A model for data ethics instruction for non-experts

Leigh Phan[1], Stephanie Labou[2], Erin Foster[3], Ibraheem Ali[4]

Abstract

The dramatic increase in use of technological and algorithmic-based solutions for research, economic, and policy decisions has led to a number of high-profile ethical and privacy violations in the last decade. Current disparities in academic curriculum for data and computational science result in significant gaps regarding ethics training in the next generation of data-intensive researchers. Libraries are often called to fill the curricular gaps in data science training for non-data science disciplines, including within the University of California (UC) system. We found that in addition to incomplete computational training, ethics training is almost completely absent in the standard course curricula. In this report, we highlight the experiences of library data services providers in attempting to meet the need for additional training, by designing and running two workshops: Ethical Considerations in Data (2021) and its sequel Data Ethics & Justice (2022). We discuss our interdisciplinary workshop approach and our efforts to highlight resources that can be used by non-experts to engage productively with these topics. Finally, we report a set of recommendations for librarians and data science instructors to more easily incorporate data ethics concepts into curricular instruction.

Keywords

data ethics, data literacy, data science, data privacy, algorithmic bias, community engagement

Introduction

The last decade has seen an unprecedented increase in the availability and prevalence of technological solutions used to simplify complex decision-making processes. Importantly, use of these technologies is so pervasive that it influences nearly every aspect of modern life. These solutions span from algorithms that regulate content and marketing that a typical person sees in their social media feed (Orlowski, 2020), to resumé reading software for job applicants (Borsellino, 2018), to recidivism predictors for former criminals (Angwin et al., 2016; O'Neil, 2016), to insurance eligibility algorithms that approve funding for vulnerable populations (Obermeyer et al., 2019). Certain models have yielded substantial benefits; for example, pandemic prediction models for informing public policy decisions, among others (Alzahrani et al., 2020). However, other models are filled with explicit or implicit biases that exacerbate racially biased decision-making, amplify stereotypes or misinformation, and make it extremely difficult to identify those accountable for failing to properly train these technologies to identify bias (O'Neil, 2016; Noble, 2018 Benjamin, 2019).

While the need for thoughtful exploration of algorithms within an ethical framework is increasingly recognized (Holmes et al., 2021), academic curriculum has lagged in incorporating these aspects as requirements. This issue is similar to the increased need for data literacy training, more generally, across academic domains. Disciplines relating to Science, Technology, Engineering, and Math (STEM) tend to incorporate some data literacy instruction into their curriculum; however, disciplines in the social sciences and humanities generally do not have the resources to do so (Dennis et al., 2021).

As a result of this curricular gap, academic libraries are increasingly called upon to deliver data literacy education through community instructional approaches like the Carpentries (Carpentries, 2021), or consultations catered to individuals and small groups (UCLA DSC, 2021; ACRL Report, 2021). In addition to this disparity that exists in data science and digital literacy training, ethical uses in data or data ethics are generally not required for degree completion. While data scientists can technically come from any discipline, research data ethics practices are almost certainly not covered at an undergraduate level. Even curricular requirements for Computer Science majors at UC campuses, where data science is most commonly taught, reveals there is limited (UCLA Curriculum, 2021; UC Berkeley Curriculum, 2021), optional (UC Irvine Curriculum, 2021) or no formal required training (UCSD Curriculum, 2021) in data ethics as part of their curricular instruction.

Historically UC libraries' instruction - whether through the Carpentries or other models - was limited to their respective campus communities. This siloed approach to instruction changed when the 2020 COVID-19 pandemic induced a complete disruption of in-person instruction. In response to the move to virtual instruction, UC campuses pivoted rapidly to remote learning environments. Despite the challenging circumstances, remote learning presented new opportunities for collaboration across the UC campuses and allowed for a broader reach of previously campus-specific events. A compelling example of this is the UC Libraries co-hosting UC Love Data Week in February of 2021, which was comprised of a week of free events for UC affiliates on a wide range of topics including: data publishing, data repositories, data cleaning, and text mining to name a few (Universities of California, 2021).

As Love Data Week organizers from UC Berkeley, UCLA, and UC San Diego, we saw this collaborative UC Love Data Week as an opportunity to present to a wider audience across the UC system on the subject of data ethics. As data information professionals, we are familiar with some aspects of data ethics, but none of us have received formal training or education on this topic. However, we felt that the importance of the topic, in addition to our professional experience with ethical issues in data, created a unique opportunity to begin a dialogue among UC affiliates and lay the groundwork for a cross-campus community interested in continuing conversations and learning around this subject. For UC Love Data Week 2022, we broadened the scope of the initial workshop to feature student speakers from different disciplines. Importantly, taking into consideration the highly diverse body of learners across the UC system, we sought to design a workshop that allowed non-experts from a wide variety of disciplines to engage with the topic of data ethics.

Workshop approach

In this section, we will go into detail about how we structured our 2021 workshop before moving into how we incorporated lessons learned from the first iteration to inform our 2022 workshop. We grounded our initial workshop as part of UC Love Data Week 2021 with the overarching question: *Who is impacted by your research?* Due to the time allotted for this workshop to cover a topic as broad as data ethics (e.g., only 1 hour in 2021), this guiding question allowed us to help attendees consider ethical considerations in their own research, beyond the confines of the workshop. Leading with this guiding question also allowed us to limit the scope and introduce critical topics as applied to research data analysis in various disciplines (Phan et al., 2021). To inform discussions, we provided examples of ethical issues related to data that are well-described in the academic literature. We focused on transferable concepts to help attendees build context for questions of data ethics that might arise in their respective disciplines (Figure 1).

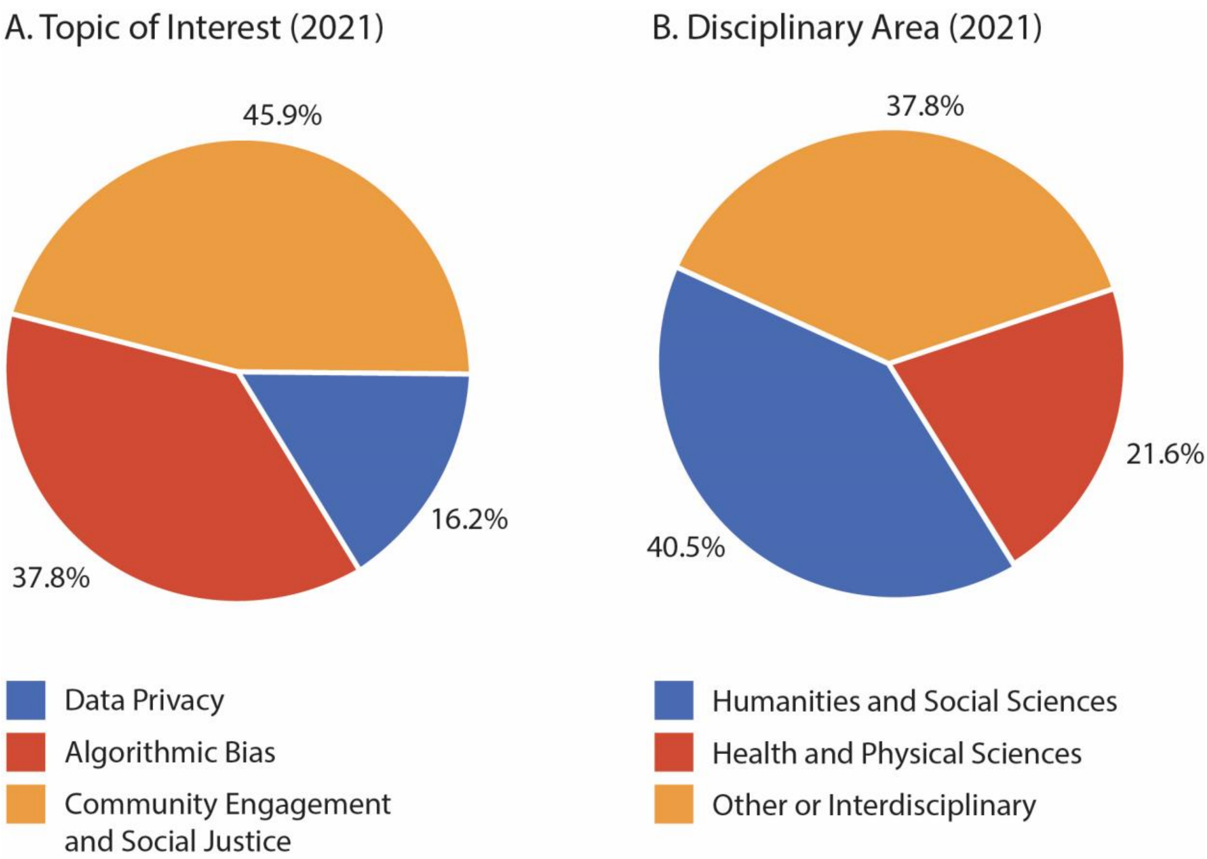

Figure 1: Background and Interests of Attendees for 2021 workshop. A. Learners were asked to indicate the topic of most interest to them in the workshop. B. Learners for the 2021 workshop were asked to indicate their primary professional background and disciplinary areas.

For the 2021 workshop, we created a pre-workshop survey to gather information about our workshop attendees' backgrounds and interests as they relate to topics in data ethics in order to best customize workshop content. The workshop attracted participants from a wide variety of disciplines, including 40.5% from Social Sciences and Humanities, 21.6% from Health or Physical Sciences disciplines, and 37.8% self-described as Other or Interdisciplinary, across many levels of experience. While preparing for the initial workshop, we focused on three aspects of data ethics: Data Privacy, Algorithmic Bias, and Engaging Communities in Research. From the pre-workshop survey, 16.2% were interested in the topic of Data Privacy, 37.8% were interested in Algorithmic Bias, and 45.9% were interested in Community Engagement and Social Justice (Figure 1) so we framed our workshop content around these three topics. Our workshop began with lecture content that described real world examples in which these three ethical topics were at play. Following the lecture portion of the workshop, we hosted breakout discussions in which we asked attendees to bring up ethics issues that they observe in their fields of study, related to one of the three data ethics themes discussed in more detail below.

Data privacy: striking a balance between privacy and transparency

There have been a number of high profile data breaches in the last several years (Larson, 2017; Barrett, 2019; Alder, 2020; Mihalcik, 2020). For this reason, we expected most attendees to be familiar with the general concept of data privacy, and that breaches of privacy have consequences. Therefore, we chose to focus more narrowly on examples that highlighted disparities in the consequences of data breaches on individuals, based on political context, economic and social status. For example, individuals experiencing poverty can be subjected to high levels of police surveillance, but have few tools to protect themselves from breaches of that information (Green and Gilman, 2018). Vulnerable populations that work with researchers, like people who identify as transgender or sex workers, may be at extreme risk in the event that their identities are revealed in the event of a data breach (Sandy et al., 2019; Sinha, 2017).

While topics such as personally identifiable information (PII) are often covered in human-subjects research trainings such as those created by Institutional Review Boards (IRB), we wanted to go beyond the technicalities described in those trainings and emphasize the importance of data privacy beyond merely a mandatory training and study design approval before research can begin. Importantly, we highlighted the study published by Rocher and colleagues in 2019 which showed that it is possible to re-identify the vast majority of the American population using only 15 demographic attributes (Rocher et al., 2019), and showcased other examples where supposedly anonymous data was reidentified.

With that context in mind, and following discussion of disparities in consequences of data disclosure, we provided a brief overview of methodologies for anonymizing data, removing personally identifiable information (PII), and aggregating data as a means of protecting the identities of individuals associated with studies. This workshop took place in February 2021 and we could safely assume every participant had heard of the decadal Census taking place in 2020. The US census serves as an example of a heavily relied upon data set used for informing public policy and determining state and federal representation. Beyond being a dataset that affects literally every US resident in a practical sense, it is also used for a wide array of interdisciplinary research and as such is applicable to many researchers. This made the 2020 Census an ideal case study to discuss how geographic scales can impact data re-identification; trade-offs between microdata and aggregated data, in terms of variables made public; and how differential privacy, specifically statistical methods for introducing noise into a dataset, can be used to preserve individual participant privacy (see Abowd, 2018).

Importantly, we did *not* offer a single one-size-fits-all conclusion for this portion of the workshop, but rather, encouraged attendees to think critically about their data, weighing the importance of data privacy vs transparency. As academic spaces grapple with a reproducibility crisis (Baker, 2016), partially due to incomplete data transparency, researchers must also be mindful that data shared is not done so at the expense of participant privacy. To connect this idea to workshop attendees' own work, we concluded this portion with the following questions:

- What are the minimum variables needed for meaningful analysis?
- Could they avoid collecting unnecessary identifiable data?

- What are the best practices or accepted de-identification methods in their research domain?
- What would be some privacy-preserving statistical methodologies that are most suitable or needed for their work?

Research models and algorithmic bias

To build on the concept of how data-driven research impacts vulnerable populations beyond aspects of privacy, we next covered predictive research models themselves, providing examples of how bias in algorithms can lead to bias in results. We began with a conceptual example that the audience would be familiar with, by explaining that researchers often use predictive models to understand complex phenomena. In this case, we used two examples of such models: infection prediction rates of COVID-19 and the trajectory of a hurricane. Importantly, we pointed out to learners that predictive models rely on specific assumptions in order for their predictions to be fulfilled correctly; similarly, algorithms follow the same principles. If the assumptions are biased, or fail to acknowledge bias in the data collection, the outcome of a predictive model will be biased. We unpacked three examples where assumptions in algorithms lead to racial or unjust bias. These examples included a variety of commonly-applied technologies, including algorithms used for prediction of recidivism, health insurance financial allocation, and facial recognition.

In our first example, we introduced the COMPAS ('Correctional Offender Management Profiling for Alternative Sanctions') recidivism algorithm, a software used to assess a defendant's risk of committing another crime within two years. As part of the COMPAS system, once defendants are booked in jail, they are required to answer a questionnaire; the algorithm incorporates the questionnaire results to predict the defendant's likelihood to reoffend. In 2016, Angwin and colleagues investigated COMPAS and found that it resulted in high false positives, predicting Black defendants to be at a higher risk of recidivism than they actually were two years later (Angwin et al., 2016). We emphasized that rather than a single flawed algorithm, COMPAS is one of numerous examples of algorithms which exacerbate systemic racial bias in the justice system by using data informed by racially biased policing to train algorithms that predict recidivism or criminality and failing to protect individuals from false-positives. (O'Neil, 2016).

Our second example focused on an algorithm commonly used to simplify funding approvals for insurance claims. In 2019, Obermeyer and colleagues uncovered that an algorithm designed to provide additional support for people with significant healthcare needs was racially biased. The algorithm designers used money spent on healthcare as a proxy for healthcare needs, assuming that the more money spent, the sicker the person. However, the designers failed to recognize that disparities in access to healthcare results in less healthcare spending for Black patients on average compared to white patients (Obermeyer et al., 2019). As a consequence, the algorithm assumed Black patients were healthier on average than white patients, even though the Black patients suffered from more chronic health conditions than the white patients at any given health score created by the algorithm. By using the number of chronic conditions, rather than money spent, to determine a health score, Obermeyer and colleagues were able to more than double the insurance approval rates for Black patients from 17.7% to 46.5% (Obermeyer et al., 2019). For this example, we highlighted to our learners that the choice of convenient, seemingly effective proxies, for difficult to access information can be an important source of algorithmic bias in contexts such as this.

In our final example, we highlighted facial recognition software, specifically use cases of facial recognition in law enforcement. The Gender Shades project spearheaded by Buolamwini and colleagues assessed the accuracy of multiple commercial facial recognition software. While many of the facial recognition software had high accuracy, all companies performed better on male than female subjects, and generally on lighter-skinned subjects overall. The authors discovered that most of the facial recognition algorithms were trained on Caucasian samples, resulting in poor accuracy, underrepresentation and misrepresentation by gender and race (Buolamwini and Gebru, 2018). In addition to Gender Shades, we introduced FairFace, a project led by Karkkainen and colleagues, who developed a more balanced racial composition compared to several other datasets used commercially for facial recognition (Karkkainen and Joo, 2021). In contrast to FairFace, we presented examples of Amazon's Rekognition software - at one point used by the U.S. Immigration and Customs Enforcement agency - which mistakenly matched U.S. Members of Congress with a mugshot database. Juxtaposing examples of improvements to facial recognition accuracy along with cases that demonstrate the negative impact of inaccurate models allowed workshop participants to critically examine the real-world consequences of algorithms on communities impacted by these models.

Community engagement and research impact

A critical component for holistically considering appropriate methods for data anonymization and creating ethical and equitable algorithms is engagement with the communities impacted by the work. Therefore, our last topic focused on providing our learners with examples of how research projects across disciplines have constructively engaged with the communities affected by their research. We highlighted three ways this engagement has taken place: (1) reporting results to study participants, (2) utilizing community expertise, and (3) enabling broader social change through policy development.

As an example for reporting results to study participants, we highlighted the For Healthy Kids! Project. This project, published by Thompson and colleagues in 2017, was an examination of pesticide exposure amongst primarily immigrant farmworkers and non-farmworkers in a region of the state of Washington, United States. The authors used a community based participatory research approach to engage with the community consistently throughout the data collection process, such as using town halls and community boards to inform the community about the dangers of chronic pesticide exposure. Following the study, the research team reached out to the individual study participants to provide their results. While the authors were unable to reach all study participants, their efforts highlight the importance of using existing community infrastructure and direct communication to make a greater impact to study participants in a vulnerable community (Thompson et al., 2017).

In our second scenario, as we had discussed the dangers of racial bias associated with predictive policing software earlier, we wanted to note a constructive example of community engagement with a historically criminalized population. In the case of the SAFE Lab led by Dr. Desmond Patton out of Columbia University, this research group studies social media communication from gang-involved and affiliated youth and how that might result in 'off-line' instances of community violence (Frey et al., 2019). In the case of this particular work, the SAFE Lab actively works with community domain experts (often former gang members) to ensure that the social media communications are not misinterpreted by the researchers and accurately analyzed (Patton et al., 2019). Additionally, to address community concerns about the use of social media data as a means of surveillance, particularly by police, the SAFE

Lab has established a code of ethics that helped the researchers explore different modalities of obtaining consent and did not allow data to be shared with groups that engage in punitive and criminalizing actions.

In our final example, we introduced learners to the Los Angeles-based Million Dollar Hoods project, led by Dr. Kelly Lytle Hernández, which is a community-based research initiative focusing on the human and fiscal costs of mass incarceration (Lee, Lytle Hernández and Tso, 2018). This initiative engages deeply with communities in and around Los Angeles County to 1) gain access to relevant data on incarceration in these regions and 2) produce outputs such as data driven reports and dashboards that show the disproportionate effect of incarceration. Additionally, this research initiative emphasizes building skills of community members to allow for them to actively engage with research focused on their needs. An example of this work is the Big Data for Justice Summer Institute, held at UCLA, which provides training on working with data and relevant tools for both university students and community members (Big Data For Justice Summer Institute, 2022). Most compellingly, the Million Dollar Hoods project has informed changes to local and state policy around issues of mass incarceration by providing data-driven evidence of disparity and discrimination. A cornerstone of this work is 'centering' the voice of the communities most affected by these issues to inform the research questions and contribute to the research process.

Together, these scenarios demonstrate that direct community engagement improves the distribution of research results, generates new pathways to build expertise amongst communities, and can help improve public policy on issues of direct relevance to the respective communities.

Incorporating data ethics into data literacy instruction

At times it may seem that ethical issues in data collection and modeling are so pervasive that engaging with the topic feels daunting, especially for those without any formal training on these issues. Fostering intentional engagement within individual sub-disciplines can lower the barrier and amplify engagement with these critical issues. In doing so, we as librarians can take ownership over how data ethics impacts us as researchers, resource providers, and data users directly, and provide a higher level of support and guidance on best practices for patrons. While we might not feel like experts as instructors on these topics, by asking critical questions (as indicated above) we were able to feel empowered to find resources that impact our own work (such as Census data, or facial recognition datasets) while presenting examples that would inspire our audience to consider ethics in their own disciplines. Furthermore, making connections between themes in data ethics and examples that are relevant to learners' backgrounds can create entry points to integrating ethics in existing library data literacy curricula.

Understanding and learning from your audience's background

In preparation for the initial workshop in 2021, we conducted a pre-workshop survey which allowed us to gauge our audience's backgrounds and their interest in the themes we planned to present, and to better prepare for facilitating breakout discussions. Being attentive to learners' backgrounds and customizing examples to match learner interest (as expressed in the pre-workshop survey) creates an environment more conducive to learners making connections between themes within data ethics and their own field(s) of study. Presenting a variety of examples which resonate with the audience not only lowers the barrier to entry for learners to engage in these topics, but also provides a framework for them to identify ethical considerations within their own research beyond the workshop setting.

With the knowledge that UC Love Data Week 2021 was well-attended among UC graduate students and staff, we planned the 2022 Data Ethics & Justice workshop with this audience in mind. Building on the lessons learned and feedback from the 2021 workshop, we designed the 2022 iteration to highlight the work of three graduate researchers from various disciplines whose work intersects with data ethics. These presentations were used to demonstrate real-world ethics consideration in both research and industry and also provided a framework for breakout room discussions, themed to the issues brought up by each speaker.

Our first speaker for the 2022 Data Ethics & Justice workshop was a PhD candidate in Gender Studies researching racism and disinformation through social media data analysis. To begin the workshop, they discussed their experience gathering Twitter data, which led them to question the ethics of mass social media data collection. They recommended protecting the privacy of social media users, and offered ethical frameworks for those conducting research using social media data. The second speaker was a Master's candidate in Statistics whose work experience included medical billing and information security at a global healthcare software company. They emphasized the importance of cybersecurity on personal and organizational levels and provided an example of a real-world healthcare system whose database was compromised. This speaker concluded by giving recommendations for protecting our own personal privacy and data as well as those in our own organizations. Our third speaker was a PhD candidate in Computer Science who focuses on human-computer interaction and investigates ways to improve privacy communication. They presented on limitations of consent for ethical data collection, providing examples from their own research in which individuals may either feel pressured to provide their personal information, unknowingly give more of their own personal information, or how social media sites can make inferences about an individual based on the information their friends provide.

Following the presentations, each speaker paired with a workshop organizer to facilitate discussion in a breakout room, focused on the data ethics aspect discussed in the respective speaker's presentation. We aimed to lower the barrier for engagement by encouraging attendees to join discussion rooms whether or not they felt ready to contribute to the discussion, highlighting the value of learning through listening. In addition, we created a notes document accessible for all participants to share discussion points. Following 25 minutes in breakout rooms, all participants reconvened for a group recap of breakout room main points and concluding thoughts before the workshop ended.

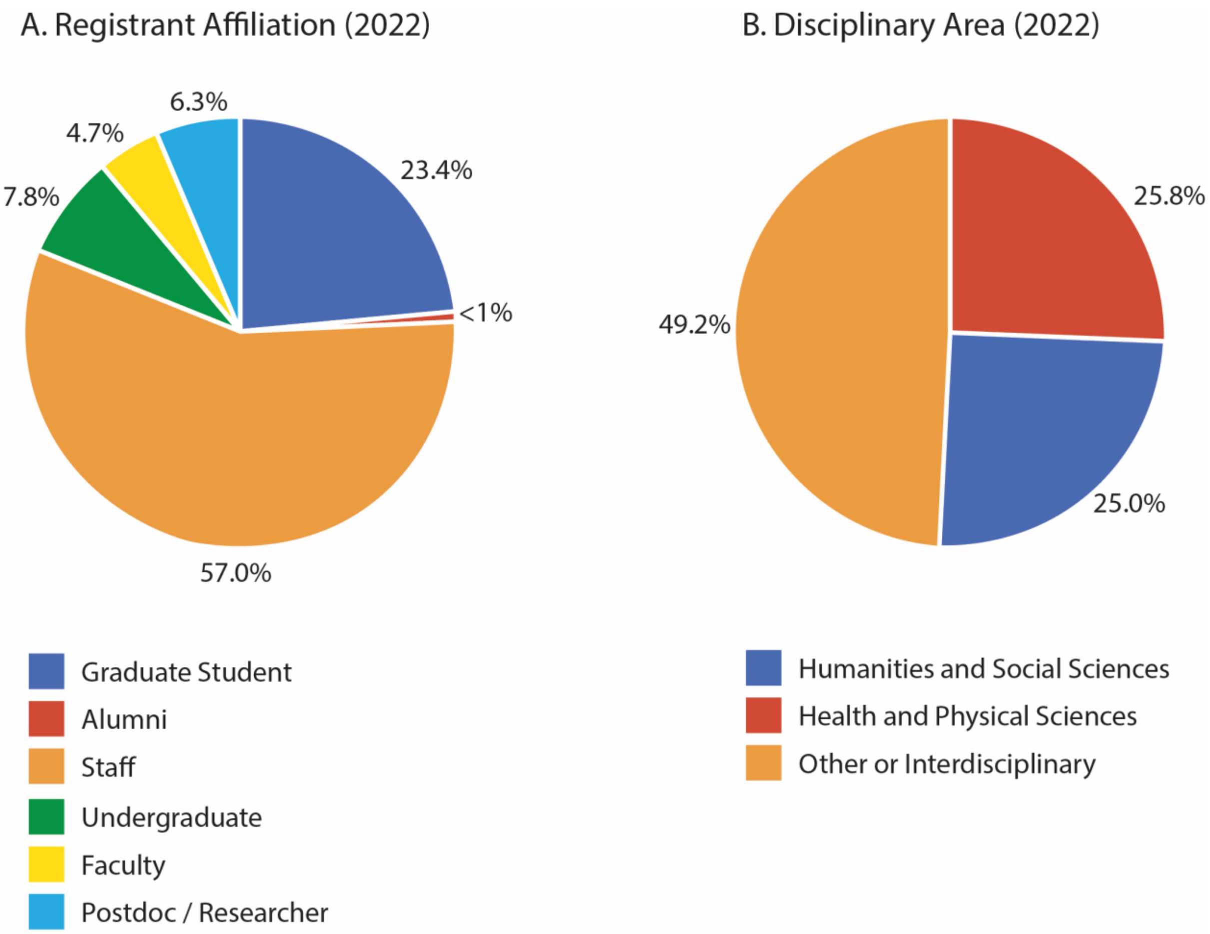

Figure 2: Background of Workshop Attendees for 2022 Data Ethics and Justice workshop. A. Learners were asked to indicate their academic status from a drop-down list. B. Learners were asked to indicate their primary professional background and disciplinary areas.

In designing both our initial and 2022 workshops, we focused on a common theme of understanding the audience's background. For example, since a large portion of 2021 attendees were graduate researchers and staff, we recognized the benefit for the audience to learn directly from peer researchers. This approach provided scholars the opportunity to share their work while creating a forum for discussion around data ethics across disciplines.The demographics of registrations for our 2022 workshop suggest that there is high interest from scholars, staff, and community members outside the university. Overall, there is evidence of increasing interest on the part of staff and interdisciplinary researchers.

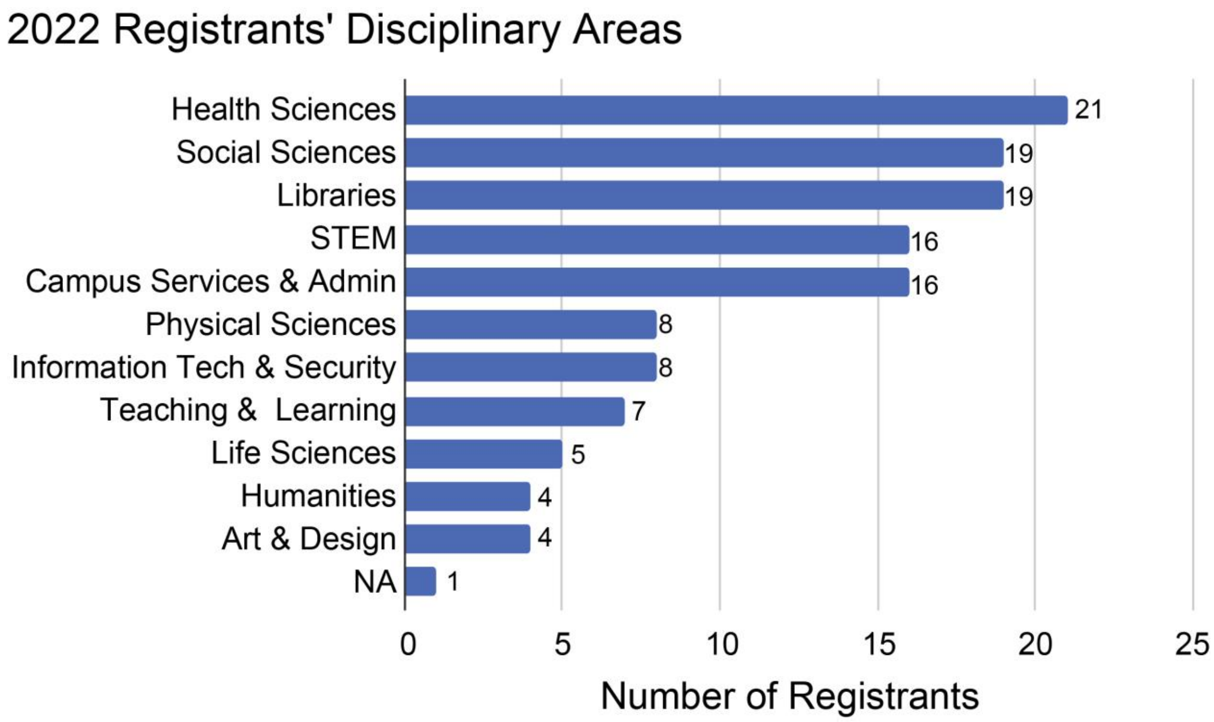

Figure 3: Registrants' Disciplinary Areas for 2022. Learners were asked to indicate their professional background and disciplinary areas. This information was collected solely for the 2022 workshop.

Work with disciplinary instructors to identify curricular gaps

Considering the growing interest, and limited required training, in data ethics and justice across disciplines and professional fields, as library professionals, we can incorporate data ethics training into data literacy curricula through partnerships with disciplinary instructors. Moving beyond one-time workshops, strategic partnerships between data science instructors, instructors that specialize in a particular discipline, and librarians can be a first step to identifying an approach to data ethics training that leverages all instructors' respective expertise. Working with disciplinary instructors can help set the scope of the audience's background, and instructors can incorporate data ethics examples in standard data literacy curricula. For example, as data science and data literacy instructors are often called on to teach research data management, instructors can incorporate topics of data privacy and examples of privacy violations (and how to avoid them) into existing training modules.

Domain-specific examples and experience from instructors, complemented by data literacy fundamentals from librarians, can also serve to emphasize for learners the idea that researchers can and should directly engage with the community they are studying. Learners get a better sense of the appropriate level of privacy required for the data being collected in their field of study, as well as any disciplinary standards and best practices for data de-identification. Furthermore, when developing models or software, it can be critical to put faces to the data, so to speak, by involving communities who will be impacted by the models themselves. These communities can also inform selection of proxies for difficult to measure parameters, as demonstrated in the examples of algorithmic bias.

These concepts apply to a wide range of disciplines and fields, and allow researchers to make more informed decisions throughout the research lifecycle.

Collaborate with offices of research

Cross-institutional collaborations may present opportunities to further expand the reach of data ethics curricula and data literacy overall. Such collaborations can develop by promoting data ethics-related instructional workshops and events with cross-campus colleagues and units.

Within a single campus, forming partnerships and lines of communication with offices of research could help identify gaps unseen by individual disciplinary experts. Furthermore, working with offices of research and other campus-wide, domain-agnostic entities could help in identifying existing platforms such as seminar series and journal clubs - both within and across disciplines - which could be a useful method of bringing these workshops directly to researcher communities. While we did not directly engage offices of research in our 2021 or 2022 workshops, we identified them as potential collaborators for future workshops of this nature. While the library is a domain-agnostic campus entity suited to providing training in ethical issues in data, offices of research are also well-versed in this area and can provide additional real-world and campus-specific examples relevant to learners.

Conclusion

In response to increased demand for data science professionals and cross-disciplinary data science training (Dennis et al., 2021), community-developed and collaborative models of providing computational training have emerged (e.g., The Carpentries), but standardized training for data ethics remains scarce. The workshops described here serve as starting points for data services providers and instructors to expand their knowledge and develop critical competencies to incorporate ethics discussions in their own work with patrons as well as during formal instruction and training settings. Through teaching this workshop, we discovered that there is substantial interest across the UC system in these topics, and an increasing interest among staff and particularly among fellow librarians. We developed an approach to teaching these topics with limited formal training - which includes noting our non-expert status for learners - and found that we could build on this work as we continue to engage with our research communities.

The time for workshops such as this is now. There are a number of major societal consequences when research is conducted without deep consideration of the communities involved in the research process. Or, as we described in our workshop: *Who is impacted by your research?* As librarians and data services practitioners play an increasingly important role in data literacy instruction, data analysis guidance, and data management best practices, intentional engagement with these topics can begin to address some of the downstream consequences of the gaps in this training. We at the Universities of California libraries, expect to build on the lessons learned from teaching these workshops and work within and outside our institutions to further identify and fill curricular gaps in data ethics training.

References

Abowd, J. M. (2018). The U.S. Census Bureau Adopts Differential Privacy. Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2867. Available at: https://doi.org/10.1145/3219819.3226070

ACRL Report. (2021). ARL/CARL Joint Task Force on Research Data Services Releases Final Report. Association of Research Libraries. Available at: https://www.arl.org/news/arl-carl-joint-task-force-on-research-data-services-releases-final-report/

Alder, S. (2020, August 17). Healthcare Data Leaks on GitHub: Credentials, Corporate Data and the PHI of 150,000+ Patients Exposed. HIPAA Journal. Available at: https://www.hipaajournal.com/healthcare-data-leaks-on-github-credentials-corporate-data-and-the-phi-of-150000-patients-exposed/

Alzahrani, Saleh I., Ibrahim A. Aljamaan, and Ebrahim A. Al-Fakih. "Forecasting the Spread of the COVID-19 Pandemic in Saudi Arabia Using ARIMA Prediction Model under Current Public Health Interventions." Journal of Infection and Public Health 13, no. 7 (July 2020): 914–19. Available at: https://doi.org/10.1016/j.jiph.2020.06.001

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine Bias. ProPublica. Available at: https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. Nature, 533(7604), 452–454. Available at: https://doi.org/10.1038/533452a

Barrett, D. (2019, July 29). Capital One says data breach affected 100 million credit card applications. Washington Post. Available at: https://www.washingtonpost.com/national-security/capital-one-data-breach-compromises-tens-of-millions-of-credit-card-applications-fbi-says/2019/07/29/72114cc2-b243-11e9-8f6c-7828e68cb15f_story.html

Benjamin, R. (2019). Race after technology: abolitionist tools for the new Jim code. Polity.

Big Data for Justice Summer Institute. (n.d.). UCLA Bunche Center. Available at: https://bunchecenter.ucla.edu/programs-events/thurgood-marshall-lecture-2/ (Retrieved March 10, 2022)

Borsellino, R. (2018). Get Your Resume Past the Robots and Into Human Hands. The Muse. Available at: https://www.themuse.com/advice/beat-the-robots-how-to-get-your-resume-past-the-system-into-human-hands

Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. Conference on Fairness, Accountability and Transparency, 77–91. Available at: https://proceedings.mlr.press/v81/buolamwini18a.html

Carpentries. (2021). The Carpentries. The Carpentries. Available at: https://carpentries.org/index.html

Curriculum, UCLA. (2021). 2020-2021 Computer Science Curriculum. Computer Science Curriculum, UCLA. Available at: https://www.seasoasa.ucla.edu/curric-20-21/83-compsci-cur20.html

Curriculum, UCSD. (2021). B.S. Computer Science | Computer Science. UCSD Computer Science Major Requirements. Available at: https://cse.ucsd.edu/undergraduate/bs-computer-science

Curriculum, UCB. (2021). Requirements: Upper Division | Computing, Data Science, and Society. Available at: https://data.berkeley.edu/degrees/data-science-ba/upper-division

Curriculum, UCI. (2021). Computer Science, B.S. < University of California Irvine. Available at: http://catalogue.uci.edu/donaldbrenschoolofinformationandcomputersciences/departmentofcomputerscience/computerscience_bs/#requirementstext

Frey, W. R., Patton, D. U., Gaskell, M. B., & McGregor, K. A. (2020). Artificial Intelligence and Inclusion: Formerly Gang-Involved Youth as Domain Experts for Analyzing Unstructured Twitter Data. Social Science Computer Review, 38(1), 42–56. Available at: https://doi.org/10.1177/0894439318788314

Green, R., & Gilman, M. (2018). The Surveillance Gap: The Harms of Extreme Privacy and Data Marginalization. N.Y.U. Review of Law & Social Change. Available at: https://socialchangenyu.com/review/the-surveillance-gap-the-harms-of-extreme-privacy-and-data-marginalization/

Dennis et al. 2021. From a Data Archive to Data Science: Supporting Current Research. In Herndon, J. (Eds.). Data science in the library: tools and strategies for supporting data-driven research and instruction. (1st ed., pp. 99-109). Facet Publishing.

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2021). Ethics of AI in Education: Towards a Community-Wide Framework. International Journal of Artificial Intelligence in Education. Available at: https://doi.org/10.1007/s40593-021-00239-1

Karkkainen, K., & Joo, J. (2021). FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. 1548–1558. Available at: https://openaccess.thecvf.com/content/WACV2021/html/Karkkainen_FairFace_Face_Attribute_Dataset_for_Balanced_Race_Gender_and_Age_WACV_2021_paper.html

Larson, S. (2017, July 12). Verizon customer data leaked through an online security hole. CNNMoney. Available at: https://money.cnn.com/2017/07/12/technology/verizon-data-leaked-online/index.html

Lee, E., Lytle Hernández, K., & Tso, M. (2018). Policing Transitional-Aged Youth in Culver City. The Million Dollar Hoods Project. Available at: https://www.culvercity.org/files/assets/public/documents/city-manager/public-safety-review/policing-transitional-aged-youth-in-culver-city-for-website.pdf

Mihalcik, C. (2020, March 31). Marriott discloses new data breach impacting 5.2 million guests [2020]. CNET. Available at: https://www.cnet.com/tech/services-and-software/marriott-discloses-new-data-breach-impacting-5-point-2-million-guests/

Noble, S. U. (2018). Algorithms of oppression: how search engines reinforce racism. New York University Press.

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453. Available at: https://doi.org/10.1126/science.aax2342

O'Neil, C. (2016). Weapons of math destruction: how big data increases inequality and threatens democracy (First edition). Crown.

Orlowski, J. (2020). The Social Dilemma [Documentary]. Netflix. Available at: https://www.netflix.com/title/81254224

Patton, D. U., Pyrooz, D., Decker, S., Frey, W. R., & Leonard, P. (2019). When Twitter Fingers Turn to Trigger Fingers: a Qualitative Study of Social Media-Related Gang Violence. International Journal of Bullying Prevention, 1(3), 205–217. Available at: https://doi.org/10.1007/s42380-019-00014-w

Phan, L., Labou, S., Ali, I., & Foster, E. (2021, August 27). Ethical Considerations in Data. Available at: https://doi.org/10.5281/zenodo.5297228

Rocher, L., Hendrickx, J. M., & de Montjoye, Y.-A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. Nature Communications, 10(1), 3069. Available at: https://doi.org/10.1038/s41467-019-10933-3

Sandy, J. E., Herman, J., Keisling, M., Mottet, L., & Anafi, M. (2019). 2015 U.S. Transgender Survey (USTS) (ICPSR 37229). Resource Center for Minority Data. Available at: https://doi.org/10.3886/ICPSR37229.v1

Sinha, S. (2017). Ethical and Safety Issues in Doing Sex Work Research: Reflections From a Field-Based Ethnographic Study in Kolkata, India. Qualitative Health Research, 27(6), 893–908. Available at: https://doi.org/10.1177/1049732316669338

Thompson, B., Carosso, E., Griffith, W., Workman, T., Hohl, S., & Faustman, E. (2017). Disseminating Pesticide Exposure Results to Farmworker and Nonfarmworker Families in an Agricultural Community: A Community-Based Participatory Research Approach. Journal of Occupational & Environmental Medicine, 59(10), 982–987. Available at: https://doi.org/10.1097/JOM.0000000000001107

UCLA DSC. (2021). About | UCLA Data Science Center. Data Science Center, UCLA. Available at: https://www.library.ucla.edu/about-0

Universities of California. (2021). UC Love Data Week. UC Love Data Week. Available at: https://uc-love-data-week.github.io/uc-love-data-week.github.io/

Endnotes

[1] Leigh Phan, Data Scientist, University of California Los Angeles, UCLA Library Data Science Center, Email: leighphan@ucla.edu.
[2] Stephanie Labou, Data Science Librarian, University of California San Diego, UC San Diego Library, Email: slabou@ucsd.edu.
[3] Erin Foster, Service Lead - Research Data Management Program, University of California Berkeley, UC Berkeley Library & Research IT, Email: edfoster@berkeley.edu.
[4] Ibraheem Ali, Sciences Data Librarian, University of California Los Angeles, Louise M. Darling Biomedical Library, UCLA Library Data Science Center, Email: ibraheemali@ucla.edu.