

Saving, Reusing, and Remixing Web Video: Using Attitudes and Practices to Reveal Social Norms

Catherine C. Marshall
Microsoft Research, Silicon Valley
1065 La Avenida
Mountain View, CA 94043
1-650-693-1308
cathymar@microsoft.com

Frank M. Shipman
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
1-979-862-3216
shipman@cs.tamu.edu

ABSTRACT

The growth of online videos has spurred a concomitant increase in the storage, reuse, and remix of this content. As we gain more experience with video content, social norms about ownership have evolved accordingly, spelling out what people think is appropriate use of content that is not necessarily their own. We use a series of three studies, each centering on a different genre of recordings, to probe 634 participants' attitudes toward video storage, reuse, and remix; we also question participants about their own experiences with online video. The results allow us to characterize current practice and emerging social norms and to establish the relationship between the two. Hypotheticals borrowed from legal research are used as the primary vehicle for testing attitudes, and for identifying boundaries between socially acceptable and unacceptable behavior.

Categories and Subject Descriptors

H.4.3 [Information Systems]: Communications Applications.

General Terms

Design, Experimentation, Human Factors, Legal Aspects.

Keywords

Video reuse, social media, property rights, social norms.

1. INTRODUCTION

The growth and ubiquity of video as an everyday medium raises new issues about its ownership and control. Before the advent of inexpensive digital video production and display tools, video belonged in a handful of distinct realms. On one hand, home movies were used to record special occasions and family events; on the other hand, professional video was the realm of experienced videographers who used the medium primarily for entertainment and education.

The distinction between professional video and amateur video is blurring. Home video producers can purchase sophisticated cameras and video manipulation tools that were once limited to professionals. Although not all home video enthusiasts have the talent and skill to make professional-quality videos, they do have access to the means of production and venues for distribution. A popular amateur video on YouTube or Vimeo may have cost little to produce, and may garner as many views as a movie or TV

series that cost millions to make and is available through a for-pay subscription service like Netflix or Hulu. Sites like Funny or Die blur the distinction further, as homemade videos vie with professionally-made web series for audiences' attention; in fact, this blurring is sometimes purposeful, as amateur videos aspire to be mistaken for professional videos and vice-versa.

At the same time, because modern video is a digital medium—and therefore easy to download, copy, republish, and remix—professional video makers (and some amateurs) are nervous about protecting their assets, either via law, technology, or explicit provisions of a service's terms and conditions. It is unusual for video producers and consumers to even be fully aware of the applicable laws, pivotal legal cases, policies, and provisions that restrict copying or (more infrequently) promote reuse. To further complicate matters, social media sites like Facebook and Twitter make it easy to repost videos to reach new audiences.

Many competing interests are at work. Generally video sharing services and file storage sites err on the side of caution: why provoke needless copyright infringement lawsuits? Similarly, professional media production concerns such as Sony and Disney aggressively push forward to protect their assets [27]. Meanwhile, non-professionals may have other concerns: why prohibit distribution if what you're seeking is fame? Amateur video-makers rely on a combination of sharing, remix, and reuse to achieve viral status. At the same time, personal video makers may rely on privacy through obscurity [32].

Given these tensions, developers, service providers, and policy makers are faced with a number of decisions that will influence users' behaviors; concomitantly, as scholars such as Lessig point out [13][14], social norms are emerging that guide peoples' behavior and attitudes. We are interested in characterizing these social norms, and exploring the practice of saving and reusing different recorded genres. We do this by performing three related studies that gather peoples' reactions to a series of hypotheticals (modeled after those used in legal research [24][31]) and the characteristics their video use.

We are specifically interested in three common practices—saving video on personal storage; reusing video as-is in new venues; and remixing video content to create new forms—along with questions introduced by new technologies such as cloud storage and social media as well as questions raised by concepts like permission, attribution, and community-contributed metadata.

To provide background for this study, we first summarize related work. Next we describe our method and characterize the participants and their use of online video. We then present findings about participants' attitudes toward the storage and reuse of recorded media. Finally, we discuss the implications of these findings on design and policy.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media.
WWW 2013, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2035-1/13/05.

2. RELATED WORK

Studies of video sharing have explored how people upload and share content and the communities this activity creates [25]. Two research areas more closely related to our studies investigate (1) the design of systems to support video remix and (2) the automatic identification of reuse/remix in video collections.

Support for Video Reuse. Video reuse relies on existing video content. Most video sharing sites are not designed to support remix. An exception is Metavid.org, an archive of congressional presentations intended as the raw footage that will serve as fodder for remixes [4]. Video remixing also relies on tools for collecting and editing existing video while creating a new video, e.g. [29]. Some tools focused on particular classes of video. For example, Vihavainan et al. [33][34] explore the design and use of tools for remixing videos recorded by audience members at a concert. Looking at the effect of technology design on remix practice, Diakopoulos et al. [5] explore how constraints of a remixing service influence users' creations and how users' values related to authorship affect the design of the service.

Detection of Reuse. Much of the research in the detection of video reuse explores techniques for recognizing the overlap within video collections (e.g. deduplication). A few studies go beyond the identification of reuse and discuss applications of the results of this process. San Pedro et al. [27] describe the application of tags based on identifying remix within YouTube. A more specific study explored how using the results of remix detection to automatically provide attribution compares with human attribution [21]. In their study of identifying video memes through duplicate detection, Xie et al. [35] found a considerable amount of remixed content and that the mix of content from citizen journalists and traditional news media varies depending on the topic.

3. METHOD

We conducted a series of studies to investigate emerging ethical norms and current practice associated with saving and reusing online video content. The studies used crowdsourced questionnaires to elicit this data; each questionnaire was implemented as a Mechanical Turk Human Intelligence Task (HIT) following published best practices [6][11][12]. Participants were solicited from US-based Turkers who had a history of reliable task completion and were self-reported users of the media type and genre in question. By imposing these restrictions, we hoped to ensure a background level of cultural congruity; participants were thus subject to the same legal system and were exposed to many of the same cultural touchstones.

Each of the three questionnaires we discuss was designed around a different media type or genre, and sought to characterize participants in several ways that would afford straightforward triangulation among questions and would give us a reasonably nuanced picture of who they were, what they thought, and what they did. First, we collected some standard demographic information; this information helped us understand what kind of people were filling out the questionnaires. For example, were they students? Was Mechanical Turk their primary income source? How experienced were they as Internet users and content contributors? Second, we set up familiar hypothetical situations—for example, recording a job interview over Skype—coupled with variations of each situation's details to elicit some basic responses that would reflect the participants' attitudes about saving and reusing online content. Finally, we asked participants about

their own practices, for example, what kind of video they had shared or watched themselves. Two reading comprehension questions helped ensure that participants were paying adequate attention to the scenario details.

Table 1 summarizes the three studies. The number of acceptable responses is reported along with the total (in parentheses). Length refers to the number of questions in the HIT, including the comprehension questions and a concluding question (“Would you be willing to work on more HITs like this one?”) that we used to make sure the Turkers were not unhappy with our approach (reported in the “Other” column). The studies yielded 634 valid responses out of a total of 719 responses. Throughout this paper we use the abbreviations listed in column 1 to refer to the specific studies and to label participant quotes (e.g. PC014 refers to participant 14 in the podcast study).

We approached data quality issues conservatively: we erred on the side of caution and discarded data based on a point system. Unanswered questions, violations of the going-in restrictions (e.g., participant reported themselves to be other than English-speaking), a suspiciously short work time, or wrong answers to the comprehension probes were each worth one point. If participants scored 2 or more points, their data was discarded. We paid participants according to standard Mechanical Turk rates even if we discarded their data. The point system gave us a means to retain data from participants who might have found one or two questions confusing, or who completed HITs prior to accepting the work (a common practice reported in online forums like Turker Nation). In practice, bad responses were easily detected from a participant's answers to the open-ended questions, but the point system allowed us to handle borderline cases consistently.

Our scenarios set up a series of hypothetical situations in which facts were varied to test aspects of participants' attitudes about the ownership and control of online video content. Specifically, we tested the fairness associated with four standard actions: storing video content, sharing video content, republishing video content, and removing video content. The facts characterize standard features of the situation: the type of online content, where it is stored, and the stakeholders who are taking the actions.

For example, we use a series of hypotheticals to vary whether the person taking the action owns the material, is portrayed in the material, or is peripherally involved in the creation of the material; similarly, we use hypotheticals to vary the extent of the content that is affected (e.g. does the action involve all of a work or only a short excerpt?) and to explore other ethical dilemmas having to do with digital content ownership and manipulation, including concepts such as permission, anonymity, and privacy.

We recount the podcast scenarios and hypotheticals in greater depth as an example to show how the hypotheticals work. The podcast scenarios revolve around 4 named characters: two on-air hosts, an engineer who helps them record the weekly show, and a guest who is on the particular podcast in the scenarios. First we vary who saves the podcast to his local hard drive to test the concept of ownership and storage: can any of the people involved in the podcast's production save the podcast? We use a slippery-slope construction to move from one hypothetical situation to the next. Once we have explored aspects of the connection between

Table 1. Study identifiers, focus, number of responses and questionnaire structure

Study	Media type/genre	# responses	Length	Demographic	Hypotheticals	Practice	Other
1 (PC)	Interview podcast	225 (239)	42	12	22	5	3
2 (VC)	Recorded Video	200 (229)	40	12	20	5	3
3 (ED)	Academic lecture	209 (250)	44	12	25	4	3

Table 2. Summary of scenarios used in the studies and the main concepts they test

Study	Description of central scenarios	Concepts tested
1 (PC)	Two comedians host a weekly interview-format podcast; an engineer (who records a podcast of his own) helps them with production. This week's guest, a musician, performs a parody of a popular pop song during his interview. A fan tags the interview, which enables an ex-band member to find the episode and comment on it. The podcast is re-edited and republished under a variety of circumstances.	<i>Ownership, local storage, reuse, remix, excerption, commercial v. non-commercial reuse, permission, removal of social metadata v. removal of primary content.</i>
2 (VC)	A computer programmer's job interview is conducted via Skype. He records the interview and shares it with a friend (to help a friend apply there too). When he is not offered the job, he posts the interview on YouTube in an effort to interest other employers in him. He also republishes it with a commentary track to help others interview better (using only his side of the interview). Finally he turns it into a parody for a comedy website.	<i>Ownership, local storage v. cloud storage, reuse, remix, excerption, parody, educational use, reposting on social media v. publication on a comedy website.</i>
3 (ED)	An astronaut delivers a popular commencement address at a large public university; she talks about the importance of the manned space program. The lecture is published on a free online service that requires special software be installed on the viewer's computer (the service also provides for-pay content via this delivery vehicle). A geologist critical of the manned space program blogs a rebuttal.	<i>Ownership, local storage, transcoding, permission, republication, ownership, storage, republication, and remix of social metadata.</i>

ownership and storage, we test republication (without any further modification of the podcast); this test moves into adding content before the material is republished. Then we switch the participant's attention from the creators to the guest and his actions: can the podcast's guest take the segment he appears in and modify it before republishing it? Can he excerpt a song he performed on the podcast and sell it separately? The situations gradually test the edges of the participants' ethical boundaries.

These types of scenarios and hypotheticals are borrowed from legal theory, where they are used to help legal scholars explain doctrine and explore the moral underpinnings and consequences of legal rules [15][31]. Hypotheticals are used widely in the law [16] and may be used to "present, support and attack positions (e.g., by testing the consequences of a tentative conclusion, pressing an assertion to its limits, and exploring the meaning of a concept)" and to "factor a complex situation into component parts (e.g., by exaggerating strengths, weaknesses or eliminating features)." ([24], p. 168)

Thus we borrow a technique widely used in law schools and in certain oral argument situations (e.g. the US Supreme Court makes extensive use of hypotheticals to decide cases) and apply it to elicit everyday attitudes toward the storage and reuse of digital video content. This enables the participants to reason about the ethical forces involved in the scenarios without requiring that they imagine potentially widely varying details themselves; it also allows us to make meaningful comparisons between their Likert scale responses to the questions. Table 2 summarizes each scenario and the concepts it tests.

4. PARTICIPANTS

Because they are Turkers, participants are likely to be Internet-savvy; we also require that they have some investment in video content (as users). These limitations are imposed by design; we are more interested in people who have actual experience working out ownership and reuse questions for themselves, for example as they repost videos in Facebook or create YouTube videos. They will have grappled at least implicitly with questions like, "will I get caught if I use Dropbox to share this copyrighted movie with my friend?" or "is it okay for me to post a homemade music video based on my favorite band's hit song?" Although it is unlikely participants will have encountered the specific situations we pose in the hypotheticals, they will have less trouble imagining them than someone who has little experience with video.

What type of Turker responds to this sort of questionnaire? As we might expect, the population is dominated by Turkers in their twenties and thirties (born in the 1970s and 1980s); only about 13% were born prior to 1970 and about 13% were born after 1989. Although often this type of research is reported to have an uneven female-male ratio [10], our studies run counter to this trend, attracting a fairly balanced population. 55% of the podcast study participants are male, as are 53% of the videoconferencing study participants and 50% of the educational recordings study participants. They report being well-educated; over 90% said they have attended at least some college and over 60% have a college degree. About 1/3 are currently students. This self-reported characterization is supported by participants' responses to open-ended questions.

Although the participants are completing questionnaires for pay, generally Mechanical Turk is not reported to be a major source of income; rather many participants either have free-lance jobs in the digital economy (e.g. as graphic designers, programmers, IT support, fact-checkers, writers, or editors), or they do other types of work that places them in front of a computer for much of the day (clerical, library, or paralegal work or office management).

5. PRACTICE-RELATED FINDINGS

In past studies, we have found that participants are better able to reflect on specific ethical questions if they have had to address them in their own experience; they are less apt to fall back on hyperbole and apocryphal stories. Thus we asked about video-sharing experience in a number of ways: via a check-box list of online activities (PC,VC,ED), via open-ended questions about online activities and online publishing (PC,VC,ED), via an open-ended question about video-sharing and reuse in the abstract (VC); via a multiple choice question about what they did with the last useful video they had encountered (ED); and via a yes/no question about whether they had ever shared non-music recording (PC). Their answers paint a broad-brush picture of current practice that is useful in interpreting other results and in further understanding the population that the participants represent.

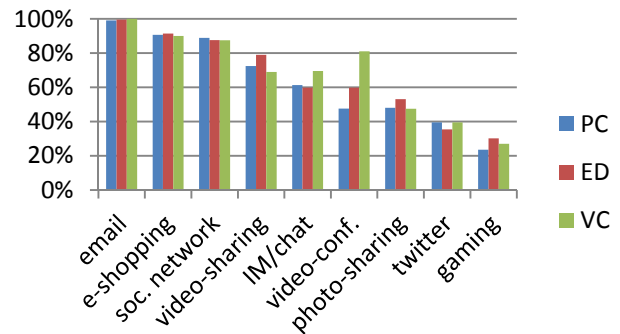


Figure 1. Participants' online activities

First, let's look at the three questions common across studies. Participants were asked to check off their online activities. Figure 1 shows that almost everyone reports using email and social networking. Video-sharing is the fourth most popular activity, more common than photo sharing (at odds with the population at large).

A parallel open-ended question supports this finding. Some participants clarified the activities they picked from the list (e.g. by shopping, they meant bidding on eBay items); others filled in gaps by citing activities we did not list (e.g. finding and listening to music). Many participants reported watching videos, curating the best to share with friends or more broadly on sites like Reddit: "I upload videos to YouTube & DailyMotion (occasionally)" [PC089] or "I've shared video/podcasts/articles on line." [ED037] Amateur and professionally produced videos are reported to be a seamless part of an online media diet: "I use [the Internet] to read up on news, watch TV Shows, upload videos on YouTube, interact with friends on Facebook, shop, and read email." [PC125]

As a secondary probe, we asked participants what they publish online; 161 out of 634 participants cited publishing video content. We did not make a bright-line distinction between sharing and publishing. Nor did we distinguish between publishing original content or republishing found media. We wanted to see if participants would make these distinctions themselves, and to see if any ownership and reuse issues arose organically before we presented the hypotheticals. Reports varied. Some referred to profile creation and status sharing as publishing, indistinguishable from other types of sharing; e.g. ED195 reported she publishes "Various status updates, some profile details such as gender and age, Pictures, videos, funny stuff, informational/education videos or stories." Under this rubric, participants said they published personal videos; e.g. ED188 said "I have a blog about my life and my dogs, I post to Facebook about my life and my dogs, I post videos on YouTube." Other participants mentioned creative efforts or cited a mix of genres: "I have shared/published articles, a few short (self-made) films, and I have re-published videos made by myself or others. I occasionally share photos." [PC089] Many participants described publishing multiple content types they had encountered; e.g. PC022 said she published: "Pictures, Basic information on myself, stories/quotes/videos I have found online."

To many, curation is tantamount to publishing. ED207 said "...I publish any interesting documentary or video related to science. And I published a few instructional videos on software usage." He acknowledges the value of his curation of science videos by reporting it in tandem with posting his instructional videos. The distinction between curation and publication may also be enforced by site (e.g. Pinterest v. LinkedIn) and genre (articles v. profiles): "I share pictures, videos, and articles with friends on Google+ or Pinterest. I have a public profile on Google+ and on LinkedIn which gives personal information about myself and my work history." [PC224] Participants feel that certain genres such as comedy and instruction invite republication without concern for content rights, e.g. ED068 said she published both humorous and educational material: "videos from funny or die or from Youtube - also share videos related to education for nurses (my former career)." Studies of encountered information (e.g. [17]) have found that sharing published material may have social motivations such as educating friends or keeping in touch: e.g. "I share videos or articles that I feel my friends would like or should know about..." [VC190]

Some responses hinted at more subtle ownership and reuse issues. VC027 said he published "...footage of gameplay from video

games." The gameplay was his, but the game content was not. Does the subject of a video own sufficient interest to publish it? VC140 reported that he published videos not just by him, but also of him: "I usually publish videos of me doing things, like DJ'ing..." How broadly are rights extended through one's friends and family? PC153 said that he "Posted video of [his] son's band on youtube..." Participants also raise a distinction between audio and video tracks; e.g. PC148 said "...Occasionally I might post a video of something from the travels but generally it's just still images. I am leary [sic] to jump into the world of audio/podcasts, I keep quiet while videotaping." We explore the separation of audio and video streams in greater depth via the hypotheticals.

A few responses imposed a linguistic separation between the two terms: one shares what one has found, and one publishes what one has recorded. For example PC141 said he "Share[s] links to interesting websites and videos. Publish videos of pets. Publish pictures of vacations, home, and garden."

Each study included one or two genre-specific practice-oriented questions. We briefly explore these responses with the thought that they reveal the participants' backgrounds, and further contextualize the quantitative portions of the studies.

Keeping podcasts. In the podcast study, we focused on *keeping*, because it is unlikely that participants create or remix podcasts. Thus we asked how many podcasts participants subscribed to and how many they retained after they had listened to them. Out of 225 responses, more than 60% (141) listen to between 1 and 10 podcasts regularly, and about ¼ listen sporadically. Only 11% listen to more than 10 podcasts regularly. Over half (57%) store at least some of podcasts with the intention of permanence. By contrast, 25% delete them after listening and 15% relegate retention to software control. These results show that the storage scenarios represent a familiar decision—can participants intentionally keep content that they've downloaded?

Watching and creating educational videos. In the educational video study, we focused on what participants had watched and whether they had ever created this type of video. Again, we were interested in discovering how far the scenarios were from their own experiences. Most commonly, they had used educational videos occasionally (44%), but it was also fairly common to have used them frequently (28%). Some participants had used them in school (18%), and about 10% had completed entire video courses.

While this made us confident that participants were familiar with educational videos, we wondered if these videos were academic (reflecting the scenario) or practical. We open-coded the responses and identified eight subgenres of instructional videos listed in Table 3, including Humanities and Social Sciences (HSS); Science, Technology, Engineering, and Math (STEM); professional development (e.g. law, EMT, criminal justice, nursing, journalism); how-to tutorials (on practical topics such as auto repair; hobbies; and self-improvement); technology use (how to use hardware, software, and video games); broad topic series (e.g. TED Talks); language-learning; and miscellaneous responses (e.g. "anything I find interesting" [ED135]).

Table 3. Genres of instructional videos viewed

Genre	#	%	Genre	#	%
HSS	68	33%	Technology Use	34	16%
STEM	64	31%	Talk Series	22	11%
Professional Dev.	58	28%	Language	14	7%
How-to Tutorials	47	22%	Misc.	10	5%

Finally, we wanted to find out if participants created or recorded educational videos. Almost one-third (31%) had recorded classes; 15% had shared these recordings online. Nearly as many (14%) had created their own educational recordings; 8% had published them on sites like YouTube. Taken together, 22% had published an educational video that they had recorded or created.

Using and recording videoconferences, and sharing videos. In the recorded videoconference study, we wanted to confirm that participants had used Skype or other videoconferencing tools. Exposure to videoconferencing and experience with video-sharing were basic qualifications for understanding the hypotheticals. Only 2% of the participants reported never having used videoconferencing, and 6% had tried it once. Most commonly, participants were occasional (42%) or weekly users (30%). Frequent or daily users made up 20% of study participants.

About two-thirds (133/200) reported that they had shared videos, although we suspect they underreported reuse of copyrighted material, since most reported sharing footage they had recorded themselves.

Table 4 shows the video-sharing breakdown. Open-coding started with a distinction between amateur, found, and professional videos; further distinctions emerged from patterns in the data. Creative efforts are videos participants created or were the subject of, and represent efforts to produce planned artistic works. Everyday documentaries are ‘point and shoot’ videos that the participant had a hand in creating (as videographer or subject). Found videos are similar to everyday documentaries, except that participants are curators rather than creators, finding the videos and forwarding them to others. Republished videos are professionally made—e.g., TV clips, movie trailers, or music videos—or are recorded from professional performances. The miscellaneous category primarily consists of responses from participants who believe they have not shared the type of videos we are asking about. A few participants justify this as a conscious decision: “I have not published any video content - the reason is because I know once you put something out there, it's there for good...” [VC161]

Each category raises different ownership and control issues, issues we anticipate that the participants have become aware of through their own experiences. For example, do subjects or performers have the same rights to content as videographers? Do members of one’s social network have more rights than strangers? Do public figures have a reduced expectation of privacy? Does genre have an effect on rights? Does the extent of distribution—or attempts to limit it—matter in determining future rights to reuse? If someone records street performers in a public place (in VC048’s case, breakdancers in the subway), does he or she need to obtain the dancers’ permission prior to posting the video on YouTube?

Copyright legislation’s fair use provisions and existing case law address many of these situations, but it is rare for people to have an accurate picture of what these provisions are, even if they are in professions that rely on fair use. For example, Aufderheide et al. have found that journalists are seldom aware of what fair use provisions dictate, and are indeed more conservative than they need to be [2]. Similarly, we might expect the study participants to reason about the video-sharing instances they report, but we might also expect this reasoning to be based on an emerging notion of fairness and social norms rather than on legal concepts.

Reuse in the abstract. We asked videoconferencing participants about reuse as a general practice. Almost without exception, they

Table 4. Video categories, frequency in data, and use examples

Category	#	Example
Creative effort (e.g. performance, film, sketch comedy)	19	"Srgt. [sic] Richard Cleener was the last video I uploaded to youtube. It was a weekend film challenge and it turned out great!" [VC135]
Everyday documentary (of or by family members)	23	"My son making fart noises with his arm. So that the family could see." [VC128]
Everyday documentary (of or by friends)	19	"I shared a memorial video to a friend who was killed in Iraq." [VC006]
Everyday documentary (vacation/travel footage)	8	"video of vacation to Egypt, to share with family and friends" [VC144]
Everyday documentary (pets)	18	"I took video of my cat rolling around on the floor and eating candy." [VC038]
Everyday documentary (instruction, diagnosis, reviews)	9	"I had shot a video of a positive PPD (Mantoux) on my arm a while ago. I posted it so that others who go through this test will have a rough idea about what will be their end result if at all they take the test." [VC093]
Everyday documentary (business-related)	8	"I produced a small seminar for work, and posted the video via DailyMotion..." [VC044]
Everyday documentary (sports/videogames)	12	"I share video highlights of my basketball team." [VC107]
Found video	12	"It was a dog balancing 36 treats on its nose. Thought it was hilarious, so decided to share." [VC136]
Professional video (specific subgenres)	8	"I shared a video that talks about the SOPA bill..." [VC016]
Professional video (movie and tv clips, commercials, music videos)	8	"The last video I shared was "baby languages" it is a part of an episode from Oprah about 5 sounds that all babies make and what they mean..." [VC047]
Performance video (recording of interview, music, standup, etc.)	7	"It was a brief clip from a sketch comedy show. I shared it because a friend of mine has a thing for Scottish men, and I thought the Glaswegian comedian would brighten her birthday." [VC193]
Miscellaneous negative responses	49	"I have not shared a video - not comfortable with it." [VC080]

interpreted the question in terms of the study’s scenarios, which involved publishing a recorded job interview. Thus, in contrast to our other practice-derived findings, participants’ answers reflect aspirational beliefs rather than normal behavior (similar to what we see happening when privacy questions are pursued [1]).

Recorded videoconferences fall under the rubric of what we’re calling ‘everyday documentary’ (or what Hill and Monroy-Hernandez refer to as functional works [9]). They take the focus off the video’s artistic merit or commercial value, and instead explore issues like fairness, privacy, and anonymity. What would participants think were central ideas, preconditions, and issues for the distribution of videoconference content? At one end of the spectrum were participants who felt there should be no publishing restrictions and at the other were participants who felt it was never acceptable. Contrast VC016’s response with VC179’s:

"I think it is ok at all times. The person is recording something that they are doing on their own so they should have complete rights to the recording. It is originally theirs." [VC016]

"it is almost never a good idea. a videoconference is almost always a "closed door" meeting, and just because it can be recorded doesn't mean that it should..." [VC179]

We open-coded participants’ responses and found that they generally reasoned about reuse starting from four situational elements: (1) the content itself and various aspects of the content (e.g. did it reveal confidential, private, or embarrassing

information? Was it libelous or malicious?); (2) the production context (e.g. were the people in the videoconference notified beforehand?); (3) the reuse context (e.g. the purpose of publishing the recording); and (4) the legal and technological circumstances (was copyright being infringed?). The dominant concepts we identified within these four overarching categories are shown Table 5. If different concepts appeared in a single response, we coded the response in multiple categories. The counts show the relative frequency of each category.

These themes and other recurring concepts (e.g. minors should be treated differently; credit should always be given; recording can be done for archival purposes) may serve as a bridge to the attitude-related findings we explored through the hypotheticals presented in the next section. As in law, the hypotheticals push participants into stating various boundaries that may not have appeared when they discussed their own practices or their attitudes in the abstract.

6. ATTITUDE-RELATED FINDINGS

Our past studies have shown that experience crucially shapes and refines peoples' attitudes toward ownership [17][19]; it is from this experience that social norms emerge. The results from the practice-oriented questions suggest we thus focus our analysis of the responses to the hypotheticals on the two actions that are common to participants' experience: saving and reusing online video. Table 2 summarizes the three scenarios we used to drive the hypotheticals—a multi-creator comedy podcast, a recorded videoconference of a job interview, and a recorded scientific talk. In this section, we delve into details of individual hypotheticals as we examine aspects of saving and reusing.

The hypotheticals related to *saving* explore four aspects of ownership: (1) the distinction between creator and subject; (2) the difference between local storage and cloud storage; (3) the distinction between primary content (i.e. the recording) and secondary content (e.g. comments and reviews); and (4) the effect of content transformations (e.g. format changes) on perceptions of ownership. Storing content on the cloud provides a venue for probing the boundary between saving and sharing; the cloud potentially exposes the content to external parties, although it is not explicit sharing. Transcoding (converting the video from one format to another) also pushes on the boundary between storage and reuse by posing potential content changes.

Next we explore concepts associated with *reuse*. In the scenarios, reuse appears in its most recognizable form: the video content is taken from one online venue and republished in another. The hypotheticals make a distinction between republication—where content is taken in its entirety and put on a different site or presented to a different audience—and *remix*, where the content is manipulated or excerpted and used in different ways (e.g. the genre may shift from instruction to comedy). Again, we look at issues that arise from reuse of associated secondary content, including social metadata such as comments or tags.

In this group of scenarios we also examine concepts from participants' responses to open-ended questions in past studies, including permission, credit, commercial v. non-commercial use, manipulation and fraud, and changing the mood of the work (e.g. from serious to funny) [18][19].

The results are organized by scenario, so the responses to the successive hypotheticals can be compared.

Table 5. Participants' reactions to reusing recordings

Situational element	Concept	#	Example
content	personal content	35	"...I think it is okay when it is informative information or impersonal items but personal events should be kept private." [VC059]
	confidential content	27	"[It's okay] when it's good for the public interest, as long as it does not violate non-disclosure and/or confidentiality agreements ..." [VC136]
production context	recording notification	40	"...I say you would need ... to let them know ahead of time that it could be shared. I don't think it is ok to post any recording if someone in it is unaware..." [VC171]
	permission or consent	104	"It's ok to share video content if it belongs to you or you get permission from the person that made the videoconference content." [VC071]
	scope	68	"Recording and sharing of videoconference content is OK when all parties have granted permission..." [VC131]
	public space or person	16	"I think it would be okay ... if it is in a public space." [VC095]
	expectation of privacy	15	"It's ok when it is done with permission or when it won't undermine someone else's legitimate expectations." [VC022]
reuse context	personal use	14	"Recording videoconferences I [am] 100% OK with, as long as the recording party intends to keep it for personal use. When publishing and sharing gets involved, things become less simple..." [VC088]
	internal use/ limited distribution	7	"...Meetings can be shared to help people who miss them catch up... Also, there should be some restrictions as to who can access these meetings." [VC189]
	good intentions/ helping others or education	20	"I think it is okay when it is intended as an archive for those who were involved in the videoconference, or for documentation. For example, I have weekly videoconferences as work with employees in other states and countries, and they recorded videos serve as proof of the decisions we made, and as resources for employees..." [VC188]
	misuse	7	"I believe it is only okay when it is for entertaining or informing purposes. I think [sic] it is a bad idea when it involves a situation like the one this survey was about because it may allow people to get jobs which they are under-qualified for." [VC091]
	libelous intentions	16	"...It is a bad idea when it is done maliciously or to cause harm by reputation or slander." [VC078]
legal and technological circumstance	imposed anonymity	24	"...If you remove names and specific content and use it as a guide it's ok by me..." [VC078] "...Only the edited down versions where the owner has claim to the video i.e.: his face and voice, should be allowed." [VC125]
	technology dictates use	3	"If the content is publish[ed] one should expect it to be available and searchable by other people." [VC178]
	legally regulated	6	"It should be subject to the same regulations as e-mails: anyone who is a party should be able to do whatever they like with it..." [VC197]
	conditional reuse is fine/ control of online content is ultimately impossible	105	"...I like that the internet is still mostly unregulated and am against anything that regulates it like SOPA." [VC191] "Be very careful and be certain it is not anything you care deeply about because once it is out there it is no longer yours to control." [VC056]

6.1 Podcast scenario

In the podcast scenario, we first explored whether it matters who saves the recording: an on-the-air creator of the podcast (host Jordan), a behind-the-scenes creator (engineer Chip), or a podcast guest, Rocky. Figure 2 shows the results. As we have seen in past studies, saving is generally uncontroversial. Participants seem to believe that if you encounter content online, you should be able to save it locally. All three alternatives trend highly positive; there is almost no difference between whether Chip or Jordan can save the podcast. On the other hand, there is mild hesitation about whether the guest (Rocky) should be able to save the recording. The difference between the responses for Rocky and the other two is statistically significant ($p < .0001$, Wilcoxon signed-rank test).




Figure 2. Responses to saving podcast hypotheticals.

A more controversial activity in prior questionnaires has been republishing part or all of an existing work. We explore this through a series of nine hypotheticals. The first six distinguish between republish as is and remix; the final three explore who has the authority to grant permission to republish. Figure 3 shows the results of the first six hypotheticals. The first two hypotheticals posit republication of the recording as-is; the variation is between creator and guest. The third hypothetical looks at partial reuse instead of full republication (Rocky excerpts a song he performed during the original recording). Hypotheticals four and five posit remix (bonus material is added and interview segments are removed). In the final hypothetical, the guest not only remixes the episode, but also republishes the comments along with his remix.




Figure 3. Responses to reusing or remixing podcast.

The results show that republishing the whole episode as-is is more acceptable than the form of remixing we described in the hypotheticals. More participants believed the behind-the-scenes creator (Chip) should be able to republish the episode on his own website than the guest (Rocky) should. This difference was statistically significant ($p < .04$). Commercial reuse was also

explored: guest Rocky excerpted his performance of a new song from the podcast so he could sell it on iTunes. This hypothetical elicited 64% positive and 24% negative responses.

When presented with a hypothetical in which Chip or Rocky published a remixed version of the whole episode, opinions were more sharply divided, with 41-42% positive responses and 45-47% negative responses respectively. The variation in who created the remix did not result in a statistically significant difference in the responses ($p < .13$). A final hypothetical explored whether Rocky could not only publish a remixed version of the podcast on his own website, but also copy some of the positive user comments attached to the original podcast to promote the remixed version. Not surprisingly, this was viewed most skeptically of all.

6.2 Videoconference

Meetings may be conducted using videoconferencing software that enables the meetings to be recorded. What ownership issues are raised by these recordings? Our second questionnaire uses scenarios in which interviewees Bill and Kyle record Skype-based job interviews.

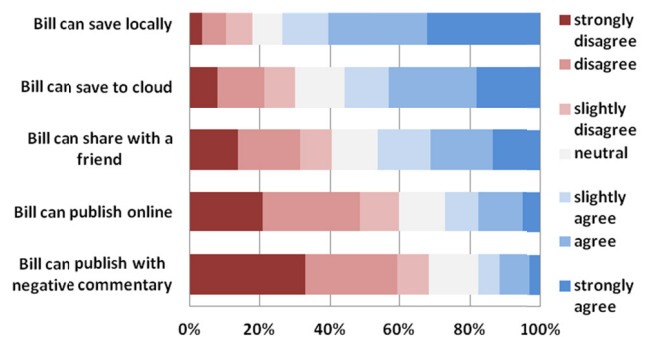


Figure 4. Responses to saving/sharing interview.

In the first set of hypotheticals, Bill records his own job interview with a software company. Saving digital media is normally uncontroversial, but since Bill recorded a conversation, we wondered if saving the recording might raise additional issues. Sure enough, while most of the participants agree that Bill can record the interview for his own use, there are more negative and neutral responses than for any other of our comparable hypotheticals about saving content. We also used this recorded interview to test the difference between saving content locally and storing it in the cloud. From there, we transitioned to a hypothetical in which Bill shares the recording with a friend, then distributes it more broadly online. A final hypothetical finds him overlaying a commentary track about the company's unfair hiring practices. Figure 4 shows the progression of responses to these hypotheticals.

The distinction between saving locally and storing the recording on the cloud provoked a significant drop in agreement ($p < .001$). Participants may have viewed cloud storage as less secure; they may have found cloud storage as a step closer to sharing; or they may have less understanding of cloud storage's implications.

Only 46% agreed that Bill could share the recorded interview with a friend who was also looking for a job; this agreement dropped to 27% when publication replaced sharing and to 18% when Bill overlaid a commentary track about the company's unfair hiring practices.

But before we assume that there are no widely acceptable uses of the recorded videoconference, we posited two modified versions of the interview video. In the first, Bill remixes the video by extracting his answers to the interview questions so neither the company nor the interviewer’s questions are included. He uses these answers to create three videos: one shows off his knowledge in hopes of getting a job, another adds commentary about “Dos and Don’ts” for job interviews, and a third is a self-deprecating comedy video in which he replaces the interviewer’s questions with recordings of his brother asking the questions. More than 65% of the participants agreed that Bill could publish or share these remixes. Figure 5 shows that the three most positive responses were elicited by remixes that showed some concern for public good, either by providing lessons or humor for others.




Figure 5. Responses to remixing interview.

In the final hypothetical of the series, Bill publishes his answers to attract another employer; the details are altered to test the effects of anonymization. In this variation, Kyle watches Bill’s “How to” video before his own interview. He records the interview and, like Bill, edits out any traces of the company from the recording. He then publishes his new video to “pay it forward,” but adds tags that identify the company and specify that it is a job interview. Participants have a significant negative response to the inclusion of the identifying tags. This poses an issue for social media since tags can be added after the content is posted – if someone deduced and tagged the company portrayed in Bill’s recording, the same situation would arise.




Figure 6. Responses to permission for interview reuse.

Figure 6 shows a last remix episode that explores whether Kyle can use Bill’s answers in his new video to create a more valuable “how to interview” video. 73% of the participants believe Kyle must get Bill’s permission before he uses the clips. 32% thought Kyle could include Bill’s answers as they were already available

online, while 43% thought acknowledgement was sufficient. While not all responses are logically consistent, they do provide a relative assessment of participants’ beliefs about permission, acknowledgement, and reuse.

6.3 Educational Recordings

Many public and institutional lectures are recorded and published online. These recordings may be of notable speeches or regular class presentations. Because the institutions that collect this content often do not have the infrastructure to provide the content online, they enter into agreements with online services such as iTunes or YouTube. This questionnaire explored the rights of different stakeholders in this situation, including presenters, their peers, institutions, service providers, and viewers.

In the study’s primary scenario, an astronaut (Sally) delivers a commencement address on the importance of manned space exploration. A video of the address is captured by the university and is made available through a web service; the video attracts numerous comments, ratings, and tags. A geologist (Cheryl) believes NASA’s resources would be better spent on unmanned missions and wants to use the video to provoke discussion.




Figure 7. Responses to saving educational video.

We start by exploring different parties’ rights to save the video and the associated metadata (comments, tags, etc.). In contrast to hypotheticals in prior studies, saving in this case explicitly involves transcoding the video from the web service’s proprietary format to a more common format. Perhaps because of this we see some controversy over saving encountered content. There is general agreement that Sally can save and transcode the video of her own presentation. There is somewhat less agreement that she can save the metadata that has been added by other users of the web service ($p < .03$). In contrast, there is substantial disagreement about whether Cheryl, Sally’s scientific peer, has the same rights. Interestingly, Cheryl’s right to save the metadata is seen as greater than her right to save the original video ($p < .04$). See Figure 7.




Figure 8. Responses to sharing/republishing educational video.

Beyond saving the video, who can share or republish it? Figure 8 shows that most participants agree that Sally should be able to share the transcoded video with friends via Facebook and to put it on her public website. There was no statistical difference between the responses to these two questions.

Given that Cheryl is seen as having less right to save the video, it is unsurprising that she is also seen as having less rights to share and publish it ($p < .0001$). Cheryl's right to share the saved video with friends via Facebook is perhaps the most divisive variation; the results are close to symmetric around neutral. The majority (55%) of participants feel that Cheryl has a diminished right to republish the video (difference from sharing, $p < .0002$).

One important characteristic of social media is the community involvement in commenting on, rating, and tagging content. Who owns this secondary content? What rights do the authors of the original content or the added content have? How about the host web service? We explored these issues through a series of hypotheticals about the reuse of reviews and permission for this reuse. Results are shown in Figure 9.




Figure 9. Responses to sharing/publishing reviews of video.

84% of the participants agreed that Sally has a right to share selected reviews of her commencement address via Facebook; 76% believe she should be able to republish selected tags and reviews from the web service when she posts the video to her own web site. The difference between the reactions to sharing and republishing is significant ($p < .001$).

As we mentioned earlier, Cheryl's rights to the reviews are viewed more positively than her rights to the video itself. 59% agree that Cheryl can share selected reviews from the web service with her friends via Facebook and 50% agree she can republish selected reviews on her own website to help make her point about unmanned space exploration. The difference between her rights and Sally's rights is significant ($p < .0001$) as is the difference between sharing and publishing ($p < .001$).

In our earlier studies, authors' permission often arose when participants were asked open-ended questions about appropriate practices for content reuse. The scenario in this study involves a number of stakeholders that might have some say over content reuse. We first asked whether Sally should have to get permission to reuse the video from either the university that recorded the lecture or the web service hosting the video. Figure 10 shows that 49% thought Sally should get the university's permission while only 32% thought it was necessary to seek the service provider's permission. The difference is significant ($p < .001$).




Figure 10. Responses to the concept of permission for educational video reuse.

Because Cheryl has no special claim to the commencement video (she neither delivered the talk nor recorded the video), we explored whether she would need to get permission from Sally, the university, or the web video service provider. Again, the web service provider's permission is the considered the least necessary (difference is significant, $p < .001$). This is in direct contrast to many interpretations of today's terms-of-use agreements and intellectual property laws. Participants instead thought Cheryl should have to get the university's and Sally's permission (the responses were statistically indistinguishable).

We also asked who needed to give Cheryl permission to republish reviews of the video found on the web service. Nearly half of the participants thought Cheryl should get permission from the review's author and the service. (See Figure 11.) The difference in responses to these two questions were not significant. Only 38% thought Cheryl should get Sally's permission to republish reviews. This is significantly different than the responses to the other two questions ($p < .01$).




Figure 11. Responses to permission for video review reuse.

While slightly over half of the participants disagreed with the need to get Sally's permission, it is interesting that her claim of authorship (she was the speaker in the video) would extend to the video's comments. Do people view the primary content creator as a proxy for the community's authority? We cannot say from this questionnaire alone, but when we combine these results with those from a similar questionnaire about product reviews [30], it is clear that rights are not limited to the content creator. People apply a more nuanced sense of ownership in such situations.

7. DESIGN IMPLICATIONS

By analyzing data gathered from 634 participants, we are able to characterize emerging social norms and current practices associated with video storage, reuse, and remix. What comes out

most strongly is that participants' attitudes and actions are constrained neither by the specific legal provisions that guide reuse, nor the terms and conditions that govern the content on most social media sites, but rather by a nuanced ethos that they have developed through experience. The question, then, is can we design technology and policy that facilitates the constructive aspects of current practice, while satisfying the more restrictive aspirational norms that participants expressed? We observed a very real (and understandable) tension between the two.

First, let's look back at the highlights of what the participants said and did. Through their account of their own experiences, we see that over 70% of the participants report sharing videos (we suspect underreporting, since it is so easy and unremarkable to share a video using a link) and it is likely that even more watch videos; many mention this specifically in an open-ended question about their online activities. The podcast study shows that many participants store downloaded material with an intention of keeping it indefinitely as a personal archive. The educational video study further demonstrates that the recordings represent a range of genres, and span professionally produced material as well as amateur videos. The videoconferencing study shows that many participants have also produced their own videos, including creative efforts, personal footage, and other types of documentary recordings. Finally, also from the videoconferencing study, we learned that participants reason about reuse from different ethical perspectives, including the nature of the content, the conditions of reuse, the details of production, and the technological and legal circumstances that they are aware of. Thus participants are judging the hypotheticals we pose from a vantage point of personal familiarity and are applying their experience.

Next let's look at the hypotheticals. From our previous studies, we believe that downloading primary content and storing it locally is essentially uncontroversial. We tested and confirmed this for online video. But we also introduced three provocative variations (1) the use of cloud storage instead of local storage; (2) the storage of transcoded content instead of duplicate content; and (3) the addition of secondary content to primary content (reviews of videos v. the videos themselves). Cloud storage and transcoding were both situations that blurred boundaries: cloud storage takes the content out of the realm of local control (a service provider is necessarily involved) and transcoding means potentially making changes to the content even if those changes are not directly observable. Secondary content introduces a notion of community ownership, or at the very least, ownership dependencies. Indeed, both cloud storage and transcoding seemed to reduce participants' comfort with saving video. Secondary content proved to be paradoxical; ownership of primary content does not extend to secondary content, but there is a more relaxed attitude toward secondary content in general; it may simply be regarded as less valuable.

Other hypotheticals brought in republishing and remixing. What did we learn there? As we have seen before, ownership comes into play more strongly when a video is shared or republished. In the least controversial case, an owner can fairly freely republish video content, even if there are co-owners. Commercial purposes add complexity, as does remix (especially if it violates abstract notions of fairness by, say, omitting negative reviews, or crossing the boundary into art theft [23]). This echoes some of our earlier findings: participants are sensitive to reuse situations that seem *per se* unfair—a recording that compromises hiring fairness, a satirical video that seems unduly mean-spirited, or not asking permission when it may be advisable.

What are the design and policy implications of what we have seen? The results clearly cannot be taken too literally: participants sometimes favor a permission-based system for reuse when fair use cases establish that none is necessary [26] (not to mention that experience suggests permission will be unobtainable [20]). They also signal certain problems with existing labeling schemes such as Creative Commons [3], because in the abstract, participants seem more sensitive to reuse than they are when confronted with particular instances and reciprocal examples (i.e. when they are the reusers). Nor is all reuse equal in the eyes of content creators and content reusers; the acceptability of reuse depends on circumstantial factors like the nature of the content (e.g., is it personal?), the differential scope of the audience (e.g., is the reused video in the process of going viral or is it now playing to an audience of 10? How different is this scope from the original?), the type of reuse (e.g., has it been included in a clip show? Will the video's original intent be distorted?), and the way the implied (or explicit) social contract between all potential owners of both the original and derived work is handled (e.g., is attribution or anonymity desired?).

Notice that only one of these factors is known at publication time (the nature of the content). Others are contingent on how the content is reused (e.g., changes in genre, audience, or publication venue). Still others are not revealed until time has passed (e.g. the differential scope of the audience). That these factors are crucial to how any labeling scheme is used makes us think that supplemental mechanisms might be desirable (hypothetical scenarios and mixed-initiative dialogs to help content creators envision reuse or decide between attribution or anonymity, or triggers that reveal when the differential scope or audience has changed). Still others depend on, say, the motivations for storing the video (past work tells us that individuals archive work that is not their own just as surely as institutions do [7][19][22]).

Thus ownership-driven questions need to be approached thoughtfully, lest we impose restrictions when none are necessary, or we do not anticipate types of reuse that will trigger the most extreme reactions when these reactions could have been averted. Our future work will continue to explore these questions on a per-genre, per-media basis, using a variety of qualitative and quantitative methods to unpack assumptions, identify social norms, and look forward to an unfolding spectrum of reuse situations.

8. ACKNOWLEDGEMENTS

This work was supported in part by National Science Foundation grants IIS-1049217 and DUE-0938074.

9. REFERENCES

- [1] Acquisti, A. and Grossklags, J. Privacy Attitudes and Privacy Behavior, in J. Camp and S. Lewis (Eds.) *The Economics of Information Security*, Kluwer, Boston, pp. 165-178.
- [2] Aufderheide, P., Jaszi, P., Boyles, J.L., and Bieze, K. Copyright, Free Speech, and the Public's Right to Know: How Journalists Think about Fair Use, (July 30, 2012). Retrieved from <http://dx.doi.org/10.2139/ssrn.2119933>
- [3] Boyle, J. *The Public Domain: Enclosing the Commons of the Mind*, Yale University Press, New Haven & London, 2008.
- [4] Dale, M., Stern, A., Deckert, M. and Sack W.. 2009. Metavid.org: a social website and open archive of congressional video. In *Proc. DGR '09*, 309-310.

- [5] Diakopoulos, N., Luther, K., Medynskiy, Y. and Essa, I. The evolution of authorship in a remix society. In *Proc. HT '07*. ACM, New York, NY, USA, 133-136.
- [6] Downs, J., Holbrook, M., Sheng, S., and Cranor, L. Are your participants gaming the system?: Screening Mechanical Turk workers. *Proc. of CHI'10*. ACM, 2399-2402.
- [7] Greengard, S. Digitally Possessed. *Communications of the ACM*, 55 (5), 2012, 14-16.
- [8] Hill, B., Monroy-Hernandez, A., and Olson, K. Responses to Remixing on a Social Media Website. *Proc. AAAI Conference on Weblogs and Social Media*. 2010. 74-81.
- [9] Hill, B. and Monroy-Hernandez, A. The Cost of Collaboration for Code and Art: Evidence from a Remixing Community. *Proc. CSCW 2013*. ACM. 1035-1046.
- [10] Ipeirotis, P. Analyzing the Amazon Mechanical Turk Marketplace. *ACM XRDS* 17, 2, Winter 2010.
- [11] Jakobsson, M. Experimenting on Mechanical Turk: 5 How Tos. *ITWorld*, September 3, 2009.
- [12] Kittur, A., Chi, E., and Suh, B. Crowdsourcing User Studies with Mechanical Turk. *Proc. of CHI'08*. ACM, 453-456.
- [13] Lessig, L. *Code, Version 2.0*, Basic Books, 2006.
- [14] Lessig, L. *Remix: Making Art and Commerce Thrive in the Hybrid Economy*, Penguin, New York, 2008.
- [15] Levi, E. *An Introduction to Legal Reasoning*. University of Chicago Press, 1949.
- [16] MacCormick D. and Summers, R. (eds.) *Interpreting Precedents*, Ashgate/Dartmouth, 1997, pp. 528-9.
- [17] Marshall, C. and Bly, S. Saving and Using Encountered Information: Implications for Electronic Periodicals. *Proc. CHI'05*, ACM, pp. 111-120.
- [18] Marshall, C. and Shipman, F. "Social Media Ownership: Using Twitter as a Window onto Current Attitudes and Beliefs", *Proc. of CHI 2011*. ACM. 1081-1090.
- [19] Marshall, C. and Shipman, F. "The Ownership and Reuse of Visual Media", *Proc. of JCDL 2011*. ACM. 157-166.
- [20] McDonough, J., Olenorf, R., Kirschenbaum, M., Kraus, K., Reside, D., Donahue, R., Phelps, A., Egert, C., Lowood, H., Rojo, S. *Preserving Virtual Worlds Final Report*. 8/ 31/2010.
- [21] Monroy-Hernandez, A., Hill, B.M., Gonzalez-Rivero, J. and boyd, d. Computers can't give credit: how automatic attribution falls short in an online remixing community. *Proc CHI '11*. ACM, New York, NY, USA, 3421-3430.
- [22] Odom, W. Sellen, A., Harper, R., and Thereska, E. Lost in Translation: Understanding the Possession of Digital Things in the Cloud. *Proc. CHI'12*. ACM. 781-790.
- [23] Perkel, D. The Art of Theft: Creativity and Property on deviantART. *Material World Blog*. 2010. Retrieved from http://blogs.nyu.edu/projects/materialworld/2010/07/the_art_of_theft_creativity_an.html.
- [24] Rissland, E.L. and Ashley, K. Hypotheticals as Heuristic Device. *Proc. Strategic Computing Natural Language*, Marina del Rey, California, May 1-2, 1986, 165-178.
- [25] Rotman, D., and Preece, J. The "WeTube"; in YouTube – creating an online community through video sharing. *Int. J. Web Based Communities* 6, 3 (June 2010), 317-333.
- [26] Sag, M. Predicting Fair Use. *Ohio State Law Journal*, 73 (1), 2012, 47-91.
- [27] Samuelson, P. Too Many Copyrights, *CACM* 54, 7 (2011).
- [28] San Pedro, J., Siersdorfer, S., and Sanderson, M. Content redundancy in YouTube and its application to video tagging. *ACM TOIS*. 29, 3 (July 2011), Article 13.
- [29] Schmitz, P., Shafton, P., Shaw, R., Tripodi, S., Williams, B., and Yang, J. International remix: video editing for the web. *Proc. MM '06*, 797-798.
- [30] Shipman, F. and Marshall, C. Are User-contributed Reviews Community Property? Exploring the Beliefs and Practices of Reviewers. *WebSci 2013*. ACM. forthcoming.
- [31] Solum, L.B. *Legal Theory Blog*, entry dated July 11, 2010. <http://lsolum.typepad.com/legaltheory/2010/07/legal-theory-lexicon-hypotheticals.html>.
- [32] Stutzman, F. Twitter and the Library of Congress. <http://fstutzman.com/2010/04/14/twitter-and-the-library-of-congress/>.
- [33] Vihavainen, S., Mate, S., Seppälä, L., Cricri, F., & Curcio, I. We want more: human-computer collaboration in mobile social video remixing of music concerts. *CHI '11*, 287-296.
- [34] Vihavainen, S., Mate, S., Liikkanen, L. and Curcio, I. Video as memorabilia: user needs for collaborative automatic mobile video production. *Proc CHI '12*. 651-654.
- [35] Xie, L., Natsev, A., Kender, J.R., Hill, M., and Smith, J.R. Visual memes in social media: tracking real-world news in YouTube videos. *Proc. ACM MM '11*, 53-62.