Bringing the National Security Agency into the Classroom: Ethical Reflections on Academia-Intelligence Agency Partnerships

Christopher Kampe, *North Carolina State University*
Gwendolynne Reid, *Emory University*
Paul Jones, *North Carolina State University*
Collen S., *North Carolina State University*
Sean S., *North Carolina State University*
Kathleen M. Vogel, *University of Maryland*

Journal Title: Science and Engineering Ethics

Publisher: Springer Link | 2018-01-09

Type of Work: Article | Post-print: After Peer Review

Publisher DOI: 10.1007/s11948-017-9938-7

Permanent URL: https://pid.emory.edu/ark:/25593/vjbgx

Final published version:
https://link.springer.com/article/10.1007/s11948-017-9938-7

Copyright information:

Bringing NSA into the Classroom: Ethical Reflections on Academic-Intelligence Partnerships

POST-PRINT JUNE 2017, DOI: 10.1007/s11948-017-9938-7

Christopher Kampe, Gwendolynne Reid, Paul Jones, Colleen S., Sean S., & Kathleen M. Vogel

Author Note

Christopher Kampe, Communication Rhetoric and Digital Media Program, North Carolina State University 27695, Raleigh, NC; Gwendolynne Reid, Communication Rhetoric and Digital Media Program, North Carolina State University, Raleigh, NC 27695; Paul Jones, Laboratory for Analytic Sciences, North Carolina State University, Raleigh, NC, 27695; Colleen S., Laboratory for Analytic Sciences, North Carolina State University, Raleigh, NC 27695; Sean S., Laboratory for Analytic Sciences, North Carolina State University, Raleigh, NC 27695; Kathleen M. Vogel,** School of Public Policy, University of Maryland, College Park, 2101 Van Munching Hall, University of Maryland, College Park, College Park, MD 20742, email: kvogel12@umd.edu

**Corresponding author

Acknowledgments: This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government. The author wishes to thank the reviewers of this paper for insightful comments and suggestions for revisions.

Bringing NSA into the Classroom: Ethical Reflections on Academic-Intelligence Partnerships

"When, you know, your professor walks into class the first day for a security course and she says, by the way, the NSA is going to monitor all your research for the next, you know, four or five months, it's a little unnerving." –Amber, graduate student (Interview, Raleigh, NC, May 3, 2016)[1]

Abstract

Academic-intelligence collaborations are on the rise for a variety of reasons. These can take many forms, one of which is in the classroom, using students to stand in for intelligence analysts. Classrooms, however, are ethically complex spaces, with students considered vulnerable populations, and become even more complex when layering multiple goals, activities, tools, and stakeholders over those traditionally present. This does not necessarily mean we must shy away from academic-intelligence partnerships in classrooms, but that these must be conducted carefully and reflexively. We hope to contribute to this conversation by describing one purposeful classroom encounter that occurred between a professor, students, and intelligence practitioners in fall 2015 at North Carolina State University: an experiment conducted as part of a graduate-level political science class that involved students working with a prototype analytic technology, a type of participatory sensing/self-tracking device, developed by the National Security Agency. This experiment opened up the following questions that this paper will explore: What social, ethical, and pedagogical considerations arise with the deployment of a prototype intelligence technology in the college classroom, and how can they be addressed? How can academic-intelligence collaboration in the classroom be conducted in ways that provide benefits to all parties, while minimizing disruptions and negative consequences? We discuss our experimental findings in the context of ethical perspectives involved in values in design and participatory/self-tracking data practices, and discuss lessons learned for the ethics of future academic-intelligence partnerships in the classroom.

Keywords: Intelligence, prototype, research ethics, participatory sensing, self-tracking, values in design

Introduction/Background

Since the mid-2000s, various government and non-government reports have urged intelligence analysts to collaborate with outside parties (especially academics) in order to further develop their craft and improve intelligence products (Nolan 2013; U.S. National Research Council 2011; Office of the Director of National Intelligence 2008; Miller 2008; Treverton 2008; The Commission on the Intelligence Capabilities of the United States 2004). In 2008, the Office of the Director of National Intelligence issued ICD 205, which states that intelligence analysts should leverage "outside expertise as part of their work" to "explore ideas and alternative perspectives, gain new insights, generate new knowledge, or obtain new information" (Office of the Director of National Intelligence 2008). In 2011, the U.S. National Research Council published a report, *Intelligence Analysis for Tomorrow: Advances from the Behavioral and Social Sciences*, encouraging the Intelligence Community to embed analysts in academic research settings in order to "participate in research and to network with [social and behavioral]

[1] Note: Pseudonyms are used for the names of all interviewees listed in this paper.

scientists who can be consulted later," while simultaneously expanding "opportunities for continuous learning that will enhance collaboration, innovation, and growth in the application of [social and behavioral science] analytical skills" (U.S. National Research Council 2011: 85;88).

In response to these calls and its own needs, the National Security Agency (NSA) established a five-year partnership with North Carolina State University (NC State) in 2013, to create the Laboratory for Analytic Sciences (LAS) (NSA 2013). The mission of LAS is to create new collaborations between the NSA, academia, and industry to address challenging big-data problems involving intelligence collection and analysis. Pursuant to this mission, LAS funds interdisciplinary research teams to design, develop, and refine cutting-edge technologies intended to assist intelligence analysts with their work. Currently, approximately 50 analysts from the NSA and other intelligence agencies have relocated from their offices on multi-year assignments to work alongside academic and industry partners in LAS teams at NC State on defined research problems. This article concerns one such team and the prototype ("Journaling application"), which was brought into a graduate political-science classroom. This collaboration provides a unique opportunity to examine a variety of ethical issues surrounding the use big data analytic prototypes involving a professor, students, and the NSA. The research questions this paper will explore are: What social, ethical, and pedagogical considerations arise with the deployment of a prototype intelligence technology in the college classroom, and how can they be addressed? How can academic-intelligence collaboration in the classroom be conducted in ways that provide benefits to all parties, while minimizing disruptions and negative consequences? We discuss our experimental findings in the context of ethical perspectives involved in values in design and participatory/self-tracking data practices, and discuss lessons learned for the ethics of future academic-intelligence partnerships in the classroom.

The Journaling Application: A Participatory Sensing Device for Intelligence Analysts

The LAS is broadly interested in technologies that could help intelligence analysts to better understand how individuals and teams make sense of information, accomplish tasks, and develop analytic workflows with big data—knowledge that could be used to improve analytic tradecraft and confidence in intelligence products. The "Journaling application" (hereafter, "Journaling") prototype was created in 2014 at LAS by computer scientists from the NSA and other intelligence agencies, as well as technical experts from RENCI, a North Carolina-based technology institute. It is a type of participatory sensing or self-tracking device, which gathers user-supplied data about their computer usage (Shilton 2009; Shilton 2010; Neff and Nafus 2016). In the case of Journaling, it works as a web-based computer application installed on laptops or desktop computers and records user work patterns (i.e. resources used, duration, frequency, periodicity); these data can be used to automatically generate activity reports about user workflows (see Figure 1). The prototype requires that a user tag a source/inquiry, with the explicit purpose of linking a resource to an operation within a larger process—with the stated intention of allowing users to see how they have allocated their attention, and quickly identify resources associated with prior work.

Fig. 1 The "recent work dashboard" displays user-viewed resources and the tags manually/automatically applied to them; users can sort by tags (i.e. stages of the project)

Beyond capturing simple time on task, the end goal of Journaling is to understand the (goal-based) activities users are engaging with, while they leverage other resources and software (Jones et al. 2016). In the future, when paired with novel recommendation software (a separate LAS R&D project), Journaling might suggest resources, e.g., datasets, articles, or different strategies for approaching a problem to help intelligence analysts produce threat assessments (Jones et al. 2016). Although Journaling currently requires a certain amount of active tagging of user activities, in the future, it is anticipated that the application will learn user work patterns through artificial intelligence means, and be able to operate and suggest resources and analytic workflows with minimal or no direct user input.

To use Journaling, the user installs a web-browser plug-in and activates it. The application then records the user's browser-based research activities and is prompted to associate relevant resources with a defined list of goals and tasks related to a specific project. Put another way, the user first selects a larger goal (e.g. Security Threat Analysis on a specific topic), and then selects a specific activity within that goal (e.g. Survey Available Data on the topic); as the user accesses online resources (e.g. reports, articles, Google documents), the application tags those resources with the activity and goal (See Figure 2).

Fig. 2 Photo of the Journaling activity/goal tree (i.e. larger activity made of up more discrete tasks/stages)

Although the application is currently used in an unclassified environment, the longer-term project goal is to develop the application for deployment in classified spaces, to help intelligence analysts work with both classified and unclassified data. Journaling could assist users in two ways: first, by providing an account of their work activities, which would enable them to see where their time/efforts have been directed and where to improve their analytic efforts; second, by enabling supervisors with access to these reports to see what their teams were working on and where they were experiencing problems/slow-downs, without requiring team progress reports.[2]

To improve Journaling functionality, the designers/programmers need users to generate data, provide usability feedback, and offer suggestions for improvement. In its earliest stage, Journaling was deployed by an LAS designer, also a PhD Candidate in Computer Science, in a graduate seminar and data collected about how computer science students solved problem sets. Though the gathered data was useful as a pilot study, the researchers were concerned about their applicability, as the activities of computer-science students might differ greatly from those of intelligence analysts, who work on more complex and diverse data sets and topics. Thus, in 2015, the Journaling development team began looking for another pool of participants at NC State on which to experiment—ideally, a class in which the students would perform research and writing activities analogous to those of intelligence analysts.

[2] In its initial design, the Journaling focused on a single user but has since been augmented to support collaborative workflows (e.g. in which two or more analysts can see each other's work, share data, and collaborate on a research problem). See (Jones et al. 2017).

Starting a Collaboration

In summer of 2015, one of us ("A")[3] was restructuring his graduate course, "Science, Technology, and National Security," which examines the roles of science, technology, and technical knowledge production as they relate to threats from conventional (e.g. missiles, bombs) and unconventional (e.g. nuclear, chemical, and biological) weapons. A hybrid of Political Science and Science and Technology Studies, the course encourages students to examine security concerns from a variety of perspectives: social, political, economic, and technological. The culminating assignment for this course (which became the focus of our study) was a Security Threat Analysis (20-30 pages in length). For this assignment, students were separated into groups; each group was required to select a contemporary security threat and examine it using the perspectives/methods taught in the course. When "A" taught this course previously (Fall 2014), he observed that students had difficulty engaging with security issues from multiple perspectives at once—rather, they tended to disregard social factors and focus exclusively on technological factors (e.g. materials, technical information, equipment) for their security assessments. Consequently, before teaching the course again, he wanted to construct a new type of assignment and pedagogical approach that would prevent, or at least mitigate, these analytic shortcomings. Through a mutual LAS colleague, "A" met with the lead technology developer of Journaling, and the two decided to collaborate on an experiment of mutual interest and benefit.[4] Ultimately, we created a six person interdisciplinary research team for the experiment involving individuals with backgrounds in computer science, engineering, science and technology studies, intelligence, and communication, rhetoric, and digital media studies.

At the beginning of this collaboration, we understood how Journaling gathered data and what data the programmers would require to further develop the application, but we did not have a clear sense of how that data gathered could be transformed into something pedagogically useful. "A" met with "B"[5] (Programmer) and "C"[6] (Usability/Documentation expert) to discuss how Journaling could be modified and meaningfully integrated into the coursework, while simultaneously gathering use data to help the programmers refine the application. Drawing from his previous experience with this class, "A" initially established two goals for the application and course:

1. Proactively encourage students to examine security threats from multiple analytic perspectives (e.g. social, economic, political, technological);
2. Assist in the identification of student learning/analytic problems (e.g. confusion, misdirected efforts, gaps) prior to submission of assignments.

Given that Journaling was still in early development, it was doubtful that we would accomplish the second goal; nonetheless, we were intrigued by its underlying concern, and started to thus orient the study. We quickly realized that Journaling, by itself, would be unable to accomplish any of this—rather, it would have to be deployed with specially tailored instructional

[3] Name removed for anonymity.
[4] Note: At the time of the experiment, "A" was involved in separate research project for LAS; during the course of the experiment LAS provided "A" with funding to support graduate student assistance on the project.
[5] Name removed for anonymity.
[6] Name removed for anonymity.

material and validated by external instruments. Consequently, we directed our efforts for the experiment towards three objectives:

1. Developing a well-documented, major assignment to methodically guide students through analysis of a security threat
2. Developing companion instruments (i.e. surveys, simplified evaluations, interview questions) to productively interpret Journaling-derived data and student experiences with the technology
3. Customizing the Journaling interface to better support/scaffold student work

Here, it is worth noting that these objectives are all instrumental in their focus—underpinned by the logic that we could develop an application that would assist students and instructors. We do not concern ourselves with the success or failure of the Journaling application in this regard; rather, in this paper, we focus on what was learned ethically during the set up and conduct of the classroom experiment.

Research Context

Reconciling Goals: Development and Education

For better or worse, the university classroom can be used as a testbed for new technologies. Although this practice can lead to enriched student experiences, it can just as easily lead to confusion and distraction (e.g. if the technology is unsuitable or unfamiliar to the instructor). These problems are compounded when students are asked to interact with prototypes undergoing revision (as bugs are found or new features are integrated). Furthermore, even stable, tested prototypes still require training/tutorial sessions (which take away from instruction time) and conscious use on the part of the students. By "conscious use," we mean not merely the physical effort required to use the application, but also the conceptual effort of thinking within the terms allowed by the application. As researchers, we respect the primacy of education in the classroom, and did our best the mitigate these problems, but we must self-reflexively acknowledge that projects such as this one create two separate agendas: student education and prototype development. As instructors, our obligation is to the students and the furthering of their education/professional development; as research collaborators with intelligence practitioners, it was to assist developers in refining a prototype purposed with helping intelligence analysts do their job more efficiently and effectively. We do not offer a solution to this tension, but we propose a method to mitigate it. This focus allows us to foreground the benefits/complications for instruction and academic-intelligence collaboration that emerged from our partnership with the LAS and the NSA.

The goal of LAS—and by extension, that of its software—is to assist intelligence analysts in doing their jobs more efficiently and effectively. To refine Journaling, the designers needed to put it in the hands of participants who were "more analogous" to intended users. "A's" course, because it concerned national security issues and required students to produce a Security Threat Analysis, proved to be one of the best local analogs for this. But it was by no means perfect: despite subject matter similarities, the lived-work of an intelligence analyst is different from the lived-work of a student. Respecting the primacy of education, we considered how and where the application might fit into the course, without imposing an undue burden on students. "A" conceptualized the course as having two parts:

I. Theory-oriented (i.e. familiarizing the students with theory, technique, methodological approaches, and historic instances relevant to the course)
II. Practice-oriented (i.e. guiding the students through the production of a Security Threat Analysis)

We concluded that Journaling, due to its limitations, would not have clear benefits for (I), as it did not conceptualize the act of "becoming familiar with a field." However, we conceived of (II), the Security Threat Analysis, implicitly contained discrete practices we could scaffold for the students which would teach students how to do a threat assessment, as well as allow for gathering data useful for the Journaling developers. Consequently, we focused our attention on (II), and worked to tailor the application for the Security Threat Analysis, which students would complete in the latter half of the semester. Therefore, we approached the design of the experiment from two ends:

I. Working with the programmers to modify the application so that it could productively guide students through the practice of an intelligence researcher
II. Working with the instructor, so that the class was constructed in a manner compatible with the recording/ordering capabilities of the Journaling application

Working with the Programmers, we used Journaling to scaffold the various steps associated with Intelligence Analysis: We drew extensively from the "Generic Analytic Workflow" developed by Dhami and Careless (2015), which breaks down an intelligence analysts' workflow into six distinct stages. The stages themselves, though presented in a set order, can be returned to or later iterated upon after the analyst moves on to a later stage. The guiding logic is that one should not skip earlier stages, though one is encouraged to return to them as understanding of the domain increases. We performed several, minor alterations to better pair it with the classwork. We note that the students did not (in their tagging habits) return to earlier stages in later weeks (perhaps the fault of our design)—their tagging progression was unidirectional. Table 1 shows the "stages" students could select (within the Journaling application).

Stage	Sub-Tasks
1. Understand Requirements	Research Customer
	Research Problem
2. Plan Response	Articulate Hypothesis
	Survey Available Data
3. Obtain Data	Exclude Data (Irrelevant)
	Include Data (Relevant)
4. Process Data	Identify Trends in Data
	Summarize Sources
5. Compose Interpretations	Establish Limitations
	Establish Reasoning
	Summarize Findings/Judgments
6. Communicate Conclusions	Review Content

Review Style/Format

Table 1 "Stages" students could select

With this framework in place, our efforts turned to integrating the application into the classroom, or—to some extent—modifying the course structure, so that the application (and its imposed scaffold) worked within the structure of the classroom, while attending to students' needs and concerns.

Theoretical and Ethical Framework for Data Collection and Analysis

In an effort to simultaneously refine the Journaling application and responsibly integrate it into the classroom, we employed Design-Based Research, which is predicated upon: rapid iteration of the prototype based on student/instructor feedback; documentation of problems and consequent interventions (e.g. new instructional materials); observation of students' work with the evolving prototype, allowing participants to suggest modifications to improve usability; and follow-up interviews, in which participants and instructors reflect on their experiences. In practice, educators have used Design-Based Research to bring "unfinished" technologies into the classroom, with the stated intention of modifying them in response to emergent concerns (Barab and Squire 2004). As a method, it embraces the natural messiness of the classroom setting and encourages continuous modification and intervention, while rigorously recording both the changes and the causes that motivated them. Although Design-Based Research cannot produce generalizable results, it allows researchers/designers to refine an educational technology, while establishing a praxis for its use and reflecting on its theoretical implications (Brown 1992).

In addition, we also took a Values in Design (VID) or reflective design approach, in which we sought to examine the social and ethical values of bringing this technology into the classroom and how we might design the experiment and use of the prototype to take these values into account (Agre 1997; Friedman 1998; Flanagan 2005; Sengers et al. 2005; Knobel and Bowker 2011; Shilton and Anderson 2016). In our case, one of the key issues that we sought to address at the very beginning of this research endeavor were student sensitivities about privacy of their data, particularly since the Journaling application was a type of participatory sensing/self-tracking device. Privacy has been identified as a key challenge surrounding participatory sensing/self-tracking data, and raises many questions such as who has access to the data, what level granularity the data has, user understanding about what the data mean, how long will the data will be retained, and what kind of longitudinal engagement will exist between the user and the researchers (Shilton et al. 2009). The ideal goal for participatory sensing is to have equitable use, meaningful community participation, local control, transparency, and empowerment between users and researchers/collectors of the data (Shilton 2010a). If these goals can be attained, Shilton argues that the collection, analysis, and use of participatory sensing data can lead to "empowering surveillance," which can be a positive outcome that gives more power and control to individuals involved in providing data for this kind of research (Shilton 2010a). This approach gives research subjects the power to have a say in how their data will be gathered, analyzed, and used, and provides them with a voice across the experimental process. In light of this literature, we will discuss below how we worked to address privacy issues and empower students as part of our experiment.

Experimental Design and Implementation
The work of our study was broken into *nine stages*:

Preparation: As described briefly above, we developed a framework for the class and the Security Threat Analysis project, performed over eight weeks and broken into multiple deliverables, according to an established structure for analytic workflow based on Dhami and Careless (2015). We then integrated that framework into Journaling and developed instructional documentation for the application.

Ethics: As instructor for the course, "A" was sensitized to the need to think about ethical and privacy concerns students might have regarding the Journaling prototype. In this way, "A" played the role of what Katie Shilton calls a "values advocate" or "values facilitator" on the research/design team (Shilton and Anderson 2016). With "A's" prompting, the team spent lengthy conversations discussing ethical and privacy considerations regarding this experiment: how the application could gather data from the students, how we could modify the Journaling application to maximize student privacy, who would have access to the raw data, how the data would need to be handled, and how we could use the data. Because the research project was funded by LAS (and by extension, the NSA), "A" thought it important to describe to students the study in depth and allow students to specify changes in use of the prototype, so as to maximize transparency and comfort. We also went through two separate IRB (and lengthy) approval processes: one at North Carolina State University, and a separate one at the National Security Agency. These processes required that for the experiment to take place, we have informed consent from each student and also ethics training by all researchers involved in the project. Ultimately, the IRB protocol was assigned Exempt status by both entities, meaning that the experiment was evaluated to be minimal or no risk by both institutions. We will discuss more details of the informed consent process below.

Introduction to students: In the first week of class, we brought the intelligence analysts and programmers into the classroom and explained the intended study to the students. We detailed the purpose of the application, the instructional framework in which it would be deployed, the data it would gather, and how that data would be used. We shared the NC State– and NSA-approved human subjects consent form for the study. Each week, over an eight week period, "A" spent a portion of class time asking students about questions regarding the experiment. Students were encouraged to voice any questions or concerns they had regarding the experiment, which we had two months to address before asking them to interact with Journaling. Students were given the opportunity to demo the prototype to see how it would collect and store data. All students were asked to use the prototype as part of the coursework, but were given the option of withholding their data, excluding it from any/all study. Furthermore, no grade or credit was associated with use of the application, so students were not rewarded for using it or penalized for not using it. Students were given nearly two months to consider and sign the informed consent form; all 10 students in the class consented to participate in the experiment.

Training: Before the group project work began, the application's designers came into the classroom and tutored the students on its use. Students were given the option to work with an LAS-provided laptop, a University-provided laptop, or to install the application on their own personal computers.

Group assignments: Students were separated into three different self-selected groups of 3–4 students each (10 students total). They were given the opportunity to choose their group topic. Students then worked in their teams over an eight week period on the security threat assessment; during this time data was gathered as part of the Journaling experiment.

Usability studies and group discussions: We repeatedly checked in with students as they used Journaling and composed their final projects. At week 2 and again at week 4 of the experiment, we distributed brief usability surveys to gather feedback regarding the application and its effects on students' work habits. In addition, the instructor reserved weekly class time to allow students to discuss their concerns regarding the application and/or the Security Threat Analysis, which were recorded via field notes.

Ongoing iteration: Using the class discussions and usability reports, we modified the application's behavior and adjusted assignment parameters. In week 5 of the experiment, the NSA analysts returned to the class, helped the students to modify the installed applications, facilitated a Q&A session, and presented a mid-course analysis of student-generated data to that date. As a more general practice, "A" and "C" met once weekly for approximately three hours and revised assignments, instructions, and expectations using student feedback.

Proximal reflections: After the final presentations (week 8 of the experiment), we held a focus group (absent the instructor and programmers), in which we asked students to share their individual and group experiences with working on the prototype and final project. Here we were interested not only in hearing feedback on usability issues, but also on any student concerns about the experiment. All 10 students provided feedback, and we recorded and transcribed their responses. Using focus-group data, individual student reflections, and student course evaluations, we identified a number of recurrent themes—some specific to the Journaling application, and others pertaining to their experience of the course, and ethical issues that emerged as part of the experiment. Using these themes, we developed a final (usability-focused) review of the application and a detailed set of follow-up interview questions.

Distal retrospectives: We conducted the semi-structured follow-up interviews six months after the course. Eight interviews were conducted individually and without the instructor/programmers present, during which we presented their group's Journaling data, offered them our interpretation, and asked them to discuss their experiences during the study. Here again we were interested in probing broader student reflections on the experiment. We recorded these interviews, transcribed them, and developed qualitative codes based on our findings.

Research Findings

"Opening" the Experiment with Students

We shift here to a discussion of our findings, with a focus on how we tried to ethically consider and involve students into the experiment. We begin with the question of surveillance and privacy—an ethically complex issue raised by Journaling's design—as a tool that inherently surveils through collection of user data—and by its association with the NSA. Because instructional technologies gather data on students and provide access to these data in asymmetrical ways, they have been regularly critiqued as participating in panoptic disciplinary regimes, while being simultaneously enlisted for their enhanced pedagogical affordances (Slade

and Prinsloo 2013; Dawson 2006; Hawisher and Selfe 1991; Janangelo 1991). Journaling's association with the NSA, however, makes what for many students has become a naturalized relationship with instructional technologies more visible and fraught. Recent scandals and controversies related to NSA surveillance of private American citizens, though not raised by our students explicitly (Edward Snowden, for example, was not specifically named), certainly made themselves felt in the initial discussions of Journaling with students. In participant "Nathan's" words, "the NSA has this record of being able to go in and see people's, you know, I guess track what people do over the Internet and with their machines" (Interview, Raleigh, NC, May 6, 2016).

We did not dismiss these concerns; rather, we did our best to create an environment of openness and accommodation into the earliest stages of the study. Before launch of the course, we spent a considerable time trying to anticipate and respond to student concerns about privacy and surveillance. In practice, "responding" to students involved: (1) getting clarifications; (2) modifying the design of the study, and (3) modifying the behavior of the application. Although we were able to anticipate some privacy concerns, there were others raised by students that we had not fully considered. To illustrate this, we discuss three issues that emerged at the early stages of this study and the steps we took to address them in order to respect and empower students.

The first issue we encountered was a simple dilemma: Should we ask students to install Journaling on their own laptops or provide them with class-specific ones? The former approach requires them to install a surveillance device on a personal machine; the latter requires them to use a separate, unfamiliar computer for classwork. From a design/researcher perspective, having students use their own laptops can better simulate real workflows and requires less planning (i.e. securing, maintaining, retrieving the loaner computers) and workarounds. The use of a laptop for the experiment, however, provided a useful "values lever" (Shilton 2012) that focused our attention on the various privacy issues associated with installing and using the application on the laptop. Over a series of discussion and deliberations about the pros/cons/tradeoffs with different laptops, we worked as a research team to respect student concerns, ultimately offering them three options for their involvement in the experiment: to use their own laptop, a university-owned laptop, or an LAS-provided laptop.

The second issue concerned the possible data that the application could gather—for instance, if a student left the application running during personal web browsing. Addressing this concern required multiple steps: The technical team installed a filter that would ignore sets of websites (e.g. pornography, file-sharing) that Journaling might record, and an edit feature that would allow students to remove recorded sites that they did not wish stored. We also allowed students to stop submitting data at any time (e.g., by disabling the plug in or switching to a different browser). There is evidence that some students chose this option during the course of the eight week experiment. Students were also given the option of changing some parameters of the prototype to suit their personal preferences (e.g., the mechanism of tagging sources). As a final precaution, we restricted data visibility. Students were only allowed to see their own raw data and not the raw data of other students. In addition, only the technical programmers, and not the instructor nor the graduate student assistants, could see a student's raw data. We made this choice so that neither the instructor nor the graduate student assistants might be biased in some

way by the individual student raw data. Students were informed of this arrangement at the beginning of the course before the experiment began. Students, the instructor, and the graduate student assistants were allowed to see visualizations and tabulations of data that had been de-identified.

Finally, students wanted to better understand what we were trying to glean from the data. In response, we provided two presentations (midpoint and endpoint) and one individualized debriefing (i.e. interviews with individual participants several months later), in which we showed participants the aggregate, de-identified data as we tried to make sense of it. This included data about the number of relevant unique URLs tagged, the top relevant domains logged, the top analyst tasks logged, the most active research team. We also provided some visualizations of the data, such as initial content analysis among the different research teams (Figure 3). Visualizations have been found to be useful in helping participants involved in self-tracking research make sense of what the data mean (Shilton et al. 2009). In an effort toward transparency about data gathering, we provided midpoint and endpoint data presentations to the students.

Fig. 3 Visualization of Initial Content Analysis of the Research Teams During the Experiment.

In addition, as noted above, this experiment went through multiple levels of IRB review and approval at NC State and at NSA headquarters; all interviews required the students' signed, informed consent. All those involved with conducting the experiment were required to take and pass CITI human subjects protection training. We spent eight weeks introducing the experiment to students, explaining the experiment, the data and how the data would be used, giving students the opportunity to raise questions and get answers about the experiment each week, and having students meet the LAS/NSA development team to better understand the experiment's intelligence merit. We made it clear that students could fully participate in the group project and security analysis without penalty to their course grade, if they chose to not have their data analyzed.

Paradoxical Relationships with Intelligence, Surveillance, and Privacy

The amount of preparation, consideration, and sensitivity we spent on the experimental design was not unwarranted. Retrospective student interviews confirm, in fact, that Journaling's design as an instrument built to collect data from users' digital activities, coupled with its NSA connection, raised considerable student apprehension. "Amber" put this particularly evocatively:

> Well, yes, when, you know, your professor walks into class the first day for a security course and she says, "by the way, the NSA is going to monitor all your research for the next," you know, "four or five months," it's a little unnerving. Especially, you know, no offense, but the NSA has gotten some bad press lately. (Interview, Raleigh, NC, May 3, 2016)

Amber's sentiments were echoed by most of her classmates. "Luc" stressed concerns about the use of the data: "I would say there was definitely some apprehension, and it seemed like, classroom-wise, we would all sort of talk about sort of a little bit of apprehension, because we

didn't know exactly how the information was going to be used" (Interview, Raleigh, NC, May 4, 2016). According to Luc, the class's initial apprehension about participating stemmed from distaste for "having this software live on our computers." This reaction is perfectly reasonable and expected, particularly coming from developmentally advanced graduate students. The lack of such reactions, in fact, would have been more worrisome. With most of these students interested in security studies and actively thinking about intelligence and related security issues, we expected hard questions from them about the project. Indeed, the first several class discussions included detailed questions about the design of the application, and the collection and use of students' data. Before the project started, however, only a couple of students had raised questions or concerns about the surveillance aspect in class; instead, they were more concerned that the NSA would sell their collected data for profit, or about intellectual property rights regarding their data.[7]

However, students' aspirations to join the Intelligence Community—while making the project a natural fit for them and equipping them with more context for critically informed consent—also raises concern that they might be vulnerable in ways other students might not. Namely, the very aspirations that made these students an excellent fit for a project like Journaling, and as proxies for intelligence analysts in the design process, also may have caused students to weigh the potential personal and professional benefits of participation more heavily than other (potentially negative) factors. This sort of surveillance and relinquishing of some degree of privacy may have been seen as the price of entry for the opportunity to interact with members of the Intelligence Community or to work on an intelligence-related project.

Retrospectively, this reasoning may have been behind the ease with which students' fears were alleviated—once initial concerns about privacy were addressed, students readily agreed to participate, attributing this (in the focus group and individual interviews) primarily to transparency about the tool and study, and the humanizing effect of meeting NSA analysts and designers. For example, "Jackson" recounted that members of the Journaling team who visited the class "explained very clearly that it was going to be, you know, very anonymous, and we could opt out," and that "the program wouldn't be observing us 24/7 and all of our web traffic; it would only [observe] within a specific browser within the work project" (Interview, Raleigh, NC, May 3, 2016). These conditions, along with the fact that he was specializing in security studies, made Jackson "really on board with it." Amber similarly pointed to transparency about the purpose and scope of the project as key, with repeated reassurance that the application was for "research purposes only," and that the data/technology "wasn't going to be sold," which "puts your fears at rest" (Interview, Raleigh, NC, May 3, 2016).

The post-semester focus group and interviews also revealed the humanizing effect of meeting the Journaling team, and its importance in alleviating students' concerns about surveillance and privacy. Luc, for example, explained that "having a face associated with this endeavor" was critical, making him "buy in a little bit too, because you knew there were people behind it" (Interview, Raleigh, NC, May 4, 2016). Students seemed relieved by the direct interaction with members of LAS and particularly glad to find them relatable. For Jackson, the

[7] With this concern, we checked with LAS management to ensure that the data would not be sold to a third party, but be used solely for NSA purposes. In addition, according to LAS management, the students would not hold intellectual property rights for their data.

analysts "weren't exactly what I expected an analyst to be like. … They were very friendly, explanatory of what they were thinking, what would be going on, what the hardware would be used for. You know it wouldn't be sold or anything like that. I have very positive impressions" (Interview, Raleigh, NC, May 3, 2016). Amber expressed her relief that "no one seemed seedy or shady" (Interview, Raleigh, NC, May 3, 2016). Notably, both these students had aspirations for joining the Intelligence Community, so their relief could be partly in finding and identifying with relatable members of their aspirational profession. Regardless, the direct human interaction with LAS members played a key role in students' comfort level and project participation.

However, the study's design—specifically, making multiple technical means of participation available—had an equally persuasive effect on students' decisions to participate. Out of 10 students, only 3 (Amber, Cooper, and Jackson) chose to install Journaling on their personal computers; the rest accepted a University-/LAS-owned laptop instead. Thus, although project transparency and the humanizing effect of interacting with LAS members played a role, most students felt more comfortable compartmentalizing class and study participation from the rest of their lives, using what "Trevor" called a "sterile computer" (Interview, Raleigh, NC, May 6, 2016). Indeed, students might have felt coerced without this option built into the study design. Trevor, an adult learner working for the Department of Defense, might not have chosen to participate without the option of separating personal data from study-related data. In his words,

> I have a crazy family, so I'll get crazy emails from my crazy sisters and my parents who are, you know, it's all sort of redneck drama. You know, I wouldn't feel as comfortable on my personal computer for reasons like that or for privacy concerns … but on a, sort of, sterile computer system, I had no problem with it. (Interview, Raleigh, NC, May 6, 2016)

"Erik," another adult learner—a naturalized US citizen with extensive defense and international experience—echoed Trevor's sentiment, adding an element of suspicion about the stated parameters of the program and study. Erik's apprehension could not be alleviated through transparency, as he was concerned that the program might include covert elements: "Once this project is done, you know, and, okay, I'm getting rid of it. Let's say I'm uninstalling it. Is it really uninstalled? Or is it just going to keep on tracking?" (Interview, Raleigh, NC, May 5, 2016). For Erik, the key to participation was the school-provided laptop, at which point his reservations "just went out the window" (Interview, Raleigh, NC, May 5, 2016). It is important to note, that although it seems like this decision was made quickly and lightly, the entire informed consent process lasted eight weeks, with several question and answer sessions, so that students would have time to consider what was important to them to assure their privacy during the experiment.

In many ways, the students held a paradoxical relationship to surveillance and privacy, one characterized by competing and simultaneously held impulses: on one hand, suspicion and resistance to surveillance; on the other, the aspiration to participate in a field that makes routine use of it, and trust in authorities to use data ethically. This paradoxical relationship is borne out in a few minor, related categories to emerge from analysis of the course that echo surveillance-related tropes in the broader culture. For example, when asked directly about these issues, several students responded that "I don't subscribe to conspiracy theories" or "I don't worry too much about that sort of thing" (Interview, Raleigh, NC, May 4, 2016). Similarly, two students

mentioned not having anything to hide, with Amber joking that “you can share my vacation pictures—they’re already on the Internet” (Interview, Raleigh, NC, May 3, 2016). Luc’s answer, however, encapsulates how many seem to cope with the paradox and conflict of surveillance, namely, resignation that “this is the world we live in” (Interview, Raleigh, NC, May 4, 2016).

In hindsight, there are aspects of the study that we could have improved to further enhance student privacy and control during the experiment. For example, although we gave students the option to delete data from their Journaling record at any time during the experiment, we did not give them the direct access to do this because the backend of the application was not user friendly; students would have needed in-depth computer training to know how to access and delete their data. Instead, the students were asked to place a request to one of the LAS computer programmers if they wanted their data removed. Because we did not provide a means for the students themselves to remove the data that was stored on the server, there may have been some reticence to ask for this assistance. This could be something that we could change for a future experiment, to give students more autonomy and direct control over their data, although there would have to be additional planning and training incorporated to accommodate this. We could have also given students more tagging options as part of the prototype function. We can unambiguously state that the students the students who received more prompts to tag found it irritating/distracting. We discovered this quickly and this led the designers to a more open approach that gave the students greater control over the automatic working of the applications -- but we did not give students the option to "partially opt out" -- and this was a lack of consideration/imagination on our part. We believe that many students would have gladly "opted in", knowing that it would help the development of the application. As a general practice, one should initially design the application so that students are given control (which assumes adequate instruction) over any automatic feature that could conceivably interrupt their work-flow. If this is not feasible, students should be given multiple study participation options.

In addition, we could have spent more time thinking about longitudinal engagement with the students and their data. We could have thought about additional ways to involve students in our analysis over time. One issue with this, however, is student fatigue with the experiment after the course was over. During the final set of individual interviews with students we noted some student fatigue/disinterest with having to reflect on the course and experiment. We also could have involved students explicitly as part of the experimental design process—although this is more difficult to accomplish in the limited time frame of a semester-long project. A year-long research project (e.g., a master’s level capstone project) would be more amenable to involving students in the design of the experiment and prototype features.

In sum, these findings suggest that when involving technologies and people of the Intelligence Community in the classroom, instructors should be aware of differing and, at times, conflicting perspectives about intelligence, and give students the opportunity to voice opinions from their perspectives. Instructors should also be proactive in designing instructional materials and assignments that recognize and protect student privacy, and involving students in this process. Since this experiment and the privacy issues that it has raised, the LAS/NSA team has become more interested in further studying how social, ethical, and privacy concerns and users’ priorities can be incorporated into the design and operation of the Journaling application.

Intelligence Technology as Opportunity, Tangent, and Obstacle for Pedagogy

Our hope in integrating Journaling into the course was that students would learn from interacting with it, while assisting in its ongoing design. Although our analysis suggests that this was mostly the case, students' experiences were more mixed and variable than we had anticipated. These learning complexities related to aspects of the tool itself and how students work, as well as to the broader context, expectations, and set of participants that constituted the learning environment in this experiment.

Positive responses to the experience, for example, fell under a few categories in our qualitative coding of interview data related to either learning about intelligence or personal cognitive benefits. Several students saw working with Journaling as an opportunity to access and experience intelligence work, a sentiment that came up during the end-of-semester focus group and retrospective student interviews. Nathan, for instance, reflected that using the tool "put me more in tune with the intelligence cycle" and gave him insight into "how real world analysts" worked with data, resulting in "a better understanding of what it would be like to work in the Intelligence Community" (Interview, Raleigh, NC, May 6, 2016). Nathan's experience illustrates well how the choice of a special topics course on threat analysis contributed to the benefits students saw in using an intelligence technology. Similarly, several students pointed to the natural fit between the tool and the course as contributing to the benefits gained from using it. Amber explained that the fact that part of the course related to technology in security studies and threat assessments created synergistic benefits between the two: "The class is a Science, Technology, Society course, and then you're bringing in technological tools…I think it fit perfectly with the class. I think it enhanced it" (Interview, Raleigh, NC, May 3, 2016). Because one of the course goals included critically assessing security risks posed by technologies like drones and chemical weapons, critically assessing an intelligence technology presented another opportunity to think analytically and evaluate technology.

At a more personal level, there were unexpected benefits of using the tool that several students raised, related to metacognition and focus. Journaling's integration in students' browsers, along with its periodic prompts to tag webpages, served as continual reminders of the course and focus of each week's assignment, particularly because the course discussions and scaffolding used the same terminology as Journaling's "goal tree." Jackson explained how the tool "would sometimes remind me to stay on focus of what portion of the project we were in at that point," which helped him prioritize tasks and information (Interview, Raleigh, NC, May 3, 2016). Several other students spoke of the metacognitive benefits of using the tool, particularly those related to their research processes. "Jenna," a student who ended up having a lot of trouble with the tool, spoke of that experience as providing her with additional self-awareness of her typical workflow and habits: "I learned a lot about how I like to work, and I learned a lot about what I'm thinking when I'm working" (Interview, Raleigh, NC, May 5, 2016). Both Amber and Erik echoed Jenna's experience, adding the insight that Journaling rendered workflows solidified over years of practice visible for inspection and reevaluation. Amber, remarking that she had "been in school for a long time" now, found that using Journaling helped her "think about the way you research, as opposed to just going out and doing it" (Interview, Raleigh, NC, May 3, 2016). Similarly, Erik noted that "old habits die hard," and that the tool "showed me that … there are things that I could improve in the way I can do research" (Interview, Raleigh, NC, May 5, 2016). These benefits were bound less to the context of the course than to the tool's interruption

of routinized, automatic operations, which afforded students conscious reflection and reshaping of those operations.

This interruption, however, was also behind a perception of Journaling as, at best, tangential to the main focus of the course or, at worst, a disruption or obstacle to students' writing, research, and thinking. Several students spoke of the tool as being tangential to the course and their learning, providing neither benefit nor detriment. "Cooper," for example, described it as an inert part of his experience: "I guess I would say I don't think it made it better or worse—it was just kind of there" (Interview, Raleigh, NC, May 4, 2016). Trevor said the application was tangential to his class goal, which was to demonstrate measurable, documented (via a grade) mastery over course material by producing quality assignments. In his words,

> Even when I did log [his research and writing activity], I was kind of just like, you know, I wasn't putting a lot of thought into what step or what task I was working on. I was just like, ah, that sounds good, fine. Because I'm more focused on getting that process done. I personally am a terrible test subject, because I'm focused on the task and especially since I'm being graded for it. If I, perhaps if this were more voluntary or incentivized, but not so much in terms of grading, maybe you would get better data on the process tracing in the task selection. (Interview, Raleigh, NC, May 6, 2016)

Trevor illustrated the multiple coexisting classroom goals, which were prioritized differently, depending on subject position. For students having trouble translating the course-based learning to further professional or academic advancement without a good grade, the graded classwork took precedence over other class-related activities, particularly when the tool was deemed tangential to the course. Interacting with Journaling, then, became simply an act of good will toward the instructor, LAS, and the Intelligence Community, one that could be deprioritized when it interfered with students' primary goals.

This experience is borne out by Journal's uneven logging activity data, which show most logging occurring during analytic stages that were cognitively more straightforward, and when user activity was easier and quicker to categorize. Collectively, students logged more activity in the first four stages, which involved primarily research-related activity and therefore categorizing results from web searches and reading. Logging dropped off dramatically in the final two stages (see Table 1: Compose Interpretations and Communicate Conclusions), which included more interpretive, evaluative, and communicative activities that were more cognitively challenging and perhaps less visible to users. This result seems to support Trevor's insight that students' class goals took precedence. Essentially, students had to continually triage between multiple goals. When some of their higher-priority goals became more challenging (i.e. producing an unfamiliar genre for a grade), many of them chose—consciously or not—to discontinue using Journaling, a trend compounded by the lack of discrete markers between tasks to be tagged during these final stages (which most activity occurred in a single Google Doc instead of the constant movement between webpages during earlier stages).

Interviews with students revealed a number of reasons why they stopped logging at various points, but the interviews seem to confirm that, understandably, use of the tool was only one priority among many (e.g. producing quality deliverables, learning course material,

managing personal responsibilities). For example, Erik explained that "when it comes to creating deliverables," the tool "becomes cumbersome" (Interview, Raleigh, NC, May 5, 2016). Jenna recounted the cognitive challenges of learning to be an effective graduate student during her first semester of graduate school, while also trying to learn about the intelligence process and how to write a completely new genre of intelligence writing. For her, being "in the process of learning those three things at once" and having "something repeatedly pinging me to ask me what I'm doing and categorize my action" was "frustrating" and produced "cognitive overload" (Interview, Raleigh, NC, May 5, 2016). Jackson recalled a mundane technical failure with the LAS MacBook he borrowed, and the two weeks it took him to install the program on it, thus subordinating it to other priorities during those weeks (interview, Raleigh, NC, May 3, 2016). And Trevor recounted a difficult medical episode with his son in November, a responsibility that he rightly prioritized, which would have added to the cognitive and logistical burden of working with two computers and logging class-related activity (Interview, Raleigh, NC, May 6, 2016).

Disruption, in fact, was one problem we anticipated and hoped to reduce through our Design-Based Research approach, which feeds research findings on users' experiences back into a fast, iterative design process. This approach led to adjustments to Journaling's settings—particularly how and when it prompted users to tag activity—thus reducing the disruptions students experienced over time. However, this did not fully address some students' experience of the Journaling project—particularly in the time it took to implement—as a distraction from the course and an obstacle to deeper learning. As one student wrote in the course evaluation, "We spent at least three classes discussing this computer program and the study, and I think that time would have been better spent in classroom discussion" (Course evaluation, December 2015). This raises another ethical issue: testing technology in a classroom may displace valuable instructional time, which has implications for student learning.

In order to address this potential ethical problem, one needs to choose contexts in which testing a technology and participating in a given study can work synergistically with course goals and the need for regular metadiscourse about those connections. It also underscores the need to be sensitive to how study participation may at times compete with students' other goals and demands, with learning necessarily being the overarching priority. In our case, we had hoped to provide students with a document based on their Journaling data, which we felt could enable reflection and enhance further learning. We used the output of the tool as a way of discussing the work that transpired within the class, but we remained critical of the tools capacities to record an accurate picture of events. For instance, when a group of three students worked together – work that was not done in a Journaling enabled browser could not be seen by the application. In this regard, we found that the short-comings of the application (i.e. the data blind spots) served as a useful vehicle for discussing the often incomplete data sets that can be produced from automated monitoring. In addition, one risk of working with a developing tool, however, is that it can take more time to develop analytical approaches to making meaningful use of generated data than allowed for within a semester. Our inability to make good on this plan before the semester ended reinforces the need for building in multiple artifacts, interactions, and experiences that support students' learning and goals.

One important pedagogical lesson that resulted from this experiment, which was a specific goal for the instructor, was to teach students how to think about conducting a technology

assessment differently. This resulted from the unique integration of course material, with Journaling, the scaffolded assignments, and the instructor feedback. This multifaceted analytic process allowed students to work with concepts covered earlier in the course, and not only incorporate them into the final group report, but also—for some students—to look at security threats differently. Jenna shared her views about the applicability of course project ideas into future security-related encounters:

> Now, going through my classes, we go on a lot of trips and we talk to a lot of people who are in the IC [Intelligence Community], and they're talking about nuclear things and satellite images and these things, and particularly tacit knowledge and how dangerous it is for people to get certain weapons, or how easily they can create them. And other classmates who I'm still in contact with and I have been in these talks [with], and one of us will ask about tacit knowledge and how easy do you really think it is? Because "A," one of his main points in our units was there's a social element to this, and it's not as easy as people think. And that person invariably will say, "Oh yeah, if they get this, it'll be really dangerous. It will be a snap for them to produce stuff." And we'll kind of give each other the side eye, like, that's not what "A" said. I think you're wrong. So that kind of has followed me. And I know at least three other people who consistently, like, we ping, it's a little like, "All tacit knowledge is important." I think that that has been very helpful in forming probably the way that all of us work now. And I know that two of them are now graduating and going to work in the IC and in, like, you know, the government, and that will continue to form how they evaluate things. So, overall, [the class] seemed very successful (Interview, Raleigh, NC, May 5, 2016).

Therefore, at least one pedagogical problem that plagued this course in the prior year seems to have been addressed through this experimental architecture. Although Journaling, by itself, was not able to reveal this example of student learning, the one-on-one evaluative interviews that followed the experiment captured this important finding.

Interacting with the Intelligence Community: Raised Stakes and Benefits

The involvement of the NSA in this project created a special kind of vulnerability and a special kind of benefit for the students involved in the experiment. Although it is difficult to untangle elements of the course, such as its scaffolding, the assignment design, Journaling, and interaction with LAS members, it seems that this last element likely played a large role in benefitting students. The involvement of intelligence practitioners led many of the students to have a qualitatively different experience of the course. By working with the prototype, they were given more direct access to intelligence analysts than would be possible otherwise—many saw that as a professional boon and led some to minimize the frustration and burden (time/privacy) associated with using the prototype---and created for them an important learning experience in and of itself. Jenna, one of the students who ultimately experienced the most frustration with the tool, spoke about how good it felt to "just be part of something" (Interview, Raleigh, NC, May 5, 2016). As someone aspiring to go into intelligence, she had heard about "how difficult it is to streamline the intelligence process and how long it takes to really get a feel for researching and gathering information." She knew "that they were trying to solve this Olympic problem" and was happy about being able to contribute. Erik echoed this sentiment, expanding it to encompass research more generally: "Ultimately, the point of, you know, coming

to a university and doing studies and research is to contribute something" (Interview, Raleigh, NC, May 5, 2016). Participating in the study seemed one way to accomplish this, something enhanced by direct interaction with analysts, which "puts in perspective how a tool like this, you know, can contribute to improve their work" (Interview, Raleigh, NC, May 5, 2016).

This direct interaction, however, also added to the stakes of the course, with several students describing it as intimidating. Suddenly, the deliverables they were producing—particularly their final report and final presentation—no longer felt like a classroom exercise, but rather like authentic threat analysis reports that could contribute to real-world intelligence work. This was particularly true for the final presentation, an event that included members of LAS (and therefore NSA) in the audience. The presentations were not scheduled for the regular classroom on the main campus, but rather a conference room in the satellite campus that houses LAS. Students were prepped to expect a few members of LAS to attend to act as "intelligence customers" for students' final reports; however, the number that did attend surprised both students and instructor. Invited by the NC State manager of LAS to show support for the LAS-funded experiment and prototype, approximately 30 NSA analysts and LAS staff attended the student presentations, including the NSA director at LAS and two visiting NSA officials.

This surprise revealed another unforeseen outcome of an externally grant-funded experiment: LAS/NSA has its own interests, concerns, and activities that the academic collaborators could not fully control, and which shaped the classroom learning experience for students. Jenna described the shift that occurred for her as NSA analysts entered the room:

> We were just kind of hanging out at this conference table and looking over our stuff, and then people start coming in. And it's just more and more. And more and more. And you know that they're important. And they're LAS people. And a lot of our class was interested in intelligence, so we were just like looking around. Like, he said there were going to be two people—there were 30 people! (Interview, Raleigh, NC, May 5, 2016)

For Luc, the class's "swirl of important folk" had made it feel "a tad more relevant than an average class," but this final class meeting was on another order altogether: "I don't know if 'intimidating' is the right word, but things took on a more decidedly official feel … the ante of the game perhaps maybe got upped just a tad" (Interview, Raleigh, NC, May 4, 2016). "A" also felt stressed before the presentations. He was concerned about: the students and how they felt about being thrown into an unexpected situation; how they would perform given the circumstances; how the NSA officials would react; how the presentations might reflect on his own instructional capabilities.

The LAS audience was not passive, so the stakes of the presentations and students' analytic work rose further as the audience participated in the Q&A portion of each group's work. As Jenna nicely described:

> The part that freaked all of us out the most is the questioning portion, because a lot of people [there] we did not know. We could tell they were informed people. And they would ask things like … "Did you think about anything in the future?" Like, "what might happen in the next six months or a year?" Or, like, "five years or ten years? And how

> would this affect your paper?" And for the purpose of the assignment "A" said was one year. So when he asked us about five years or ten years I was like, uh. You know, no. And then they were just like, uh-huh. And then I was like, oh my God. We failed. (Interview, Raleigh, NC, May 5, 2016)

Luc, who generally found the experience intellectually invigorating and stressed the feeling that "what we were doing maybe mattered a little bit," also found the questioning by LAS officials disconcerting and questioned their reasons for being there. Although he felt their work mattered, he also had doubts:

> Were the things that we were doing, did it really translate into actually … you know, was some of this information useful to people? I don't know. Was some of this stuff timely that people asked? I don't know. Were they there in a purely observatory capacity? Or *what were they there for*? (Interview, Raleigh, NC, May 4, 2016)

Luc's question is interesting in that he seemed, however briefly, to be wondering what, exactly, he had participated in. Although students understood the layered goals from integrating research and design activities in the classroom, the interest LAS officials took in students' findings—though exciting and flattering—felt also potentially unwarranted, raising the prospect that there were broader, invisible motives. The stakes, then, felt raised, but not in ways students could fully see or understand.

The intensity and pressure of this experience, however, ended up being a confidence booster for many of them, because the students were able to perform well. Nathan, in particular, who described the experience as being "thrown to the wolves," found it validating: "Even though we got grilled by the Director of LAS for a while, we actually were able to answer his questions fairly well, and so we were happy about it" (Interview, Raleigh, NC, May 6, 2016). Indeed, Nathan (and potentially others) received some direct benefits from performing so well. While chatting after the presentations, the LAS Director who had grilled Nathan asked him if he had "ever thought about coming and interning with us?" which Nathan was able to parlay into a summer internship. So, although it is doubtful that the stakes of students' threat assessments were real, students were right to perceive them as such, because this experiment did, for some, become an opportunity for professional advancement.

This outcome, in fact, raises the possibility that most of the benefits students gained from using Journaling came not from their mere interaction with it, but rather from the greater technology—the LAS/NSA apparatus *behind* the technology. In the course, Journaling was not simply a browser-based software application, but also: institutionally provided laptops; a computer programmer provided by LAS and affiliated with NSA; two additional NSA analysts/technical specialists; a usability specialist with instructional technology, documentation, technical writing, and instruction expertise; institutional scrutiny by a University and an NSA Institutional Review Board; intense weekly planning meetings between instructor and usability specialist; intelligence practitioner involvement; and carefully planned and iterative scaffolding. All of these elements constituted the "technology" of Journaling; they also shaped the student learning experience.

Conclusions

Recommendations

Although there were some bumps along the way, this experiment proved to be a successful collaboration that implemented important ethical considerations into the prototype use for the class experiment, with unique aspects of student learning. Although this project benefitted from having collaborators with compatible goals, like-minded iterative research agendas, and interest in ethics and situated knowledge, we have learned a number of lessons from this unique experiment regarding academic-intelligence partnerships in the classroom that might be beneficial to researchers contemplating similar engagements.

For students

Whenever possible, use research methods and experimental design that consider ethical perspectives to benefit student participants and other stakeholders. In our case, we purposefully included some variations in methods (e.g. use of LAS laptops, reporting of collected de-identified data, option to delete data) to provide students with additional privacy protections, if desired. In this way, we were able to design a prototype and experiment that embodied values important to the individuals involved in the experiment (Shilton and Anderson 2016) (See Figure 4). We intentionally anticipated ways to protect their privacy stringently, even if they did not voice privacy concerns. We felt this extra care necessary because of the potential sensitivity of having an NSA partner involved. We also aimed to maintain transparency about the experiment's research goals, methods, and data analysis throughout, from initial consent through final presentations. This was not always maintained (e.g. the unexpected NSA audience during final presentations), but we strove to achieve transparency to the greatest extent possible. We also provided students with alternative means of meeting course requirements if they were not comfortable participating in the Journaling experiment, to fully protect their rights and prevent subsequent feelings of coercion. However, to our surprise, no students chose this alternative.

Fig 4. From Privacy values to Journaling/experimental design features

For instructors

Time and significant effort are required to successfully develop and deploy a prototype (intelligence or otherwise) into a classroom setting in order to maximize student learning and meet research objectives. In our case, the instructor spent a significant amount of time in the months leading up to the experiment to design appropriate instructional materials and assignments to make the technology understandable to students and also to address student privacy concerns. In addition, the instructor specifically took on the role within the team of "values advocate" to promote privacy considerations throughout the experiment and was an integral part of the research team; it is important to have someone on the design team adopt this role as an insider working within the team (Shilton 2010b; van Wynsberghe and Robbins 2014). Previous literature has shown the problems of having a values advocate who is not an accepted and integral part of the research team (Guston and Sarewitz 2002; Manders-Huits and Zimmer 2009; Rabinow and Bennett 2012). The instructor also worked closely with a usability expert and technology developers to adapt the prototype to course objectives in order to enhance student learning, as well as allow sufficient time to develop appropriate human-participant instruments and obtain all necessary IRB approvals. Specifically, it required additional IRB review by NSA headquarters, which involved the submission of additional documentation beyond that of the NC

State IRB. In anticipating and addressing student concerns as part of our experimental design and IRB applications, our team went above and beyond current codified ethical standards in thinking through the ethical implications of this sort of study in a classroom setting. Thus, although ethics oversights bodies are important and helpful, instructors/practitioners should also be proactive to think about ethical implications of bringing technologies into the classroom, particularly with intelligence collaborations, especially since the issues posed by these sorts of technologies and relationships are emergent. As this study illustrates, new academic-intelligence collaborations can offer a new set of ethical and procedural challenges that stakeholders need to take into consideration for a classroom-based research experiment.

Beyond specific Journaling data outputs, this experiment provided students with the ability to reflect on their analytic process. This reveals the importance of building in opportunities and resources to support metacognition as one of the principal pathways for learning. This course also benefited from having a close synergy with intelligence analysts. Therefore, it is beneficial to partner intelligence practitioners with courses whose learning outcomes can be readily augmented by research and interaction with the Intelligence Community. In this way, experiments and collaborations can provide new and powerful ways for students to learn about complex problems. Instructors need to develop multidisciplinary qualitative and quantitative evaluation instruments to assess student learning and technology experience over the course of the experiment.

For technology developers/intelligence practitioners

One key lesson we learned was to be sensitive to your user population, and the social and ethical concerns they may have with how the technology (and its data) might be used. In our experiment, the technology developers were very open to considering potential student concerns and adapting the experiment/technology to address them. The developers were purposefully very transparent in their explanations to students of how the technology worked, how the data would be used and stored, and who would have access to them. This required developers to devote several hours of their time (beyond setup and data gathering) to create a trusting environment for the experiment. In addition, it required that some of these individuals—intelligence analysts—reveal their identities; not all of them had done so in public, unclassified settings before. Some of the positive student comments on the course experiment regarded their opportunity to interact closely with members of the Intelligence Community (e.g. providing weekly feedback on the prototype, giving the final presentation to NSA analysts). This requires that intelligence practitioners be open and willing to interact with students outside their regular workload and be reciprocal in their commitment to transparency and engagement. The LAS is a particularly suited site for promoting interdisciplinary research that takes into account values in technology design. What might be surprising to some, the NSA leadership at LAS has been a strong supporter of a values in technology design approach and is sensitive and reflexive to public concern over surveillance---this has created a very positive and open environment in which to explore issues of privacy and surveillance as part of prototype development at LAS. For example, since we ran this experiment, LAS has created a “privacy research team,” a new thrust area of research for the lab, that involves a collection of academic, industry, and intelligence participants to work on various ethics and privacy issues that are expected to emerge with different big data prototypes that the LAS is developing.

More broadly, this project also provides an opportunity to illustrate the importance of co-construction of a technology with its users. NSA analysts have said that they had previous experiences in which new technologies were foisted upon them in their work—technologies that had not been developed together with the analytic community and thus often proved unsuitable (e.g. too cumbersome, impractical, little value added) for analytic work, or sometimes merely rejected outright by analysts. By developing a process and practice of close engagement between technologists and users, it is possible to develop a new analytic technology, enabling analysts to work more efficiently and accurately, that also reflects social values important to users. This experiment was one step forward in that direction and has stimulated additional research at LAS on privacy issues. Future experiments with Journaling will involve NSA analysts, and we are in the process of probing the social and ethical issues that might arise with that set of users working within an intelligence organization.

In this experiment, we provide an example of how one can create "empowering surveillance" even when working with the NSA (Shilton 2012). We understand that many find any connection to or collaboration with the Intelligence Community unsettling, if not absolutely indefensible; these concerns reflect a variety of political and intellectual commitments that we understand. We do not advocate that everyone should jump on the intelligence bandwagon, but for those who do have an interest in exploring academic-intelligence collaboration in the classroom, this experiment provided one lens through which we hope to clarify the various social, ethical, and pedagogical matters that can shape academic collaborations involving intelligence prototypes.

Acknowledgments: To be added.

References

Agre, P. (1997). Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In G.C. Bowker, L. Gasser, S. Leigh Star, and B. Turner (Eds.), *Social Science, Technical Systems and Cooperative Work: The Great Divide*, (pp. 131-158). Hillsdale, NJ: Lawrence Erlbaum.

Barab, S., & Squire, K. (2004). Design-based research: Putting a stake in the ground. *Journal of the Learning Sciences*, *13*(1), 1–14.

Bødker, S. (1987). Through the interface—A human activity approach to user interface design. *DAIMI Report Series*, *16*(224), 1-167.

Brown, A. (1992). Design experiments: theoretical and methodological challenges in creating complex interventions in classroom settings. *Journal of the Learning Sciences*, *2*(2), 141–178.

Dawson, S. P. (2006). The impact of institutional surveillance technologies on student behaviour. *Surveillance & Society*, *4*(1/2), 69–84.

Dhami, M., & Careless, K. (2015). Ordinal structure of the generic analytic workflow: A survey of intelligence analysis. Paper presented at the European Intelligence and Security Informatics Conference 2015, Manchester, UK.

Engeström, Y. (2001). Expansive learning at work: Toward an activity theoretical reconceptualization. *Journal of Education and Work, 14*(1): 133-156.

Flanagan, M., How, D., and Nissenbaum, H. *Values in Design: Theory and Practice*. Working Paper, 2005.

Friedman, B. (ed) (1998). *Human values and the design of computer technology*. Chicago: University of Chicago Press.

Guston, D.H. and Sarewitz, D. (2002). Real-time technology assessment. *Technology in Society 24* (1-2), 93-109.

Hawisher, G. E., & Selfe, C. L. (1991). The rhetoric of technology and the electronic writing class. *College Composition and Communication*, 42(1), 55–65.

Janangelo, J. (1991). Technopower and technoppression: Some abuses of power and control in computer-assisted writing environments. *Computers and Composition*, *9*(1), 47–64.

Jones, P., Sharma, S., Moon, C. & Samatova, N. (2017). A Network-Fusion Guided Dashboard Interface for Task-Centric Document Curation. *The 22nd Annual ACM Conference on Intelligent User Interfaces (IUI),* Limassol, Cyprus.

Jones, P., Thakur, S., Matthews, M., & Cox, S. (2016). A Versatile Platform for Instrumentation of Knowledge Worker's Computers to Improve Information Analysis. *The Second IEEE International Conference on Big Data Computing Service and Applications (BDS),* Exeter College, Oxford.

Jones, P., Thakur, S., Matthews, M., Cox, S., Streck, S., Kampe, C., Srinath, P., & Samatova, N. (2016). Journaling Interfaces to Support Knowledge Workers in their Collaborative Tasks and Goals. *The 17th International Conference on Collaboration Technologies and Systems (CTS),*Orlando, Florida.

Knobel, C. & Bowker, G.C. (2011). Computing Ethics: Values in Design. *Communications of the ACM*, 54(7): 26-28.

Manders-Huits, N. & Zimmer, M. (2009). Values and pragmatic action: the challenges of introducing ethical intelligence in technical and design communities. *International Review of Information Ethics*, *10,* 37-44.

Miller, B.H. (2008). Improving All-Source Intelligence Analysis: Elevate Knowledge in the Equation. *The International Journal of Intelligence and Counterintelligence* 21(2), 337-354.

Neff, G. & Nafus, D. (2016). *Self-tracking*. Cambridge, MA: MIT Press.

Nolan, B.R (2013). Information Sharing and Collaboration in the United States Intelligence

Community: An Ethnographic Study of the National Counterterrorism Center, Ph.D. Dissertation (Philadelphia: University of Pennsylvania).

Office of the Director of National Intelligence (July 16, 2008). Intelligence Community Directive 205: Analytic Outreach. https://fas.org/irp/dni/icd/icd-205.pdf. Accessed March 11, 2017.

Rabinow, P. & Bennett, G (2012). *Designing Human Practices: An Experiment with Synthetic Biology.* Chicago: University of Chicago Press.

Sengers, P. Boehner, K. David, S, and Kaye, J. (2005). Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility* (pp. 49-58). New York: ACM, Retried from http://dl.acm.org/citation.cfm?id=1094569

Shilton, K. and Anderson, S. (2016). Blended, not Bossy: Ethics Roles, Responsibilities, and Expertise in Design. *Interacting with Computers*, *29*(1), 71-79.

Shilton, K. (2014). This is an Intervention: Foregrounding and Operationalizing Ethics During Technology Design. In K.D. Pimple (ed.), *Emerging Pervasive Information and Communication Technologies (PICT).* Dordrecht: Springer (pp. 177-192).

Shilton, K. (2012).Values levers: Building Ethics into Design. *Science, Technology, & Human Values, 38*(3): 374-397.

Shilton, K. (2010a). Participatory Sensing: Building Empowering Surveillance. *Surveillance & Society*, *8*(2): 131-150.

Shilton, K. (2010b). Technology Development with an Agenda: Interventions to Emphasize Values in Design. Retrieved from http://escholarship.org/uc/item/72n146cj

Shilton, K. Burke, J.A., Estrin, D, Govindan, R., Hansen, M. Kang, J. and Mun, M. (2009). Designing the Personal Data Stream: Enabling Participatory Privacy in Mobile Personal Sensing. Retrieved from http://escholarship.org/uc/item/4sn741ns

Shilton, K. (2009). Four Billon Little Brothers?: Privacy, mobile, phones, and ubiquitous data collection. *Communications of the ACM, 52* (11), 48-53.

Slade, S., & Prinsloo, P. (2013). Learning analytics ethical issues and dilemmas. *American Behavioral Scientist*, *57*(10), 1510–1529.

Spinuzzi, C. (2003). *Tracing genres through organizations: A sociocultural approach to information design*. MIT Press.

The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction, *Report to the President of the United States*. (March 31, 2005). https://fas.org/irp/offdocs/wmd_report.pdf. Access March 11, 2017.

Treverton, G.F. (2008). *Assessing the Tradecraft of Intelligence Analysis* (Santa Monica, CA: RAND Corporation.

U.S. National Research Council (2011). *Intelligence analysis for tomorrow: Advances from the behavioral and social sciences.* Washington, DC: U.S. National Academies Press.

Van Wynsberghe, A. and Robbins, S. (2014). Ethicist as Designer: A Pragmatic Approach to Ethics in the Lab. *Sci. Eng. Ethics*, *20*, 947-961.

Fig. 1

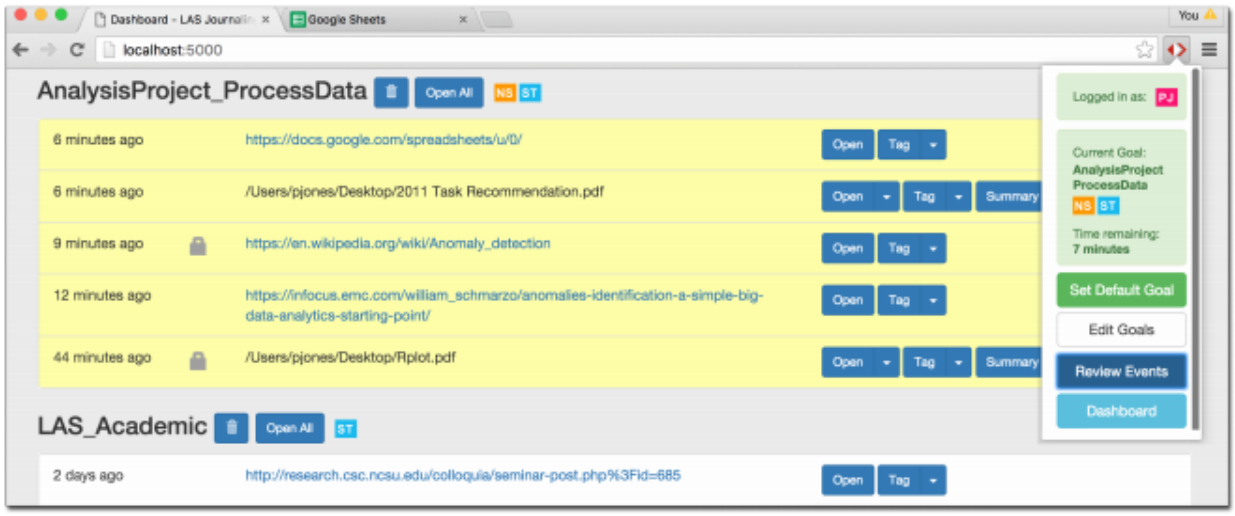
Dashboard - LAS Journalin
Google Sheets
You
localhost:5000
AnalysisProject_ProcessData
Open All
NS ST
6 minutes ago
https://docs.google.com/spreadsheets/u/0/
Open
Tag
6 minutes ago
/Users/pjones/Desktop/2011 Task Recommendation.pdf
Summary
9 minutes ago
https://en.wikipedia.org/wiki/Anomaly_detection
12 minutes ago
https://infocus.emc.com/william_schmarzo/anomalies-identification-a-simple-big-data-analytics-starting-point/
44 minutes ago
/Users/pjones/Desktop/Rplot.pdf
LAS_Academic
2 days ago
http://research.csc.ncsu.edu/colloquia/seminar-post.php%3Fid=685
Logged in as:
PJ
Current Goal:
AnalysisProject ProcessData
Time remaining:
7 minutes
Set Default Goal
Edit Goals
Review Events
Dashboard

Fig. 2

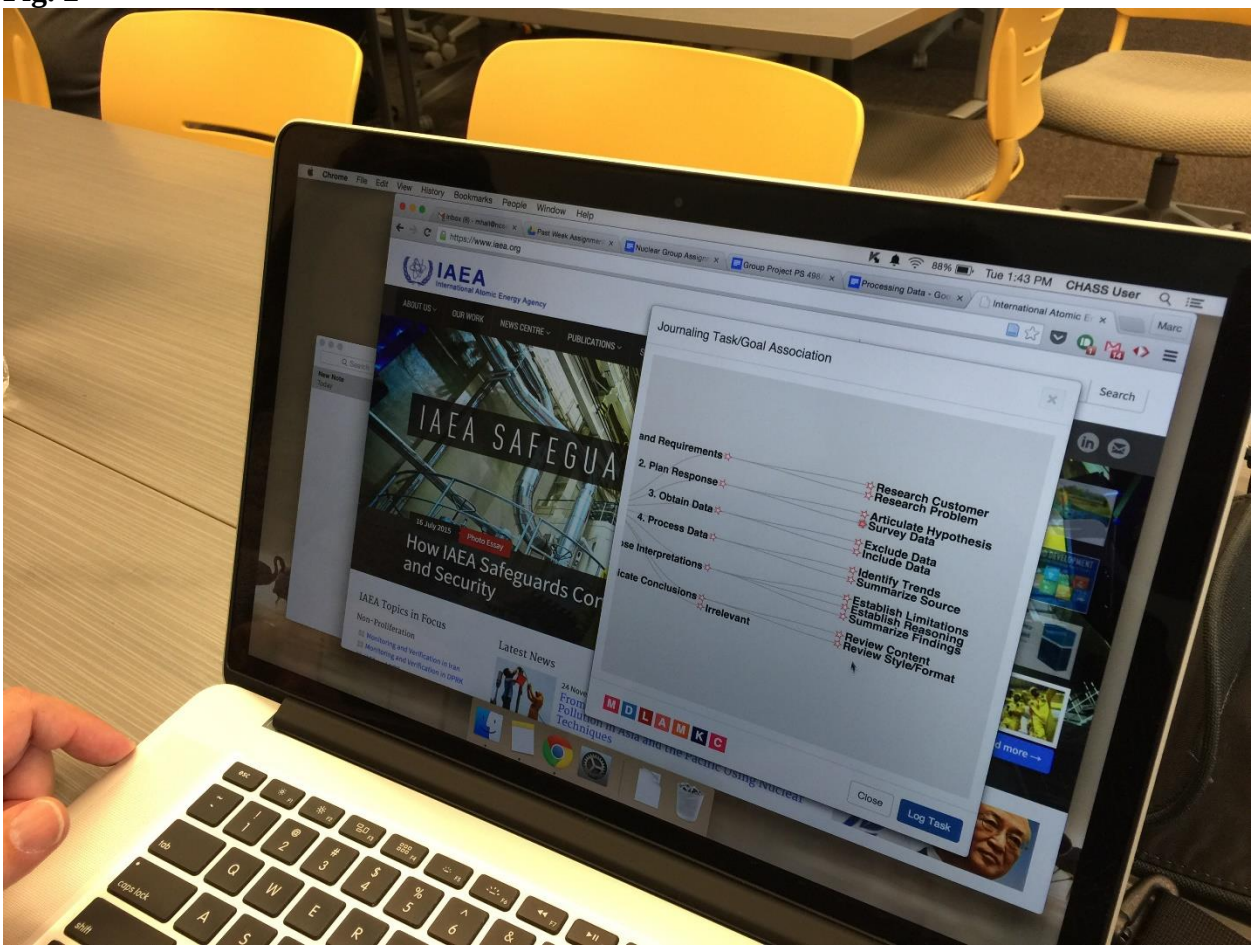
IAEA
Journaling Task/Goal Association
Research Customer
Research Problem
Articulate Hypothesis
Survey Data
Exclude Data
Include Data
Identify Trends
Summarize Source
Establish Limitations
Establish Reasoning
Summarize Findings
Review Content
Review Style/Format
Close
Log Task
How IAEA Safeguards
and Security
IAEA Topics in Focus
Latest News

Fig. 3

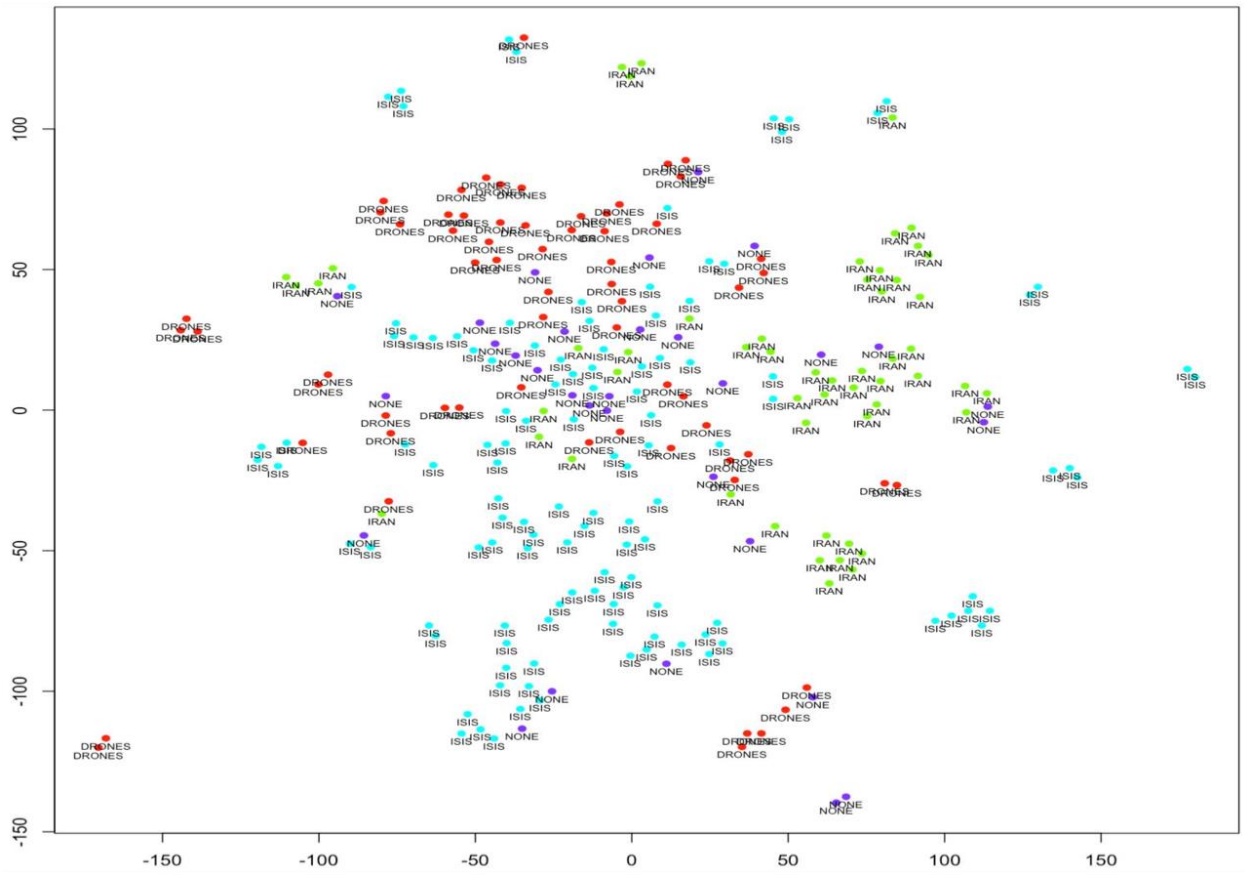
100
50
0
-50
-100
-150
-150
-100
-50
0
50
100
150
DRONES
IRAN
ISIS
NONE

Fig 4.

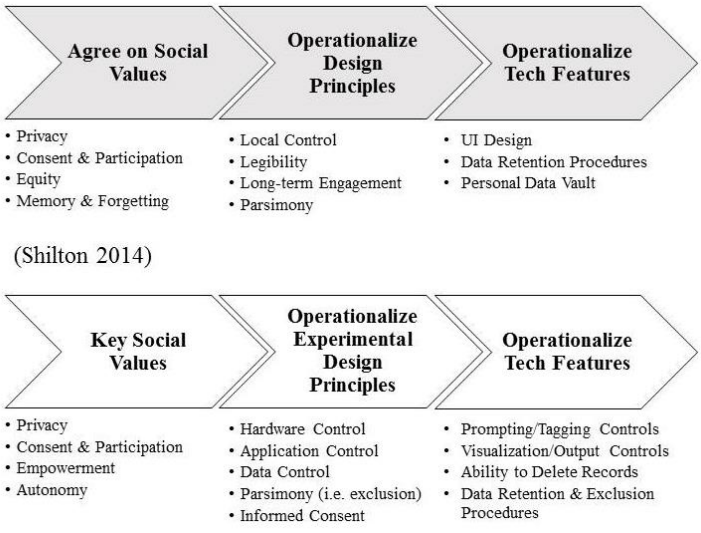
Agree on Social Values
Operationalize Design Principles
Operationalize Tech Features
• Privacy
• Consent & Participation
• Equity
• Memory & Forgetting
• Local Control
• Legibility
• Long-term Engagement
• Parsimony
• UI Design
• Data Retention Procedures
• Personal Data Vault
(Shilton 2014)
Key Social Values
Operationalize Experimental Design Principles
Operationalize Tech Features
• Privacy
• Consent & Participation
• Empowerment
• Autonomy
• Hardware Control
• Application Control
• Data Control
• Parsimony (i.e. exclusion)
• Informed Consent
• Prompting/Tagging Controls
• Visualization/Output Controls
• Ability to Delete Records
• Data Retention & Exclusion Procedures