Persona: A Tool for Managing Identities

William Broniec
Bryan Bennett
Peter McAughan
Ashok Krishna
wbroniec3@gatech.edu
bbennett37@gatech.edu
pmcaughan6@gatech.edu
akrishna43@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

ABSTRACT

We present *Persona*, a browser extension which – much like password managers – tracks accounts and user's passwords. Additionally – in an attempt to help users avoid contextual collapse common to the modern Internet – *Persona* tracks the top-level identity to which these accounts and passwords belong and manages an "active" persona, which serves to remind the user into thinking carefully about identity and its relationship with the current logged-in web session. Further, we discuss design principles and present the results of a qualitative user study suggesting the advantages of such a tool.

CCS CONCEPTS

• **Security and privacy** → **Social network security and privacy**; **Privacy protections**; **Usability in security and privacy**; *Social aspects of security and privacy*;

KEYWORDS

Privacy, Pseudonymity, Usability

INTRODUCTION

As of June 2018, roughly 4.2 billion people possess access to the Internet around the world [4]. This radical uptick in online inter-connectivity has created unforeseen possibilities for interpersonal communication and spurred the existence of communities such as Facebook, Reddit, and Twitter with global user bases and sizes inconceivable prior to the Internet.

Online communities possess many unique affordances and challenges, one of the more controversial being identity. Though true anonymity in the Internet is quite difficult, if not impossible, to achieve with even tools such as Tor being ultimately traceable [10], the Internet does provide some veil of protection against other engaged users and immediate accountability differs drastically from face-to-face interactions. Many dominant organizations on the Internet such as Google and Facebook have fought to create a one to one correspondence between physical people and virtual personas – often referred to as the "real name" Internet – in which there exists a single, unique link between physical persons, online personas and legal name usage. Here we use the term persona to refer to a linked collection of online accounts, either implicit or explicit.

Social networks often provide these means of linking accounts, and the advantages of such a de-anonymized system have been described at length, such as ease in locating other users, accountability, and parallelism to typical offline social interaction. However, such a system has also concerned privacy advocates since the inception of the Internet as exploitable and undermining the genuine value in the pseudo-anonymity and pseudonymity provided online. We stand with these privacy advocates in recognizing the value of partial anonymity online, but as more entities on the Internet have pushed for systems realizing the "real name" Internet, current solutions of online privacy are inadequate in providing user-friendly means for maintaining and managing different personas.

RELATED WORK

Though both with histories long predating the Internet [13], pseudonymity and anonymity have long been recognized as contributing unique value to online communication. Ma, Hancock, and Naaman [9] demonstrated that given a variety of contexts, users disclosed the most information when they were most anonymous, second only to contexts of users with close social ties. Both Ribiero [12] as well as van der Nagel, Frith [15] likewise illustrate that virtual identities empower users to pursue new avenues of personal development. Part of this is a result of protection from identification and thus a lack of social pressures found in the offline world. Hillier, Horsley, and Kurdas [5] explored this avenue specifically in the context of SSA (Same-Sex Attracted) youth, finding that 70% of those interviewed found the Internet as reducing their social isolation and providing support. Though this is a small demographic, the results are eminently generalizable to the broader population. A study from 2015 in Britain showed that 63% of people are using the Internet to search for personal medical information

[1]. Traced back to the original user, this information could be used to profile and discriminate against real people in unlawful or unethical manners.

Thus, public divulgence of certain aspects (medical conditions, socially stigmatized traits) tied back to an identifiable person can have permanent and meaningful consequences to users across a variety of scenarios. This type of context collapse is less emphasized in the existing literature compared to collapse between different circles of the same identity (e.g., friends versus coworkers), but it is evident that the obfuscation of a single online persona and scaffolding of multiple have both naturally arisen in society as a means for users to express themselves and divulge sensitive information. The tenuity of existing systems to deliver fully on perceived private actions has been well documented. In 2006, AOL released "anonymized" search results which researchers were easily able to correlate back to end users [2]. Likewise, in 2017 a similar scenario occurred when CNN traced an online Reddit account back to its owner via analyzing post history [6]. In either of these cases, the use of different personas could protect a user from identification. While modern research has involved machine learning of syntactic style as a countermeasure linking personas accounts to a single identity [3], research has still recognized pseudonymity as a valid defense strategy coupled with others to avoid context collapse [11]. Past research has recognized the potential for login systems to specifically support for pseudonymous accounts [8], but we see a tool existing outside of any one system as far more flexible and useful.

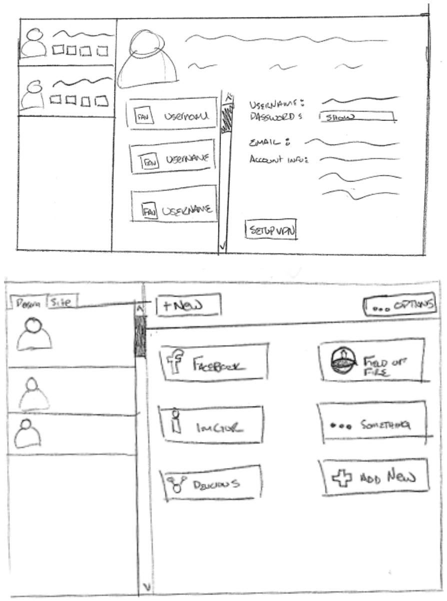

Figure 1: Two early iterations on *Persona*'s potential interface

TOOL DESIGN

Initially, *Persona* was conceived and partially developed as a desktop application targeting users of similar tools, such as LastPass. However, feedback led the team to develop a Google Chrome plugin using the same data schema and general approach as the desktop version. This refocus allows greater flexibility and hones in on a larger user base.

Persona's top-level interface presents a user's extant *Identities*. *Identities* represent a grouped set of accounts and are given a title, an email, and (optionally) an expiry time stamp. Clicking on an *Identity* presents the user with her *Accounts*, which function identically to accounts in a password manager. Clicking an account allows the user to edit the account, or copy the account's password to the clipboard.

The novelty in *Persona* is in the manner in which Identities are tracked. Clicking on an Identity presents the user with the ability to mark the Identity as active. Marking an Identity as active will cause any other Identity to become logged out and allow the user to sign-in as the active Identity, which makes it much more difficult to stay logged in as an incorrect Identity. This feature is intended to prevent the user from inadvertently giving away their "main" account by forgetting to log out of an Identity manually.

Being a prototype, however, there are many ways in which *Persona* can stand to improve. The user interface is in need of a redesign, as noted by at least one participant ("Why not check boxes? I mean...instead of having to click into the Identity?", noted Participant 5.) The development team also decided to focus on the utility of the idea of Identity management rather than implement top-tier security. To that end, passwords are stored in plain-text in the prototype and there exists no way to lock down the system with a master password.

EVALUATION

There are a variety of ways to test the success of our identity management system. Ideally, we would be able to conduct a long-term diary study in which we determine if the identity manager gets organically adopted by our subject pool as part of their regular Internet use. However, given the time constraints of the semester schedule, we have decided to conduct sessions that consist of interview, observation, and think-aloud elements using an initial UI prototype. The interview captures information as to the subject's perceptions of online pseudonymity and anonymity, as well as their current practices of identity management. After asking users to accomplish some tasks with our tool, a think-aloud session captured additional thoughts the user had during use. Given more time we would have liked to incorporate quantitative measures into these interactions with our tool, but given constraints we utilized qualitative measures over quantitative ones to determine whether our system meets user needs and conforms to their existing behavior.

To begin our assessment, we launched a needs-finding survey with a primary subject pool of masters students in the Georgia Tech Human and Computer Interaction Master's program. This needs-finding survey focuses on the familiarity of users with throwaway accounts as well as user sentiment regarding the usefulness of throwaway accounts. We received eleven responses to this survey and used this information in order to refine our prototype.

After building our prototype, we recruited 6 participants (4 males, 2 females, ages 19-35, mean age of 23.3) from the Georgia Tech student body. Though this sample is not representative of the general population, this is not problematic because *Persona* is aimed at people who actively use multiple online properties and Georgia Tech students fall within that group. We did our best to ensure an even gender distribution in our sample even though we did not utilize gender in our analysis. This is to ensure that our sample is diverse on at least one dimension.

Each session took roughly 60 minutes to complete. First, we asked participants about their perceptions of context collapse or linkage of their in-person identity with their online activity. This introduction both gave us information about how our participants valued this context separation, and also introduced the idea of a 'persona' to the participant. Responses to these and further questions were audio recorded for analysis.

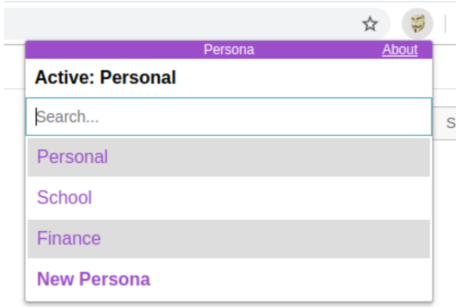

Figure 2: *Persona's Identity* view

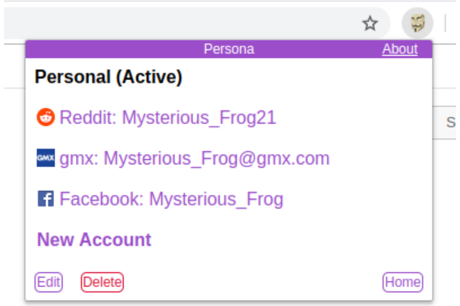
Persona
About
Personal (Active)
Reddit: Mysterious_Frog21
gmx: Mysterious_Frog@gmx.com
Facebook: Mysterious_Frog
New Account
Edit
Delete
Home

Figure 3: *Persona's Account* view

Participants were first walked through a number of pre-existing personas in the tool; one for personal use and one for questions on a finance forum. The user was then given username and password data for an online account and were asked to use the tool to create a new persona for private use, and connect this account data to the persona. Afterwards, the participant was asked to set this persona as the "active" persona to change browser interactions. As a secondary task, users were given a scenario where they were motivated to ask a personal question on a public forum in a way where personal information could not be connected to their personal identity. Users then used our tool to select the persona they wanted to ask the question under, use the auto-fill function of our tool to login to the connected account, ask a question, and delete the persona information from the computer. These tasks gave users a practical understanding of the fundamental functionality of our identity manager.

After completing these tasks, users were asked about their perceptions of the tool; particularly if they found the tool useful and usable. Criticism of the tool and opinions about the importance of the tool were discussed and recorded.

After the think-aloud session, we asked the subject to perform the System Usability Scale for the identity-manager interface. This was a very quick wrap-up process that took around five minutes. Using this scale provides us with an externally validated measure of usability, which will complement our think-aloud and interview observations.

RESULTS

We asked our participants to rate how concerned they were about their online activities being linked to their real-world identity on a seven point scale (1 being least concerned, 4 being neutral, and 7 being most concerned). Our mean score for this rating was 4.7, with a minimum rating of 3 and a maximum rating of 6. This indicates that are user group was somewhat concerned about the scenario *Persona* is meant to protect against.

In using the tool, participants required very little to no interaction from the researchers to accomplish the required interactions with our tool, and after completing the tasks user feedback towards *Persona* was largely positive. Four of our six subjects said they would use the tool if it were available on the market. The two who said they would not use the tool based their comments on a lack of personal need for it rather than a flaw in the tool's functionality, and described the importance of such a product. One participant (S4) said he would use the tool even if he had to pay to use it.

Negative user feedback was primarily targeted at the "Mark as active" feature. However, the negative feedback was due to user confusion as to what the text meant as opposed to the actual functionality of the feature. This indicates that the feature needs to be renamed or reworked rather than removed entirely. One possibility is to change the wording of "active" to "on/off", as "on/off" is the language used by many existing applications. However, further research must be performed to determine the wording that is most intuitive to users.

Our SUS data indicates are tool is highly usable. We obtained a mean SUS score of 91.7 (minimum: 85, maximum: 97.5) which corresponds to an "excellent" usability rating. [14]. Though we are pleased with the SUS scores, we recognize that there still usability improvements that can be made. One participant (S6) suggested incorporating a short tutorial similar to that of existing mobile apps. This could enable users to understand the intention behind the Persona tool and allow users to apply to tool for the use cases they deem important.

DISCUSSION

The scope of our user study was only that of a pilot evaluation, but we can draw some initial conclusions from the results. The high usability scores and perceptions of our tool likely stem from a simple design. Users never have to go through more than a couple clicks to achieve a desired function, and the form of a browser extension results in an easily accessible location whenever a user is engaged in online activity. These generally validate our design approach to move away from a desktop application and towards a browser extension, trying to make it as easy as possible to facilitate a separation of online contexts. There are common criticisms across the user feedback that are worth implementing, but our initial prototype satisfies the foundations of our initial vision.

There was an interesting discrepancy between a user's perceptions of the necessity of our tool and their own disposition towards using it. Subjects as a whole felt somewhat neutral to the importance of maintaining multiple separate identities in their own lives, but most subjects said they'd use our tool that specifically accomplishes this function. We hypothesize that this is because as users are exposed to a concrete tool and use case, they begin to realize the potential and importance of online context separation in their own lives. To further explore the validity of this hypothesis and the previously mentioned results, we'd like to conduct a larger scale study.

During our subjective feedback collection process, a small number of users raised concerns of accountability and potential negative actions that could be carried out easier through their usage of *Persona*. Online trolls often rely on fake accounts to conduct behavior without repercussion [7] and providing users greater convenience in managing transient accounts might expedite this process. While our literature survey highlighted the variety of positive opportunities created by this situation, this negative viewpoint must also be considered.

However, while *Persona* protects against identification by other users of a given site, it provides no protection against privileged users, administrators, nor against automated systems such as free-text comparison tools which analyze a user's writing style. This could very well be seen as a weakness, as it could lead to unintended divulgence of information, but the authors contend that this problem simply falls outside the scope of *Persona* and that this problem has already been explored by existing tools such as Virtual Private Networks.

CONCLUSION

We believe our study is an important contribution to the growing study of online context collapse and pseudonymity, and that our developed tool is a benefit for privacy-concerned users. As Internet communication grows, debates in the field of online identity and the "real name" Internet will continue. Oftentimes negative behavior such as hacking and trolling are cited as major drawbacks to tools encouraging Internet anonymity and should be recognized as serious consequences. However, previous studies have clearly demonstrated real value in the unique and less personal interactions which online pseudonymity enables. Our tool and study serve to explore how internet users perceive and use online identities and provide them a usable software for managing these personas in a manner unlike any other existing tool.

REFERENCES

[1] Aviva. *Aviva Health Check UK Report*. Nov 2015.

[2] Barbaro, M., and Zeller Jr, T. A face is exposed for aol searcher no. 4417749. *The New York Times* (Aug 2006).

[3] Braunlin, J. Using nlp to identify redditors who control multiple accounts.

[4] Group, M. M. Internet world stats - usage and population statistics, 2019.

[5] Hillier, L., Kurdas, C., and Horsley, P. “it’s just easier” the internet as a safety-net for same sex attracted young people.

[6] Kaczynski, A. How cnn found the reddit user behind the trump wrestling gif, Jul 2017.

[7] Klempka, A., and Stimson, A. Anonymous communication on the internet and trolling. *Concordia Journal of Communication Research* (2014).

[8] Leavitt, A. This is a throwaway account. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW ’15* (2015).

[9] Ma, X., Hancock, J., and Naaman, M. Anonymity, intimacy and self-disclosure in social media. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI ’16* (2016).

[10] Murdoch, S., and Danezis, G. Low-cost traffic analysis of tor. *2005 IEEE Symposium on Security and Privacy (S&P ’05)* (2005).

[11] Rao, J., and Rohatgi, P. Can pseudonymity really guarantee privacy? *Proceedings of the 9th conference on USENIX Security Symposium - Volume 9* (2000).

[12] Ribiero, J. C. The increase of the experiences of the self through the practice of multiple virtual identities. *PsychNo Journal 7*, 3 (2009), 291–302.

[13] Stang, C. M. *âĂIJNo longer IâĂİ: Paul, Dionysius the Areopagite, and the apophasis of the self*. Harvard Divinity School, 2008.

[14] User Experience Magazine. Determining what individual sus scores mean: Adding an adjective rating scalejus, 2009.

[15] van Der Nagel, E., and Frith, J. Anonymity, pseudonymity, and the agency of online identity: Examining the social practices of r/gonewild. *First Monday 20*, 3 (2015).