1-1-2020

Contextual healing: Privacy through interpretation management

Fatma Outay

Rula Sayaf

Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 177 (2020) 308–317

Procedia Computer Science

www.elsevier.com/locate/procedia

The 11th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2020)
November 2-5, 2020, Madeira, Portugal

Contextual Healing: Privacy through Interpretation Management

Fatma Outay[a,*], Rula Sayaf[b]

[a]*College of Technological Innovation, Zayed University, DXB, UAE*
[b]*PSI Algebra, Leuven 3000, Belgium*

Abstract

Contextual privacy is an essential concept in social software communication. Managing privacy of data disclosed in social software dependence strongly on the context the data is disclosed in. The sheer amount of posts and audiences may lead to context ambiguity. Ambiguity can affect contextual privacy management and effective communication. Current contextual privacy management approaches can be either too complex to use, or too simple to offer fine-grained control. In many cases, it is challenging to strike a balance between effective control and ease-of-use. In this article, we analyse contextual privacy by in relation to context and communication. We examine a relevant contextual privacy management framework based on the maintenance of the interpretation of data. We propose an architecture based on the utilisation of intelligent mechanisms. We conceptually analyse the usability aspect of the proposed architecture. We conclude by arguing how our conceptual framework can enhance communication and privacy in private and public spaces.

Keywords: *Contextual privacy; social software; context; communication; identity; artificial intelligence.*

* Corresponding author. Tel.: +971 4 402 1789 ; fax: +0-000-000-0000 .
E-mail address: fatma.outay@zu.ac.ae

10.1016/j.procs.2020.10.092

1. Introduction

In the era of social software, privacy management becomes fundamental to facilitate communication in private and public spaces. Social software communication can be achieved through disclosures of multimedia posts (e.g., images, videos) rather than the exchange of explicit messages. Through a disclosure, a user expresses a particular communicative message and builds a desired identity. A post can be disclosed in a private space to a particular set of recipients; alternatively, it may be posted publicly for a large, and a priori unrestricted audience with a specific communicative message. Inappropriate dissemination can affect the communicative message, and the user's identity [35]. Privacy as self-determination is the key to protect data and identity [31]. It is also a means to facilitate communication in private and public spaces.

To facilitate communication, privacy management mechanisms should offer control on contexts to preserve the communicative message. These mechanisms should guard against misappropriation of data such as the following scenario (**Scenario 1**): The proud mother and carefree Cathie shares a photo of her young daughter bathing and shares it with public audience. The intended message is that her baby is cute. However, she did not anticipate that pervert Pete could also have access and he disseminates it in an 'pornographic' context. Cathie finds out about the inappropriate dissemination, and reports the abuse to the social software provider as it is a violation due to the identity damage she and her daughter have experienced.

Avoiding misappropriation of posts is complicated and is not often possible by most security and privacy management mechanisms. Avoiding misappropriations requires controlling dissemination contexts. Due to the high dimensionality of context and its complex nature [40], controlling it can be complicated. To avoid such complexity, security and privacy management mechanisms focus on simplifying the control offered to users. Offering users usable privacy management mechanisms requires a shift in the conceptualization and architecture design of privacy management mechanisms. Privacy management mechanisms should offer context control to facilitate identity management and communication with ease-of-use for average users. Proposing such contextual privacy mechanisms requires examining the role of context in communication and privacy management. It also requires a shift in designing such mechanisms towards utilizing artificial intelligence to assist users in managing their privacy [21].

In this article, we examine contextual privacy, and propose an architecture design for contextual privacy management. The examination validates a previously proposed framework for contextual privacy for social software called CPS^2 [35]. CPS^2 is a framework for contextual privacy management based on managing the interpretation of a post. This paper explores possible effective contextual privacy mechanisms by contributing the following:

1. Analyzing the concepts relevant to contextual privacy and communication (Section 2).
2. Elaborating on the issue of controlling context, and analyzing the causes of context ambiguity (Section 3)
3. Validating the previous conceptualization of contextual privacy, proposing an architecture design for CPS^2, and discussing a deep learning approach for the possible implementation. Also, analyzing the usability aspects of the proposed architecture design CPS^2 and comparing it to the well-known theory Contextual Integrity [30], and discussing the implications of this design for users experience (Section 4).
4. Analyzing how CPS^2 can address context ambiguity, enhance privacy in private and public spaces; facilitates cooperative communication and helps avoiding adversarial communication (Section 5).

2. Contextual Privacy Concepts

In this section, we illustrate contextual privacy as a means to managing the communication context to protect data and guard identity.

Context is "any information that can be used to characterise the situation" [1]. An online context is any information that can be used to characterise an online situation. The context implies the topic of communication and possibly some characteristics of the interlocutors. For instance, in a health-related communication context, the context implies the topic is health-related and that some of the interlocutors are doctors, nurses or patients. A context can be approximated by the set of available informational parameters in the situation. We refer to those parameters as the context-approximation parameters (CAP). Whenever these parameters are not readily accessible to the observer, it is challenging to approximate the context, and the context is said to be ambiguous.

A situation in social software can be characterised by a post and a context wherein the post is put (Figure 1(a)). The communication context includes the data surrounding the post, the principal owner and an audience. The principal owner is the user who discloses the post in the original context. The original context is the situation in which the post is originally disclosed via the software (Figure 1(b)). The audience in a context can be potential or actual audience. The potential audience are the users that can view the post. The actual audience are the members of the potential audience who have already viewed the post. When the actual audience contribute to the communication (e.g., comments or likes), they become subordinate owners. Another class is extended audience that is a feature of Facebook (Jan. 2014). It occurs when the friends of a friend become part of the audience of public posts. When a friend of the owner likes a public post, the post is displayed in the newsfeed of the friends of this subordinate owner, as if the principal owner shared the post with these friends. The data about the post, the owner and the audience in a social software context are the CAP. These CAP may differ according to the design of the software and the information available.

By being able to control the context, one can affect the approximation of context. Adding or removing data to context may affect the context-approximation parameters, and as result the context changes or transitions. Whenever the observer—a member of the audience—is unaware of this transition, there can be discrepancy between the perceived and the actual context causing ambiguity (Figure 1(b)). Approximating context is required to interpret posts [11]. The interpretation is a set of meanings, or the set of values of parameters [5]. Based on the context, the relevant interpretation can be disambiguated [4]. When the post is put in a certain context (‘contextualised’ [27]). Decontextualization is the process of taking a post out of the current context, to where the interpretation is unavailable [27]. We identify the third process of moving posts between two contexts as recontextualization.

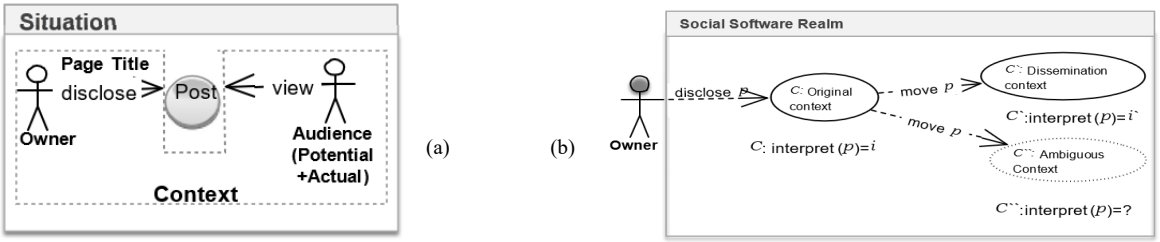

Figure 1: (a) A simplistic representation of a communication situation. Context includes all the parameters in the situation except the post. (b) The post p is disclosed in the original context C, where it has the interpretation i. When p is put in C0, the dissemination context, the interpretation is i0. When p is in the ambiguous context C00, an observer may not be able to interpret p.

Communication via Social Software In most communication, people aim to convey certain messages. In social software, posts can convey communicative messages to the audience. The audience may be well-known to the principal owner, or they may be strangers. Familiarity between the interlocutors affects the kind of communication in a continuum of trust between adversarial and cooperative communication. We focus on these two extreme ends of the communication spectrum to emphasize the varying roles of context and privacy.

In cooperative communication, interlocutors work together to understand the meaning of the communicated message. According to Grice, facilitating the inference of a message requires the interlocutors to cooperate and put an effort to clarify the communication [20]. Grice stated that providing true and sufficient amount of information that is relevant while avoiding ambiguity are key to unambiguous context. In cooperative communication, privacy concerns are relatively low while the importance of context is relatively high. The interlocutors trust each other to behave properly and clarify the context if needed.

Adversarial Communication is characterised by the manipulation of the communicated message. An interlocutor—the adversary—acts maliciously and misleads others into misinterpreting the message to disrupt the communication or force others to reveal certain information [15]. When context is ambiguous, communication can be adversarial [38]. In this communication, users may protect privacy by providing less information [41], with detrimental consequences on the clarity of context and the inference of the communicative message, thereby resulting in a misinterpretation and hence an unintentional adversarial communication (Scenario 7).

In summary, preserved privacy and clear context facilitate cooperative communication and play a role in avoiding adversarial communication.

Identity and Privacy are related to communication through which a desired identity is expressed [18, 14, 32]. By controlling context, a user establishes the appropriate situation to express a desired identity.

Privacy as control is demonstrated as informational self-determination and facilitates identity management [31]. Through controlling context, privacy facilitates the maintenance of more than one identity. This control aims at ensuring that others would infer the communicated message appropriately.

3. Issues Related to Context

In this section, we discuss the issue of controlling context and the issues of ambiguity that hinder privacy and identity management.

3.1. Controlling Context

Most existing contextual privacy mechanisms offer a low degree of context control [34]. Access control mechanisms offer owners the possibility to control and constrain access to posts by audience. These mechanisms offer control over the audience or the disclosure context. A recent survey of access control mechanisms in social software has identified issues in fine-granular context control [34]. Most mechanisms simplify context by capturing it by means of a single parameter [34]. Some approaches represent context by only roles of users [7], location or time [3]. This simplification may fail to actually capture contexts that users may want to control. In Scenario 1, limiting access to the photo to users with the role 'mother' is not enough to avoid 'pornographic' comments about the photo. Context may change to a different context by the evolution of the communication and this evolution requires the owner to change the access control policy. To avoid inappropriate dissemination and prohibit the evolution of context to inappropriate ones, the owner should formulate policies about all possible appropriate or inappropriate contexts, which is beyond what a user can do [35].

3.2. Context Ambiguity

The causes of context ambiguity may vary and are challenging to specify. The literature has identified three sources of context ambiguity [29]: 1) audience invisibility and the obscured viewing of owners' post, 2) contexts collapse due to lack of boundaries in social software situations, 3) and blurred boundaries between private and public and how posts can be accessed. We argue in the following that the invisibility of context parameters is the main cause of context ambiguity.

Invisible Owner: Scenario 2 Bill is an activist against gentrification. He anonymously posts a photo of a recent outrage with violent protesters and shares it with his wider friends. Some of the actual audience are unable to infer the reason of the violence and its relation to the people in the photo. The anonymous post may disrupt the audience's ability to infer that the photo is to report violence in Bill's neighborhood. Knowing Bill is an activist, makes possible approximating that the context as social uprising.

Invisible Subordinate Owner: Scenario 3 Sam is a police officer and comments on Bill's photo, saying that he was attacked and injured. Dean, being unable to see that Sam posted the commented, assumes a fellow protester was injured. He comments back saying that the police are brutal. The invisibility of the subordinate owner Sam, affects the approximation of context by Dean. Dean's comment is inappropriate because misrepresents the police as brutal.

Invisible Potential Audience: Scenario 4 Rex, a friend of Bill who works for the secret service, is a member of the potential audience. If he views the communication context after Dean has commented, he might disseminate the comment in a page about people who encourage violence against the police. Had Dean been aware that the potential audience include Rex—the intelligence employee—he could have been able to approximate the context more accurately and reason that his comment was inappropriate to be disclosed or that it may be assigned an interpretation that is not right. The possible act of taking Dean's comment and putting it in another context is a recontextualization of his post. The recontextualized comment is interpreted differently than intended. Recontextualizing the comment is a privacy violation to Dean, and does not contribute to the identity he is expressing.

Invisible Actual Audience: In the previous scenario, if the actual audience were visible,

Dean would have been able to detect the context change— Rex becoming a member of the actual audience, and assuming he has his profession accessible in the social software. To Dean, Rex is an adversary who may (mis)interpret his message. Being aware of this transition, Dean could have removed his comment, or taken precautions.

Invisible Extended Audience: Scenario 5 (scenario 3 cont.) After Sam has commented on Bill's post, Sam's colleagues, who are also police officers, become part of the audience and can see the interaction. They all think that Dean is the one who urged the crowd to attack the police.

When Sam becomes a subordinate owner, his friends become part of the audience. Such an extension of the audience changes the context. The invisibility of the extended audience challenges Dean to approximate the context to reason about how the new audience may perceive his message. The extension of audience is a case of blurred boundaries between private and public. The new audience can be total strangers for Dean, and sharing his comment with them takes it out of the private space he thought his comment will only be viewed in. Moreover, to the extended audience, the new communication context can be ambiguous, and this is a problem because the audience who get accidental access to personal data of others do not have the same ethical obligations and responsibilities towards respecting the privacy of the owners [33].

3.2.1. Ambiguity and Privacy

Ambiguity negatively affects privacy, as well as communication. A communicated message and its interpretation contribute to one's identity [19, 42]. When it is possible to correctly interpret a post, communication is effective, the expressed identity, and hence privacy are maintained—unless other violations occur. In other words, effective communication contributes to preserving one's privacy [35].

4. Contextual Privacy Management

The interdependence of privacy and communication validates the previous conceptualization of contextual privacy by Sayaf et al. [35]. The previous scenarios show that when the interpretation changes, privacy might be negatively affected. By using the interpretation to manage contextual privacy, we argue in the following, that the complexity of controlling context could be overcome.

4.1. CPS^2: Contextual Privacy Framework for Social Software

The previously proposed framework CPS^2 is a conceptual framework to manage contextual privacy without overloading users and without sacrificing the richness the context offers [35]. In CPS^2, the focus is on the online context only. It assumes that the interpretation in a specific context is appropriate if the owner allows the disclosure of a post in this context. CPS^2 states that an owner can identify the appropriate interpretation of her post, and based on this interpretation, dissemination and context changes can be managed. For example, a user discloses a disease-related post and specifies the appropriate interpretation to be 'disease'. By automatically checking this interpretation in any context, recontextualizations can be controlled to avoid interpretation changes. The aim of CPS^2 is not to control the interpretation the audience infer, but to facilitate the inference of the intended interpretation of the owner and avoid misappropriating posts within social software contexts. The framework is designed based on the concept that the surrounding context helps interpreting posts to prevent dissemination in inappropriate contexts.

4.2. An Architecture Design for Contextual Privacy Management

In this section, we elaborate on the architecture design for CPS^2 proposed in [35]. We extend the description of each layer, propose the interaction between these layers (Figure 2), and investigate techniques of machine deep learning to implement inference layers.

***Context Inference Layer*:** Responsible for processing the communication situation (e.g., through a Facebook page), to approximate the current context within the social software realm.

***Interpretation Inference Layer*:** Responsible for inferring the interpretation of data whether textual or visual, based on

the context inferred by the previous layer. Inferring the interpretation is similar to how a search engine matches a search query to a document: the document is the context and the query is the post. The query has a specific interpretation in a document. Based on the appropriateness of this interpretation, only transitions and recontextualizations that maintain the interpretation are allowed.

***Contextual Privacy Management Layer (CPML)*:** Responsible for facilitating contextual privacy management by maintaining the appropriateness of interpretation. Following access control, CPML allows users to specify the appropriate interpretation of their posts. CPML verifies any action or transition of context to maintain the appropriateness of the interpretation. Alternatively, without specifying the appropriate interpretation, CPML notifies the owner when the interpretation changes from the interpretation in the original disclosure context, following accountability and auditing approaches. The owner judges the appropriateness of the new interpretation, and accordingly the change of context is allowed or prohibited.

These layers can be implemented by machine learning models, especially deep machine learning generative models. Deep learning focuses on computational models for complex information representation [6]. Generative models are useful for unsupervised learning with a high number of parameters [22]. Generative models can learn a joint probability distribution over observable data and labels to estimate the conditional probability P(O|L) and P(L|O), where L is a label and O is a set of observable data variable. In CPS^2, the observable data is the CAP, and labels are information about context names and interpretations. The Multimodal Learning with Deep Boltzman Machine [39] could be applied to learn a multimodal data representation. The model classifies images and tags them; and it can also retrieve images corresponding to a set of tags. This model can be applied for context and interpretation inference and context retrieval. On top of such a model, CPML can be implemented as an access control or accountability and audit approach.

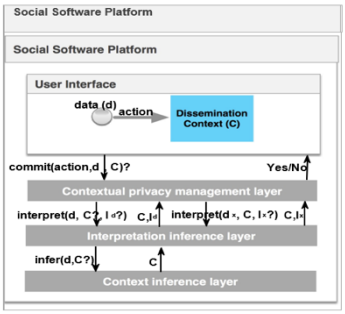

Figure 2: The interaction between layers. Upon submitting a request to add a post d, CPML checks whether the action can be committed by consulting the interpretation inference layer. To infer the interpretation, the context inference layer is consulted

4.3. Conceptual Analysis of Usability

In this section, we present a comparative assessment of the usability of CPS2. We assess the most relevant and widely accepted conceptual framework for contextual privacy 'Privacy as Contextual Integrity' (CI) proposed by Nissenbaum [30]. CI is an essential work that incorporates context relevant aspect to manage contextual privacy. It mainly addresses the issue of limiting recontextualization of posts by controlling four parameters: contexts, actors, attributes, and transmission principles. CI requires the specification of the norms including: terms of information flow; the prevailing contexts and possible sub- and super-contexts; subjects, senders, recipients; and transmission principles.

Usability is an important aspect in achieving the objectives of security and privacy management mechanisms [10]. If a mechanism is not easy to use, average users would fail to preserve their security or privacy regardless of the effectiveness of the mechanism. Assessment of usability requires significant engineering effort. To avoid such time-consuming tasks, many works have focused on assessing usability at the design phase [37, 10]. We use the usability of CPS2 and CI using the 'Security Usability Model' proposed by Braz et al. [10]. In their model, they select metrics from the Quality in Use Integrated Measurement model for usability standards [37]. This model is considered to provide the best usability standard.

In principle, CI requires more effort of users and may pose challenges to usability in contrast to CPS^2. CI requires

specifying parameters that may be challenging to specify in advance, for instance, users may not be aware of the terms of information flow in the system, or they not be able to predict how the terms may change over time. The most challenging aspect of CI is that it is based on the prohibitive requirement of specifying appropriate contexts. On the practical level, when CI is deployed in formal access control models or technical mechanisms [7, 24], users are still required to specify the same number of parameters stated in CI. In contrast, CPS^2 limits the number of parameters users need to specify to the interpretation of their data. And, it requires the incorporation of intelligent mechanisms to overcome the burden of handling context. These two aspects make CPS^2 satisfy the most of the metrics to a higher degree than CI (Table 1).

Implications for User Interface Design and Engineering We foresee three main design aims that enhance the user interaction experience of social software, enhance usability, and compliance with privacy requirements [2]:

Context change alerts. Besides alerts of inappropriate interpretations, users can be alerted when CAP change.

Awareness tools. More generally, users will be made more aware of how their communication evolves.

Feedback loops. Users will have the opportunity to provide feedback to the system (e.g., rate alerts or confirm blocked interlocutors), thereby generating labels that become part of the training data and gradually improve the system recommendations.

5. Applying CPS^2

5.1.1. In Private Contexts

Private contexts are the communication contexts in which owners constrain access to data to protect their privacy. CPS^2 offers contextual privacy management and requires minimal involvement of users during the following phases.

Table 1: Usability metrics. Given that this is an estimate of the performance of the designed system, we only use two degrees ‘high’ and ‘low’ to indicate the estimated degree. In UM6, we assess the time required to load the whole application in social software. CPS^2 requires loading the data into the layers and performing inferences about the current contexts in the systems. CI does not require such inferences for all contexts, only for contexts in which items have been recontextualised. Given the limited

Usability Metric	Description	CPS^2	CI
UM1- Minimal Action	the amount of action required to achieve the task	low	high
UM2- Minimal Memory Load	the amount of information the user should have in mind to complete the task	low	high
UM3-Operability	amount of effort required to operate an application	low	high
UM4-Privacy	whether users’ personal information is protected	Yes	Yes
UM5-Security	whether of the application protects information in the system against security threats		depends on the hosting system
UM6-Load Time	time required for the application to load	high	low

Disclosure of a Post: The owner provides values for the various CAP, such as post attributes and the potential audience. The context inference layer infers the context of the situation. The interpretation layer infers possibly a set of interpretations of the post. CPML prompts the owner with the set of possible interpretations to select or add the appropriate interpretation of the post—in case it follows an access control approach. In case it follows an accountability and auditing approach, CPML saves the inferred interpretations from the original context.

Context Evolution: CPML checks changes in contexts and allows only those that continue to preserve the appropriateness of interpretation. CPML either allows the change, or prohibits it or notifies the owner about it.

Recontextualization: When a post is added to a situation, CPML interacts with the interpretation layer to infer the post interpretation. If the new interpretation has not been specified as appropriate, the recontextualization is prohibited.

5.1.2. In Public Contexts

CPS^2 offers contextual privacy management in public spaces. As an example, some Facebook users suffered from privacy violations by the misappropriation of their profile photos—that are by default public—in the incident of 'prostitutes of Antwerp' [13]. Profile photos of some girls were put in a 'prostitutes of city of Antwerp' context. The possible interpretation in the new context negatively affected the identity of the girls and counted as a privacy violation for the girls and was reported to the police information and Facebook [13]. With CPS^2, users can have a certain degree of control when posts are public to avoid inappropriate dissemination. Although, it is not possible to practice total freedom of speech, e.g., sexists and racist comments are legally prohibited comments and are not socially accepted. From an HCI point of view, technologies are required to support freedom of speech as well as practicing one's right to privacy. Our work seeks to explore a more delicate approach to reach freedom of data communication while respecting boundaries of others. In the setting of an accountability and audit approach, our work does not affect freedom of speech. It offers keeping data subjects informed about the usage of their data. Legal rules that specify what the subject's rights are can be integrated into our proposed framework. E.g., the EU General Data protection Regulation 2016/679 allows subjects to rectify inappropriate usage of their data [36].

5.2. Enhancing Privacy in Ambiguous Contexts

CPS^2 can be applied in situations with ambiguous contexts. Consider scenarios of Section 3.

Invisible Owner: In Scenario 2, had Bill specified a protest-related interpretation, despite the anonymity of the post, only comments about the protest would be allowed and the integrity of the intended communication is preserved.

Invisible Subordinate Owner: In Scenario 3, after adding the subordinate owner's comment, the interpretation of the post clearly indicates a protest in which protestors attacked the police. CPS^2 will prohibit Dean's comment.

Invisible Potential Audience: In Scenario 4, had Dean added an interpretation of his comment that the protestors are victims, the act of recontextualising his comment into a context where its interpreted as an encouragement for violation would not have been allowed.

Invisible Actual Audience: In Scenario 4, even if Dean is unaware of when Rex becomes part of the actual audience, by specifying the appropriate interpretation of his message, Dean could avoid the inappropriate recontextualization.

Invisible Extended Audience: In Scenario 5, even if Dean had specified the interpretation of his message, the possible misinterpretation by the extended audience may have happened. In this scenario the audience view the comment and interpret it without performing any specific action on it. In this case, CPS^2 has no effect outside the social software.

5.3. Enhancing Communication

CPS^2 can enhance communication where the interpretation is essential. This type of adversarial communication in scenario 1 can be mitigated by CPS^2. Consider another form of adversarial communication as that of this reported by Boyd [9] (**Scenario 7**): Carmen was sad because she broke up with her boyfriend. She wanted to express that to her friends but not to her mother so that she would not worry. Carmen posted lyrics from "Always Look on the Bright Side of Life" from the film "Life of Brian", where the main character is about to be crucified. She knew that some of her friends would infer her exact communicated message, while her mother would infer a literal meaning of the post. The interpretation is disambiguated based on the knowledge of the audience with the film and not only the online context. This approach is referred to as social steganography [9]. But it is insufficient for contextual privacy management: any of Carmen's friends could comment in a way that reveals an interpretation that Carmen does not want to make explicit. CPS^2 allows the audience to communicate without being concerned that they might reveal an inappropriate interpretation, and become unintentionally adversarial.

6. Related Work

Various works realize the importance of context in privacy management and focus on context-based privacy management approaches. However, most approaches lack the dynamic adaptivity to changes in context [34]. Moreover, most works on context-based privacy management address the complexity of controlling context by

simplifying the representation of context, rather than capturing the essence of the complexity. The simplification of context representation can be seen in various models. In the access control model proposed by Fong [16], relationships between the audience and owner represent contexts. Current social software such as Facebook and Google+ adopt models similar to Fong's. In their implementations, users can pre-specify contexts by specifying groups of friends. Users can disclose their posts to a specific group or context. The contextual privacy approaches based on Nissenbaum's work [30] also require the prohibitive complexity of specifying details in advance. To overcome such complexity, they also simplify contexts and replaces them with roles [7, 24].

There are multiple shortcomings with the simplistic context representations, mentioned above. Firstly, grouping contacts is a time-consuming process [26]. Secondly, changes in friends cannot be reflected easily in the grouping. One model addresses the challenge of the manual grouping of friends by utilizing clustering algorithms to group friends [17]. However, it does not adapt the groupings over time. Thirdly, such a model does not offer protection against recontextualization. Fourthly, empirical studies with Facebook users showed that grouping friends is not relevant to privacy management [23]. In contrast to the models discussed above, our conceptualization of contextual privacy reduces the parameters needed to be controlled without simplifying the representation of context.

7. Conclusion and Future Work

Context is an essential ingredient for communication and privacy management. This article emphasises its role in interpreting posts and for privacy management. By conceptualising contextual privacy as a means to maintain the interpretation of posts, users could manage their privacy in a context-based manner without being faced with the complexity of controlling context through traditional mechanisms. The proposed architecture design using intelligent mechanisms is promising for addressing the complexity of controlling context, overcoming context ambiguity, and enhancing communication. Our future work aims at conducting a user study to test the hypotheses of our concept of contextual privacy and the usability of the framework. Such a study would assist in guiding the realisation of the proposed design.

References

[1] Gregory D. Abowd, Anind K. Dey, Peter J. Brown, Nigel Davies, Mark Smith, and Pete Steggles. 1999.

Towards a Better Understanding of Context and Context-Awareness. In Proceedings of the 1st international symposium on Handheld and Ubiquitous Computing. Springer-Verlag, 304–307.

[2] A. Acquisti, E. Balsa, B. Berendt, D. Clarke, W. De Groef, R. De Wolf, C. Diaz, B. Gao, S Gurses, J.¨ Pierson, F. Piessens, R. Sayaf, T. Schellens, F. Stutzman,

B. Van Alsenoy, and E. Vanderhoven. 2011. SPION Deliverable 2.2–Requirements and Conceptual Framework. Technical Report. KULeuven.

[3] Nabil Ajam, Nora Cuppens-Boulahia, and Frederic´ Cuppens. 2010. Contextual privacy management in extended role based access control model. In

Proceedings of the 4th workshop, and 2d conference on Data Privacy Management and Autonomous Spontaneous Security. Springer, 121–135.

[4] V. Akman. 2000. Rethinking context as a social construct. Journal of Pragmatics 32, 6 (2000), 743–759.

[5] A. Analyti, M. Theodorakis, N. Spyratos, and P. Constantopoulos. 2007. Contextualization as an independent abstraction mechanism for conceptual modeling. In Information Systems, Vol. 32. Elsevier, 24–60.

[6] Itamar Arel, Derek C Rose, and Thomas P Karnowski. 2010. Deep machine learning-a new frontier in artificial intelligence research. Computational Intelligence Magazine, IEEE 5, 4 (2010), 13–18.

[7] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum.

2006. Privacy and contextual integrity: Framework and applications. In IEEE S& P'6. IEEE Computer Society, 184–198.

[8] D.M. Boyd. 2008. Taken out of context: American teen sociality in networked publics. Ph.D. Dissertation. PhD Dissertation. University of California-Berkeley, School of Information.

[9] Danah Boyd and Alice Marwick. 2011. Social steganography: Privacy in networked publics. International Communication Association, Boston, MA (2011).

[10] C. Braz, A. Seffah, and D. M'Raihi. 2007. Designing a trade-off between usability and security: a metrics based-model. In HCI–INTERACT 2007. Springer, 114–126.

[11] P. Brezillon. 1999. Context in problem solving: a survey.´ In The Knowledge Engineering Review, Vol. 14. Cambridge University Press, 1–34.

[12] Lorrie Faith Cranor. 2005. Security and usability: designing secure systems that people can use. " O'Reilly Media, Inc.".

[13] R. De Wolf. 2013. Over ‘spotted’, ‘hoeren’ en ‘failed’-pagina’s. Electronic article. (2013).

http://www.knack.be/nieuws/belgie/ dader-antwerpse-hoeren-foto-geklist/ article-4000230766578.htm

[14] J.M. Dimicco and D.R. Millen. 2007. Identity management: multiple presentations of self in Facebook. International ACM conference on Supporting group work (2007), 383–386.

[15] Marta Dynel. 2008. There is method in the humorous speaker’s madness: Humour and Grice’s model. Lodz Papers in Pragmatics 4, 1 (2008), 159–185.

[16] Philip W L Fong. 2011. Relationship-based access control: protection model and policy language. In Proceedings of the first ACM conference on Data and application security and privacy (CODASPY 11). ACM, New York, NY, USA, 191–202.

[17] B. Gao, B. Berendt, D. Clarke, R. De Wolf, T. Peetz, J. Pierson, and R. Sayaf. 2012. Interactive Grouping of Friends in OSN: Towards Online Context Management. International Workshop on Privacy in Social Data (PinSoDa) (2012).

[18] Erving Goffman. 1959a. The presentation of self in everyday life. Garden City, NY Double Day (1959).

[19] Erving Goffman. 1959b. The presentation of self in everyday life. Garden City, NY (1959).

[20] Herbert P Grice. 1975. Logic and conversation. In The Logic of Grammar, Donald Davidson and Gilbert Harman (Eds.). Harvard Univ., 64–75.

[21] Seda Gurses and Claudia Diaz. 2013. Two tales of˜ privacy in online social networks. IEEE Security & Privacy 11, 3 (2013), 29–37.

[22] Geoffrey Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. Neural computation 18, 7 (2006), 1527–1554.

[23] Patrick Gage Kelley, Robin Brewer, Yael Mayer, Lorrie Faith Cranor, and Norman Sadeh. 2011. An investigation into facebook friend grouping. In HCI–INTERACT 2011. Springer, 216–233.

[24] Yann Krupa and Laurent Vercouter. 2012. Handling privacy as contextual integrity in decentralized virtual communities: The PrivaCIAS framework. Web Intelligence and Agent Systems (2012).

[25] Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. 2011. We’re in it together: interpersonal management of disclosure in social network services. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, 3217–3226.

[26] Heather Richter Lipford, Andrew Besmer, and Jason Watson. 2008. Understanding privacy settings in facebook with an audience view. In Proceedings of the 1st Conference on Usability, Psychology, and Security. USENIX Association, Berkeley, CA, USA, Article 2, 8 pages.

[27] J. McCarthy. 1993. Formalizing context (expanded notes). In Computing Natural Language, A. Aliseda, R. J. van Glabbeek, and D. Westerstahl (Eds.). CSLI Publications, 13–50.

[28] Ian Mcculloh. 2009. Detecting changes in a dynamic social network. Ph.D. Dissertation. Carnegie Mellon University. Advisor(s) Carley, Kathleen M.

[29] J. Meyrowitz. 1985. No sense of place: The impact of electronic media on social behavior. Oxford University Press New York.

[30] H.F. Nissenbaum. 2010. Privacy in context: Technology, policy, and the integrity of social life. Stanford Law & Politics.

[31] Leysia Palen and Paul Dourish. 2003. Unpacking

”privacy” for a networked world. In Proceedings of the SIGCHI conference on Human factors in computing systems. ACM, 129–136.

[32] Zizi Papacharissi. 2011. A networked self. A networked self-identity, community, and culture on social network sites (2011), 304–318.

[33] S.S. Petronio. 2002. Boundaries of privacy: Dialectics of disclosure. SUNY Press.

[34] R. Sayaf and D. Clarke. 2012. Access control models for online social networks. Social Network Engineering for Secure Web Data and Services (2012), 32–65.

[35] Rula Sayaf, Dave Clarke, and Richard Harper. 2014. CPS 2: a Contextual Privacy Framework for Social Software. In SECURECOMM 2014. Springer.

[36] R. Sayaf, J. B. Rule, and D. Clarke. 2013. Can Users Control their Data in Social Software? An Ethical Analysis of Data Control Approaches. In IEEE S& P Workshops (SPW). 1–4.

[37] A. Seffah, M Donyaee, R. Kline, and H.K. Padda. 2006. Usability Metrics: A Roadmap for a Consolidated Model. Journal of Software Quality (2006).

[38] Brian Skyrms. 2011. Pragmatics, logic and information processing. In Language, games, and evolution. Springer, 177–187.

[39] N. Srivastava and R. Salakhutdinov. 2012. Multimodal learning with deep boltzmann machines. In Advances in neural information processing systems. 2222–2230.

[40] Teun A Van Dijk. 2008. Discourse and context. A Sociocognitive Approach, Cambridge University (2008).

[41] Rineke Verbrugge and Lisette Mol. 2008. Learning to apply theory of mind. Journal of Logic, Language and Information 17, 4 (2008), 489–511.

[42] Andrew F Wood and Matthew J Smith. 2004. Online communication: Linking technology, identity, & culture. Routledge.