockups - www.balsamiq.com

E.7.1. Venue overview


High resolution version of Figure 5.22.



created with Balsamiq Mockups - www.balsamiq.com

E.8. Top videos

High resolution version of Figure 5.23.



created with Balsamiq Mockups - www.balsamiq.com

E.9. Top photos

High resolution version of Figure 5.24.

