

Personal vs. Commercial Content: The Similarities Between Consumer Use of Photos and Music

Frank Bentley, Crysta Metcalf, Gunnar Harboe

Motorola Labs

Schaumburg, IL USA

{ f.bentley, crysta.metcalf, gunnar.harboe } @motorola.com

ABSTRACT

We describe the results of two ethnographic-style studies that investigated consumer use of photos and music respectively. Although the studies were designed, executed, and analyzed separately, in our findings we discovered striking similarities between the ways in which our participants used personally captured photos and commercially purchased music. These findings have implications for the design of future systems with respect to handling and sharing content in photo or music form. We discuss making allowances for satisficing behavior, sharing media as a way to reminisce or to communicate an experience (tell a story), getting sidetracked while browsing, and similarities in organizing behaviors.

Author Keywords

Consumer, Music, Photo, Sharing, Ethnographic Study.

ACM Classification Keywords

H.5.1 Multimedia Information Systems: Evaluation/methodology, H.5.2 User Interfaces: User-centered design, H.5.5 Sound and Music Computing: Methodologies and techniques.

INTRODUCTION

In previously published work, researchers have studied personal photos and personal music collections separately. In our experience, the consumer electronics industry also emphasizes the differences between “personal” content that is captured by the user and “commercial” content that is produced and purchased. Both operate under very different business models and creation mechanisms which have contributed to the different ways different types of media are handled. For example, the personal content industry is focused on selling capture devices for consumers to capture and create personal content, like photos, and on providing

services for sharing that content with others. In contrast, the business model of commercial content like music is often in the selling of the content itself. The end user is not involved in the content creation process and the lyrics and music are produced by people who are not personally known to the typical user. Digital commercial content arrives at the user with strict DRM regulations and metadata provided by the content supplier.

We approached these studies with this mindset, assuming that the differences between media captured/created by a user affords different uses than media that is created by others (most often strangers) and purchased/downloaded. While there were certainly some differences, we were surprised by the many similarities that we observed.

In the winter of 2002-2003, we conducted an ethnographic study of personal photo use and sharing (“Photo Sharing Study”) which investigated personal photo use in both printed and digital form. We wanted to understand how photos were shared, and the circumstances around the sharing as well as how the use of photos affected their organization (or lack thereof). This study built upon previous work in our lab, including research in people’s homes, examining how participants label, organize, and remember how to find photos.

Two years later, in the winter of 2004-2005, we conducted a separate study looking into how context affects the ways that consumers use music (“Music Context Study”). After the analysis of this study was complete, we noticed many similarities between the ways in which our participants used music and the ways that our previous participants had used their personal photos. We believe that these similarities have implications for the design of future music and photo applications.

PREVIOUS WORK

Much existing research has gone into studying photo usage or music usage separately. For the most part, our findings back up those results; however, we add additional insights and the perspective of combining results of two separate studies on the two media types. While there is some current research on the combination of photos and music, the results presented have been abstracted data that lack the details of use necessary for invention. Our goals in both studies were to innovate in the field, requiring an in-depth

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2006, April 22-27, 2006, Montréal, Québec, Canada.
Copyright 2006 ACM 1-59593-178-3/06/0004...\$5.00.

look at the micro-behaviors surrounding the use of each media type. The resulting analysis led us to notice similarities in use that we had not seen in the previously published work.

Personal Photographs

Over the past few years, a number of researchers have been interested in the use of personal photographs. Many focus on aspects of sharing photos with others. One of the most comprehensive observational studies was conducted by Frohlich et al. in their look into user requirements for Photoware [17]. They observed the communication that surrounds photo sharing and concluded that photos are really taken to “communicate experiences with others.” They saw a need for photo sharing systems that afforded rich communication and instant photo sharing. Frohlich also investigated the ability to combine sound with photos in order to create richer connections to the past [16]. In his study, users explicitly chose ambient sounds, music clips, or voice-overs to accompany photos to enrich the experience of viewing this personal content.

Balabanovic et al. [2], building on the work of Chalfen [9], who provides considerable evidence that experiences of sharing are one of the most common uses of photos, created a system that allowed for sharing and storytelling around photos on a handheld device.

Schiano et al. [32] looked at photo use among teens and saw many of the same behaviors. They focused on photos as being used to reminisce and relive past events. They also talk about photo organization ranging from event or chronological albums to haphazard shoeboxes.

Viégas [37] looked at the ways in which people describe photo sets differently to different people, and created the Collections system in order to allow people the ability to customize their display for their audience. Crabtree et al. [11] also realize the importance of recipient design in ‘photo-talk.’

All of these studies showed the importance of sharing experiences through sharing photos and the importance of photos as a memory aid and springboard for stories surrounding their particular context. Our findings agree with these results and emphasize the importance of sharing and communicating around photos. We add additional observations into how people satisfice in choosing their photos and how these same experiences can be evoked with music.

Other researchers focused on ways that people organize their photo collections and ways that this organization can be accomplished automatically. Graham et al [22] focused on events as a way of grouping photographs, allowing users to find groups of photos based on the time they were taken. Gargi, Deng and Tretter [20] also found that users tend to group their photos based on time-based “events.”

Systems like PhotoFinder [24], iPhoto, and Flickr support photo tagging to add metadata to photos that can later be searched for. Systems like Mobile Media Metadata [13] assume a smaller amount of user involvement in tagging, by sharing annotations among users. Our study and others [22, 31, 28] show that manual tagging rarely occurs due to the added effort that it requires and the simple ease of use of browsing large numbers of thumbnails by events or time. Indeed, Mills et al. suggest that “It is quite possible that users may not be willing to annotate images and may never even wish to perform a search” [28]. Our research sheds light on why that is.

Commercial Music

We found that the social science and HCI research on music has focused on a few key areas (resembling those studied for photos): music information retrieval [12, 14], musical categorization [25], and influences on music preference [10]. There has also been a significant amount of work done in music recommendation prototyping [15]. But again, we found a dearth of behavioral detail on music organization and searching that could help us invent new music experiences.

Voida et al. [38] looked at music usage and sharing with iTunes. Relevant to our work, they looked at how the music that is shared helps to create an image of the person sharing the music. While often seeing other’s music in an office setting helped to reinforce images of coworkers, they found that often people would learn more about their coworkers and change previously held opinions after browsing their music collection. This reminds us of Viégas’ Collections [37] in that users wanted the ability to control who could see which parts of their collections and thus maintain different views of their identity.

Cunningham, Reeves and Britland [12] report that finding music in libraries and stores is a collaborative experience, with browsing music, having conversations about the music, and looking at cover art together being used in affirming relationships and providing social bonding experiences.

Brown, Sellen and Geelhoed studied buying, listening, and organizing music use [6, 7]. They discuss the collaborative nature of music sharing and using music to learn more about friends’ identities. We saw several examples of this in our study, and feel that our methodology of task-based interviews and touring participants’ music collections complements their research with more concrete examples and detailed descriptions of use, especially in the area of selecting music to play.

Kim and Belkin [25] suggest that people associate music with certain events and identify music with certain emotions. The authors infer that “people’s musical information needs are often related with certain uses of music such as for a party, relaxation at day’s end, ceremonies and dancing.”

Downie and Cunningham [14] make a similar argument based on the analysis of over 160 music newsgroup information requests. In almost 20% of the postings the authors found the importance of a social connection or event that brought the music to mind in the first place, or an emotional/social memory associated with the music. While the direction of the research is leading toward an association between music and reminiscing, there is no direct association made, and no descriptions of how music results in reminiscences in daily life.

Crozier [10] takes this line of reasoning a step further and makes the argument that “enjoyment of music is essentially a social experience.” The author argues that music is a part of significant events in people’s lives, and in many ways part of the social context of living. In understanding the effect of social influence on music preferences Crozier also describes how issues of identity and self-presentation may affect the public expression of musical choice. However, he does not explore the functions of music in telling stories or recalling past experiences.

All of these authors indicate that the social context of music selection and listening is vitally important for understanding music use, and we took the information from these sources as clues for what to look for in our study. However, despite these articles, there was little in the literature to suggest how to design new and unique tools that facilitate social music use within and between the different contexts in which people work, play, and otherwise live their lives. Indeed, Hargreaves and North [23] note the dearth of research examining the social influence on the use of music.

METHODOLOGY

The two studies described herein were conducted using rapid assessment procedures [3] for data collection, what Randall, Harper and Rouncefield call a “practical approach” to ethnographic-style fieldwork for design [30]. In large part the two studies were of very similar nature. Both relied heavily on ethnographic techniques such as semi-structured interviewing, object descriptions, and direct examination of collections and organization/storage methods. Similarly, the analysis techniques in both studies relied heavily on team-based grounded theory affinity diagramming. While the models in the photo study differed slightly from the models used in the music context study, the models in both studies were adapted from the Beyer and Holtzblatt Contextual Inquiry models [5]. The photo study was conducted first, as a stand-alone study. The music study, conducted over a year later was also intended as a stand-alone study. It was only after analyzing the music data, that we began to see the similarities in photo and music usage patterns.

Photo Sharing Study

For the Photo Sharing study, the 6 interview participants were Motorolans. Three were female, 3 were male; they ranged in age from their 20’s to their 50’s, and ranged in

educational level from having a high school diploma to a Master’s degree. Participants ranged in occupation from office manager to marketing to engineering, and none were known to us prior to their participation in the study.

Self-reports of sending and receiving photos varied from “daily” to “4 or 5 times a year,” and in-person sharing varied from “monthly” to “4 or 5 times a year.” We asked the interview participants not to modify their digital collections prior to the interviews. The interviews were conducted at work where many participants had told us they sent and received, and otherwise shared and displayed, photos. We requested that the participants bring hard copies of photos they had sent or received recently, as well as any associated letters or e-mail messages. All participants also gave us access to their computers, and during the interviews they showed us the photos they sent and received through e-mail and put on websites or saved from websites. Some participants also had photos stored on cameras or camera phones and on CDs, which they brought with them to the interviews. We videotaped and audio taped the one-hour interview sessions.

We had selected people who reported sharing photos in the last week or last month, so we asked them to recall specifically the occasion of their recent sending or receiving events, and to show us the photos that were sent or received. We probed them for intents, relations with the people they shared with, and the occasions for sharing. We had also defined sharing broadly, so we encouraged our participants to talk about showing people photos as well as giving people photos or sending photos.

Analysis of the data was conducted by a multidisciplinary team and included selecting and abstracting the data while watching the videos of the interviews as a team. We extracted direct quotes and observations of behaviors and we created an affinity diagram, workflow diagrams, and task sequence models to use in design, and to inform new application development.

The affinity diagram was used as a modified version of grounded theory analysis [4], building up, inductively, a hierarchy of themes, to highlight general patterns in photo sharing behavior. The item level of analysis preceded the identification of patterns, and from this we could draw conclusions about our participants’ goals and intentions .

Workflow diagrams and sequence models were used to discover patterns of problem areas in current user tasks: places where the currently available techniques and technologies make attaining the user’s ultimate goal difficult. They were also used to find patterns of communication flow, and patterns of steps used in photo sharing. This data drove design by helping us conceptualize what the user wanted to do, and the types of tasks the users considered necessary while doing what they wanted to do. We used the data to create use case scenarios that describe possible solutions to the problems we saw.

Music Context Study

For the Music Context Study, we recruited 13 adult volunteers for detailed semi-structured and task-based interviewing along with home-tours. We worked to ensure diversity in participants and diversity in formats of music collections. Only two of the 13 participants were Motorolans. Females comprised 7 of our 13 participants, and ages in our sample ranged from 22 to 60, with a multitude of different occupations, but 12 out of 13 had a bachelor's degree or higher.

One of the essential factors in the final choice of participants was recency and frequency in reported music listening behavior. Another essential requirement was for a music collection equivalent to at least 100 audio CDs, and 10 out of our 13 participants had extensive music collections that were the equivalent of over 300 CDs. We asked the interview participants to leave their music collections exactly as they usually maintain it between the screening and the final interviews.

The bulk of the data collection took place during a single, 2-3 hour interview session at the participants' homes, consisting of two parts. During the first phase, participants were engaged in semi-structured interviews about storing music, finding music, and sharing music and music information. Again sharing was defined broadly and we asked them to recall specific occasions of their recent listening events, and describe to us the circumstances, music, and people involved.

The second phase consisted of task-based interviews, where the participants selected music from their personal collections based on different scenarios created from background information about real-life situations that we collected over the phone prior to the in-home interviews. The participants were asked a number of questions after each scenario in order to uncover their reasoning for the various music selections.

Participants took us on a tour of the various places where they store and listen to music, including living and family rooms, bedrooms, basements, kitchens, patios and cars. All participants showed us their music collections, and we were able to examine sub-collections of music in trunks of cars, in car doors, in closets, on dressers, and various other places where our participants put music. We videotaped and audio taped the two to three-hour interview sessions.

Methods for analysis of the interview data were similar to those used in the photo study and were drawn from Beebe [3], Beyer and Holtzblatt [5], and LeCompte [26]. As in the photo study, the first part of analysis was the Item Level of analysis: utilizing the raw data and identifying "events, behaviors, statements, or activities that stand out" [26]. This was followed by identifying patterns in the quotations and observed behaviors, and creating an affinity diagram, looking at the data together and extracting the connections as a team. We then utilized the themes (higher levels of the

affinity taxonomy) to answer the research questions, draw conclusions and explore interrelationships.

We also created abstracted models that showed relevant aspects of the participant's physical and social environment, and indicated the flow of information and artifacts

SIMILARITIES

In the analysis phase of the music study, when we compared our initial hypotheses with the emergent themes, we began to recognize similarities between photo usage and music usage.

Search Behaviors: Satisficing and Sidetracking

In our studies, we saw certain patterns emerge in how people look through their music and photos. These patterns were similar for both media types, though we also found interesting variations. The high-level finding is that people are rarely looking for one specific item. Rather, they have a certain "kind of" thing in mind, and select the first item that matches adequately. In one example, Gina¹ was looking for alternative rock among her cable TV music channels, but when she came across a song she thought was OK on a different channel she left it at that station.

In technical terms, we might say that users are "satisficing" rather than "optimizing" their selection of music or pictures. Formally, satisficing is a stopping rule for sequential search, where an aspiration level is fixed in advance, and the search is terminated as soon as an alternative exceeds that level [33, 34]. More generally, we can use it to describe any method that only attempts to find an alternative that is "good enough", without regard for whether it is the best possible [35].

The theories of satisficing and bounded rationality are well established [1, 21], and this kind of behavior has been observed in web browsing [1, 8], job hiring, and fighting fires [21], among others. However, we find no studies that describe this behavior specifically for searching for music or for photos.

The reason people satisfice seems to be that they have so many things to choose from that they cannot evaluate them all. In the Music Context study all our participants had large music collections, but we also saw this problem with cable TV music channels and digital radio. Debbie: "It takes me forever to figure out what channel I want 'cause I have the Sirius." We found the same thing in the Photo Sharing study, where it might take many minutes for people to look through their photos.

Although we say that users satisfice, this is only a convenient approximation of how they behave. Users are not rigorously executing any algorithm, but are often distracted, change their minds, get bored, etc. Our discussion here is necessarily a somewhat generalized

¹ All names are pseudonyms

description, but we will outline briefly some different forms of satisficing we saw.

One variation that occurred both for photo finding and music finding was a two-pass method, where the user first conducted a satisficing search to pick out a number of candidates, and then made his or her final selection from those. For instance, participants who created scrap books used this method to pick out the photos to include, and Joyce chose which CDs to put in her car this way. This method is similar to the *reduction* technique described by Agosto [1] for web browsing. However, in our studies the users selected a subset small enough that they didn't need to use satisficing to select within it.

While the classical description of satisficing assumes a static aspiration level, and therefore a fixed criterion for accepting alternatives, we saw some evidence that the level changed during search, especially as people got increasingly impatient with the process. (Todd and Miller describe adaptive aspiration levels for satisficing-like algorithms [36].) This might mean making a selection they told us was a compromise, or going back to an alternative they had previously rejected.

When our participants were listening to music on the computer, they would sometimes play a mix or random shuffle of their music, and skip songs they didn't like. Chris: "I don't like this song, I don't like this song. I'll just flip through it." This approach appears to be very similar to standard satisficing, except in this case music has to be actively rejected (pick the songs to skip), instead of actively accepted (pick the songs to play). However, we noticed that although our participants might let many songs play before skipping one, once they skipped one song they often continued skipping several more times before they let a song play again. Our explanation for this behavior is that they lowered their aspiration level in order to avoid having to go to the effort of skipping a song. Once they had started skipping it was easy to keep doing so, and they raised their aspiration to a more discriminating level. This gives us a two-state model with a high and a low aspiration level, depending on which state the user is in.

We also saw cases where people gave up searching without having found anything, what Agosto calls *termination* [1]. Several of our participants kept a default selection to fall back on in case they couldn't find anything within a reasonable time. A participant in the Photo Sharing Study kept pictures from her Florida vacation with her boyfriend on her camera: "when I want to see it, I can pick up my camera and boom." For one participant in the Music Context Study it was the radio that was her default selection Gina: "I started on the first disc and listened to a track or two off of it, then went to the second disc. I got to the fourth one before I was like 'Forget it!' and put the radio on." (See also Brown et al. [7].) It should be noted that falling back to a default in this way implies a breakdown in

the search process, and is usually associated with some level of dissatisfaction.

Many times, our participants declared that they were looking for a particular photo or piece of music, but in the end selected something else entirely. One person told us she was looking for a photo of "all of my family, in the field across the street," but the picture she ended up picking was from a totally different time and place. We explain this by speculating that people are actually just looking for certain features (the aspiration level is the degree of match with those features), and that the item they mention is one they remember having those features. To some extent they are searching for the item they remember, and to some extent they are browsing for something that matches what they want. Satisficing is a matter of degree.

However, we must also acknowledge that sometimes people end up with something other than what they intended when they started because they actually change their mind during the search process. In our studies, the participants were often sidetracked while looking for media. Through associations suggested by items they came across, they would start looking for other things they had not originally intended. These were not negative events, but allowed our users to rediscover media and reminisce about their past.

In photo research conducted prior to the Photo Sharing Study, participants unintentionally started to browse through their photos during purposeful tasks. When participants were asked to find certain photos they had described to us, they would often stop looking for the photo altogether as they got sidetracked and rediscovered old photos. They would then describe the event to us, or express delight in seeing the old pictures. Said one: "Wow, look at these pictures, my son's getting big...I haven't seen these in a while." One participant even likened it to "going through an old photo album" that he forgot he had.

In a very similar way, participants in the music study were inspired by music playing to listen to other music in their collection. For instance, if a song by Weezer came up, they might decide that they wanted to play the entire album. Lamar: "Sometimes when I listen to things it triggers [...] me that I want to listen to one of those [...] CDs." Brown et al. emphasize the importance of serendipitous *browsing* over goal-directed *searching* when it comes to buying music [7]. We find that the same is true when people go through their own collections of music and photos. These sorts of discoveries led to some users getting on "kicks" and listening to the music that they found frequently while for others it led to rediscovering past events and the circumstances around the music itself.

Our participants tended to choose something that was convenient when they were picking music to play, just as they picked the most easily available photos to send or share with others. In fact, in our observations this was one of the main influences on what they picked. This is consistent with Agosto's finding that physical constraints

are an important factor in bounded rationality [1]. Participants often just played what was already in the player. Gina: “[I’ll] check what’s in here first. If there’s anything cool in [the player].” (See also Brown et al. [7].) In the same way, on the computer users would just play the current play list or they would pick music that was lying out already. For photos, we found that the most recently captured photos were the ones most likely to be sent out. We also found that the photo that is “good enough” and found quickly is more often sent than “the best possible” photo that is hard to find. This tendency is readily explained by the concept of satisficing. People go with the first good thing they find, and that means they are more likely to play music or share photos that are in one of the first places they look. Martignon and Hoffrage [27] emphasize the *ecological rationality* of methods and heuristics like satisficing: they take advantage of the structure of the environment in order to be effective and efficient. In this case, the fact that the photos and music people are likely to want tend to be convenient to access helps make satisficing a generally successful strategy.

Many of these behaviors can be described as forms of satisficing, but we also saw people being sidetracked, and browsing relatively aimlessly. Searching for specific items, though it happens, appears to be a fairly rare activity.

Ties between media and events

Throughout both of our studies we saw how music and photos are interleaved in people’s lives, as, simply, part of life. But it is not enough to say that these forms of content are fundamental parts of the social context of living. Music and photos, as we found, are a fundamental aspect of how we communicate with each other and how we communicate with ourselves. They are used when we are alone as much as they are used when we are together with others. And we found that both personal and commercial content are used to reminisce, to tell stories, and to provide a continuing communication with ourselves as well as others.

Our initial study into photo usage had shown a strong tie between photos and the event in which they were captured. The photo was not merely a snapshot in time, but for the people who captured them, a relatively small number of photos represented the entire context around that particular instant in time. During the literature search prior to the photo sharing study we found Balabanovic [2], Chalfen [9], and others concentrating on the centrality of sharing, and how the conversational context is intimately and inextricably associated with sharing photos. Photos were being used to reminisce and tell stories; the photos were essentially a mechanism for enabling conversation, and it was the memories themselves that were the focus of conversation, not the photos. Often in the Photo Sharing Study, a participant would go on to tell us a long story about a photo that they happened to be browsing past, thus bringing back the context around the photo.

The participant who kept her Florida vacation photos on her camera also had a photo from that trip on her screen-saver, and she said it “reminded” her of her boyfriend who was then living far away from her. Another participant found himself reminiscing about special times in his son’s life as he created a photo CD to share at his son’s birthday party.

Beyond using media to connect with past events and reminisce, we observed users using media itself to tell these stories to others. This can be seen in photos when users send particular snapshots of things that they come across in everyday life that they think others would enjoy, or creating custom slideshows that illustrate a child growing up throughout the years.

When sending photos to others, there is an associated story that is either implied, if the recipient is expected to know the story, or explicitly stated if the recipient was felt to need the contextual information. For example, one participant used her camera phone to send a single picture of a Christmas tree in the town square to her friend, with the message “Hey look at this, we’ve got a huge tree in the neighborhood.” She believed her friend would understand that this was the same town square they had visited the previous year on vacation, but without the tree. The photo needed no other explanation. But another participant sent a group of photos of her new house to relatives and friends, and the e-mail told the story, enhanced by the explicit detail of the labels: “the back yard.” This group of photos needed the story told with them. One CD created by a participant was, itself, the story-telling mechanism. The participant put together a CD of pictures and video from his son’s first haircut to give to relatives.

Story-telling with photos is also, obviously, conducted in person. One participant in the Photo Sharing Study described looking through her entire website collection with a close friend. Another participant mentioned sharing an entire “slideshow” of her husband running a marathon at their party – and explained that her guests, who did not know the context, asked questions and prompted discussions of the event on “display”. All of our participants mentioned re-telling the stories of events using the photos to “update” friends and family and to brag about children or grandchildren (or their husbands).

We did not expect to see this strong tie between music and event-based story-telling when looking into the contexts around music use. However, participants consistently mentioned stories and times in their lives that revolved around their music. Music was very important in their lives. Joyce: “I always say you could get to know me very well by the music I listen to.”

Often, music embodied strong personal links not just to times or event, but also to other people in their lives. Music, in these cases, was used to reminisce. Sometimes this reminiscing was done alone, but sometimes it was done with others – where the music took second stage as the story of a time gone by became front and center. For

example, music reminded listeners of the times they had listened to it before. Kara vividly remembers listening to particular music: “Me and my sisters sitting on the porch and talking in our Nike sweatshirts. We used to play this song over and over.” There are other songs that she remembers the same way. During our interview, she tied particular songs and CDs to people and events in her life and recalled the associated stories without even listening to the music, just by looking at the CD covers as she flipped through her stack. Emma told us that she often tells these stories to others: “If a song brings out a story [...] I usually share that...” Clearly music is putting people in a reminiscent frame of mind.

Other times, music provided a way for our participants to express what they were feeling to others. Music was being used as a story-telling device in a number of different ways. As with photos, our users often made special music mixes in order to say something special to someone or convey parts of their own life to others. Chris made a CD for his brother with songs of importance throughout his brother’s life for his birthday. “I pulled together music that he used to listen to when we were kids in high school [and other music].” He then made a photo collage of his brother through the years and sent it as a part of the CD. Here, very explicitly, music is being used as a part of a story, the story of Chris’ brother’s life. In another episode she described to us, Kara and her husband together listened to a song that resonated with him because of his experiences in the military.

When Isabella was dating someone just after college, she created a series of “Life of [Isabella]” CDs that told her life story. Some of the music was just music that she really liked at different points in her life, while other music was music she later related to based on the lyrics and past experiences. Her boyfriend at the time made a similar set of CDs for Isabella relating the story of *his* life, and they used these CDs to learn more about each other as part of the “getting to know you” conversation.

In this example, it was not the music that was the important topic; it was the story of their lives. They wanted each other to know something about them, and used the music to express those some-things. Music is being used as a shared artifact in storytelling. All these instances show the power of music in relating to people’s lives and the powerful things that can be said through music.

As with photos, sharing music, it seems, is part of an ongoing conversation between giver and recipient – an emergent communication interaction that may take place over many days, weeks, months, or even years (as with exchanges of mixed CDs). These “conversations” involve different forms of dialog depending on the medium used, but they illustrate how commercial content (music) can trigger associations so personal that it takes on the characteristics of personal content (photos).

This behavior suggests that commercial content ties in with the events and places people received or experienced the content. Systems could use this contextual information to suggest music or tie in photos or other media from events to music this is playing. For both photos and music, the building of social relationships is a significant part of how they are being used, outside of “buying” or “browsing” together. [12]

Organization

Music organization styles and photo organization styles are also similar: some people organize meticulously, and others just throw things wherever they can fit them or wherever is convenient at the time. Music, like photos, gets “archived,” and both end up in dusty collections that are rarely revisited. This can be represented by photos in shoeboxes under the bed or in closets, or likewise cases of CDs stored away or kept dusty on shelves. Photos that are “close to the heart” are kept close at hand, and “favorite” music (or music that is currently in favor) is kept convenient to the player. Labels vary from precise (date, time, event, and people for photos; descriptive title and track list for music) to cryptic (a photo titled “the kids” included pictures of nephews as well as the participant’s own children; a CD titled “MP3 mix 1”).

As did Kim and Belkin [25], we found evidence that people use their own personal categorization scheme for the music in their collections, identifying certain music with events, such as “wedding songs” or emotions such as “sad,” but we also saw that they had identifiers for their music that were theme-based, such as “I hate men” or “feel good.” Here, again, we find a similarity in participants’ photo organization schemes, with folder labels created based on categories such as “E-mailed photos” and “my new house” and photos themselves labeled with descriptions such as “Ho Ho Yawn” – a photo of one participant’s grandson, yawning as he sat next to Santa.

Keeping poor quality music is also similar to how people make decisions about keeping photos: people seem to keep both music and photos based more on their personal connections with the media than on evaluations of quality. Sometimes the quality of the pictures are poor just as the quality of music is poor (as with the collection of college tapes Joyce had), but they are kept anyway. We found participants in the photo study keeping pictures that were so dark that the subject matter was unrecognizable, but the photo still brought back memories of the Ice-Capades and was kept because it was the only one they took of her daughter’s favorite character Cinderella.

The ways in which people organize their collections show several themes that could be used in creating new systems. Especially interesting are ideas to support active collections and media “on the shelf.”

IMPLICATIONS

We believe that this work has implications for the design of systems that interface to personal photo or music content, or both. Having found many similarities between the use of music and the use of photos, one implication is that good ideas from one kind of system should be examined to see if they could be applied to the other. For instance, applications that allow people to browse their photos chronologically or by event could fruitfully be adapted for music, given support for usage history and other metadata.

We believe that many of the social practices around photo sharing and reminiscing also apply to music listening and sharing behaviors. For example, the work at HP into storytelling about photo events [15] can be adapted to creating systems for storytelling around music. One particular design, which came directly from our data, was to create a system that automatically creates music mixes that are representative of certain events or people. Looking at the usage history of the content played by the user at certain times over their life, it is possible to automatically create mixes that contain the most frequently played music from certain periods given the right metadata. For example, the system could pick a song from each five years of a user's life. Or it could automatically create a CD for a particular vacation or event for a particular person based on music most frequently played at that event or with that person. A system like this would allow people to tell stories of their lives as well as reminisce about music that was once meaningful to them much in the way people talk about the stories around the individual photos when showing someone photos from an event. Specifically, this idea came from users like Isabella, who as we reported created CDs of her life to share with her boyfriend as part of a getting-to-know-you process. This is similar to sharing photo albums (i.e. the most looked at pictures) to learn about one's past.

This same implication has led to the creation of a prototype for a different use case. The idea came from the observation of users commenting on each other's music combined with people telling stories about themselves through the mixes that they create. In short, the prototype allows users to see on the idle screen of their mobile phones a short history of the artists and titles of the music that their friends are playing and for them to post lightweight messages about the songs. In a brief user study, we recruited four teenage friends to use this system for five days. Their music players were outfitted with plugins that would allow the data to be sent to their friend's phones and users were told of the functionality of the system in seeing what their friends were playing. During the study, two of the users played music specifically to send a message to other users. One purposefully played a song that was given to her by a friend so that he could see that she was playing it. At one point, she played only songs that had her first name in the titles for a period of time until everyone noticed. Another user played a song that had humorous in-group context specifically "for" a friend, and although the friend could not

hear the song, they received the metadata on their phone. The user who played the song wished that the other friends would notice too and realize the humor in song choice.

The other possibility suggested by the similarity of photo and music use is to combine the two. By binding together the experience of music and pictures, we could create "experience packages" that evoke more memories of past events and support richer storytelling than either one alone.

The current common practice of storing metadata for photos and music in separate systems using different formats is unfortunate because connections cannot be made. While we believe in the importance of applications that perform their particular functions well (e.g. a photo browser or music player), we believe they should rely on a shared underlying database of contextual metadata and usage history (such as we present in [19]). For example, users like Kara, quoted above, explicitly remembered particular events where she had listened to particular music. If all content were tagged with sufficient metadata, connections could automatically be made across media to create rich experiences for remembering and reminiscing about past events. Currently, applications like Google Desktop Sidebar show an ambient slideshow of random personal photos on the side of the user's screen. If music and photos shared a relational metadata repository, they could show photos from events when the currently playing music was played or songs from that particular month or year. This would allow users to relive those past events in a richer form than just hearing the music without additional explicit effort. Frohlich [16] also saw the possibilities of combining audio and photos to create rich experiences in his book *Audiophotography*; however, his systems required users to attach a specific ambient sound or music clip to a song. Systems based on relational metadata allow these connections to be made by the application itself, thus freeing the user from elaborate tagging and creation of explicit multimedia works. This is not a mere combination of music and photos in one application, but a common infrastructure that allows applications to make use of the meaningful connections between the two media types.

Another implication of this work is that systems should explicitly tie context into multiple media experiences. We have found that for both personal content and music, users often relate the content to the time of capture/purchase or to times when the content was experienced. Not merely a feature for convenience, this will enable systems that more prominently support reminiscing and storytelling.

Finding that users often select music and photos by satisficing, we feel that systems should be designed to support and enhance this behavior. Even in computer applications, users were often reduced to going sequentially through each item in the whole collection, or a significant subset of it. We believe the keyword search available in modern photo and music applications does not work well with a satisficing strategy, because of the overhead

involved in iterating on the search, as described by Information Foraging Theory [29]. Rather, applications should take advantage of the flexibility they have in dynamically organizing data, in order to present the user with a quite small subset of items likely to include something they will like. Making recent items conveniently available is a good first step, but we feel more contextual selections can also be generated.

People's search behaviors also imply a need for systems that support much more serendipitous browsing over explicit tagging and search. Such browsing experiences, whether exploring music through types of randomized play, or scanning through photo albums on the way to an event that the user is looking for, allow for exploration of content that would not be seen in a direct search. This exploration often leads users to discover media that they had forgotten and brings back a rush of memories of certain events. For music, these experiences can start users on a "kick" of listening to music that they had temporarily forgotten about or remember back to times in their lives when they had listened to a particular song or album. These rich experiences should still be supported in new digital systems for managing this content. As content libraries grow, it becomes increasingly difficult for people to remember everything that they have. Therefore, experiences that allow for discovery of content and the ability to view similar content will lead to users finding unexpected but "good-enough" results to satisfy their needs.

Instead of designing systems based on distinctions between personal content and commercial content, it is important to note the similarities in uses of different kinds of content and design with this in mind. It is possible that the insights outlined above can contribute to the creation of applications that help people combine their photos and music more easily and in creative ways for story-telling and reminiscing. Creating interfaces that explicitly tie this context to the media could allow for very rich experiences. Imagine being able to browse both songs and photos from a particular day such as a graduation or wedding. Tying photos, music, and events together can allow for much richer multimedia experiences than just happening on a song or picture in separate interfaces.

While social experiences bring certain music to mind, we found that music brings social experiences to mind, and this can be a crucial distinction when designing new music use applications. Just as with photos, music can be used intentionally as a trigger to help people relive past experiences in more meaningful ways.

Finally, we feel that it is important to create systems that allow for people to communicate using media. The saying claims that a picture is worth a thousand words, and people indeed use photos as well as music to convey complex thoughts, feelings, or portrayals of their identity to others. Systems that can exploit the rich meanings behind photos and music, and allow for easy sharing of content, can help

people better communicate with those who are close to them.

ACKNOWLEDGMENTS

We thank all those who were also involved in the original Photo Sharing Study including Emilee Patrick, Krishna Avasarala, and Mark Lill. We also thank those involved in the Music Context Study including Vivek Thakkar as well as Drew Harry for his work on the Music Presence system and evaluation.

REFERENCES

1. Agosto, D. Bounded Rationality and Satisficing in Young People's Web-Based Decision Making. *Journal of the American Society for Information Science and Technology* 53:1 (2001), 16-27.
2. Balabanovic, M., Chu, L.L. and Wolff, G.J. (2000). Storytelling with digital photographs. Proc. Conference on Human Factors and Computing Systems. The Hague, The Netherlands.
3. Beebe, J. (2001). Rapid Assessment Process: An Introduction. Walnut Creek, CA: Altamira Press.
4. Bernard, H.R. and Ryan, G.W. (1998) "Text Analysis: Qualitative and Quantitative Methods" in Handbook of Methods in Cultural Anthropology, H.R. Bernard ed., Altamira Press, 595-646.
5. Beyer, H. and Holtzblatt, K. (1998). Contextual Design: Defining Customer-Centered Systems.
6. Brown, B., Geelhoed, E. and Sellen, A. (2001) Music Sharing as a Computer Supported Collaborative Application. in Prinz, W. ed. Proceedings of Ecscw 2001, Bonn, Germany., Kluwer Academic Publishers, 2001, 179-198.
7. Brown, B., Geelhoed, E. and Sellen, A. (2001) The Use of Conventional and New Music Media: Implications for Future Technologies. In Hirose and M., editors, Proc. Interact'2001, 67-75
8. Card, S.K. et al. Information Scent as a Driver of Web Behavior Graphs: Results of a Protocol Analysis Method for Web Usability. *Proc. CHI 2001*, ACM Press (2001), 498-505.
9. Chalfen, R. Snapshot versions of life. Bowling Green State University Press, Bowling Green OH, 1987.
10. Crozier, R. (1997) Music and social influence. In: *The Social Psychology of Music* ed. by D. Hargreaves and A. North. Oxford University Press, pp. 67-83.
11. Crabtree, A., Rodden, T., and Mariani, J. (2004) Collaborating around collections: informing the continued development of photoware. Proc. CSCW '04. ACM Press, New York, NY, 396-405.
12. Cunningham, S., Reeves, N., and Britland, M. (2003) An Ethnographic Study of Music Information Seeking: Implications for the Design of a Music Digital Library.

- Proc. ACM/IEEE-CS joint conference on Digital libraries, Houston, Texas..
13. Davis, M., King, S., Good, N., and Sarvas, R.. (2004) "From Context to Content: Leveraging Context to Infer Media Metadata." Proc. International Conference on Multimedia (MM 2004). New York, 188-195.
 14. Downie, J. S., and Cunningham, S.J. (2002) Toward a Theory of Music Information Retrieval Queries: System Design Implications. Proc. 3rd International Conference on Music Information Retrieval.
 15. Dunne, J., et al. (2002) mpME!: Music Recommendation and Exploration. Proc IUI 2002, San Francisco, California. ACM Press publishers.
 16. Frohlich, D.M. (2004) Audiophotography: Bringing photos to life with sounds. Kluwer Academic Publishers.
 17. Frohlich, D., Kuchinsky, A., Pering, C., Don, A. and Ariss, S. (2002). Requirements for PhotoWare. In the Proc. CSCW 2002. New Orleans, Louisiana.
 18. Futrelle, J., and Downie, J.S. (2002) Interdisciplinary Communities and Research Issues in Music Information Retrieval. Proc. 3rd International Conference on Music Information Retrieval, Paris, France.
 19. Gandhi, B.; Martinez, A.; Bentley, F. (2004) Intelligent multimedia content management on mobile devices. Proc. ICME '04. 1703 – 1706.
 20. Gargi, U.; Deng, Y.; Tretter, D. R. (2002) Managing and Searching Personal Photo Collections. HP Labs Technical Report HPL-2002-67, 2002.
 21. Gigerenzer, G., Todd, P.M., and the ABC Research Group. *Simple Heuristics That Make Us Smart*. Oxford University Press, New York, NY, USA, 1999.
 22. Graham, A., Garcia-Molina, H., Paepcke, A. and Winograd, T. (2002). Time as essence for photo browsing through personal digital libraries. In Proc. International Conference on Digital Libraries.
 23. Hargreaves, D. J., and North, A.C. (1997) The social psychology of music. In: *The Social Psychology of Music* ed. by David J. Hargreaves and Adrian C. North. Oxford University Press, 1-21.
 24. Kang, H. and Shneiderman, B. (2000). Visualization Methods for Personal Photo Collections Browsing and Searching in the PhotoFinder. In Proc. ICME 2000.
 25. Kim, J.Y. and Belkin, N. (2002) Categories of Music Description and Search Terms and Phrases Used by Non-Music Experts. In: Proceedings of the 3rd International Conference on Music Information Retrieval, Paris, France.
 26. LeCompte, M.D. and Schensul, J.J. (1999) *Designing and Conducting Ethnographic Research. Ethnographer's Toolkit 1*. Walnut Creek, California. Altamira Press, Sage publications.
 27. Martignon, L., and Hoffrage, U. Why Does One-Reason Decision Making Work? A Case Study in Ecological Rationality. In *Simple Heuristics That Make Us Smart* (Gigerenzer, G., Todd, P.M., and the ABC Research Group). Oxford University Press, New York, NY, USA, 1999.
 28. Mills, T., Pye, D., Sinclair, D, Wood, K. (2000). ShoeBox: A Digital Photo Management System. AT&T Laboratories Cambridge.
 29. Pirolli, P., and Card, S.K. Information Foraging Models of Browsers for Very Large Document Spaces. *Proc. Working Conf. on Advanced Visual Interfaces 1998*, ACM Press (1998), 83-93.
 30. Randall, D., Harper, R., and Rouncefield, M. (2005). Fieldwork and Ethnography: A Perspective from CSCW. Proc. EPIC 2005. 81-99.
 31. Rodden, K. and Wood, K. R. (2003) How do people manage their digital photographs? In Proc CHI 2003. 409-416.
 32. Schiano, D., Chen, C.P. and Isaacs, E. (2002) How Teens Take, View, Share and Store Photos. Proc. CSCW 2002. New Orleans, Louisiana.
 33. Simon, H.A. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics* 69 (1955), 99-118.
 34. Simon, H.A. Rational Choice and the Structure of the Environment. *Psychological Review* 63 (1956), 129-138.
 35. Simon, H.A. Invariants of Human Behavior. *Annual Review of Psychology* 41 (1990), 1-19.
 36. Todd, P.M., and Miller, G.F. From Pride and Prejudice to Persuasion: Satisficing in Mate Search. In *Simple Heuristics That Make Us Smart* (Gigerenzer, G., Todd, P.M., and the ABC Research Group). Oxford University Press, New York, NY, USA, 1999.
 37. Viégas, F.B. (2000). Collections: Adapting the display of Personal Objects for Different Audiences. Master of Science Thesis, Massachusetts Institute of Technology.
 38. Volda, A., Grinter, R.E., Ducheneaut, N., Edwards, W.K., and Newman, M.W. (2005). "Listening in: Practices surrounding iTunes music sharing." Proc. CHI 2005. ACM Press, pp. 191-200.