7

Privacy and Audiovisual Content: Protecting Users as Big Multimedia Data Grows Bigger

Martha Larson, Jaeyoung Choi, Manel Slokom, Zekeriya Erkin, Gerald Friedland and Arjen P. de Vries

7.1 Introduction

In this chapter we discuss the relationship between privacy and algorithms that make use of large amounts of multimedia data. As users continue to post their audiovisual content online, and as companies continue to collect user profiles and interaction data, concerns about privacy are becoming increasingly urgent. With much of their behavior, users unwittingly share information about themselves that could compromise their privacy. As big multimedia data continues to grow bigger, it is essential to understand not only the risks for users, but also the potential of using multimedia technology to protect users. In this chapter we look at online multimedia sharing and discuss a selection of multimedia algorithms that can help to protect user privacy.

The success of privacy protection algorithms depends on their connection to the human and social aspects of privacy. Effective privacy solutions are situated in the intersection of:

1) education: educating users on the danger of privacy and providing them with the information they need to make informed privacy decisions
2) legal: developing legal frameworks that protect users
3) technical: creating technologies that are able to protect data from unauthorized access or misuse.

This chapter focuses on multimedia algorithms, but looks beyond a purely technical approach to privacy. We connect to the human and social aspects of privacy by focusing on algorithms that put control in the hands of users, i.e., of the people who produced the data in the first place and whose privacy needs to be protected. Solutions that give users control help to free them from dependence on external parties such as service providers. Our focus on user control should not be interpreted as detracting from the responsibility of external parties to protect user privacy. To the contrary, parties that collect, hold, and process user data have an enormous responsibility to protect privacy. It is important to understand that we focus on user control because this perspective on privacy is conventionally underrepresented in previous treatments of privacy and big data from a technical perspective.

Big Data Analytics for Large-Scale Multimedia Search, First Edition.
Edited by Stefanos Vrochidis, Benoit Huet, Edward Y. Chang, and Ioannis Kompatsiaris.

Companion website: www.wiley.com/go/vrochidis/bigdata

The chapter is organized as follows. In the remainder of section 7.1 we motivate and define privacy. Next, in section 7.2, we turn to the discussion of what must be done to protect users' privacy. Section 7.3 discusses the particular privacy challenges raised by multimedia, and specifically by big multimedia data. Section 7.4 presents example techniques and algorithms, and finally section 7.6 provides an outlook for the next steps for multimedia privacy research.

7.1.1 The Dark Side of Big Multimedia Data

Anyone who has learned a new skill by watching a video on YouTube, delighted in exploring a faraway place by browsing collections on Flickr, or reconnected with a long-lost friend by sharing photos on Facebook knows how the rise of online multimedia has added functionality and joy to our daily lives. However, along with the appeal and benefits of sharing videos and photos, sharing multimedia is accompanied by risks. An important risk is *cybercasing*, the use of information available online to mount real-world attacks. This term was introduced in 2010 in [1] in order to help raise awareness of rapidly emerging privacy threats of big multimedia data online. Here, we present cybercasing as a motivating example in order to illustrate the importance of privacy.

The discussion of cybercasing in [2] focused on the inference of users' geo-location by people with criminal intent, e.g., people aiming to rob a house when the owner is away. The danger represented by unwittingly revealed geo-location information can be extended to other sensitive information such as family status and health state, which can be inferred by automatic methods using information shared or otherwise available online. The findings of [1] have serious implications. While users typically realize that sharing locations compromises their privacy, they (i) are unaware of the full scope of the threat they face when doing so and (ii) often do not even realize when they publish such information. The threat is elevated by developments that make systematic search for specific geo-located data and inference from multiple sources easier than ever before. The authors of [1] base these insights on a summary of the state of geo-tagging as of 2010, an estimate of the amount of geo-information available on several major sites, including YouTube, Twitter, and Craigslist, and an examination of its programmatic accessibility through public APIs. The magnitude of the threat of cybercasing is illustrated with set of scenarios demonstrating how easy it is to correlate geo-tagged data with corresponding publicly available information for compromising a victim's privacy. The investigations in [1] were able to find, for example, private addresses of celebrities as well as the origins of otherwise anonymized Craigslist postings. Additionally, further work [3, 4] has shown that it is possible to predict users' home, work and vacation locations using geo-tagged posts on Flickr, Twitter and YouTube.

Protecting users' privacy defends users against cybercasing and against other unintended ill-effects of information misuse. Now that we have highlighted the importance of protecting user privacy with the example of cybercasing, we turn to examine in more detail how privacy is defined.

7.1.2 Defining Multimedia Privacy

Opening a dictionary provides a definition of privacy common in every day usage. Merriam Webster,[1] for example, lists the following two definitions of privacy:

[1] https://www.merriam-webster.com/dictionary/privacy.

A: *the quality or state of being apart from company or observation: seclusion*
B: *freedom from unauthorized intrusion one's right to privacy*

The dictionary definition provides a natural connection to our everyday experience of the importance of privacy. As human beings, we need sleep as a physical necessity. In addition to sleep, we also need seclusion. The need for seclusion is reflected in the fact that we seek moments away from the broader circle of other people around us. As humans, we seclude ourselves to relieve ourselves or while engaging in intimacy. Seclusion provides us with time away from observers.

A possible explanation for the physiological and psychological importance of seclusion to humans is the following. Independently of whether observers pose a threat or not, observers require us to maintain a readiness state. Even if the readiness state involves only low-level readiness to comply with social conventions, for example to give a polite response to a greeting, perpetual readiness can wear us down. This view is echoed by the book *Understanding Privacy*, which states, "Privacy protections are responses to problems caused by friction in society" (p. 76). If seclusion is respite from the burden of social interaction, the parallel between sleep and seclusion can be extended further. Like sleep, privacy by way of seclusion is needed for the body and mind to be able to restore themselves. However, the parallel also hints at reasons why privacy is difficult to understand. Like sleep, seclusion is a state that is desirable in the right quantity, but becomes undesirable, or even dangerous, in excessive quantities. Further, each individual has different needs for both sleep and seclusion, and sometimes these needs are only apparent with long-term deprivation. Finally, like sleep, seclusion is a phenomenon that is currently not yet well understood, despite its obvious critical importance for human health and well being.

We are not experts on either sleep or seclusion, and make no claims on the underlying nature of either here. Rather the point that we make is that privacy is an inherent human need, and the fact that we do not understand it completely should not discourage us, but rather provide further motivation to study it. In order to gain insight on how to study privacy related to big multimedia data, it is necessary to go beyond the dictionary definition of privacy and look at more formal characterizations.

The authors of the book *Privacy and Big Data* [5] identify three basic types of privacy, which we use as the basis for the following definitions.

- Physical privacy: the state of being free from intrusion into your personal space, including your possessions and your own body.
- Information privacy: the state of control over information about you that is collected, stored, processed, or shared.
- Organization privacy: the state of secrecy used by companies and governments to hide their activities from competitors and enemies.

When privacy is discussed in the context of big data, the focus is on *information privacy*. For example, the distinction between information privacy and other sorts of privacy is made explicit by [6]. In the context of multimedia, information privacy can also be viewed as a sort of seclusion. Instead of being out of range of observers, information privacy requires withdrawing from the view of the camera. Unlike observers, however, the camera records, allowing a person to be observed and re-observed endlessly. The persistence of multimedia creates a new type of friction: a person needs to maintain a state of readiness to clear up misunderstandings that arise when recordings are viewed out of context or juxtaposed in a way that creates a misleading perspective. Information privacy must protect users from the myriad of mirrors that is online information. These mirrors confront the user with distorted, disorganized, and contradictory views of themselves. Without information privacy, nothing protects the user from the wear and tear of the effort of maintaining their identity and reputation in the face of these fragmented views.

Information available online can also lead to damage that is more focused than wear and tear. Specifically, information used in the wrong way by the wrong person can lead to dangers such as the cybercasing attacks discussed above. Safeguarding users' information privacy in cyberspace means protecting users from real-world harm.

Here we cite some other helpful definitions of information privacy as a starting point for further discussion of the complexity of privacy. The multimedia privacy tutorial at ACM Multimedia 2016 [7] mentioned [8], who define information privacy as "The claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others." Further, work in the area of privacy for recommender systems [9] follows [10] in defining information privacy as "an individual's claim to control the terms under which personal information—information identifiable to the individual—is acquired, disclosed, and used."

There are two aspects where these definitions fall short. First, they fail to emphasize that there are two types of information that users have claim to control: explicit and implicit information. Explicit information is consciously known by the user. For example, the user must have the ability to control images of his children. Implicit information is information that is not consciously known, but is derivable from multimedia content in some manner, i.e., by an intelligent system (trained on aggregated data) or by a human expert. For example, the user must have the ability to control information on the health state of his children. He himself might not see that an image reveals that his child has curvature of the spine. However, such a condition is noticeable to an expert. A definition of information privacy must encompass the claim to control both explicit and implicit information. Second, these definitions fail to consider the issue of inaccurate information. The danger of inaccurate information is acute with intelligent systems. An intelligent system might analyze the people who a user associates with, as depicted in the photos he posts online, determine that they have committed traffic violations, and conclude that the user represents a risk and should not be offered a job requiring driving. There is no external reference by which to judge if this decision is accurate or fair to the user. The aggregation might have included an error, and it might also be wrong to assume that someone is a bad driver because their friends are. A definition of information privacy must encompass the claim to the control of purported information about the user, independently of whether this information is correct.

In order to address these two aspects, a broader definition of information privacy is needed, such as the one provided in [11], which defines information privacy as "Practically securing the implications of communication". In other words, we can recognize that privacy has been protected when we have protected users from unintended consequences that result from sharing or interacting with multimedia online.

The fact that this definition uses the word "securing" raises the question of the relationship between privacy and security. Wikipedia defines information security as "...the practice of preventing unauthorized access, use, disclosure, disruption, modification, inspection, recording or destruction of information" [12]. From this definition we can see that the focus of information security is on protecting information. Although protecting privacy is also about protecting information, it does so with a different focus. Protecting privacy is about protecting people, not only information. The authors of [11] point out explicitly that the goal of their privacy research is not providing a secure communication line between communicating parties, but rather ensuring that information that is available to the public conveys only what the person producing that information intended it convey. The focus on people also differentiates privacy from confidentiality. As formulated by [9], confidentiality is related to keeping specific pieces of information secret, whereas information privacy focuses on protecting the individual who the information is about.

Achieving the goal of successfully protecting people requires a collaboration between the people-focused view of privacy and the information-focused view of security. Here, we list several aspects of the people-focused view of privacy that are easy to overlook if too much emphasis is placed on the information focused-view.

- First, protecting a person's information privacy has two facets: it must prevent unintended implications of information sharing by that person or about that person, and it must ensure that the person has the perception that his privacy is protected. For multimedia privacy, the topic of user perceptions of privacy is addressed by [13].
- Second, privacy has a subjective component. For any given piece of information, one user may be comfortable with having that piece of information public and another may not. In the case of multimedia privacy, it has been observed that users differ in the kinds of photos that they would like to keep private [14]. Users' individual judgments may or may not be consistent with more objective assessments of risks. Subjective judgments have a different status depending on whether users are able to understand and reason about risks.
- Third, privacy is contextual. Under one set of circumstances sharing of certain information may be seen as acceptable, but under another it might constitute a privacy violation.
- Fourth, privacy protection must be robust in the face of information accumulation and improvement of information inference technologies (i.e., intelligent systems). A user's privacy may be protected for today, but that protection must extend into arbitrary futures. In the case of cybercasing, this means that information that a user shares today must be prevented from contributing to future prediction of his geo-location by malicious parties.
- Finally, as already mentioned above, information privacy must be protected whether or not the user himself is aware of the information to be protected.

We also mention another important people-focused concept in privacy, namely, *intentional privacy*. Here, we follow the definition used in [15]: "intentional privacy – the right of the individual to prevent or forbid further communication of observed events of exposed features (e.g., publishing photos or video footage)". This definition can be subsumed under the "implications of communication" mentioned by the definition of informational privacy above. However, discussing it separately allows us to correct a frequent misunderstanding. Often privacy is approached as a "barn door" problem. In other words, people assume that once a user has shared or given permission for the use of information "anything goes" and there is nothing to be done, just as it is senseless to close the barn door after the cow has already escaped the barn. However, "intentional privacy" clearly spells out that it is the intention of the user that is important. The user has no obligation to state intent explicitly ahead of time. Instead, as pointed out by [13], the implicit assumptions of users concerning how their data will be used must be respected. The importance of the scope of sharing is underlined by [9], who state, “Privacy involves keeping a piece of information in its intended scope” (p. 269). This scope comprises the user’s intended audience, the types of usage that the user intended to allow, and the intended lifetime of the information.

This section has introduced a definition of privacy and discussed its intricacies. We have seen that the successful protection of privacy requires protecting many things, in many ways, under many circumstances, and under changing conditions. It is inevitable that some researchers will view privacy as impossible or too complex. However, instead we wish to encourage the view that the complexities of privacy are exactly what makes privacy such an interesting research field, and that they should have the effect of inspiring new and creative solutions. Researchers who feel that there is “no cure” for privacy threats should reflect on the medical world. Medical researchers do not abandon research on a disease because it seems difficult or impossible to cure. Instead, they try to understand it and also to prevent its spread. As multimedia researchers, we need to understand privacy and prevent the spread of technologies that threaten it. In the remainder of the chapter we turn to discussing technical solutions for protecting users.

7.2 Protecting User Privacy

In this section we provide a more concrete description of what we need to protect when we are protecting user privacy.

7.2.1 What to Protect

What we protect about users breaks down into two categories, which we define following [15]:

1) Personal identifiable information: personal information that makes it possible to identify someone.
2) Personal information: any information relating to a person.

The distinction is also formulated as identity vs. attribute disclosure. Here, we provide definitions following the discussion of [16]. Disclosure takes place when available data is used to accumulate knowledge that can be potentially misused. Identity disclosure happens when a link is made between knowledge and a particular person. Attribute disclosure takes place when new knowledge is accumulated. We mention this distinction for completeness, and do not discuss it further here.

Instead, we take a closer look at personal identifiable information and personal information. Personal identifiable information is described as a set of identifiers. Specifically, [15] lists three types of identifiers:

1) biometric identifiers: permanent characteristics that differentiate a person from other people (including voice and gait)
2) soft biometric identifiers: differentiating characteristics that are less permanent (e.g., height and weight) and
3) non-biometric identifiers (including hairstyle).

A great deal of research has been carried out on eliminating identifiers from multimedia content, a field which is referred to as de-identification. The topic of multimedia de-identification is covered in depth in [15]. Key areas include the de-identification of faces, gaits, and body silhouettes, as well as gender, age, race, and ethnicity. Open challenges include the visual acceptability of the resulting multimedia content after the information has been removed, and also the ability of systems to operate in real time, as is necessary for surveillance systems. For de-identification algorithms it is important to remember the context of the judgment. For example, a person's face could be de-identified, but that person would still be identifiable in a photo to a friend who was present at the moment that the photo was taken. Further, [15] mentions the problem of "pairwise constraint identification", which means that if a person's face is de-identified in a photo, another photo taken in the same context can serve to identify the person due to other matching factors such as hairstyle and dress.

In the rest of the chapter we concentrate on personal information. As discussed in section 7.1.2, information falls into two categories: information of which the user is aware and information that the user is not aware of. We further highlight the difference between information that is inferable by inspection of a single piece of multimedia content, i.e., a person has visited a particular bar, vs. information that requires aggregation of information across multiple pieces of content, i.e., a person visits a particular bar every evening. Further, multimedia content can convey information by virtue of something that it does not show. For example, someone might be shown driving a car without their glasses on, which is illegal if their driver's license require them to wear glasses.

Personal information that must be protected is referred to as *sensitive information*. Although the exact nature of which information is sensitive has a certain dependence on context, certain information is always considered sensitive. From the legal perspective, the General Data Protection Regulation (GDPR) of the European Union lays out a set of rules to protect data. Here, personal data "... can be anything from a name, a photo, an email address, bank details, your posts on social networking websites, your medical information, or your computer's IP address."[2] In general, deciding what needs to be protected is part of the overall problem of developing privacy protection. Researchers must pro-actively seek to understand what must be protected, rather than expect to be handed a list of sensitive personal information.

7.2.2 How to Protect

We begin our discussion on how to protect users with a broad view of the connection between privacy and multimedia research. Recall that section 7.1 mentioned the importance of developing privacy solutions in the overlap of three areas: (i) education, (ii) legal, and (iii) technical. The broad view reminds us that not only do our solutions exist in this intersection, we ourselves also inhabit it. Concretely, this means that we should understand ourselves not only as researchers developing technology, but also as actors promoting education and upholding legal protection. We now discuss our education and legal roles briefly in turn.

Many researchers active in the area of multimedia research teach, supervise students, mentor junior colleagues, and lead labs. The students and early-career researchers that

[2] http://europa.eu/rapid/press-release_IP-12-46_en.htm?locale=en.

we work with today are the computer scientists who design the technologies of tomorrow. The success of future privacy protection technology starts with the education of those who will create it. Looking backwards, [17] points out that if privacy had been given serious thought when the Internet was first developed, today we would be facing less of a serious problem. Teaching students about the importance of privacy early will give them the motivation and the skills necessary to ensure that the technological developments that they contribute to incorporate privacy in their design, rather than as an afterthought. Further, we point out that the contribution of computer scientists with technical knowledge of data science, multimedia, and computer networks are needed to further the development of curriculum materials in order to teach students about privacy. A key example is the teaching privacy initiative, described in [18].

Researchers also support the legal system in protecting users' privacy by conscientiously following the procedures of informed consent. When we collect data from users for the purpose of developing a new algorithm, it is important to ensure that users are informed and agree to their data being collected, and to the purpose for which we are using it. Sensitive user data should be stored safely, should not be shared beyond the group designated in the informed consent, and should be destroyed after use. Destruction of data has two purposes. First, it prevents data from accidentally falling into the wrong hands due to neglect or a data breach. Second, it prevents what is known as "function creep" (cf. [15]), the effect that slowly over time, due to forgetfulness or the pressure of producing research results, data are put to uses for which they were not originally intended.

It is important to keep in mind that protecting users' privacy might be the most time-consuming aspect of any particular research project. An exemplary paper is [19], which carries out a study of young people's nighttime habits by studying photos that they take while going out in the evening. The participants who contributed photos are recruited via a recruitment campaign, which was approved by the ethics advisory board of the city. The purpose of the study is explained to the participants, and they provide informed consent. The data collected from the study is used only by the researchers for the particular purpose of the specific research, and it is not shared beyond the research group. When using data collected online, multimedia researchers should keep in mind the concept of "intentional privacy". Users shared their data online with a specific intent, and had no way of anticipating the use that it would be put to by researchers. Researchers should not assume that it is acceptable to carry out any experiment that they can conceive of using data found online. Rather, they should consider the harm that could come to individuals by being singled out by multimedia algorithms that predict information that those individuals could have no way of imagining was possible.

Although purely legal aspects of privacy are out of the scope of this chapter, we do find that it is important for researchers to be aware that privacy holds the status of a fundamental human right. The Universal Declaration of Human Rights (UDHR),[3] proclaimed by the the United Nations General Assembly in 1948, includes privacy as a right to be protected along with other fundamental human rights. The right to privacy is encoded in the legal systems of different countries in different ways. However, in general, researchers should realize that their contributions to protecting users' privacy are part

[3] http://www.un.org/en/universal-declaration-human-rights/.

of a larger picture. By protecting users' rights to privacy, they strengthen the protection of their own rights. Next, we turn to technical solutions for protecting users' privacy.

7.2.3 Threat Models

In order to develop technical solutions, we must first be able to formulate privacy problems. Upon first consideration, many researchers feel that protecing privacy is "mission impossible". The following two questions illustrate the tension between researchers, perceptions of what types of problems can be considered "solvable" and what is necessary in order to protect privacy.

How is it possible to protect users who do not seem willing to stop sharing multimedia online and publicly? Research in privacy requires not laying the blame at the door of the user, but rather providing tools and technologies to protect users as they engage in their natural behaviors and go about their activities. A frequently cited parallel to online sharing is transportation. Asking people to stay home is not a viable solution to protect them from automobile accidents. Likewise, asking people to stop sharing is not a viable solution to protect them from loss of privacy. Just as automobile safety took years to develop, and is also an ongoing target of much research attention (i.e., in the area of self-driving cars), researchers must direct concerted attention to privacy in order to develop creative and effective solutions.

How is it possible to demonstrate progress in protecting privacy, when the privacy problem is never really solved? Research in privacy requires realizing that privacy solutions are not absolute. Instead, privacy problems are like security problems: they can also be solved with respect to a definition of a model. A frequently cited parallel is protecting a house. A house can be protected with locks and bars on the windows and doors. These protections will prevent a person with a standard set of implements from entering. However, someone with a backhoe can easily make a hole in the wall and walk right in. The house is considered secure because of the very small probability that someone with a backhoe will want to enter. When we secure a house, we secure it with respect to a model that does not include the possibility of a backhoe.

These examples illustrate the importance of being able to formulate a concrete privacy problem in order to be able to move forward with developing solutions. The basis for the description of privacy protection problems is a *threat model*. A threat model is a detailed description of the situation in which the solution is intended to provide protection. The model characterizes the source of the threat (adversaries) and the opportunities for threat (vulnerabilities). A classic definition of a threat model from the security literature was presented by [20]. Here, we touch on the main points and highlight some respects in which a model for privacy must extend the standard threat model.

The threat model starts with a characterization of the *adversary*: the person or entity who poses the threat. The adversary has multiple dimensions that must be taken into account. First, the objective: what is the goal or the purpose of the adversary? Second, the degree to which the adversary has or can obtain access to critical data. Third, the resources at the adversary's disposal (e.g., algorithms and computational resources). Fourth, how willing is the adversary to invest those resources given the likelihood of success or failure. Next, the threat model characterizes vulnerabilities. Here, it is critical not to consider only the immediate system and/or data, but the surrounding ecosystem of data sources. Finally, the threat model describes countermeasures. The only countermeasures that must be considered are the ones that correspond to the objectives of the adversary, and the type of attack the adversary is willing and able to mount. In order to be viable, a countermeasure must fulfill a series of non-functional requirements, including being feasible to implement and easy to use. While outlining these steps in the threat model, [20] also mentions that there are three steps to any attack: (i) plan the type of attack, (ii) gain the necessary access, and (iii) execute the attack. Countermeasures can be effective at any of the three steps.

This threat model was designed for security applications. As previous mentioned, information security is oriented to protecting information and is not focused on protecting users. For this reason, the threat model does not directly transfer to the case of privacy. Here, we discuss some additional factors that must be taken into account to study privacy.

Part of securing a system is determining who to trust. Security conventionally involves one or more trusted parties. Privacy, on the other hand, must be protected in cases where no one can be trusted. The user sharing information might be irresponsible, the friends he shares it with might thoughtlessly share it on, the system that he uses to share might be attempting to target him with unwanted advertising or might be breached. In short, there may not be a single trusted entity in the entire setup, including the user himself. The threat model approach dictates that each of these possible issues be enumerated and their significance estimated. If a privacy solution does not address all of them, it will have still made an informed choice on what not to address, and its limitations will also be explicitly stated, helping to ensure that it is not inappropriately applied.

The threat model states explicitly that countermeasures must be viable, with ease-of-use being a key criterion. Security conventionally involves expert users who are already convinced of the need for security. Privacy protection measures must be easy to use for non-specialist users. Moreover, if the user is not convinced about the importance of protecting privacy, the measures must contribute to educating the user and motivating their use. If a privacy protection algorithm is created, but users do not adopt it, then it might as well not exist. The threat model approach dictates that countermeasures be assessed for the potential for user adoption before they are even developed.

In this section, we have discussed what we must protect when protecting privacy and how we should proceed. Specifically, we discussed the broad view of how multimedia researchers contribute to privacy and also the details of how to formulate a privacy problem. Next, we move on to discussing privacy and privacy threats specifically with respect to big multimedia data.

7.3 Multimedia Privacy

Protecting multimedia privacy involves securing the implications of users' multimedia sharing and consumption behavior. Users should be free to produce, share, and consume media without suffering or fearing unintended consequences.

7.3.1 Privacy and Multimedia Big Data

Tackling the challenges of multimedia privacy involves understanding why not just multimedia data, but, specifically, *big* multimedia data pose a threat to privacy. The authors of [5] clarify the nature of the debate on privacy by stating that it is "…a debate about the collection and use of our personal information from a commercial and political standpoint" (p. 2). They highlight two particular negative implications of big data for privacy. First, big data is not only big, but it is also drawn from multiple sources. By automatically matching bits of data from different sources, algorithms can infer information about people. This information was not available as long as the sources were scattered and not easy to access. Second, big data has the ability to build a complete picture of an individual. Privacy violations continue to also involve the revelation of individual sensitive facts. However, in the age of big data, it is now possible that a privacy violation reveals a complete picture of an individual's life and activities. In *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World* [21], it is argued that today's technologies enable nothing short of mass surveillance.

Other negative implications are related to the fact that standard procedures and checks become costly, if not impossible, to carry out. From a legal standpoint, [22] points out that standard procedures of gathering consent become unmanageable for big data. Consent is often reduced to whether or not the users from whom data is being gathered clicked "I agree" rather than whether they genuinely understood the purposes to which they are allowing their data to be put. This situation is different from a standard data collection context, for example during a medical trial, in which informed consent is gathered individually, and opportunities exist for clarifications and questions.

Further, as a data sets grows bigger and bigger, it becomes less and less feasible to check it for accuracy. In fact, most big data applications use data in the state in which it was collected, and make no attempt at quality control. Multimedia collections are particular difficult to check for quality. Researchers regularly work with collections of video so large that it would be humanly impossible to verify them by linear watching. Further, multimedia is often accompanied by metadata. Even if the media content can be verified, the association with the metadata might be wrong (i.e., a YouTube video that has been uploaded with the wrong title).

Big data applications prefer to use data sets that contain every existing, available data point, the so-called $N = all$ of big data. If every single data point is used, no data sources are left in which the reliability of the data can be checked. It becomes very easy to treat the data set as "reality", since nothing is known of reality beyond the data set.

Independently of the verification of big data, the question of verification of the *output* of big data algorithms arises. The discussion in [6] highlights two important privacy concerns: first, the reliability of correlations discovered in big data, and second, the trustworthiness of the interpretation of inferences.

Finally, it is important to realize that the commercial value of big data gives rise to forces that make it increasingly difficult to protect privacy. Collection and exploitation of user data generates profits. With profit comes the ability to buy the infrastructure needed for large-scale data processing, which in turn generates more profit. Protecting users' privacy does not generate profit. However, research developing privacy-protection technology also requires expensive infrastructure. The picture is different from the "small data" era in which data fit on a standard desk-top hard drive. Unless business models can also be linked to privacy research, the playing field will remain uneven.

In the field of multimedia, this means that entities that can afford infrastructure will continue to process video, images, and audio for profit, with little regard for protecting those people who produced them.

On top of the force created by profit, another force that makes it difficult to protect privacy is industry lock-in: the fact that big data needs big infrastructure locks the advantages in to players that have the resources; those players can, in turn, offer services that lock in users. They have little to no incentive to develop new business models that would embrace privacy since they face no serious competition. The result is a reinforcement of the dominant idea that a commercial company can collect and store users' data indefinitely. There is no burden of proof that the value provided to users in turn is near the value that users would claim, were they fully aware of the privacy risks. The effects of lock-in also make it difficult for researchers interested in studying privacy to be able to access the data that is necessary to develop and test algorithms effective at large scale.

The picture that emerges is that big multimedia data poses unique and pressing challenges for multimedia privacy. At some time in the past, multimedia privacy could mean adding a black bar over the eyes of a person depicted in a picture published in a magazine. Clearly, it is not enough to simply scale-up techniques for privacy protection effective with small-scale data. Instead, new technologies must be developed that make possible the protection of users from new threats that arise from big multimedia data.

7.3.2 Privacy Threats of Multimedia Data

We have seen that the information contained in big data can either be explicit (users are aware of the information) or implicit (users are unaware of the information). The two categories of information hold for general data, as well as multimedia data, in particular. Here, we turn to the question of the particular potential of multimedia data to pose privacy threats from which users must be protected. Previously, multimedia data has been described as having two privacy levels [13]. The primary privacy level encompasses the ostensive information in the multimedia content, and especially the information that is intended to be communicated. The primary privacy level of a lecture video, for example, is the topic of the lecture. The secondary privacy level encompasses the social/psychological characteristics of the data, which are the necessary, but not necessarily intended, by-product of its production. Following the discussion in [13], we list the following examples:

1) textual cues: style, vocabulary size, mastery of grammar
2) verbal cues: tone, accent, enunciation, projection of confidence
3) visual cues: manner, gestures, dress, projection of confidence.

These examples support the argument that multimedia data is more privacy sensitive than other forms of data (documents) since the secondary privacy level can more easily escape the notice of the user. Next we turn to discuss a set of examples of current technologies that represent privacy threats in various sorts of multimedia data.

7.3.2.1 Audio Data

Current open-source technology can be easily repurposed to deanonymize consumer-produced videos just based on their audio tracks. This ability was demonstrated in 2011 by [23]. The authors describe an approach that utilizes audio to link the uploaders of videos based on the audio track of the videos. Using a subset of the MediaEval 2010 Placing Task's[4] Flickr video set, which is labeled with the uploader's name, they conducted an experiment with a similar setup to a speaker identification experiment. Based on the assumption that the audio might be matched in various ways (speaker, channel, environmental noise, etc.), they trained the open-source speaker ID software on the audio tracks of the Flickr videos. Note that since the selection of videos was essentially random, the audio track can contain any types of sound. They obtain an equal error rate of 36.7% on 312 videos with 11,550 trials. The result is interesting for several reasons. It first shows that even highly tuned systems, like current speaker ID systems, are generic enough to be "abused" for a different task. Second, it shows that random Internet data is not nearly as random as one might think. A speaker ID system can be used to link independent personas. In other words, it is not safe to use different user names to keep sets of videos distinct.

7.3.2.2 Visual Data

With the rapid advancement of visual analysis technology, it has become possible to extract much useful information from images and videos such as objects (what is there) and events (what is happening). Face recognition technology has matured to the point that it is included in commercial software. Multimedia data is often shared online with location information. However, recently people have become more conscious and careful about the level of information released in each image. More people are consciously switching on and off geo-tag labeling when posting to social media and do not turn on the GPS geo-tag function when they are near everyday places such as home or work [4]. However, turning off the GPS will not necessarily solve the problem. What most people are not aware of is that when a group or stream of images are collected even without geo-tag, attackers can still learn about the target much more than expected. Visual geo-location estimation can automatically predict the location at which an image was taken, and recent years have seen it becoming more mature [24]. For instance, a search-based algorithm [25] can geo-locate 7.5% of the MediaEval Placing Task 2016 dataset within 1 km from the ground truth location. With the sheer amount of data shared online that can be easily crawled every minute, this level of accuracy is already good enough to find real-world targets among them. This information reveals the location of the photographer at the time at which the photo was taken, and also, indirectly reveals where the photographer was *not* at that moment. With a stream of geo-data, which gives a history of a user's location, much more can be revealed: habits, interests, hang outs, where the user works, and where the user lives.

7.3.2.3 Multimodal Threats

Although less work has been carried out on multimodal information inference, research in that direction is developing rapidly. In the future, we can also expect privacy threats to be elevated by the combined use of indicators from different media. The potential *and* danger of multimodal location estimation was pointed out by [2]. The ability of multimodal inference to compromise privacy will be magnified many times as machine learning approaches continue to mature. Imagine a future where multimedia query

[4] http://www.multimediaeval.org/mediaeval2010/placing.

engines *just work*. It would be possible to search by topic, location, person, camera identity, and time, even if the uploader did not explicitly include such information. If someone would like to steal an expensive piece of sound equipment, he could query for photos depicting that piece of equipment. Then, he could query for other photos that were taken by the same person. Finally, he could query for the geo-location of these photos. If these photos reveal a pattern, a time can be pinpointed at which the owner is unlikely to be home. As described earlier, current technology already makes cybercasing attacks possible. The point that we would like to make here is that with a sophisticated multimedia query engine, simple methods of anonymizing posts and suppressing metadata will no longer be enough to protect privacy.

For all types of threats, audio, video, and multimodal, users are most endangered by technologies whose performance is so advanced that it is no longer intuitive. In [26], a series of experiments was carried out on multimodal geo-locations to compare human with machine performance. In the cases in which machines can outperform humans, they do so in a way that would be nearly impossible to guess for a user not trained in multimedia analysis. For example, using a very large number of images, the automatic approach was able to discriminate mountain scenes in Asia, North America, and South America. A human expert might also have this ability, but a general user would have no reason to believe that mountains in different reasons are particularly easy to distinguish. Another example is the role of apparently harmless tags such as ‘iphone’ or ‘3g’. Such tags may be associated with photos taken in any region, the distribution of the tags itself is able to discriminate regions. Finally, it was found that videos of the Shinkansen train leaving Tokyo station could be automatically geo-located due to the distinctiveness of the sound produced. A user listening to the Shinkansen might realize that it makes a particular sound, but has no way of knowing it is that distinctive enough to pick out an exact location.

In the next section we turn our attention to multimedia algorithms that provide users with control to protect their privacy in the face of these threats.

7.4 Privacy-Related Multimedia Analysis Research

First, we look at examples of machine learning techniques that are able to infer information about what is depicted in multimedia content. In contrast to conventional multimedia analysis algorithms, which focus on the literally depicted content of images and videos, these approaches focus on less-commonly studied problems with a connection to educating users and supporting the law in protecting them.

7.4.1 Multimedia Analysis Filters

Recent years have seen the development of multimedia analysis algorithms that are able to infer the privacy status of user photos. There are several possible use scenarios for such algorithms. One prominent scenario is that the social networking application would warn a user that a photo might be private before the user makes the decision to upload a photo. This scenario is interesting because it not only helps to protect the user in the moment (avoiding an uploading mistake of this particular photo) but it also educates the user (raises awareness of what would be considered a privacy-sensitive photo).

Other possible use scenarios include smart cameras that warn a user that a photo might be "too private" before clicking the shutter, and automatic sorting applications.

The first work, to our knowledge, that studied the classification of photos by privacy status was [27] and the accompanying demonstration, PicAlert! [28]. This work conceptualized the privacy of a photo as a generic notion: the privacy status of photos was judged independently of the personal view of the user who took the photo. The photos were crawled from Flickr, and the judges were provided with the following description of what constitutes a private photo: "*Private* are photos which have to do with the private sphere (like self portraits, family, friends, your home) or contain objects that you would not share with the entire world (like a private email)". It is notable that this definition covers both aspects related to the identity of the people pictured in the photo, but also related to the subject material that the user photographs. In other words, it is clear from this definition that a private photo can be a photo other than one that depicts the user himself in a state, location, or with company that he does not wish to make public. At the time that this work was published, there was a surprising reaction of incredulity that judges make stable judgments of privacy, or that it was possible at all to carry out such classification. However, the classifier performed surprisingly well, both using visual features and using text features. A combination of the two achieved an F-measure of 0.8 (data set size 9402 photos evenly split between public and private images). The PicAlert! work appears to have had the function of allowing the multimedia community to accustom itself to the idea that it is possible to build a classifier to make classification decisions about the characteristics of images going above and beyond their literally depicted content.

Work that followed addressed two shortcomings of the PicAlert! work. First, PicAlert! used data collected from Flickr, so the chance was high that the users who originally took the photos did not consider them private (since they were posted to Flickr). Also, PicAlert! did not go beyond the generic notion of privacy. In [29], a study was carried out that investigated private photos of users together with the users' own privacy status classification. The data set used was tiny (150 photos), but the work is nonetheless important. The work addresses the scenario of classifying the photos on a user's device to prevent them from inadvertently being viewed. An accuracy of 80% was achieved using easily-available metadata and image features. The best set of features was quite diverse: latitude, longitude, elevation, Unix minutes, calendar week, weekday, day of the month, important hue, ISO-speed, local acutance, number of faces, and resolution. The paper includes a discussion of the underlying trust model and also analyzes the results in terms of their impact on the user. Private photos misclassified as non-private lead to a privacy violation, whereas non-private photos misclassified as private only lead to inconvenience. Understanding the effectiveness of multimedia filters for privacy requires not only reporting standard evaluation metrics, but also analyzing the implications of these metrics for users using the filters in practice.

The most recent work on the classification of images with respect to their privacy status is [14]. This work used photos taken by users (a total of 1511 photos collected from 27 users) and judged by users. The work neatly circumvents the issue of collecting private photos by only collecting features. An important aspect of this study is that it looked at how a generic privacy model can be adapted to cover a user's individual privacy needs. The findings verified that users have individual preference for privacy. The finding showed that very few photos labeled by users between 5 and 35 were enough to substantially improve the performance of the generic model. This result shows promise that a personalized privacy filter can be trained for individual users, requiring very little feedback. The findings also showed that it is important to find the appropriate weight to balance the contribution of the generic model and the user. An important aspect of this work is its emphasis of explainability. Although the best visual classifier directly used CNN features, competitive performance was achieved by a classifier using so-called *semantic features*, the output of a concept detector that detected a large number of different classes, such as "child", "erotic", "seaside", and "groom". The benefit of using semantic features is that they help the user to understand the decision of the classifier. The result is better informed privacy decisions on the part of the user, and, as already mentioned, a potential educational effect about which types of images can lead to privacy violations.

We finish the discussion of multimedia analysis for privacy by mentioning the connection to the research area of multimedia *user intent* [30]. Specifically, we point to work in the area of predicting *uploader intent* [31], defined as "the reasons that motivate users to upload videos to the Internet". A set of YouTube videos (1677 total) was labeled by mechanical turk workers, who judged whether the video best could be considered "personal" (uploaded for family or friends), "social" (uploaded for people with a common interest, but not everyone), or "public" (uploaded for a broad audience). Using visual features it was possible to build a weak predictor for these classes. Text features derived from the video metadata outperformed the visual classifier, achieving a weighted F-measure across the three classes of 0.53 (for comparison the dominant class baseline achieved 0.32).

In short, multimedia research has shown that it is viable to create classifiers that can estimate the privacy status of multimedia content and support users in making informed choices for sharing. However, before deploying these classifiers in products, it is important to understand how to teach users to use them wisely. A user should never be tempted to suspend their own judgment when making the decision of whether or not to share multimedia content.

7.4.2 Multimedia Content Masking

Masking approaches to privacy protect users by changing data in order to reduce the privacy risk that it represents. In [16], masking methods are defined as including perturbative and non-perturbative methods (p. 62). Masking via perturbation involves distortion of the original data, while masking without perturbation involves replacing the values of the original data with less specific values. We note that [16] points out that perturbative methods can introduce error, while non-perturbative methods just reduce the level of detail of the data. Here, we focus on perturbative masking methods, especially obfuscation techniques. We follow the definition of [32], who state, "Obfuscation is the deliberate addition of ambiguous, confusing, or misleading information to interfere with surveillance and data collection" (p. 1). The key idea of obfuscation is to produce data modeled on existing data to make the collection of data more difficult to analyze and act on. Obfuscation techniques are discussed in detail by [32], who argue that obfuscation provides people who are not in a position to exercise control over their own data with ways to evade or even sabotage the surveillance. Obfuscation is a two-edged sword, and [32] provides concrete examples of cases in which it has been used to muddy essential communication during elections, as well as protect people working against a repressive government. It is important to differentiate between obfuscation used destructively and obfuscation used as self-defense. Here, we discuss obfuscation in the service of self-defense since we are focused on multimedia algorithms that provide users with greater control to protect themselves from privacy threats.

Recent work demonstrates the potential of obfuscation to protect information about user location by foiling the content-based geo-location estimation methods discussed above. The study in [33] investigates the impact of common user image enhancements (filters and cropping) on the accuracy of geo-location estimation algorithms. Because the enhancements remove some information from images, the expected result is that filtered photos will be less geographically distinctive and more difficult to associate with the location at which they were taken. The paper reveals that filters have a small, but measurable effect. When image enhancements are applied, performance of both the search-based [25] and classification-based approaches [34] degraded.

The benefit of obfuscation-related strategies can be understood by referring back to the discussion of the viability of privacy protection measures taken to counter the privacy threats encoded in threat models. There we mentioned the importance of measures being easy to implement and also maintaining their protective function in the face of the lack of trustworthy parties. Obfuscation is an interesting technique because it does both. It is easy to understand, and, also, obfuscated data minimizes the danger of data falling into the wrong hands, i.e., due to a data breach. Further, obfuscation also provides users with the possibility of plausible deniability: any of the events that occur in their photo streams, for example, are possibly synthetic events added for the sake of allowing users to blend in.

7.5 The Larger Research Picture

The focus of this chapter is on multimedia algorithms that help users to stay in control. It is important, however, to realize that the research presented here is only part of a much larger picture involving many types of research related to privacy and big data. In this section, we point out some of the major areas that multimedia privacy researchers should be familiar with.

7.5.1 Multimedia Security and Trust

Multimedia security, which seeks to protect multimedia data, has been an important topic in the research community for the last two decades. The development of the field of multimedia security has been largely driven by the needs of industry to protect intellectual property (IP) rights. Work on protecting IP has led to research on specific areas of "organization privacy" (mentioned in section 7.1.2). Many of the technologies developed are currently in commercial use.

Specifically, IP protection has been studied by researchers working in the area of fingerprinting (e.g., [35]), watermarking (e.g., [36]), and "secure watermark detection" (e.g., [37]), namely, proving the existence of a fingerprint or watermark without revealing it. The impact of this line of research was greater than expected; a new research field, secure signal processing, was born.

In a system where security and privacy are needed, research starts with a definition of the capabilities of the attacker, his goals, and the security assumptions. In other words, a threat model, discussed in section 7.2.3, is needed. For example, in cryptography, cryptographic protocols are assumed to be secure in the presence of computationally bounded adversaries, meaning that the attacker has limited polynomial time computational power. Alternatively, we assume that the protocol is perfectly secure; no matter how much computational power the attacker has, the system is not breakable [38] since all possible solutions are equiprobable and the attacker cannot distinguish any of these. The security assumptions must also include a description of the data that should be protected from the adversary.

Security research distinguishes two types of adversary. The first one is an outsider: the traditional attacker who listens to the communication line passively or tries to modify the interaction between the involved parties actively. Traditional security measures like deploying encryption for secure communication, hashes, MACs and digital signatures for integrity and authenticity are straightforward approaches to defend against outsider adversary. The second type of adversary is an insider: an adversary directly in contact with the data. Insider adversaries are associated with cases where the sensitive data on the service provider side are leaked, e.g., due to a malicious or careless employee. This second model is the basis for a growing amount of research since it captures the notion of the untrustworthy service provider: the customers would like to use the services but do not trust the service provider with their sensitive data.

As already discussed in section 7.2.3, privacy must be protected in cases in which no one can be trusted. This concern is shared across the board by researchers oriented towards security and researchers oriented towards privacy. Even if the service provider does not use users' information inappropriately, it is not possible to assume that the information will never fall into the wrong hands, for example due to a systems breach or to a takeover by another company or a untrustworthy government. Research that looks at different types of trust issues is carried out under the umbrella of privacy-enhancing technologies (PETs) and deploys methods such as anonymization [39, 40], differential privacy [41], and computational privacy, which relies on processing encrypted data [42, 43].

7.5.2 Data Privacy

Multimedia privacy can be understood to be part of a much larger research area on data privacy. Data privacy research is described by [16] as encompassing three research communities: statistical disclosure control, privacy-preserving data mining, and PETs.

In turn, issues of data privacy are embedded into larger issues of protecting the entire pipeline of a system that makes use of user data. The pipeline of such a information system requires privacy protection at different stages:

- Ingoing information (queries): Users queries and query patterns must be protected. For example, for the problem of search in encrypted databases scenarios are addressed in which the service provider as the owner of the data should not be able to see which entry is being queried.

- Outgoing information (delivery): Information about the use of data must be protected. For example, in a content delivery system, the data owner should not be able to see who accessed or consumed which content.
- Stored information (data): Data that contains explicit information about users or data from which knowledge about users can be inferred must be protected while it is at rest. Stored information is the main challenge addressed by data privacy. It must be possible to exploit data, yet still maintain the privacy protection of users.

Two areas of research which look at data privacy in the context of systems are private information retrieval (PIR) [44] and privacy for recommender systems [45]. Here, we focus on recommender systems research. We point out that from the beginning, recommender systems research has looked at ways to keep users in control and has also made use of threat models, two aspects that we cover here. Specific examples are [46, 47], which combined secure multi-party computation and homomorphic encryption in order to create privacy-preserving recommender system protocols under the assumption of an untrusted service provider. Users collaborate to compute intermediate values without a central server, making it possible to do away with the need for a trusted service provider.

In the remainder of this section we look at recommender systems research that has yielded interesting examples of the application of masking methods. Applying masking to recommender system data involves introducing perturbations that protect privacy without affecting recommendation effectiveness. In other words, the prediction accuracy is maintained. We take a closer look at examples of two research directions.

The first research direction is interested in masking the rating data, where ratings are the attribute that needs to be protected. As examples of this type of research, we mention [48], where the authors showed that one of the solutions to the privacy issue in collaborative filtering recommender systems is to apply masking, which modifies the ratings stored in user profiles. The evaluation showed that users can mask large parts of the user profiles (ratings) without significantly decreasing the accuracy of the predictions.

The second direction of research defines specific privacy-sensitive attributes (e.g., gender, health state) that should be protected and demonstrates that prediction accuracy can be maintained even when masking is introduced. An example of this type of research is given in [49, 50]. In [49] a framework is proposed for privacy-preserving recommendations using data obfuscation. It is demonstrated that system performance on the obfuscated data does not necessarily impact the accuracy of the prediction. This type of protection would enable, for example, e-commerce vendors to share information about their users without violating their privacy. In [50], the authors proposed a new method called *BlurMe* that adds ratings to a user's profile with the goal of making it hard to infer the user's gender, while causing a minimal impact on recommendation quality.

This brief discussion highlights opportunities for multimedia privacy research to take a cue from data privacy research, which devotes effort to considering privacy at multiple stages of the pipeline, and also has yielded interesting demonstrations of how masking can be used to protect privacy and for which sorts of situations it is appropriate.

7.6 Outlook on Multimedia Privacy Challenges

In this section we look at the research challenges in multimedia privacy that remain to be tackled, as well as ways in which research must reorient itself in order to tackle them effectively.

7.6.1 Research Challenges

We expect that the future will bring further work on a wide range of areas, for example de-identification of multimedia content, which was briefly covered in section 7.2.1. Here, we focus on research challenges related to keeping control in the hands of the users who produce the multimedia content, which has been the focus of this chapter.

7.6.1.1 Multimedia Analysis

Techniques related to multimedia analysis have a bright future. Here, we mention several areas that would be well served by more research attention than they are currently receiving. First, researchers should take a closer look at mitigation. For example, [17] points out the need for approaches that are able to infer for a given photo which information the user intends to convey (foreground information) and which information is conveyed unintentionally (side information). These approaches should focus on metadata (Does the user know that the photo he is posting also contains information about the identity of his camera?) and also on the semantic content of the photo (Does the user realize that the wallpaper of the hotel is visually unique?) In addition to educating the user about danger of possible information leaks of this sort, and informing the user of potential leaks before they happen, algorithms should be developed that are able to control the damage of leaks once they have occurred.

Techniques that seek to thwart multimedia analysis by blocking inference of sensitive information are nearly as important as the multimedia analysis themselves. Here, we point to the possibility of developing multimedia storytelling techniques that would be capable of enhancing social network feeds. As mentioned in [32], in 2012 a technique was proposed and given the name Bayesian Flooding.[5] Its purpose is to hide real-life events in a stream of creative fiction, with the aim of shifting the prior probability that a user belongs to a certain demographic, as calculated by advertisers interested in targeting particular market segments. A key point of this technique is that it needs to be fun. The obfuscated feed should be engaging and not confusing for the user's followers. We will comment further on the importance of user engagement below. The user may choose not to obfuscate certain characteristics, such as his gender or his music tastes, but he might choose to obfuscate subtle information such as his psychological state as a reaction to concerns about inference of psychological fragility.[6]

7.6.1.2 Data

Conventionally, researchers assume that big multimedia data will only become bigger, and that the growth will last indefinitely and without control. Here we point out that this assumption may inadvertently be leading to a lack of interest in research topics focused

[5] http://www.kevinludlow.com/blog/1610/Bayesian_Flooding_and_Facebook_Manipulation_RD/.

[6] https://www.theguardian.com/technology/2017/may/01/facebook-advertising-data-insecure-teens.

on deriving more use from smaller amounts of data. Developing algorithms that need only a small amount of user data is an important step in helping to protect user privacy. Less data can be collected in order to achieve the same aim.

Recently, the area of recommender systems has seen growing interest in data minimization. For example, in [51] data requirements analysis is proposed as best practice. The argument this work advances is straightforward: just as we avoid developing algorithms with unnecessary computational complexity, we should also avoid developing algorithms that need unnecessary data. Evaluation of a new algorithm should include an analysis which demonstrates that the algorithm uses the minimum amount of data necessary in order to achieve its goal. The goal of a recommender system is defined in terms of one or more metrics that capture the effectiveness of the prediction that it generates. Once this goal is met, the use of additional data is no longer justifiable and should be avoided. The future could see the commercial demand rise for data minimization algorithms. First, as users grow more aware of the dangers of sharing, they will be drawn to services that set themselves apart by their responsible handling of data. Second, new laws (cf. the EU General Data Projection Regulation,[7] which comes into force in 2018) are coming into force which impose strict data requirements. Companies will find themselves in a position that they need to do more with less data. A small nudge might be enough to push industry to working with less data. After all, algorithms that use less data may also require less computation, saving on computational resources, including energy.

Another opening for researchers is the area of synthesizing realistic data. It has been pointed out in the literature, e.g., by [14], that privacy-related data sets are difficult to create due to the nature of the material. Even if it is not possible to synthesize data representing individual users, data synthesis might be helpful at the ecosystem level. Synthetic or semi-synthetic data could make it possible to better understand which types of data are the biggest risk to user privacy and which levels of aggregation are most dangerous. Further, with sufficiently advanced data synthesis, the amount of user information that must be collected could be reduced—algorithms would only need a minimal amount of data to seed the synthesis process, and the multimedia and interaction data needed for training systems could be home grown from there.

7.6.1.3 Users

In this chapter we have emphasized that protecting privacy is related to protecting information, but that its main focus is protecting users. For this reason, future research will necessitate the investment of time, energy, and resources in the detailed user studies needed in order to ensure that privacy-protection algorithms meet user needs. During the presentation of the threat model in section 7.2.3, we mentioned the importance of privacy-protection measures being viable. Critically, viability includes ease of use, and, we argued, effective privacy-protection measures may actually need to educate users, or provide incentives in order to ensure that they are adopted. It is important that users' sense of fun and feel for aesthetics are not disturbed by technology that protects their privacy. Recall that [33] investigated the ability of naturally occurring user practices of photo enhancement to reduce the accuracy of content-based geo-location estimation. Instead of developing a privacy-protection measure and hoping that users would

[7] http://ec.europa.eu/justice/data-protection.

adopt it, this work starts with a behavior that users already engage in and investigates its potential as the basis for a privacy-protecting technology. We hope that future research will be able to also productively adopt such a tactic. How to employ obfuscation techniques for multimedia privacy without compromising users' quality of experience is an important research question that must be addressed in the future.

7.6.2 Research Reorientation

We have seen that multimedia privacy holds significant challenges. It is clear that the research necessary can be measured at the scale of conferences, projects, and products, rather than at the scale of a handful of papers. In order to make real progress in privacy, it is necessary that multimedia research reorients in order to direct more resources to privacy challenges. With our call for reorientation of research we echo the sentiment of [14] that structuring of an active research community in the area of multimedia privacy is an important next step. Here, we discuss three changes that would support such a reorientation.

7.6.2.1 Professional Paranoia

The security world cultivates an attitude of "professional paranoia" towards scientific problems. This attitude predisposes researchers to look at a system or situation from the perspective of what could possibly go wrong. This concept is lacking in today's multimedia research environment, where the focus is on reasonable or probable scenarios that assume cooperative behavior from all users. The threat model is the key tool, adopted from the security discipline, that allows researchers to make headway on developing algorithms that protect users' privacy in face of the complexity and messiness of the topic. Multimedia researchers would be well served by working more closely with information security researchers, adopting both their tools and their attitudes. One hurdle to overcome before we can achieve such collaboration is the traditional reluctance of security researchers to work with detailed user requirements or to carry out user studies (e.g., as pointed out by [13]). However, security researchers are skilled in projecting themselves into the minds of the adversaries interested in compromising systems, and with enough cooperation with user-oriented multimedia researchers they will naturally also become better at projecting themselves into the minds of the users who need their privacy protected. In addition to professional paranoia an attitude of professional empathy must be developed: whatever seemingly illogical behaviors and practices users might pursue, the multimedia privacy researcher must be willing to devote time and energy to the task of protecting that user's privacy.

7.6.2.2 Privacy as a Priority

Protecting user privacy may well be the most challenging problem currently faced by multimedia researchers. Making privacy a priority requires adapting our assumption that the information world exists on a separate plane and does not interact with the real world. As computer scientists, we are used to being able to abstract away from physical reality. In the digital world, people do not suffer from hunger, disease or cold the way they would in the physical world. We feel confident that we can ignore the physical world. However, this confidence can get us into trouble when we study big multimedia data. When we create multimedia systems, we jump directly to the assumption that we can treat online users as entities abstracted away from physical-world needs. In our rush to abstraction, we forget that privacy concerns information, and this issue follows us from the physical to the digital world in ways that temperature does not. It is easy to assume that privacy violations, occurring as they do in the realm of information, cannot be dangerous. Although information effects clearly have non-physical properties, we should not assume that an information-related phenomenon like privacy cannot result in physical harm.

The assumption that the information world does not impact the physical world is reinforced by the lack of research on the long-term harmful effects of privacy-deprivation on human well-being. Given the lack of such research, it might first seem logical to assume that privacy research can wait until scientists get around to confirming the harmful effects of privacy deprivation. However, studies on such questions may actually be impossible to formulate in an ethical manner. For this reason, it is important to act now on the basis of our natural understanding of the importance of seclusion for the well-being of human beings, and not wait until it is too late.

It is interesting to note that not all researchers will be equally attuned to the importance of privacy. Anyone who has fallen victim to an attack such as cybercasing will obviously feel a more acute sense of urgency for developing algorithms to address the problem. Also, the chance of encountering a problem with privacy violation increases with age: people who are older simply have longer life stories, more that can be revealed, and more time in which the privacy violation could have occurred. Older people may be underrepresented in the research and development workforce (most obviously because older people retire). It is important that multimedia researchers do not rely exclusively on their own experiences to inspire them to carry out privacy research, but also talk to a large range of other people who have first-hand experience of the ill-effects.

7.6.2.3 Privacy in Parallel

A popular position that researchers assume without reflecting particularly carefully is that geo-location prediction is the "main task" and that privacy protection of users' photos is only an interesting side branch of the research. This position characterizes the status quo and should not be taken as a research vision. Formulating a research vision is difficult since we are limited in our abilities to predict what the future will bring. However, we do know something about the realities that give rise to the big multimedia data collections that we study: the data are created by users and the systems are profitable to the extent that they can attract and retain users. If our multimedia systems fail to protect users, we lose not only the source of our data, but the entire motivation for creating multimedia technology in the first place. If we seek to be prepared for any future, then we need to be prepared for a future in which the problem of multimedia privacy stands on an equal footing with the problem of multimedia inference.

References

1 Friedland, G. and Sommer, R. (2010) Cybercasing the joint: On the privacy implications of geo-tagging, in *Proceedings of the 5th USENIX Conference on Hot Topics in Security (HotSec'10)*, HotSec'10, pp. 1–8.

2 Friedland, G., Vinyals, O., and Darrell, T. (2010) Multimodal location estimation, in *Proceedings of the 18th ACM International Conference on Multimedia (MM '10)*, pp. 1245–1252.

3 Zheng, D., Hu, T., You, Q., Kautz, H.A., and Luo, J. (2015) Towards lifestyle understanding: Predicting home and vacation locations from user's online photo collections, in *Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM '15)*, pp. 553–561.

4 Tasse, D., Sciuto, A., and Hong, J.I. (2016) Our house, in the middle of our tweets, in *Proceedings of the 10th International AAAI Conference on Web and Social Media (ICWSM '16)*, pp. 691–694.

5 Craig, T. and Ludloff, M. (2011) *Privacy and Big Data*, O'Reilly.

6 Wacks, R. (2015) *Privacy: A very short introduction*, OUP Oxford.

7 Friedland, G., Papadopoulos, S., Bernd, J., and Kompatsiaris, Y. (2016) Multimedia privacy, in *Proceedings of the 2016 ACM on Multimedia Conference (MM '16)*, pp. 1479–1480.

8 Westin, A.F. (1968) Privacy and freedom. *Washington and Lee Law Review*, 24 (1).

9 Jeckmans, A.J., Beye, M., Erkin, Z., Hartel, P., Lagendijk, R.L., and Tang, Q. (2013) Privacy in recommender systems, in *Social Media Retrieval*, Springer, pp. 263–281.

10 Kang, J. (1998) Information privacy in cyberspace transactions. *Stanford Law Review*, pp. 1193–1294.

11 Friedland, G., Janin, A., Lei, H., Choi, J., and Sommer, R. (2015) Content-based privacy for consumer-produced multimedia, in *Multimedia Data Mining and Analytics*, Springer, pp. 157–173.

12 Information Security, Wikipedia, https://en.wikipedia.org/wiki/Information_security (accessed 1 December 2017).

13 Adams, A. (2000) Multimedia information changes the whole privacy ballgame, in *Proceedings of the 10th Conference on Computers, Freedom and Privacy: Challenging the Assumptions (CFP '00)*, pp. 25–32.

14 Spyromitros-Xioufis, E., Papadopoulos, S., Popescu, A., and Kompatsiaris, Y. (2016) Personalized privacy-aware image classification, in *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval (ICMR '16)*, pp. 71–78.

15 Ribaric, S., Ariyaeeinia, A., and Pavesic, N. (2016) De-identification for privacy protection in multimedia content: A survey. *Signal Processing: Image Communication*, 47, 131–151.

16 Torra, V. (2017) *Data Privacy: Foundations, New Developments and the Big Data Challenge*, Springer International Publishing.

17 Friedland, G. and Tschantz, M. Privacy Concerns of Multimodal Sensor Systems, book chapter, to appear.

18 Bernd, J., Gordo, B., Choi, J., Morgan, B., Henderson, N., Egelman, S., Garcia, D.D., and Friedland, G. (2015) Teaching privacy: Multimedia making a difference. *IEEE MultiMedia*, 22 (1), 12–19.

19 Santani, D., Biel, J.I., Labhart, F., Truong, J., Landolt, S., Kuntsche, E., and Gatica-Perez, D. (2016) The night is young: Urban crowdsourcing of nightlife patterns, in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*, pp. 427–438.

20 Salter, C., Saydjari, O.S., Schneier, B., and Wallner, J. (1998) Toward a secure system engineering methodology, in *Proceedings of the 1998 Workshop on New Security Paradigms (NSPW '98)*, pp. 2–10.

21 Schneier, B. (2015) *Data and Goliath: The hidden battles to collect your data and control your world*, W.W. Norton & Company.

22 Strandburg, K.J. (2014) *Monitoring, Datification, and Consent: Legal Approaches to Privacy in the Big Data Context*, Cambridge University Press, pp. 5–43.

23 Lei, H., Choi, J., Janin, A., and Friedland, G. (2011) User verification: Matching the uploaders of videos across accounts, in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2404–2407.

24 Larson, M., Kelm, P., Rae, A., Hauff, C., Thomee, B., Trevisiol, M., Choi, J., Van Laere, O., Schockaert, S., Jones, G.J.F., Serdyukov, P., Murdock, V., and Friedland, G. (2015) The benchmark as a research catalyst: Charting the progress of geo-prediction for social multimedia, in *Multimodal Location Estimation of Videos and Images*, Springer, pp. 5–40.

25 Li, X., Larson, M., and Hanjalic, A. (2017) Geo-distinctive visual element matching for location estimation of images. *IEEE Transactions on Multimedia*, 20 (5), 1179–1194.

26 Choi, J., Lei, H., Ekambaram, V., Kelm, P., Gottlieb, L., Sikora, T., Ramchandran, K., and Friedland, G. (2013) Human vs machine: Establishing a human baseline for multimodal location estimation, in *Proceedings of the 21st ACM International Conference on Multimedia (MM '13)*, pp. 867–876.

27 Zerr, S., Siersdorfer, S., Hare, J., and Demidova, E. (2012) Privacy-aware image classification and search, in *Proceedings of the 35th International ACM Conference on Research and Development in Information Retrieval (SIGIR '12)*, pp. 35–44.

28 Zerr, S., Siersdorfer, S., and Hare, J. (2012) PicAlert!: A system for privacy-aware image classification and retrieval, in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM '12)*, pp. 2710–2712.

29 Buschek, D., Bader, M., von Zezschwitz, E., and De Luca, A. (2015) Automatic privacy classification of personal photos, in *Proceedings of the 15th IFIP TC.13 International Conference on Human-Computer Interaction (INTERACT '15)*, pp. 428–435.

30 Kofler, C., Larson, M., and Hanjalic, A. (2016) User intent in multimedia search: A survey of the state of the art and future challenges. ACM *Computing Surveys*, 49 (2), 36:1–36:37.

31 Kofler, C., Bhattacharya, S., Larson, M., Chen, T., Hanjalic, A., and Chang, S.F. (2015) Uploader intent for online video: Typology, inference, and applications. *IEEE Transactions on Multimedia*, 17 (8), 1200–1212.

32 Brunton, F. and Nissenbaum, H. (2015) *Obfuscation: A user's guide for privacy and protest*, MIT Press.

33 Choi, J., Larson, M., Li, X., Li, K., Friedland, G., and Hanjalic, A. (2017) The geo-privacy bonus of popular photo enhancements, in *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval (ICMR '17)*, pp. 84–92.

34 Weyand, T., Kostrikov, I., and Philbin, J. (2016) PlaNet–photo geolocation with convolutional neural networks, in *Proceedings of the European Conference on Computer Vision (ECCV '16)*, pp. 37–55.

35 Caldwell, A.E., Choi, H.J., Kahng, A.B., Mantik, S., Potkonjak, M., Qu, G., and Wong, J.L. (1999) Effective iterative techniques for fingerprinting design IP, in *Proceedings of the 36th Annual ACM/IEEE Design Automation Conference*, pp. 843–848.

36 Bianchi, T. and Piva, A. (2013) Secure watermarking for multimedia content protection: A review of its benefits and open issues. *IEEE Signal Processing Magazine*, 30 (2), 87–96.

37 Prins, J.P., Erkin, Z., and Lagendijk, R.L. (2007) Anonymous fingerprinting with robust QIM watermarking techniques. *EURASIP Journal on Information Security*, 2007, 20.

38 Katz, J. and Lindell, Y. (2014) *Introduction to Modern Cryptography*, CRC Press.

39 Weber, R.H. and Heinrich, U.I. (2012) *Anonymization*, Springer Science & Business Media.

40 Raghunathan, B. (2013) *The Complete Book of Data Anonymization: From planning to implementation*, CRC Press.

41 Dwork, C. and Roth, A. (2014) The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9 (3–4), 211–407.

42 Lagendijk, R.L., Erkin, Z., and Barni, M. (2013) Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation. *IEEE Signal Processing Magazine*, 30 (1), 82–105.

43 Dara, S. (2013) Cryptography challenges for computational privacy in public clouds, in *2013 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, IEEE, pp. 1–5.

44 Gasarch, W. (2004) A survey on private information retrieval. The Bulletin of the European Association for *Theoretical Computer Science*, 82 (72-107), 1.

45 Knijnenburg, B.P. and Berkovsky, S. (2017) Privacy for recommender systems: Tutorial abstract, in *Proceedings of the 11th ACM Conference on Recommender Systems (RecSys '17)*, pp. 394–395.

46 Canny, J. (2002) Collaborative filtering with privacy, in *2002 Proceedings of the IEEE Symposium on Security and Privacy*, IEEE, pp. 45–57.

47 Canny, J. (2002) Collaborative filtering with privacy via factor analysis. *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '02)*, Tampere, Finland, pp. 238–245, ACM, New York.

48 Berkovsky, S., Kuflik, T., and Ricci, F. (2012) The impact of data obfuscation on the accuracy of collaborative filtering. *Expert Systems with Applications*, 39 (5), 5033–5042.

49 Parameswaran, R. and Blough, D.M. (2007) Privacy preserving collaborative filtering using data obfuscation, in *Proceedings of the IEEE International Conference on Granular Computing (GRC '07)*, pp. 380–380.

50 Weinsberg, U., Bhagat, S., Ioannidis, S., and Taft, N. (2012) Blurme: Inferring and obfuscating user gender based on ratings, in *Proceedings of the 6th ACM Conference on Recommender systems (RecSys '12)*, pp. 195–202.

51 Larson, M., Zito, A., Loni, B., and Cremonesi, P. (2017) Towards minimal necessary data: The case for analyzing training data requirements of recommender algorithms, in *ACM RecSys 2017 Workshop on Responsible Recommendation (FATRec '17)*, pp. 1–6.