Academic Advisory Committee on Social and Ethical Implications of Computing and Data Science

Year 2 Annual Report

July 15, 2025

Advisory committee members:

David Harding, *Professor of Sociology (Chair)*
Catherine Albiston, *Professor of Law, Director of Center for the Study of Law and Society*
Matilde Bombardini, *Professor of Business*
Sandrine Dudoit, *Executive Associate Dean, College of CDSS (ex officio)*
William Fithian, *Associate Professor of Statistics*
Joseph Gonzalez, *Associate Professor of Electrical Engineering and Computer Science*
Jennifer Johnson-Hanks, *Executive Dean, College of Letters and Science (ex officio)*
Ayesha Mahmud, *Assistant Professor of Demography*
Massimo Mazzotti, *Professor of History*
Rasmus Nielsen, *Professor of Integrative Biology*
Dara O'Rourke, *Professor of Environmental Science, Policy, and Management*
Alex Saum-Pascual, *Associate Professor of Contemporary Spanish Literature and New Media*

Table of Contents

1. Introduction

The “and society” advisory committee was appointed by – and reports to – the EVCP, with the following charge: “to advise the EVCP on how CDSS can fulfill its commitment to the study, understanding and teaching of the social and ethical implications of computing and data science. The advisory council will make recommendations to the EVCP concerning an integrated approach to faculty hiring, research, graduate curriculum, and undergraduate teaching programs that it believes CDSS should implement individually or jointly with faculty in other schools and colleges, as well as advise the EVCP on collaborations and synergies that it believes CDSS and other colleges and schools should pursue…the committee is not limited to considering the social and ethical implications of computing and data science solely in the context of CDSS. If it envisions opportunities that serve Berkeley’s research and teaching in those areas outside of CDSS, it is free, as well, to provide advice about those opportunities to the EVCP.”

Because the above charge is very broad, this committee has focused on general recommendations of strategies, structures, and initiatives regarding research and education in the social and ethical implications of computing and data science within and outside of CDSS. We leave to those charged with implementation the details of exactly how any particular recommendation would be carried out. We emphasize that, ultimately, efforts to further incorporate “and society” into CDSS and to build out capacity across campus at the intersection of computing and data science and the social sciences and humanities will be the collaborative work of faculty both in CDSS and across campus.

During the first year of its work, this committee focused on (i) developing a conceptual framework specifying what it believes “and society” entails, (ii) identifying challenges and making recommendations for undergraduate teaching, and (iii) cataloging both existing structures on campus that facilitate research at the intersection of data science and the social sciences and humanities, as well as what peer universities are doing at this intersection.

During this, the second year of work, this committee focused on three goals:

1. Undergraduate Education: Developing recommendations for integrating technical and social/ethical education in undergraduate programs, including (a) expanding opportunities for CDSS students to encounter social and ethical reasoning as an integral part of their technical education, and (b) expanding access to computer science, statistics, and data science education for students across campus.
2. Public Engagement: Recommending structures for managing and supporting public engagement around the social and ethical implications of computing and data science.
3. Interdisciplinary Research: Identifying strategies to catalyze interdisciplinary research on the social and ethical implications of computing and data science at UC Berkeley.

We have deliberately chosen the scope of our work to avoid overlap with the work of other committees or college processes. For example, because there is a Provost’s Advisory Council on Artificial Intelligence, we have not engaged directly with important questions about the

harms, risks, and opportunities surrounding the use of artificial intelligence. In addition, the Academic Senate Work Group on Generative AI is already "assess[ing] the opportunities and challenges of generative AI in the classroom and develop[ing] guidance for faculty." Nor have we weighed in on internal college processes to establish new departments or other types of units for Human Technology Futures or Interdisciplinary Data Science that are already underway, although we do believe that these units will be critical to implementing many of our recommendations and continuing to build bridging ties between CDSS and other units on campus that can contribute to teaching, service, and research related to the social and ethical implications of computing and data science.

On March 27, 2025, this committee received written comments from multiple Academic Senate committees on its year one report. We thank those committees for their serious and thoughtful engagement with our work, and look forward to future dialogue around these issues. Unfortunately due to shifts in the focus of our efforts from Year 1 to Year 2 and the timing of our receipt of those comments, we are not able to respond to all of them in this report. However, where there is a fit with our goals for this year, we have used the comments to inform our work.

2. Undergraduate Education

The committee continued to discuss current and potential ways to promote both i) exposure of CDSS students to social and ethical reasoning and ii) access for non-CDSS students to courses in data science, statistics and computer science. In this section we assess practices that are already under way and offer our views on possible ways forward, with the understanding that the ultimate decision regarding the CDSS curriculum rests with the CDSS faculty. It is important that such efforts reflect a diversity of approaches to social and ethical reasoning from multiple disciplines and perspectives.

Our discussion and recommendations will focus on three directions:

1. Creating "Society" integrated units in existing CDSS courses
2. Expanding course offerings for the HSD (Human and Social Dynamics) requirement within and outside the college
3. Expanding access of non-CDSS students to computer science, statistics and data science courses across campus

2.1 Creating "Society" integrated units in existing CDSS courses

In the current curriculum students are exposed to social and ethical reasoning in two broad ways:

- Dedicated courses: The Data C4AC, CDSS 94,[1] and Data C104 courses in the Data Science major have social and ethical instruction as their main curricular focus.

[1] CDSS 94 is a Special Topics course number. Planned offerings include "Anthropology of Science, Data, and Technology" and "Anthropological Study of Computational Thinking and Social Dynamics."

- Guest lectures/modules: Core courses, such as Data C8, Data 100, Data 102, and STAT 20 include some elements related to society and ethics. Examples of Fall 2024 course offerings that include a guest lecture are: fairness in housing appraisal in Data 100, and a guest lecture on air quality data in Data 6. Courses like Data 6 and Data 8 include applied examples on social and ethical topics like discrimination and racial bias in jury selection integrated into technical instruction.

We recommend strengthening this second approach with an expansion of these modules into more robust integrated units. Weaving social and ethical content into technical courses can serve as a complement to dedicated courses that establish foundational understanding, by allowing students to recognize these concepts when they naturally arise in applying technical ideas. It also has the advantage of building on existing courses rather than requiring the creation of new courses. Similarly, statistical and computing content could be integrated into social science and humanities courses through integrated units focusing on introducing technical concepts, perhaps by rebuilding and expanding the Data Science Modules program with greater involvement from senate faculty in CDSS.

These integrated units would be developed collaboratively by the Stat or EECS faculty teaching the course and faculty inside or outside of CDSS with expertise in related fields in both the specific domain (e.g. environmental toxins, racial segregation) and in societal and ethical considerations related to the question at hand. In addition to helping students see the connections between the technical course content and its social implications, well-chosen applied examples can be particularly valuable in helping students to understand technical concepts in more concrete terms, and to learn to think critically about their application.

The nature of these as integrated units, rather than stand-alone modules, would align with the comments by DIVCO about the Year 1 Report that “integrating a segment of other curricula into an existing syllabus was not pragmatic” and instead recommended “ensuring that societal and ethical considerations are thoroughly integrated into the curriculum, rather than superficially added." Feedback from one CDSS faculty member also cautioned against the “tendency to to *add,* but *not subtract,”* noting that “our undergraduates are under-served, under-resourced and over-burdened. A concern is curricula updates just add to the overall workload of students and faculty.“ Taken together, this feedback indicates that the best way to weave material into technical content may be to have it arise naturally in applications that also have technical interest, so that technical and social/ethical instruction are taught together. Ideally these integrated units should be developed to align with the main questions of the course and developed in consultation with an expert in the relevant scientific, social, or ethical domain.[2]

How to concretely promote these collaborations is an important issue that we discussed in our Year 1 Report and that requires resources, a discussion that we return to later in this section.

[2] DS has recently featured a news story about this kind of integrative work in Data 100 and Data 102 in collaboration with faculty from Data 104. This work brings social science concepts anchored in 104 into technical courses in the context of their learning objectives and examples.

2.2 Expanding courses that satisfy requirement for “Human and Social Dynamics of Data and Technology” (HSD)

The second recommendation of the committee is to continue expanding the set of courses, both within and outside the college, that satisfy the college-level Human and Social Dynamics requirement. This likely involves a combination of identifying existing courses and creating new courses, which would also benefit students from outside CDSS.

The purpose of the Human and Social Dynamics of Data and Technology requirement (abbreviated HSD), is “to ensure that students graduate with a basic understanding of how technology and data interact with human and societal contexts.” It is in the process of being implemented by the Undergraduate Study Committee of CDSS, who are converging on criteria while evaluating courses. The Undergraduate Study Committee includes a faculty member representing “Society” domains, and faculty teaching these courses across campus have been engaged by the committee in proposing syllabi. It has been suggested that the subcommittee reviewing HSD proposals for CDSS could include additional faculty outside CDSS, as is done for the Data Science Human Context and Ethics (HCE) requirement.

The HSD requirement overlaps with the related one-course Human Contexts and Ethics requirement for the Data Science (DS) major, but the HCE requirement can only be satisfied with an upper-division course, and there is a desire to have lower-division offerings for students to complete in their first year in CDSS. Currently all HCE courses are approved for the HSD requirement, and DS majors can use one course to count for both requirements.

Name of requirement	HSD (Human and Social Dynamics of Data and Technology)	HCE (Human Contexts and Ethics)
Required for which students?	All CDSS students	DS Major and Minor students
Satisfied w/ which courses?	Upper- and lower-division, with a goal of offering plenty of lower-division offerings to meet demand	Upper-division only
Body for approving courses	CDSS Undergraduate Study Committee (UGSC)	Data Science Governance Committee
Standards for judging courses	Under development by UGSC	Full description at DS major website The requirement aims to expand students’ ability to: - Analyze common assumptions of data-related

		practices and consider implications and alternatives; - Deepen their readiness to engage with non-technical forms of argumentation, perspectives rooted in social and humanistic inquiry, and diverse participants and publics; - Work meaningfully with human, societal, and ethical complexity; and - Strengthen crucial skills, centrally reading, writing, engaged presentation, and responsive discussion.
Courses most often taken to satisfy requirement?	TBD	Data C104

In recent years, 65% of DS majors satisfied their HCE requirement with Data C104 (along with many DS minors), in part because the other course offerings have limited space for DS majors and are not offered every year. Since the HSD requirement will apply to more students, more courses will need to be approved and the Undergraduate Study Committee is in the process of doing this.

Within the college, CDSS is piloting new courses which will satisfy the HSD requirement next year under the Special Topics course code, CDSS 94, including Artificial Humanities: AI, Language, and Fiction, Anthropology of Science, Data, and Technology, and Anthropological Study of Computational Thinking and Social Dynamics.

An additional course is planned for Spring 2026, in collaboration with Legal Studies, on the topic of “Cybersecurity in Context.” The course was previously piloted as a topics course in Legal Studies, but will have a new Data Science listing. Collaborating with Data Science on hiring and other operations seems to have been instrumental for the course to be able to scale from 30 seats to 100 seats, with an eventual 250 seats planned. There seems to be interest in partnering with CDSS among faculty in several other units, including the Goldman School, Political Science, and the Haas School of Business, so this may be a promising model for expanding course options.

This committee helped to identify several of the courses below that were subsequently approved by the CDSS Undergraduate Study Committee, and has identified others that are promising:

Approved HSD Courses

Approved initially:

Lower-division

- CDSS 94: Special Topics
- Data C4AC: Data and Justice
- L&S 22: Sense and Sensibility and Science

Upper-division

- American Studies / African American Studies C134: Information, Technology, and Society
- Data C104: Human Contexts and Ethics of Data
- Bioengineering 100: Ethics in Science and Engineering
- City Planning 101: Introduction to Urban Data Analytics
- Digital Humanities 100: Theory and Method in the Digital Humanities
- ESPM C167: Environmental Health and Development
- Info 188: Behind the Data: Humans and Values
- ISF 100G: Introduction to Science, Technology, and Society
- New Media 151AC: Transforming Tech: Issues and Interventions in STEM and Silicon Valley
- Philosophy 121: Moral Questions of Data Science

Approved in Spring 2025:

Lower-division:

- History 30: Science and Society (taught in Fall 2025)
- Philosophy 5: Science and Human Understanding (taught in Spring 2026)
- ISF 60: Technology and Values
- L&S 25: Thinking through arts & design

Upper-division:

- Anthro 168: Anthropology of Science, Technology, and Data
- New Media 133: Tech Wars: Security, Geopolitics, and Resilience of Digital Infrastructures
- Info 101: Introduction to Information Studies
- Info 103: History of Information
- Info 134: Information Technology Economics, Strategy, and Policy
- ISF 100D: Introduction to Technology, Society, and Culture
- ISF C100G / STS C100 / History C182C: Intro to Science, Technology and Society
- Hist 100S: Fossil Fuels and Climate History (this topic only)
- Hist 133D, Calculating Americans: Big Histories of Small Data
- Hist 183B: History of Science, Technology and Medicine in Modern East Asia

- Legal Studies 190, Cybersecurity in Context (this topic only)
- Political Science 132C: Algorithms, Public Policy, and Ethics
- Public Policy 138A: AI and Public Policy
- Public Policy 145: War? Politics, Security, and Emerging Technology
- Public Policy 147: Progress? An Introduction to Technology and Public Policy
- Public Policy C151: Space and National Security Policy
- Political Economy 156: Silicon Valley and the Global Economy
- Rhetoric 173: Science and the Public

The HSD option is currently being implemented in a way that emphasizes a variety of options for students, both in the social contexts and the type of data and technology involved. This is partly by necessity due to the very large number of student credit hours needed. An alternative model that might be worth pursuing, however, would be to come up with a more settled ethical curriculum that could even be used as a prerequisite for upper-division courses. This would be a significant undertaking that would ideally involve broad input from multiple units on campus, including all of the units in CDSS, and could possibly be implemented via multiple courses that differ thematically but share a core curriculum.

Collaborative development of upper-division “connectors”

The HSD requirement is formally an essential skill requirement focused on courses that teach foundational skills, such as lower-division courses or upper-division courses with few prerequisites. Upper-division “connector” courses that make use of technical prerequisites provide another promising model for undergraduate instruction that fully integrates advanced technical material with social and ethical issues. Legal Studies 126: Responsible AI, Law, Ethics, and Society is a potentially promising example: the course is aimed at data science / CS students as well as legal studies students, and pairs them in collaborative groups that leverage both types of skills. At a still more advanced technical level, CS 194: Understanding Large Language Models: Foundations and Safety is a course taught by EECS with a deep learning prerequisite that focuses on addressing risks from AI models.

2.3. Expanding access to computer science, statistics, and data science across campus

A third key priority is to ensure that students from outside CDSS have opportunities to learn key technical skills that will enable them to succeed in a world permeated by data and technology. There are at least three possible models for how these students would attain this education: (1) by taking service courses in CDSS, (2) by taking courses in their home departments or aligned departments, and (3) by introducing integrated units covering statistical and computing concepts into non-technical courses. We believe that different students will be better served by different models. As faculty on campus work to meet these needs, it will be important to build on what already exists and to avoid unnecessary duplication.

The service-course model is partially implemented already through courses such as Data 6 and 8, Stat 2 and 20, and CS 10 and 61A, which have few prerequisites; see next section.

- Advantages:
    - Allows for economies of scale in teaching many students from smaller majors
    - Ensures technical material stays current because the CDSS faculty is closer to research frontier
- Downsides:
    - It may be excessively technical relative to the needs of the students coming from a specific discipline
    - Does not necessarily expose students to how technical tools and methods are applied in their own disciplines
    - It may not expose students to the tools and methods currently being used in their own disciplines
    - It may be hard for students to learn technical material outside of the context of an application with which they are already somewhat familiar

The home-course model is the way many students in the social sciences learn how to apply statistics in their disciplines.

- Advantages:
    - Giving students exposure to how the methods are used in their own disciplines
    - Ensuring that courses are taught at an appropriate technical level so students don't need to take multiple prerequisite courses in CDSS. Berkeley students vary widely in their technical preparation. Some of our students enter Berkeley with a strong technical background and breeze through technical courses; for many others, even understanding basic statistical concepts like p-values and linear regression is difficult. For this latter (large) group of students, understanding modern ML/AI methods, like large language models or computer vision, might realistically require that they take a multi-course sequence. Alternative curricular goals could center on how to use AI rather than how it works, for example a course on how to use LLMs and how not to use them. Collaboration between CDSS and non-CDSS faculty could be particularly productive here. Berkeley students whose own major requirements do not overlap with the prerequisites for technical service courses may also not have room in their course schedules for an entire sequence of technical courses. This problem may be especially acute for transfer students.
- Disadvantages:
    - students in smaller departments, or in the Arts and Humanities, may not have access to such courses
    - not all departments may include faculty with the expertise to teach technical courses
    - There is a risk of variation in quality and consistency when statistics and computing courses are taught by faculty outside of CDSS.

A possible intermediate solution for courses outside of CS/DS could have one or more of the following features:

- They would be less technical and more applied
- They would have few prerequisite (some of the classes listed below, like Data 6, fit this profile)
- They could be service courses taught in CDSS or they could be taught by a pool of faculty outside CDSS units with a common curriculum, designed in collaboration with CDSS faculty
- They could be aimed at groups of students from adjacent majors, such as those in the social sciences division
- They could be revised versions of courses already being taught outside of CDSS that are updated to include ML/AI methods and their application in a particular domain
- Multiple courses could be coordinated such that parts of the courses are shared common technical material taught by CDSS faculty across courses and parts are specific to a discipline

The integrated unit model would be similar to the integrated units discussed in section 2.1, except that the integrated units would introduce statistical and computing skills into non-technical domain courses.

- Advantages:
    - Relatively lightweight and scalable
    - Would allow for reciprocal exchange of faculty effort when paired with other integrated units
- Disadvantages:
    - There may be time for only the most basic of technical skills
    - Training of ASEs may be required

Existing large service courses in CDSS

A number of large lower-division courses in CDSS currently teach DS, CS, and stats to students from a wide range of backgrounds and disciplines.

Examples of courses that satisfy requirements in CDSS majors, but which serve a large number of students are:

- Data C8: Foundations of data science (1472 students in Fall 24)
- Stat 20: Introduction to probability and statistics (~800 students in Fall 24)
- CS 61A,B,C: Intro computer science sequence (1210 in Fall 24 offering of 61A, 165 open seats)

In addition, there are less technical versions of these courses which are not taken by CDSS majors because they do not satisfy the lower-division prerequisites:

- Data 6: Introduction to computational thinking with data (50 students in Fall 24)
- Stat 2: Introduction to statistics (503 in Fall 24)
- CS 10: The beauty and joy of computing (173 seats in Fall 24, not full)

- Another course with fewer cumbersome prerequisites that teaches some machine learning is Data 100, although it has two prerequisites and a math co-requisite, which itself has two prerequisites.

It can be very difficult for students outside of CDSS to enroll in upper division courses. For example, CS 189 (Introduction to Machine learning) is perpetually difficult to get into; the Data Science major pays for additional sections so their students can take it.

It may be difficult to offer many more service courses, or to open up many more seats in these courses, because CS, DS, and Statistics faculty and TAS budgets are currently stretched very thin (see next section). If the Interdisciplinary Data Science Institute proposal is adopted, then faculty in that institute may be able to teach more courses of this type.

An important question, that we lack data to answer at the moment, is whether we can quantify the number of students who are not currently majoring or minoring in a CDSS degree who have an appetite for taking upper-division technical courses in ML or AI that may require significant prerequisite coursework.

2.4 Implementation challenges: teaching across units

The three directions we have discussed in this section all require, to different extent, faculty teaching students that are not in their unit. The current system makes it difficult to properly compensate for this cross-unit teaching and we recommend actions be taken to relax this constraint.

In Section 2.1 we recommended the creation of “Society” integrated units. This may require engaging non-CDSS faculty members in a time-consuming but potentially impactful re-thinking of core curriculum and thoughtful organic integration of social and ethical discussion. This time would have to be compensated through a partial course credit/course release, for example, or could involve an exchange of teaching across courses between different units.

In Section 2.2 we discussed expansion of HSD course offerings. There are many faculty across Berkeley with the expertise to teach courses about the intersection of data, technology and society, and the social impacts of data and technology may well be at the top of the intellectual agenda for many of Berkeley’s humanities and social sciences departments and professional schools. Moreover, we have found that many faculty are enthusiastic about teaching CDSS students. It should be possible for Berkeley to offer a wide variety of options so that students can choose a course on a topic that interests or inspires them.

One practical issue that can limit Berkeley students’ opportunities to take courses in other units is a general shortage of TAS funds across campus, which tends to disincentivize units from teaching students from outside their own unit. Partly as a result of this, many courses have few seats for students from other disciplines. Strategies for addressing a shortage of TAS resources is beyond the scope of this report, however.

3. Public Engagement

Berkeley has several strengths in terms of public engagement on data science and society. One strength is Berkeley's public mission, which includes a focus on public service, civic engagement, and research that serves the community. This strength distinguishes Berkeley from private competitors. Second, Berkeley has a strong tradition of interdisciplinarity, including collaboration across disciplinary boundaries, low fences among disciplines and schools on campus, and faculty who enjoy and are adept at working across these divides. Conversations about accountable AI repeatedly emphasize that interdisciplinary expertise is necessary for accountable technology and functional guardrails. Third, Berkeley leads in entrepreneurship of all kinds, including nonprofit public service start ups as well as more traditional start ups. Finally, Berkeley offers a diversity of views on technology regulation and governance, ensuring vigorous debate and vetting of policy proposals, and avoiding industry capture.

Several dimensions of public engagement emerged in our discussion. Some forms of public engagement will involve interactions with federal and state executive and legislative branches, with local government, and with the press. Other forms of public engagement include conducting community-engaged scholarship, sharing scholarly findings with the California public, interacting with community groups and members of the public, and fostering technology "for good" or public interest technology.

For example, UC Berkeley is actively collaborating with the California government to guide the responsible development and deployment of artificial intelligence (AI) technologies. Governor Gavin Newsom's Executive Order N-12-23, acknowledges UC Berkeley's College of Computing, Data Science, and Society as a leading institution in generative AI (GenAI) research, positioning it as a key partner in shaping state AI policy and risk assessment. The CDSS Dean's Office, alongside Stanford's Institute for Human-Centered AI and State government partners, designed and convened a Joint California Summit on GenAI in May 2024–bringing together leaders from numerous disciplines and sectors to discuss how the state can best use GenAI to better serve the people of California. The CDSS Dean's office also facilitated input from multiple colleges, schools, and institutes to various California state government reports related to GenAI's use, public benefit, and impact. A key goal of this committee is to ensure that cross cutting proactive efforts like this continue.

Looking to the future, we identified several key areas for Berkeley to have a leading voice in the public dialog about emerging technology and society. These include:

- Informing the legislature and broader public about the implications of possible regulatory approaches regarding computing and data science at the development stage, through appropriate channels. Monitoring the legislative process will be key for this area of public engagement.

    - Conducting, encouraging, and rewarding community engaged scholarship regarding data science and society. For example, this would include recognizing community-engaged scholarship in tenure and merit review decisions, and providing support for community-engaged scholarship. (see BMAP guidance on assessing and crediting community-engaged research in AP cases).

    - Building off the success of the Data Discovery Program, expanding opportunities and support for CDSS students to become involved in policy work, community-engaged scholarship, and public interest oriented research and efforts. This could include programs that facilitate policy engagement (in D.C. or Sacramento), fellowships for specific public engagement activities, course units for community-engaged or public interest projects, and research support for graduate students whose work engages with policy, communities and/or the public interest.

    - Identify opportunities to inform policy discussions and offer expertise on best practices regarding technology, data science, and society. This forward looking proactive role will help people on campus keep up with the pace of policy discussions, which move faster than research. It will also provide timely input when foundational policies about technology and data science are being set.

- Chancellor's Advisory Committee on Computing, Technology, and Society.
    - We propose establishing a standing Chancellor's Advisory Committee on Computing, Technology, and Society (CACCTS). The committee would be explicitly designed to be interdisciplinary and representative of faculty engaged in data science and society issues from across campus. It would include at least one appropriate member from each of the following: (1) CDSS, (2) social sciences including but not limited to sociology, history, political science, economics etc.; (3) humanities, including but not limited to philosophy, literature, and especially members with ethics expertise; (4) engineering, including areas outside of computer science; (5) biological and physical sciences; (6) natural resources and environmental design; (7) the professional schools including law, public policy, education, social welfare, public health, business, journalism, and information. We suggest that the Vice Chancellor for Research appoint members of the committee and provide administrative support for the committee's work. This committee would be similar to existing Research Advisory Groups and Initiatives, but would report to the Chancellor as well.

    - CACCTS Mission/Tasks. CACCTS will serve a proactive role in public engagement, tracking each of the key areas identified in section 1. The committee will also help identify major policy or legislative issues and emerging public interest concerns regarding technology, data science, and society. The Chancellor may also call on CACCTS to advise on questions about technology, data science, and society.

- Public engagement working with existing Berkeley structures and organizations.

    - There are many existing entities at Berkeley that are working on public engagement that could be helpful partners on public engagement. These include Goverment and Community Relations (GCR); Vice Chancellor for Research; Berkeley's Institutional Change Initiative focused on Community Engagement; the Berkeley Center for Law and Technology (BCLT) (especially on legislation and policy); the Institute for Governmental Studies **(**IGA) which provides some fellowship support for students working in D.C.; the CDSS Data Discovery Program.

    - CACCTS should collaborate and meet regularly with these existing entities at Berkeley to partner on issues related to technology, data science and society. In addition, several emerging entities could be well-aligned with the proposed CACCTS and this committee's recommendations. Efforts over the last year (May 2024-ongoing) across Berkeley include plans for a campus-wide Tech Policy Network and/or Tech Policy Initiative. Deans from several UC Berkeley colleges and schools including Business, CDSS, COE, Information, Law, and Public Policy, as well as the Director of CITRIS, appointed a working group to assess the landscape and identify opportunities related to technology and society. The group's report and recommendations include suggestions for fostering timely tech policy research, establishing Berkeley as a resource for policymakers, and serving as a central resource to offer expertise for government agencies and legislative briefings.

    - We recommend that there be a point person or entity to track emerging legislation and policy issues and consult with CACCTS and GCR, as well as other entities as appropriate. Consistent with our recommendation, the working group mentioned in 3.2 above recommended that to optimize interdisciplinary collaboration and autonomy, the suggested staffing of an Executive Director, Policy Director, Communications/Events, Development, Administration, Academic/Research Advisor and student assistants be placed within VCR, EVCP, or another campus-wide office rather than a single school or college.

    - We recommend that CDSS expand meaningful opportunities and support for students to become involved in community-engaged scholarship, and public interest oriented research and efforts including collaborations to inform policy and to learn about public service. Collaboration with IGA, BCLT, and the Berkeley Institutional Change Initiative can facilitate this effort.

    - We recommend that the office of the Vice Chancellor for Research work with CACCTS to develop a comprehensive web presence highlighting faculty and student research, and student opportunities related to policy, legislation,

community needs and engagement, and the public interest. This information should include work from both CDSS and outside of CDSS on campus. The office of the Vice Chancellor for Research could update and maintain this web presence. A public engagement web page could highlight faculty and student research, and student opportunities related to policy, legislation, community needs and engagement, and the public interest. The web presence will help the public find useful information and help students become more involved in community-engaged research and work related to policy and the public interest.

- We recommend that a working group to better understand and inform emerging legislation/policies be created, structured as a collaborative effort among CACCTS, GCR, and the emerging multi-school/college Tech Policy Network/Initiative that includes engagement from the Deans of Business, CDSS, COE, Information, Law, and Public Policy, as well as the Director of CITRIS. Student involvement in this work where appropriate would be desirable

4. Catalyzing Interdisciplinary Research

It is both a strength and challenge at Berkeley that in new issue areas, we often see the rapid emergence of multiple research initiatives approaching the topic from diverse theoretical, empirical, and methodological perspectives. This interdisciplinarity has led to cutting edge research and rapid public engagement on these topics. We are experiencing this now in the broad area of the “Social and Ethical Implications of Computing and Data Science,” and in particular in researching the rapidly changing applications and implications of Artificial Intelligence. Interdisciplinarity is essential in this domain due to the scale and complexity of both the opportunities and challenges it presents. Technical, social scientific, and humanistic perspectives are necessary. Fortunately, Berkeley is well-positioned to be a leader in the careful interdisciplinary examination of the implementation and use of AI in social life and social, economic, and cultural institutions by virtue of its long history of scholarship on the interconnections between society and technology.

Berkeley’s strength in the study of society and technological change

The expansive use of the term “AI”, and its increased relevance in expert as well as public discourses, has fueled a number of initiatives that aim to link forms of generative artificial intelligence with other campus priorities and scholarly traditions. While AI presents new and specific challenges and opportunities, we believe that these are best understood in a long-term perspective. There are rich and lively research traditions that, in the last few decades, have engaged with technological change and its implications. It is their tools and methods that promise, once again, to be a good starting guide for our understanding and action.

Berkeley is especially well positioned to lead in this area, as it has long standing research programs that have explored the nature of scientific and technological change, starting from

Thomas Kuhn's seminal work on the structure of scientific revolutions.[3] Since then, a number of initiatives and units have pushed forward the research agenda, adding new dimensions and creating connections across the entire campus.

The technological shift that the words "and AI" try to capture is part of the history of mechanization and automation and, as such, many of its features are familiar to the specialist. Yet, it is also a process that has distinctive technical traits and unfolds within new social and political conditions, which call for new models and explanations. We believe that Berkeley should mobilize its existing research capacity across campus to best respond to this challenge.

A necessary condition for the success of this synergistic effort is shared leadership. Given that the relevant faculty belong to different departments and divisions, we believe that any "and AI" initiatives should be supported and co-led by the relevant deans. In particular, we envision an essential role for the Dean of Social Sciences and the Dean of Arts and Humanities. Building trust across units and divisions starts from faculty collaboration, but it needs the committed support of deans as well.

Next Steps

The Committee's Year 1 report, highlights the many programs across UC Berkeley currently conducting related research, including (but not limited to):

- Berkeley Program in Science, Technology, and Society
- Human Technology Futures
- Public Interest Technology
- Human Contexts and Ethics
- BAIR Responsible AI Initiative
- Center for Human-Compatible Artificial Intelligence
- AI, Platforms, and Society Center
- Center for Technology, Justice, and the Courts
- Technology and Work
- Algorithmic Fairness and Opacity Group
- Kavli Center for Ethics, Science, and the Public
- Berkeley Initiative for Transparency in the Social Sciences
- UC Berkeley AI & Society
- Berkeley Institute for Data Science

We thus face a challenge and an opportunity to better coordinate, support, find synergies, and leverage these distributed research initiatives for greater impact. We believe there is interest across campus for better collaboration. However, historically it has rarely worked at Berkeley to simply centralize a research area.

[3] Kuhn, Thomas S. *The structure of scientific revolutions.* Chicago: University of Chicago Press, 1962.

Mechanisms for Collaboration

We believe there are benefits to creating fairly lightweight mechanisms for sharing, inclusion, and potential collaboration on research, so that faculty and other researchers can be engaged beyond those who will be part of the new CDSS units in Human Technology Futures and Interdisciplinary Data Science. The goal is to connect people, centers/institutes/labs, and ideas that do not currently interact without creating a new time burden on faculty and other researchers. The Berkeley Institute for Data Science (BIDS) would be a natural catalyst for such efforts given its prior history of drawing together scholars across disciplines.

To be successful, such a mechanism would need to have the following properties:

- Promote real inclusion of both STEM faculty and faculty from the humanities and social sciences;
- Provide repeated positive interactions to build trust in a new intellectual community and build understanding of the intellectual frameworks and practices from other disciplines;
- Provide an intellectual (not just logistical or administrative) space for “bottom-up” generation of projects, collaborations, and proposals;
- Create structures that can support collaboration and “organic” matching of ethical and social implications of AI projects with computing and data science researchers as intellectual involvement deepens over time;
- Minimal time commitment for people with many existing commitments; and,
- Opt-in, but some commitment to sustained participation over time (e.g. a semester or year).

Lessons from other Interdisciplinary Berkeley Research Collaborations

In this section we begin to survey some existing structures on campus that might be useful models of interdisciplinary collaboration, whether or not they are focused on computing or data science. There are a number of examples at Berkeley where efforts have been made to support interdisciplinary collaborative research across campus. The below is not intended to be a complete list, but rather a set of models from which we might learn about effective structures and practices.

- Human Technology Futures (HTF) group draws together around thirty faculty from a dozen departments who work to shape ethical, just, and livable technological futures. It crosses boundaries between technical disciplines and the humanities, interpretive social sciences, and the arts. Current transdisciplinary research areas include: conceptual foundations of data science and computing, their broader ramifications in society, power dynamics in creating and implementing technologies, data aesthetics, and the materialization and humanization of data through art. Many members of the group are active within the Center for Science Technology, Medicine and Society (CSTMS), an established research unit that administers the Berkeley Program in Science, Technology, and Society, which offers a PhD Designated Emphasis and undergraduate minor.

- Berkeley Climate Change Network (BCCN) connects 300+ people and 80+ programs — institutes, centers and labs — focused on climate change research across the campus and at Lawrence Berkeley National Lab. BCCN works with UDAR, VCRO, BRDO, and SkyDeck to help researchers find funding and investment opportunities from government, foundations, donors, and venture capital. BCCN publishes a twice-monthly email newsletter and hosts disciplinary working groups - such as a "Berkeley Climate Policy Group" - to discuss cross-campus collaboration. https://bccn.berkeley.edu/

- Berkeley Economy & Society Initiative "forges collaborations across fields to reframe the study of political and economic life. Hard questions are our starting point." Participants are based in a wide range of departments and programs, including political science, sociology, economics, history, law, public policy, geography, and environmental and data science. "BESI sponsors innovative faculty and graduate student research, convenes event series and international conferences, and enhances graduate student training through the Designated Emphasis in Political Economy, a new minor field option for students across Berkeley's PhD programs." BESI organizes its work around three research clusters: Capitalism and Democracy, Climate, and Technology, and hosts research talks, book panels, and working group meetings. https://besi.berkeley.edu/

- Innovation & Entrepreneurship (I&E) at Berkeley coordinates applied research across campus, including 90 leaders of diverse programs that meet monthly as the I&E Council with the goal of knitting together Berkeley's innovation ecosystem, as well as helping connect Berkeley researchers to outside funders and industry partners. (https://iande.berkeley.edu/). BEGIN (Berkeley Gateway to Innovation) then serves as a directory of innovation and entrepreneurship resources, including grants, competitions, and connections to research partnerships, as well as a hub for events and news. https://begin.berkeley.edu/

- The Social Science Matrix Research Teams program provides interdisciplinary groups of faculty and graduate students with yearlong or semester-long support for meetings, speakers, and pilot research projects to "explore or develop a novel question of significance in the social sciences." "Matrix is especially interested in original and emerging approaches that explore new theoretical and empirical questions, and that combine research at different scales and from different methodologies." https://matrix.berkeley.edu/initiatives/matrix-teams/

- "The Townsend Center Working Groups program supports research in the humanities and the humanities-related social sciences. The Townsend Center funds approximately 70 groups per year, ranging across a rich array of subjects and disciplines. The aim of these working groups is to bring together faculty and graduate students at Berkeley to create scholarly dialogue and foster the free exchange of ideas on shared research interests." https://townsendcenter.berkeley.edu/funding/working-groups

Thus far, we have observed that the practices below have often led to productive interdisciplinary collaboration, dialogue, and prolonged engagement:

- Explicit framing of the importance of interdisciplinarity in group goals, including intellectual humility and a willingness to learn from faculty from other fields.
- Sustained involvement to provide time for building relationships and understanding between faculty from different fields.
- Pilot grants to support new work at the boundaries between disciplines that require collaborators from different disciplines.
- Staff with expertise in community-engaged scholarship to facilitate faculty engagement with community partners.
- Access to development staff who can assist with fundraising for interdisciplinary research teams, labs, or centers.
- Enthusiastic graduate students, who often play a role in linking faculty across disciplines, and funding to support them.
- Multiple avenues for participation: from intermittent seminar participation to co-equal research collaboration (this “concentric circles” model allows for various levels of frequency and intensity of involvement over time and across participants).
- Joint meetings of existing seminar groups to promote interaction between them.

Recommendation

To jumpstart collaborative, interdisciplinary research on the social and ethical implications of computing and data science, we recommend that campus and CDSS initiate a program to fund interdisciplinary pilot grants, working groups, and seminars. To ensure interdisciplinarity, this initiative would be administered by the VCR and funded projects would be required to have faculty from both computing or statistics and the social sciences or humanities. Initial grant calls might emphasize areas of recognized priority, such as the environmental impact of AI or its role in learning and education.

There are multiple options for funding such an initiative, including philanthropy and state funds. One other option would be to allocate some gift funds provided by technology industry companies to CDSS labs for technical computing research to research the social and ethical implications of computing and data science. Such a scheme has the potential to provide a novel source of funding and to explicitly link technological development to the real-time study of its social and ethical implications. Clearly, gift funds would be a viable option only if they do not impose limitations on the research methods or the breadth of investigation. An interdisciplinary faculty committee should be created to evaluate and review applications for funding for small grants. We recommend that this committee include faculty from across campus. In this way, it would function as a bridge connecting different colleges and units, fostering communication and coordination. Ultimately, collaborative “and Society” research based in CDSS will succeed if it can effectively mobilize, connect, and amplify the unique collective expertise of Berkeley scholars researching the interplay between technology, humanities, and society.

5. Conclusion

This Year 2 report of the Academic Advisory Committee on Social and Ethical Implications of Computing and Data Science has endeavored to provide realistic recommendations to the ECVP regarding ways to better integrate social scientific and humanistic intellectual perspectives, approaches, and knowledge with undergraduate education, public engagement, and research in computing and data science, both within CDSS and across the campus. We conclude by summarizing those recommendations.

With regard to undergraduate education, our recommendations fall into three categories:

- Integrated units: In order to integrate "society" into existing CDSS courses and technical material into social sciences and humanities courses, we recommend strengthening the current standalone "module" approach by expanding the existing modules into integrated units and creating new integrated units that weave together social and ethical content with technical content through collaboration between the primary instructors for the course and other faculty with relevant expertise.
- Expanding HSD course offerings: In order to increase the scale and range of opportunities for CDSS students to learn social and ethical content, we recommend increasing the number and size of courses both inside and outside CDSS that meet the CDSS college-wide Human and Social Dynamics of Data and Technology Requirement (HSD). Upper division "connector" courses are one strategy among many that may be worth exploring further.
- Expanding access to CS, Statistics, and DS across campus: In order to improve access to data science methods and tools among students outside of CDSS, we recommend a two-pronged approach in which students with the time and technical background be given greater access to existing CDSS courses while others be given access to new courses in their majors or related fields that are more tailored to their background and applications in their major, perhaps requiring only Data 6 or Data 8 as a prerequisite.

With regard to public engagement, we provide a suite of related recommendations that would leverage Berkeley's long standing strengths at the intersection between technology and society and facilitate nimbler, broader, and deeper public engagement on computing, technology and society. These recommendations include a centralized website for public and policymaker access to Berkeley expertise, greater opportunities for student engagement, and more support for community engaged scholarship in this domain. Our core recommendation is the formation of a well-staffed Chancellor's Advisory Committee on Computing, Technology, and Society to lead in this space, proactively prepare for upcoming policy engagement, and broaden the range of campus experts who are publicly engaged around computing, technology, and society.

Our third type of recommendation aims to catalyze interdisciplinary research at the intersection of computing, data science, and society. We recommend an interdisciplinary pilot grant program administered by the VCR that will spur truly interdisciplinary collaboration and leverage Berkeley's strengths in both the technical fields and the study of society and technological change.