



HAL
open science

Content Management for the Live Music Industry in Virtual Worlds: Challenges and Opportunities

Marco Lüthy, Jean-Julien Aucouturier

► **To cite this version:**

Marco Lüthy, Jean-Julien Aucouturier. Content Management for the Live Music Industry in Virtual Worlds: Challenges and Opportunities. Journal of virtual worlds research, 2013. hal-01546802

HAL Id: hal-01546802

<https://hal.science/hal-01546802>

Submitted on 26 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Journal of
• Virtual Worlds Research

jvwresearch.org ISSN: 1941-8477

Arts

June 2013
Volume 6, No. 2



Art by: Catarina Carneiro de Sousa

Volume 6, Number 2

Arts

June 2013

Managing Editor

Yesha Sivan,
Metaverse Labs Ltd.
Tel Aviv-Yaffo Academic College, Israel

Guest Editors

Celeste Lovette Guichard,
Savannah College of Art and Design, GA, USA

Laura Salciuviene,
Lancaster University Management School, UK

Gary Hardee,
University of Texas at Dallas, TX, USA

Coordinating Editor

Tzafnat Shpak



The JVWR is an academic journal. As such, it is dedicated to the open exchange of information. For this reason, JVWR is freely available to individuals and institutions. Copies of this journal or articles in this journal may be distributed for research or educational purposes only free of charge and without permission. However, the JVWR does not grant permission for use of any content in advertisements or advertising supplements or in any manner that would imply an endorsement of any product or service. All uses beyond research or educational purposes require the written permission of the JVWR. Authors who publish in the Journal of Virtual Worlds Research will release their articles under the Creative Commons Attribution No Derivative Works 3.0 United States (cc-by-nd) license. The Journal of Virtual Worlds Research is funded by its sponsors and contributions from readers.



Volume 6, Number 2

Arts

June 2013

Content Management for the Live Music Industry in Virtual Worlds: Challenges and Opportunities

Marco Lüthy

Robotag Media K.K., Tokyo, Japan

Jean-Julien Aucouturier

IRCAM (CNRS UMR9912), Paris, France

Abstract

The real-world music industry is undergoing a transition away from the retailing and distribution of fixed objects (records, files) to the consumption of live, interactive events (concerts, happenings). This development is paralleled by the recent flourishing of live music in virtual worlds, which in many ways could become the epitome of its real-world counterpart. For the artists, virtual concerts are cheap and easy to organize, and can therefore be a viable alternative to performing in the real world. For the music promoter and marketer, virtual concert attendance can be traced and analyzed more easily than in the real world. For the virtual concertgoer, attending concerts that are happening a (virtual) world away is possible with a single click. Taking insights from both a survey among the Second-Life music practitioners and from our own prototype of a live music recommendation system built on top of Second-Life, this article shows that the technical infrastructure of current virtual worlds is not well-suited to the development of the content management tools needed to support this opportunity. We propose several new ways to address these problems, and advocate for their recognition both by the artistic and the technical community.

1. Introduction

Since the year 2000, when it saw its revenues decrease for the first time in 15 years, the Music Industry has been in crisis. CD shipments in the U.S. were down 25% in 2008, and a cumulative 37% from 2008 to 2011 - something trade groups such as the RIAA readily attribute to illegal file-sharing on the internet, although some academic researchers disagree (Oberholzer-Gee & Strumpf, 2007). Beyond a “Napster” generation effect, what is under question here is the value of music fixed on a medium. In July 2007, American Funk musician Prince’s Planet Earth LP was distributed as a freebie with UK’s tabloid *Mail on Sunday* - shipping three million copies in one day (Time, 2007). In May 2009, British alternative rock band Coldplay announced that it would give away copies of its new recording to those attending its upcoming tour as “a thank you to our fans” (Rolling Stone, 2009). Increasingly, the recording is slipping out of the business model for artists and their labels. That transition is not in favor of video, either: in January 2009, American Industrial Rock band Nine Inch Nails shared with their fans more than 400GB of unedited raw concert videos in HD format via their website. A few years ago, such costly material would be edited into a highly profitable DVD, selling several 100,000 units globally (Aucouturier, Fujita & Sumikura, 2012).

So where do music revenues now come from? First, live performances. After his *Sunday Mail* campaign, Prince went on to sell out a series of 21 shows in London in mere minutes. Live performance cannot easily be pirated either because it involves tons of lighting and stage rigging, or simply because of the value of the feeling of being there, of interacting with the artist - not something that can easily be triggered from an mp3 file (Earl, 2001). Second, beyond ticket and merchandising revenues commonly associated with concert production, artists increasingly build wealth on the more general idea of personal “connection” with their fans (Kelly, 2008). Connections are generated by fans’ finding USB keys with free mp3s in some random spots of concert venues, deciphering the text on the band’s T-shirts to find secret website URLs, seeing your own geographical position displayed in real-time on the band’s website as you (and thousands of others) download their songs. Being a fan of a band that cultivates such a connection, like Nine Inch Nails, therefore increasingly resembles the playing of a video game - and not one that consists in buying an mp3 file from a database for \$0.99.

Because music can be accessed for free, bands have to give fans a reason to purchase. That reason is, increasingly, a feeling of attachment to the band - a fan may decide to buy a given track because of the strong experience he or she had with one of the band's projects, happenings or games that involved that track. Buying the track is a thank-you note.

Such trends are supported not only by Grammy-winners, but by the whole sector. For instance, one Japanese musician that we followed for the present study, Komuso Tokugawa, is also a live powerhouse. Since his acclaimed appearance at the nation’s premier blues festival, he has been playing constantly at more than ten shows per week, and selling out every one of them. His fan base is uniquely devoted, often coming to attend shows from all over the world several times in the same week. Many fans have followed him since the early days, when he could still be found singing in the small pub around the corner. Between shows, fans gather on online forums and blogs to review and comment on each of the artist’s appearances with astonishing enthusiasm. During his concerts, tips fly onto the stage constantly, money thrown by a crowd that cheers and shouts in perfect rapture. Tokugawa enjoys the kind of acclaim every pop musician would like to have.

What is most fascinating, of course, is that Tokugawa doesn't *really* exist.

Tokugawa is in fact the 3D avatar in the virtual world of Second-Life of Tokyo-based musician Paul Cohen. Introduced in 2003 by California-based Linden Labs Inc., Second-Life (hereafter SL) is a computer-based recreation of the real-world, entirely built by its users - at the time of our study, a thriving community of over 1,400,000 people. Real-life users of SL appear in virtual reality as 3D avatars and they can interact with other users' avatars and their environment. Like a small but increasing number of real-life musicians, Cohen, aka Tokugawa, has made performing in Second-Life his quasi-exclusive source of revenue (Figure 1 - see also Aucouturier, 2009).

"Performing virtually has a number of advantages over real world (RL)", says Cohen, "instant setup, no travel from home studio to gig, and less wear and tear on gear"¹. Over the past 3 years, Cohen/Tokugawa has played over 800 concerts in SL, topping at more than 40 a month. None of this music is recorded, but rather streamed live from Cohen's studio. For him, going from one concert venue to the next is as simple as clicking on a "teleport" link (a feature of SL, where avatars can fly, duplicate and break numerous other physical laws). The audience can be recruited by a computer-based instant message, and they too can teleport to the concert in seconds. In a mixed-reality event we organized in April 2009 in Tokyo, Cohen/Tokugawa performed live from SL, projected onto a screen in a real-world venue. With no advertising, his SL studio filled with 50 avatar concert-goers in seconds - word of mouth at the millisecond timescale (Figure 1).



Figure 1: Paul Cohen/Komuso Tokugawa

Transposing this to the real-world music industry is vertiginous. Leonie Smith, a jazz singer in real-world and a SL television host by the name of Paisley Beebe (Figure 2), puts it simply, "In the real world, you might never get to see your favorite musician live. This is why you buy CDs. In SL however,

¹ All interviews reported hereafter were collected by the authors in 2008-2009 as part of a "Digital Anthropology" project in Temple University Japan. Following the methodology of Boellstorff, T. (2008), interviewees were contacted, and interviews conducted via avatars inworld, with some follow-up by email where needed. See also Aucouturier, 2009.

you can teleport out anytime if you want to catch that musician maybe once a week, maybe 8 times a week”. If seeing a given artist live becomes a commodity, who needs to acquire a recording?



Figure 2: Leonie Smith/Paisley Beebe

One may be skeptical about Second-Life. Experiencing a virtual concert conjures visions of uninspiring 3D polygons, awkward navigation with a mouse and keyboard, and finding one’s avatar mistakenly rendered half-way through a pillar or a wall. All this is true, but irrelevant. First, technology is constantly improving, and we’ll soon be able to teleport at will into the middle of an ongoing music video of perfect HD quality. Second, many of the recent RL music industry successes, such as OK Go’s treadmill series of music videos, explicitly traded technical quality for imperfect, do-it-yourself aesthetics, which yielded a direct benefit of immediacy and immersion in the moment (Marshall, 2010). Beyond all the glitches, what matters most in virtual music performances, therefore, is the feeling of being there -- of watching a concert with your screen flush with a constant flow of text comments from the audience, hearing the sound of clapping and fans calling out to the band with their real voice as well as seeing the performer obviously reacting to all this, addressing people by their names, taking song requests. All this can leave one wondering if there is actually anything missing from the real-life experience of music. “It’s replaced by a different sort of energy”, says Brad Reason, a DJ in RL, and in SL under the avatar name of Doubledown Tandino, “It’s about the shared musical experience, sitting home, blasting the tunes and being locked in great conversation that fills up the night with a great experience for everyone. I can read a virtual crowd by the way it texts and the way it tips” (Figure 3).



Figure 3: Brad Reason / Doubledown Tandino

The question is therefore not whether virtual worlds like SL offer conditions for music that are better or worse than, say, the radio or the mp3 before them. The key question is whether the emerging micro-economy of virtual live music is a trend so strong as to shape the future of the music industry. In short, will ubiquitous real-time, live-only events replace the archived recording, at least for a segment of popular music? If so, this is an opportunity that technology will have to follow, rather than create, just as music content management practices have had to adapt to electronic music distribution in the past fifteen years (Pachet, 2003).

This article examines the technological road ahead, i.e. the challenges and opportunities created by the technical architecture underlying current virtual worlds like SL, in view of building scalable content management systems for live music. To do so, we first present a prototype of a music recommendation system that we built on top of Second-Life to recommend ongoing concerts to users based on previous patterns of concert attendance. We will then describe the difficulties that were revealed by our effort to build the system, and how we think these can be addressed in the future.

2. Live Music Recommendation in SL (Part I. Technical Description)

Many of the modern modes of media consumption (movies, books, news, etc.) are built around the relatively novel paradigm of automatic recommendation, an information filtering technique aimed at selecting goods that are likely to be of interest to a given user (Resnick & Varian, 1997). In particular, musical items are a natural target for automatic recommendation systems: one can recommend physical records for purchase ("people who bought this CD on Amazon also bought..."), MP3 files for download ("Genius: Top songs you're missing"), or more generally artists to discover ("You listen. Last.fm

recommends"). Recently, several commercial services (e.g. bandsintown.com, songkick.com, gigulate.com), were even launched to recommend music concerts near a user's location, based on the user's internet listening profile. For a review of music recommendation technologies, see Celma & Lamere (2011).

We describe here a system able to recommend musical performances happening in Second-Life (SL) to SL users, based on their patterns of attendance at previous virtual concerts (In typical Amazon lingua, "people who have been to this concert also went to..."). Recommending SL concerts is useful since there is no centralized search engine to discover such events unless they are actively published by a promoter. On a typical day in SL, dozens of concerts may be organized, yet many of them are one-time happenings or unscheduled events such as open-mic or karaoke bars. Discovering concerts in Second Life is further made difficult by the short life-time of the events: it is no use learning from a friend about a great performance which happened yesterday.

The system as we built it stayed live in Second-Life for several consecutive months in 2010. It consists of two components executed in two concurrent locations: a scripted virtual object (in the virtual world) and a server application (on the internet, outside the virtual world). The former exports data from the virtual world, and presents incoming recommendations to the user's avatar; the latter does all of the data processing and executes all recommender logic. While the following technical description will inform our subsequent discussion in Sections 4 & 5, the non-technical reader may want to focus attention on the user scenario in Section 3, which illustrates the same concepts in a less technical format.

2.1 Server Application

The server application cycles through all currently active concerts, and finds users to whom the concerts should be recommended. At each iteration the recommender builds a list of candidate users for a given concert by utilizing a combination of one content-based (CB) and two collaborative-filtering (CF) methods (Figure 5). The CB method finds old concerts by the artist giving the current one and retrieves the list of users of those earlier concerts ("You've been to a concert by Paul. Paul is playing again now. Want to go?"). The first CF method is said to be item-based: it finds other concerts that have had a similar set of users in the past, and recommends the new concert to those of these users who have not been to the new concert yet ("You've been to a concert by John. Many of those who were there with you are now in a concert by Paul. Want to go?"). The second CF method is user-based: for each user of the current concert, the system finds other users that have attended similar concerts in the past and recommends the new concert to those who have not been to the new concert yet ("You've been to concerts by John, Paul and Ringo. Mary has too, and she's now in a concert by George. Want to go?"). Using three concurrent recommendation methods (CB and CFs) allows the system to operate even when active concerts do not yet have attending users (a problem called cold-start, see Section 4). This technique is referred to as "mixed hybridization" in the recommendation literature (Burke, 2002).

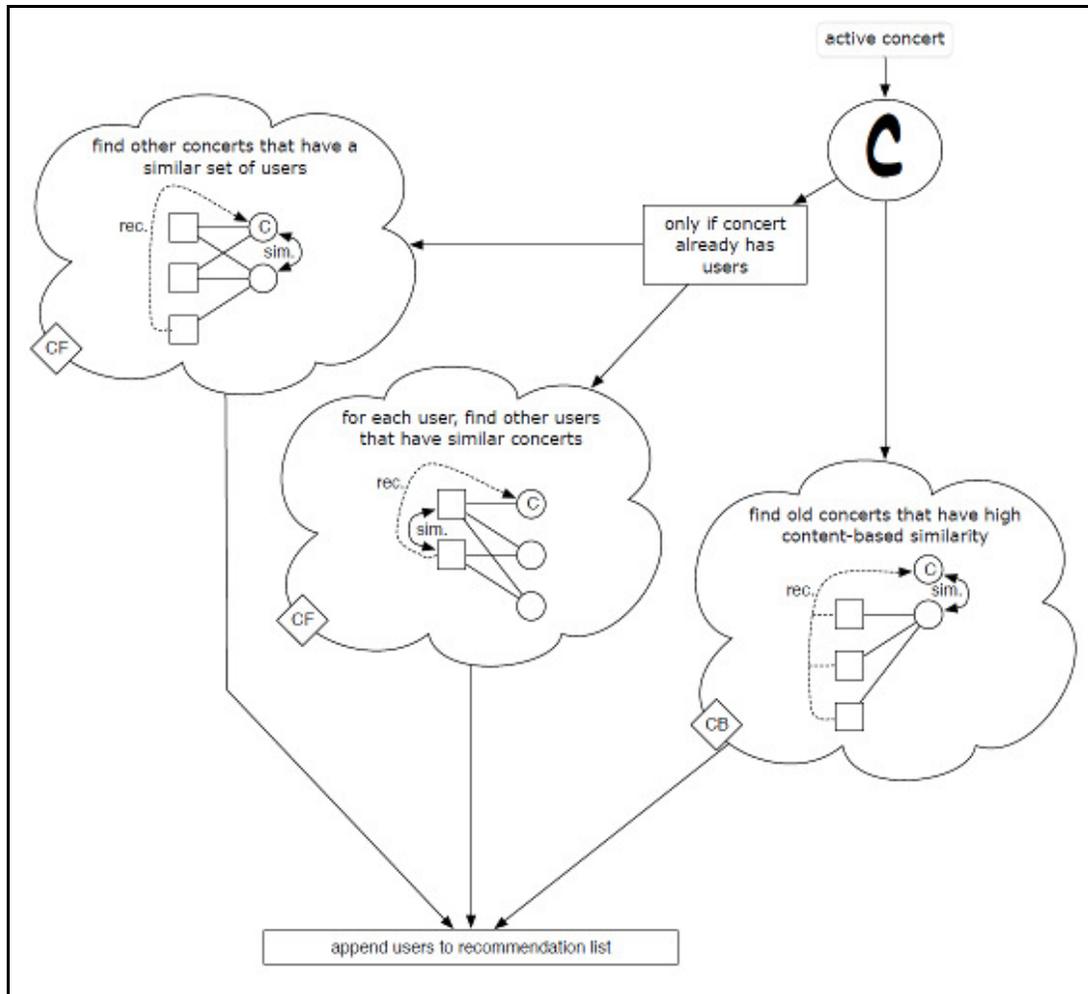


Figure 4: Architecture of the recommender system: the server application cycles through all currently active concerts and finds users to which the concerts should be recommended. The algorithm is a hybrid system, combining two collaborative filtering methods and one content-based method.

2.2 Virtual Object

Each subscriber to the service carries a virtual object, which is scripted with our system's program, and has the capacity to collect data about its "owner" avatar and to inform him/her of incoming recommendations.

More precisely, as a user moves through the virtual world, the virtual object sends regular updates to the server application (via an open HTTP connection) about the avatar's XYZ position in the virtual world. If the application determines that the object currently resides at the location of an active concert, it then calculates or updates a rating for that particular user and that particular concert: the longer the user stays in the concert (i.e. the more successive updates are received from the object at the same concert's location), the more the corresponding rating increases; if the user types text during the concert (i.e. engaging with the artist or others in the audience), the rating is also increased; conversely, if the user teleports away from the concert within the first few minutes (a negative reaction), the rating is

decreased; if the user falls inactive during the concert (i.e. in SL, the avatar is tagged as "away from keyboard"), the rating is also decreased. After each new update of the ratings, for any active user or any active concert, the whole recommendation routine is triggered again, and new recommendations are issued to any active user in-world, if needed (Figure 6).

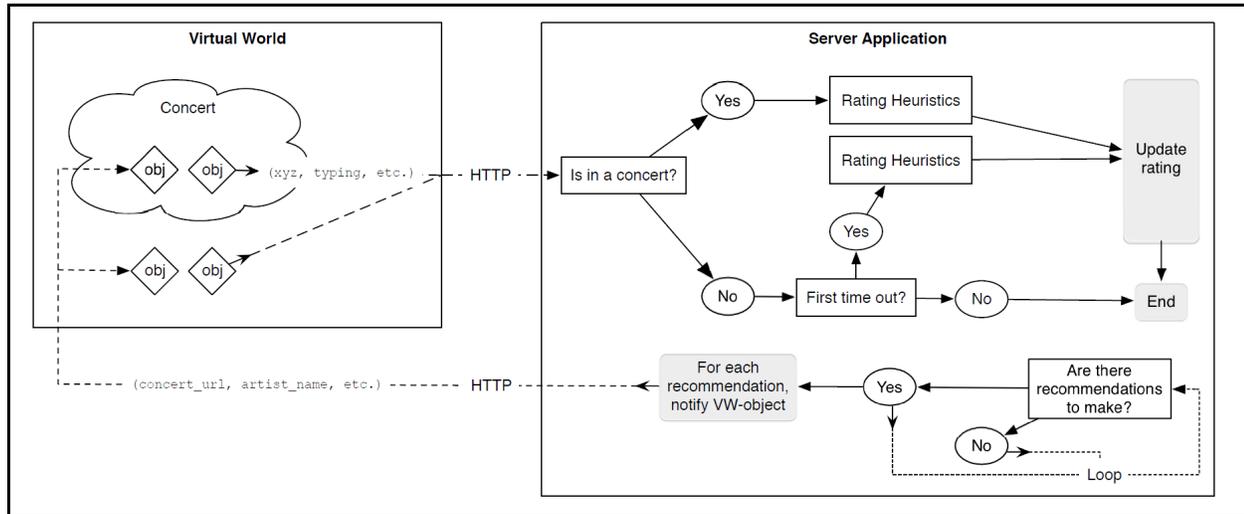


Figure 5: Inputs and outputs of the system: as a user moves through the virtual world, his/her virtual object sends information about that user's behavior and location to a server application. This information is analyzed to produce ratings for the concerts that the users are attending, which in turn inform the recommendation algorithm. When new recommendations of concerts are found, they are sent back to the corresponding users via their object in the virtual world.

3. Live Music Recommendation in SL (part II. Use Scenario)

We illustrate here the system described in Section 2 with a typical use scenario, inspired by a series of user studies realized between January and April 2010. A video demonstration of a real sequence of events, as implemented in Second-Life, can also be found online (<http://youtu.be/qJ3HpVXMvSg>).

Our scenario involves four (fictional) SL artists John, Paul, George and Ringo, whose concerts are recommended to a group of 5 SL users (Alice, Bob, Carol, Dave and Erin). The scenario starts with the system history described in Figure 6-top. Among the concerts registered with the system's server, two are currently active, in separate locations, by artists Paul and Ringo. The concert by Paul currently has no registered users attending it. The concert by Ringo has had two ratings so far: user Erin has been to the concert, but she did not like it and teleported away in the first few minutes; the virtual "recommender" object she carries detected the teleport action, notified the server, which in turn updated her rating to -3. A second user for Ringo's concert, Alice, has been there for only one song so far; her object has notified the recommendation server twice (it does so every minute), updating her rating to +2.

Paul's active concert (Paul3) is scrutinized by our system to recommend users. Searching in the database of previous concerts by the same artist, the system finds that Paul has performed twice recently (Paul1 and Paul2). User Alice enjoyed Paul2 (she had a positive rating for it); user Bob enjoyed Paul1; both therefore receive a notification in the virtual world that a new concert is happening now, which they might enjoy.

The recommendation reaches Bob just as he was thinking of taking a short break from work - “why, of course I remember that concert by Paul, and yes, why not, let’s see where and what he is playing now”; Bob teleports straight away (Figure 6-middle). Within 60 seconds, his object sends a position update to the server, which compares it to the positions of all the active concerts, and detects that he is attending Paul3. A new rating is created for the (Bob, Paul3) pair. Alice receives the recommendation while she is still attending Ringo’s concert: “Alice, you saw Paul’s concert yesterday, and we believe you enjoyed it. Paul is playing again right now in the Lizard Lounge - do you want to teleport and have a look?” Sure, she did like Paul’s last concert, but she also likes what she hears now; she ponders for a few minutes, during which her object sends another position update to the server, increasing her current rating for Ringo. She finally decides to give Paul3 a try - if she does not like it, she can always go back to Ringo’s using her recommendation history (Figure 6-bottom).

Just as she materializes at Paul3, Alice is greeted by a cheerful “Hi Alice, thanks for visiting the Lizard Lounge”, emanating from the club manager’s avatar, “I hope you enjoy your time with us!” It’s her first time there, but she likes the design of the room, and Paul is just starting one of her favorite songs - she notes that he is playing the same song list than as in Paul2, which she thought was much better than the ones in his previous set, Paul1. “Looks fantastic! Hey, Paul! I love this song!!” she types in the local chat channel. Her object detects the IS_TYPING status of her avatar and sends the information to the server, along with her position in the club which has not changed since the latest update. This raises her rating to 3. (Figure 7-top). Unbeknownst to Alice, on another side of the room, Bob is also enjoying his time. The concert has the same raw energy that he liked in Paul1, without all the annoying visuals he so disliked in Paul2. It is rare that Bob and Alice like the same type of music - in fact, the only time this happened before was at that concert by John, last week. This connection is appropriately identified by the system, which matches the two concerts, John and Paul3, and searches for users that attended the former but not the latter - user Dave is one of them. “Hi Dave! Last week, you were in that concert by John. Alice and Bob, who were there too, are currently enjoying what appears to be a fine concert by Paul in the Lizard Lounge. Do you want to teleport and have a look?” Dave naturally cannot resist the pitch (Figure 7-bottom).

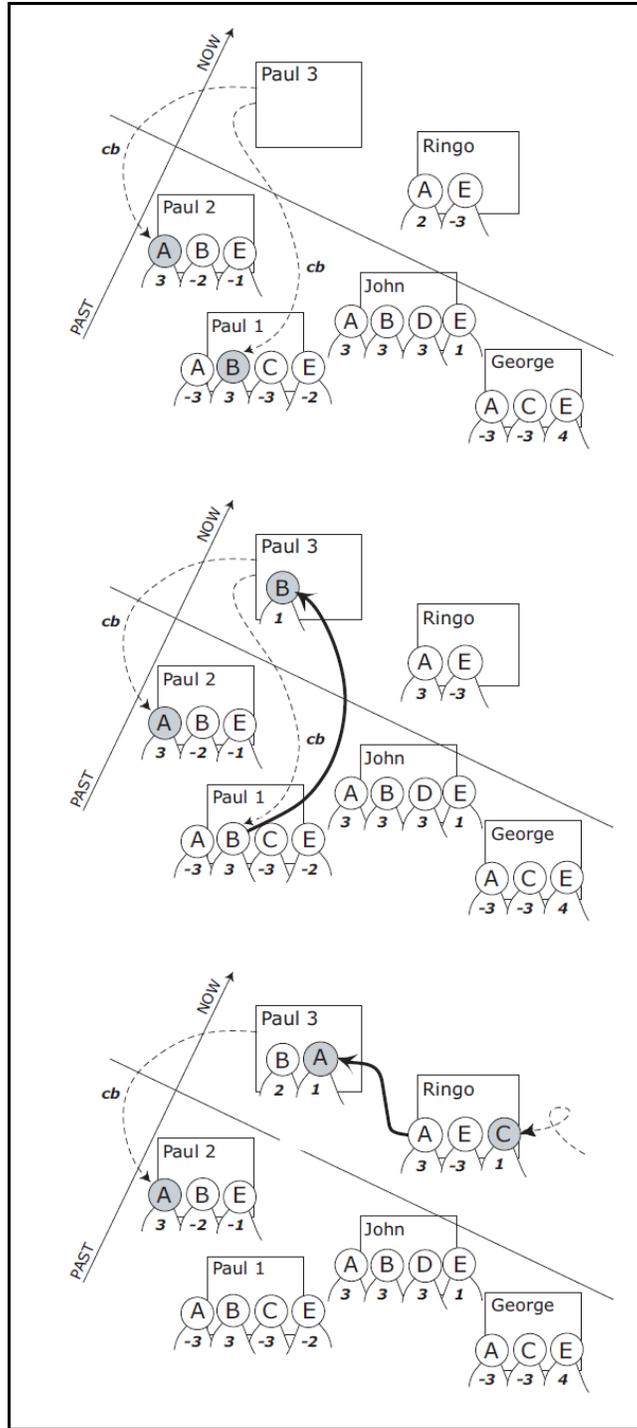


Figure 6: use scenario: users Alice and Bob are recommended concert "Paul3", because they have been to previous concerts by the same artist

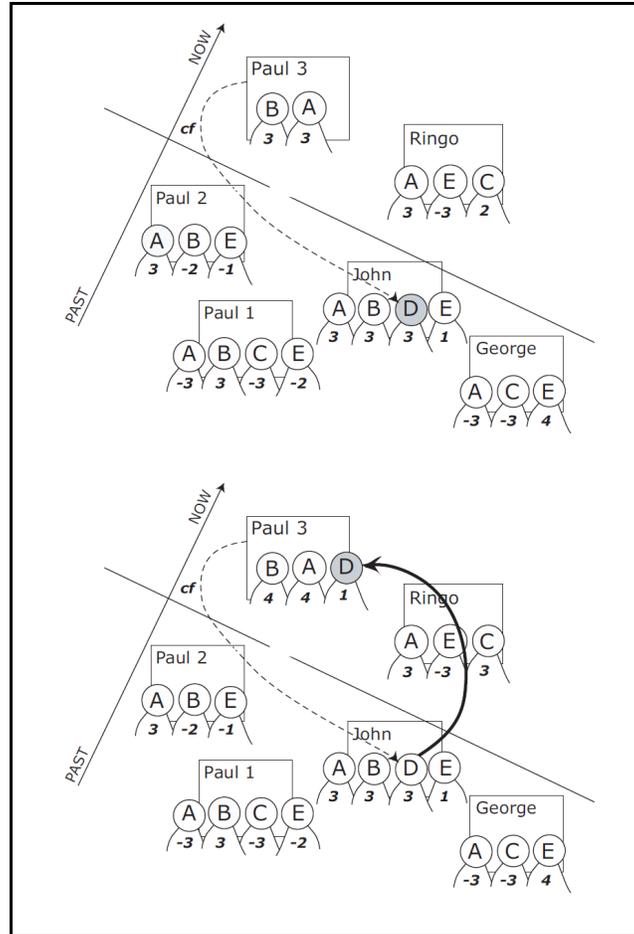


Figure 7: use scenario (continued): user Dave is recommended concert "Paul3", because he has been to a previous concert ("John") where many of those currently attending "Paul3" were also attending.

Meet Carol, the cynic. It's very rare that Carol likes a concert at all. Paul1 - she hated. George - teleported away in minutes. But the moment she saw artist Ringo takes the stage (Figure 6-bottom), she was mesmerized. Song after song, position update after position update (Figure 7), she keeps standing there, immersed in the music. When her rating reaches 4, the system determines that Carol's tastes are actually quite similar to Alice's – they both liked Ringo, but most importantly they also tend to hate the same music. As Ringo's concert ends, a recommendation for Paul3 is sent to Carol, just in time for her to join the show for the finale (Figure 8- bottom).

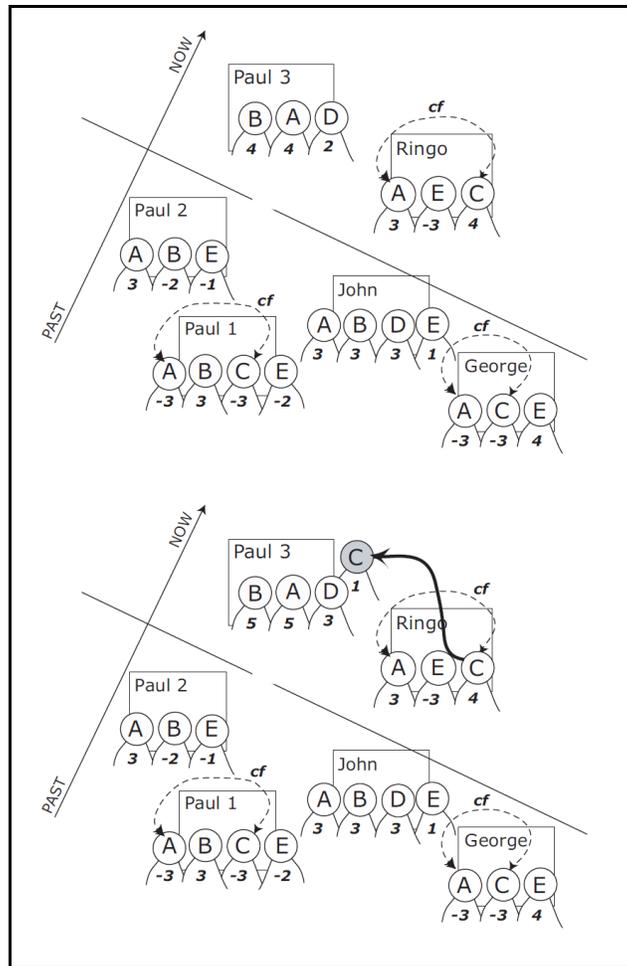


Figure 8: use scenario (continued): user Carol is recommended "Paul3", because she has a history of liking the same concerts as Alice and Alice is currently attending Paul3.

4. Challenges for Content Management

4.1 Concerts Are Not API-identified Entities

Contrary to avatars and virtual objects, concerts are not an identified part of the SL world model: they are defined extrinsically by users, as the combination of a location (technically, a URL) and a time for the event. The virtual world's application programming interface cannot serve the list of upcoming concerts, nor can it even identify if "something" is a concert.

This of course is a major obstacle to building content management technology such as recommendation systems. It would be akin to building an electronic music distribution system such as iTunes, with the constraint that it does not have access to a list of the available musical files, but only to its users and their listening behavior (*user A has clicked on her iPod and listened to a thing called "My songs/Beatles-BestOf-Yesterday.mp3" for a continuous three-minutes*). As a consequence, the recommender described above cannot monitor interactions between users and concerts in any direct way, as for example Amazon recommendations are able to operate on purchase actions of items by

users. Rather, it has to piggy-back on the virtual world and data-mine it to discover the items as well as the users' behavior with the items.

In our prototype, we circumvented this issue by relying on users to create concert items in our system: artists or promoters must register their upcoming (or just started) concert with our service, so we can start tracking potential users and issue recommendations. This makes strategic sense, as artists are easily motivated to advertise their events in order to attract a larger audience. The mechanism is currently implemented as a simple web form, with location, start and end time of the event. Other possible implementations could include registration through a custom SL object carried by the artist ("switch on the lights at the venue to register the event") or letting users define what they consider an event, in the way, for example, real-world geo-location services like Foursquare (<http://www.foursquare.com>) let their users define places that do not already exist in the system.

4.2 Concerts Lack Content - Description

A second consequence of the lack of programmatic representation for virtual world events like concerts is that there is ambiguity of what a concert even *is*. In order to track user attendance at concerts, the system needs to adopt a heuristic definition for the space and time span of an event. As described above, we relied on users to document when a concert starts and when it ends (the point after which an ongoing concert should stop being recommended to new users). Space however is more difficult: in what radius around the location should users be considered to be attending the event? This obviously depends on the physical properties of the venue and on surrounding events, but also on the psychological circumstances of the attending subject: does engaging in animated conversation with someone at the back of a crowded café count as attending the show on stage? In our implementation, we simply defined as "attending" the event any user within the same SL "region" as the registered event, between the registered start and end time of the event.

More generally, virtual world events do not come with associated content descriptions (they lack *metadata*, because they're not considered data in the first place). In our system, we assembled the essential metadata of who the performing artist is (by asking artists to register their concerts with our system), which gave us useful information to recommend concerts to users who had a history of attending concerts by the same artist (note that we could technically create recommendations without knowing the artist, by only considering patterns of co-location of users at unidentified concerts - see Figure 4). Other important metadata that could be assembled similarly are: the concert location, the concert's tracklist, the list of supporting musicians in a pop music context, etc. - all information which is a commodity for modern music content-management platforms, but is currently lacking from the data model of virtual world platforms.

4.3 Concerts Lack Identification - Technology

This current state of affairs resembles the situation of music distribution systems at the end of the 1990s, when it was realized that what was then the most common media for music consumption, the compact disc, did not carry information about its identity. For the requirements of playback on hifi systems, the CD did contain the number and duration of each of the musical tracks it contained, but not the information about, for example, what album was playing; that information was only available on the CD's physical packaging. What has been described as "the great CD bug" (Pachet, 2003) soon became a

major obstacle to linking the physical media to any type of computer information, such as inserting the CD in a computer, and seeing information about the artist, or suggestions of similar music - the same obstacle we face here with virtual concerts. This obstacle received a technical work-around with *identification technology* such as the CDDB initiative, which matched a given CD track-list to a user-assembled database of previously identified tracklists: if a CD's first track is exactly 3'27, the second track exactly 4'52, then the CD is very likely *Abbey Road*, by *The Beatles* (Pachet, 2003). In the early 2000s, a similar roadblock was found in the identification of individual digital files: how could one be certain that a file called *The Beatles - Yesterday.mp3*, that could have been mistakenly named by any user, indeed contained the media for that exact song? This problem motivated further technical innovation, namely signal processing algorithms able to identify songs directly from the "sound" of music, which were quickly industrialized with companies such as Gracenote (Cano et al., 2005). The constant underlying use of such identification technology is easily overlooked in our everyday seamless interaction with music content management platforms, by which we expect not to mistake one song for another song, but it is only after these were introduced that a whole industry could develop. Similar tools that can, for example, automatically recognize a concert's performing artist, detect when a concert starts and stops and decide if a given concert is within a user's auditory reach, are likely needed before the live virtual music industry can properly develop.

4.4 A Limit-Case for Current Recommendation Paradigms

There are several intrinsic properties of virtual concerts that make them unusual targets for recommendation systems, compared to traditional music files. First, while a given song may be available for decades, collecting massive amounts of associated user listening history, the typical length of a SL performance is only one hour. As in the real-world, virtual concerts can only be recommended before they end. Unlike the real-world, they cannot be recommended much before they start either, as they are rarely announced more than a day in advance and sometimes not even planned. Second, user data for our system is significantly smaller than for traditional web stores. Typical attendance for SL concerts is 20-30 avatars, and typical SL regions (i.e. bandwidth on hosted server space) do not allow more than 50 active avatars the same time for technical limitations².

The implication for recommender systems is that they have to operate mostly in the so-called "cold start" regime, with only a few user ratings available for an item before it goes inactive. Cold-start requires hybrid solutions that do not rely solely on collaborative filtering, but also on content-based alternatives when enough user-ratings are not available yet (see Figure 4). This functioning is a limit-case for the most-studied systems, and its behavior is still not considered optimal (Schein, Popescul, Ungar & Pennock, 2010). Similarly, little is known about the behavior of recommendation systems for which items are routinely deleted (i.e. at the end of a concert) from the system, and what is the best way to maintain/update user history for such items (Lathia, Hailes, Capra, & Amatriain, 2010). Finally, with short-lived items, offline computations are impossible. While typical systems can process recommendations as a daily scheduled task (Linden, Smith, & York, 2003), user ratings for active concerts have to be processed while the concert is active, within minutes. While our recommendation prototype proved the concept valid, much work is therefore needed to adapt recommendation algorithms to these novel constraints.

² More recent VW architectures, such as Avatar Reality's "Blue Mars" (www.bluemarsonline.com), arguably allow up to 1000 avatars simultaneously at an event

5. Opportunities

5.1 Rich Semantic Analysis of User Behavior

While inferring user intention behind any type of online behavior is obviously difficult, user behavior in virtual worlds is quite rich which provides a lot more data for an algorithm to analyze. Contrary to internet users, who mainly interact with web pages via scrolling and mouse clicks, VW avatars are tagged with complex information about status and actions, all of which are directly accessible by API. Figure 9 shows a selection of typical SL avatar behaviors found in virtual concerts. All of these correspond to actions taken in the SL client software, (e.g. typing "\bored" in the local chat channel to trigger a "bored" gesture as in Figure 9-7), and are possible to detect algorithmically with a script embedded, for example, in a virtual object. In our current implementation, we only detect four types of behaviors, corresponding to Figure 9-1,2,3,4: being in the concert (determined by comparing the avatar XYZ position to that of the concert), being away from keyboard (thanks to an AGENT_AWAY tag in the SL API), typing text during the concert (_TYPING tag) and teleporting away from the concert (thanks to the _TELEPORT listener). But further cues, such as those corresponding to Figure 9-5,6,7,8, can all be detected and analyzed computationally to yield unprecedented amounts of user information. For instance, text typed by the avatar can be automatically analyzed to detect positive or negative attitudes to the concert (O'Connor, Balasubramanyan, Routledge, & Smith, 2010).



Figure 9: Examples of user behavior when their SL avatars are in a concert. (1) Attending the concert, i.e. being located in a spatial neighborhood of the concert venue, and within reach of the live audio stream. (2) Typing, i.e. sending text in the local chat channel, visible to all in the audience and to the artist. During a concert, typical texts include cheers to the artist, song requests and factual information about the concert. (3) Being away from keyboard, i.e. a RL user leaves computer without logging off, so that the avatar remains; after about three minutes, the avatar head bows down and the word "away" appears over it. (4) Leaving the concert by teleporting away from it, i.e. sending one's avatar to another virtual location, leaving the concert venue and disconnecting from the associated audio stream. (5) Tipping, i.e. sending an electronic money transaction to a scripted object owned by the artist or the concert venue, with a mouse click in the SL client software. SL has its own virtual currency, the Linden dollar (L\$), which is exchangeable for US dollars or other currencies on market-based currency exchanges. A typical tip at a concert is about L\$200 (c. 1 USD). (6) Clicking for more information, i.e. clicking on scripted objects found in the concert venue which provide additional information on the artist, such as song list, links to web pages or (as seen here) automatic fan-group subscriptions. Information usually opens in a semi-transparent, browser-like window in the SL client software, in the manner of a head-up display. (7) Gestures, i.e. animating one's avatar with a preloaded animation, using the SL client software interface. Animations are usually emotional expressions such as looking bored, happy, excited, as well as integrated movements such as clapping, playing an instrument or dancing. Custom animations can be designed manually or bought. (8) Dancing, i.e. animating one's avatar specifically with a dance animation. Concert venues or artists typically offer a selection of dance animations for their audience which fit the music. These are usually triggered by right-clicking on a scripted object found in the venue (a so-called "pose-ball").

5.2 Real-Time Audio Content Analysis

The content-based recommendations made by the prototype described above are based on a concert's editorial metadata: its artist (concerts by the same artist) or its location (concerts in the same venue). However, virtual live music can also be recommended on the basis of its audio signal: how it sounds, what genre it is, and what type of emotion is involved ("there's great energy in this 90s rock medley being played right now at venue A, want to teleport there?"). Technically, a stream of music broadcasted from a musician's web server can be piped by a recording algorithm, which can extract its audio characteristics in real-time and use them to inform pre-trained pattern recognition algorithms. While such algorithms have been much researched in the past ten years in the Music Information Retrieval (MIR) community (Casey et al., 2008), their application in industrial recommenders has been slow compared to alternative methods based on collaborative-filtering (Barrington, Oda & Lanckriet, 2008). It has been argued that, while CF methods may work best for popular music, for which much user information exists, CB methods are most appropriate for lesser known items, which do not share large common listening history with popular songs but may sound similar (Celma & Lamere, 2011). With its many, one-off performances by small independent artists and little to no marketing, virtual music therefore seems an ideal application for such technologies.

5.3 Cross-Domain Recommendations

Typical virtual world users have diverse and multiple interactions with their avatars, which they may use variously to experience live music, meet people, shop for virtual or actual goods, play games, learn from a virtual campus, or do business (Castranova, 2005). In traditional web-based recommendations, all these activities would occur with distinct, proprietary websites (last.fm for music, amazon.com for shopping, facebook to meet people, etc.), all of which may have their own recommendation service, but with little inter-operability. Virtual worlds, on the contrary, offer multiple opportunities for cross-domain content-management, and recommendation in particular: users may be compared based on similar behavior with different types of items ("in concerts and in lecture halls, you both like to stay at the front"), or recommended items of one domain, based on similarities in another ("people who wear the same gothic outfit are attending this concert"). Such opportunities call for semantic description for events, items, actions, that can inter-operate ("attending" a course and "attending" a concert) - a project largely taken on by the Semantic Web initiative and which could find in content-management in virtual world an ideal field of application (Lesko & Hollingsworth, 2010).



Figure 10: An example of RL promotion of a virtual performance: SL musician Komuso Tokugawa tweets about an upcoming/improptu performance in SL.

5.4 Link With RL Music Industry

A further opportunity offered by the development of content management for live virtual music is in its linking with the already existing music content in "real-life". Musicians performing in SL have not waited for this opportunity to emerge: Figure 10 shows a message posted on the "RL" service Twitter by SL musician Komuso Tokugawa, advertising his upcoming, impromptu performance in SL (which can be linked to with a standard URL). However, one could envision that such performances could also be recommended to users based on, for example, what music they have recently bought online or what person they have just befriended on social networks. Another intriguing possibility would be to double all RL musical interaction with a virtual counterpart: buying concert merchandising in RL could give access to virtual copies of the same items for one's avatar; attending a RL concert gives private access to a virtual rendering of the same concert, if one wants to live it again and again (Figure 11).

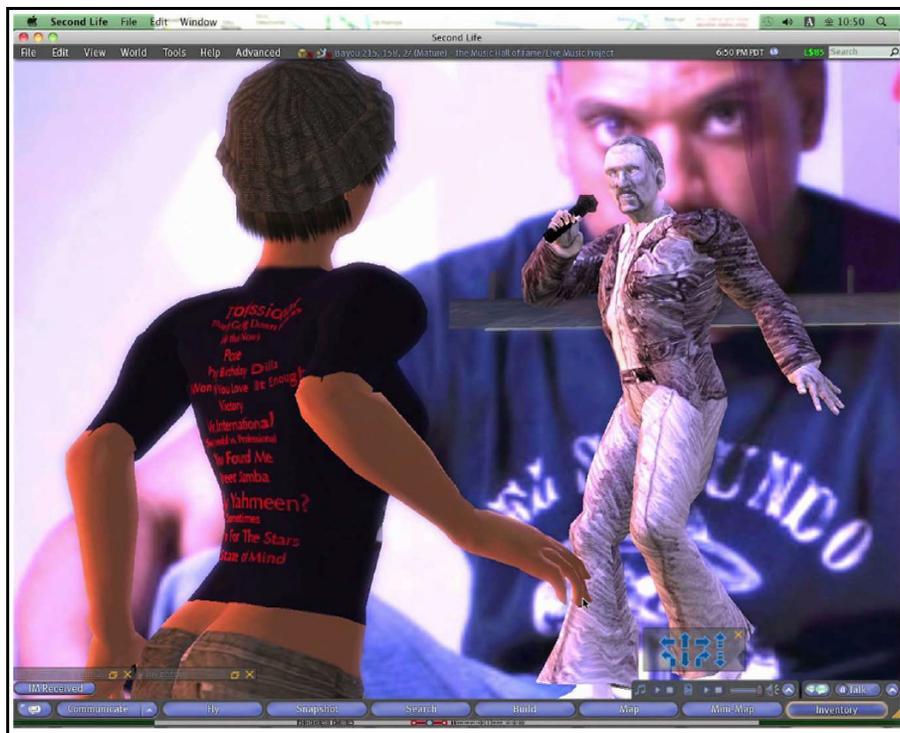


Figure 11: Real-world hip hop musician Legrand (background) performing live in the virtual world of Second-Life, as a 3D avatar (middle ground), to an audience of avatars (foreground), as promotion for his new RL release. The avatar in the foreground wears a virtual copy of Legrand's real-life merchandising (a tour tee-shirt).

6. Conclusion

In this article, we have shown that it is possible to build a recommender system on top of an existing virtual world (Second-Life), in order to both data-mine the behavior of its users/avatars to determine the level of their appreciation for live concerts, and to issue recommendations in real-time for items that are both short-lived and sparsely attended. Our prototype highlights a series of shortcomings in current virtual world architectures that limit the development of content-management technology for the live music industry, but also points at several opportunities to develop semantically-rich, cross-

domain and cross-world content management. While the phenomenon of virtual concerts may possibly never scale up to become a mass consumer reality, we argue that they epitomize a trend in the consumption of music, which, year after year, becomes less of a catalog product (e.g. a CD recording) and more of a live, instantaneous and mobile experience (e.g. a live music broadcast, streamed on a mobile phone). We believe that future recommendation systems for music, and more generally for any cultural goods, will face some of the challenges and opportunities that we have encountered in the context of this study.

Acknowledgements

We thank the many Second-Life musicians we got to interview and interact with over the past years. Even indirectly, they have provided inspiration and feedback for the system described here. In particular, we wish to thank Komuso Tokugawa (Paul Cohen), Legrand, Doubledown Tandino (Brad Reason), Paisley Beebe (Leonie Smith) and Teddi Shamrock (Teddi Davis) for their warm support. M.L. was supported for this work by a Temple University CARAS fellowship for undergraduate research.

References

- Aucouturier, J.-J. (2009). *Real music for a virtual world*. Japan Inc., n. 86, Summer 09.
- Aucouturier, J.-J., Fujita, M. & Sumikura, M. (2012). *Psychological effect of participatory music consumption: Case-study of Nine Inch Nails' Lights in the Sky*, Submitted manuscript, J. Consumer Behaviour
- Barrington, L., Oda, R. & Lanckriet, G. (2009). *Smarter than Genius? Human evaluation of Music recommender Systems*, In Proc. Retrieved from <http://cosmal.ucsd.edu/cal/pubs/Barrington-Genius-ISMIR09.pdf> , The International Conference on Music Information Retrieval, 2009, Kobe, Japan.
- Boellstorff, T. (2008). *Coming of Age in Second Life: An Anthropologist Explores the Virtually Human*. Princeton: Princeton University Press.
- Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4), 331-370.
- Cano, P. Batlle, E. Kalker, T. and Haitsma, J. A review of audio fingerprinting. *Journal of VLSI Signal Processing* 41, 271–284, 2005
- Casey, M. et al. (2008) Content-based music information retrieval: Current directions and future challenges. *Proc IEEE* 96:668–696.
- Castranova, E. (2005). *Synthetic worlds: The business and culture of online worlds*. Chicago:University of Chicago Press.
- Celma, Òscar and Lamere, Paul, If You Like Radiohead, You Might Like This Article. *AI Magazine* 32(3): 57-66 (2011)
- Earl, Peter. Simon's travel theorem and the demand for live music. *Journal of Economic Psychology*, Vol. 22(3) , June 2001, Pages 335–358
- Farouky, J. (2007, July 18). Why Prince's Free CD Ploy Worked, *TIME Magazine*, USA.
- Ferreiro, L. (2010). Second life's thriving music scene. *Los Angeles Times*, June 09 2010.
- Kelly, Kevin. (2008 March). "A Thousand True Fans". Retrieved from [http://kk.org/thetechnium/archives/2008/03/1000 true fans.php](http://kk.org/thetechnium/archives/2008/03/1000_true_fans.php), June 19 2009.
- Knees, P., Schedl, M., Pohle, T., & Widmer, G. (2007). Exploring music collections in virtual landscapes. *IEEE Multimedia*, 14(3), 46-54.
- Kreps, D. (2009, May 1). Coldplay Reward Fans With Free Live LP, *Rolling Stone Magazine*.
- Lathia, N., Hailes, S., Capra, L., & Amatriain, X. (2010). *Temporal diversity in recommender systems*. In Proceedings of ACM sigir 2010, July 2010, Geneva, Switzerland.
- Lesko, C. J. and Hollingsworth, Y. A. (2010). Compounding the Results: The Integration of Virtual Worlds with the Semantic Web. *Journal of Virtual Worlds Research*, Vol. 2(5), pp. 1-20.
- Linden, G., Smith, B., & York, J. (2003, January-February). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 75-80.

- Marshall, Wayne (2010). The Unstable Platforms and Uneasy Peers of Brave New World Music. *Berkman Center for Internet and Society luncheon series*, Dec. 2010. Retrieved from: <http://cyber.law.harvard.edu/events/luncheon/2010/12/marshall>
- Oberholzer-Gee, F. & Strumpf, K. (2007 February). The Effect of File Sharing on Record Sales: An Empirical Analysis, *Journal of Political Economy*, Vol 115(1): 1-42. The University of Chicago Press, Chicago, USA.
- O'Connor, B., Balasubramanyan, R., Routledge, B. R., & Smith, N. A. (2010). From tweets to polls: Linking text sentiment to public opinion time series. In Proc. *AAAI Conference on weblogs and social media*, Washington DC, may 2010.
- Pachet, F. (2003). Content management for electronic music distribution: the real issues. *Communications of the ACM*, vol. 46(4), pp. 71-75.
- Pan, Z., Xu, B., Yang, H., & Zhang, M. (2006). Content-based personalized recommendation in virtual shopping environment. *International Journal of Business Intelligence and Data Mining*, 1:4 , 430 - 449.
- Resnick, P., & Varian, H. R. (1997). Recommender systems. *Commun. ACM*, 40 (3), 56-58.
- RIAA 2011 Year-End Shipment Statistics, (n.d). Retrieved November 2012 from: <http://www.riaa.com/keystatistics.php> .
- Schein, A. I., Popescul, A., Ungar, L. H., & Pennock, D. M. (2002). Methods and metrics for cold-start recommendations. In Proc. of ACM SIGIR'02.
- Second life economic statistics. (2010). Retrieved June 2010 from: <http://secondlife.com/whatis/economy stats.php>.
- Shah, F., Bell, P., & Sukthankar, G. (2010). A destination recommendation system for virtual worlds. In Proceedings of the twenty-third international Florida Artificial Intelligence Research Society Conference.