



BUILDING DATA AND AI ETHICS COMMITTEES



Organizations face a difficult challenge when it comes to ethically-informed data collection, sharing, and use. On the one hand, there is increased sensitivity to ethical issues and desire for responsible stewardship of people’s information, both internal and external to organizations. On the other hand, outside of mere legal compliance, there is very little guidance for organizations about what being ethically responsible involves, let alone how to incorporate ethical consideration into product and service design, particularly at scale. Even among organizations that have adopted institutional value statements for data/AI ethics platforms, there has been limited success translating those statements to organizational practices, decision-making, and products.

This situation is problematic for stakeholders from within and outside the organization. For individuals, privacy, data control, security, and fairness are at stake. For organizations, there are new risks concerning data misuse and insecurities, with potential for loss of trust from users and consumers. Conversely, organizations that develop capacity for ethical collection, use, and sharing of information will be the ones most trusted by customers and the people they serve, which is increasingly crucial to success.¹

One strategy to develop capacity for data and AI ethics that is often proposed, and being explored by some organizations, is instituting ethics committees.² However, there is virtually no guidance on what building such a committee involves or on how such a committee would function once created.

The objectives of this report are to:



DISCUSS

the advantages of a committee-based approach to data and AI ethics



DESCRIBE

the components of a committee-based approach to data and AI ethics



IDENTIFY

the questions that an organization would need to answer in the process of developing an effective ethics oversight committee

An ethics committee is a potentially valuable component of accomplishing responsible collection, sharing, and use of data, machine learning, and AI within and between organizations. However, to be effective, such a committee must be thoughtfully designed, adequately resourced, clearly charged, sufficiently empowered, and appropriately situated within the organization.

CONTENTS

PART ONE: DATA ETHICS AND ETHICS COMMITTEES	4
01 Why the need for Data and AI Ethics?	4
02 Approaches to Building Organizational Capacity in Data and AI Ethics	5
03 Why Use an Ethics Committee?	6
PART TWO: BUILDING AN ETHICS COMMITTEE	8
01 Organizational Function: Why is the Committee Being Created?	9
02 The Ethical Content: What Should Guide the Committee's Activities and Decision-making?	9
Foundation values: What are the basic values that the committee is meant to protect and promote?	9
Core principles: What are the primary guiding principles in support of the values?	11
Substantive content and deliberative resources: What is required in practice to satisfy the core principles?	11
03 Committee Composition: Who Should be on a Data and AI Ethics Committee?	14
What are the types of expertise needed?	14
How can the committee avoid bias and conflicts of interest?	16
How are committee members selected and how long do they serve?	16
How large should a committee be?	16
04 Organizational Position and Powers: Where Should an Ethics Committee be Situated within an Organization and what Should be its Powers?	17
Where is the committee housed within the organization?	17
What is the purview of the committee?	17
When should the committee be consulted?	18
Who is required/permitted to consult with the committee?	18
What authority does the committee have?	19
05 Procedures and Governance: What are the Processes and Procedures that Govern the Work of the Ethics Committee?	20
What is presented to the committee?	20
Is there an expedited review process?	21
What are the procedures and rules for external consultation?	21
What are the standards by which committees make judgments?	21
How do committee members collectively issue a decision?	22
What is the timeline for committee review?	23
How are committees audited and evaluated?	23
CONCLUSION: GETTING STARTED	24

PART ONE:

DATA ETHICS AND ETHICS COMMITTEES

01 Why the need for Data and AI Ethics?

Legal compliance, especially with privacy regulations, is currently the dominant framing for issues concerning data collection and use. However, data and AI ethics involves much more than mere legal compliance. Not using people's information illegally is the **minimum responsibility** that organizations have, and it is not sufficient for maintaining trust, managing risk, realizing organizational values, or being responsive to public expectations. It is now clear that neither the public nor employees view merely complying with the law as good enough.

Moreover, legal guidance regarding data collection and management, including its use in AI and machine learning, lags well behind technological innovation and organizational practices. As a result, the ethical issues and questions confronted by organizations and the people working in them often arise **prior** to the development of an adequate legal regime.

Building sufficient data and AI ethics capacity can serve a key role in **maintaining trust** between organizations and the people they serve, whether clients, partners, employees, or the general public.³ This is particularly so given the current context of frequent, high-profile cases of negligent and irresponsible use of information, which has sensitized people to ethical concerns. Data ethics capacity can help organizations **manage risks** and liabilities associated with such data misuse and negligence.⁴ It can also help organizations clarify and make actionable **mission and organizational values**, such as responsibilities to and respect for the people and communities they serve.⁵ Finally, data ethics capability offers organizations a path to address the transformational power of data-driven AI and machine learning decision-making in an anticipatory way, allowing for **proactive responsible development** and use that can help organizations **shape good governance**, rather than inviting strict oversight.⁶

Thus, the organizational need for developing data and AI ethics capacity is grounded in the realities of operating in a complex environment — one with novel risks and liabilities where expectations for responsible use exceed legal requirements. People increasingly want organizations to respect values such as autonomy, fairness, justice, and rights in their data and AI practices. Overall, ethics capacity is necessary because the collection, use, and sharing of information, particularly in combination with machine learning and AI, impacts people's lives and reshapes the social, economic, and political landscape.

Approaches to Building Organizational Capacity in Data and AI Ethics

Options for building organizational capacity to anticipate and address ethical issues associated with data collection and use, as well as to identify opportunities for using data in ways that support organizational mission and values, include:

- Appointing **chief data/AI officers** with ethics as part of their responsibilities⁷
- Assembling organizationally high-level **ethics advisory groups**⁸
- Incorporating privacy and **ethics-oriented risk and liability assessments** into decision-making or governance structures
- Providing **trainings** and **guidelines on responsible data practices** for employees⁹
- **Developing tools, organizational practices/structures, or incentives** to encourage employees to identify potentially problematic data practices or uses¹⁰
- Using a **data certification system** or **AI auditing system** that assesses data sourcing and AI use according to clear standards¹¹
- Including members responsible for representing **legal, ethical, and social perspectives** on technology research and project teams¹²
- Creating **ethics committees** that can provide guidance not only on data policy, but also on concrete decisions regarding collection, sharing, and use of data and AI¹³

These approaches are not mutually exclusive. In fact, they are complementary. A chief ethics officer might chair an ethics oversight committee, some of whose members serve on an ethics advisory council or develop and deliver ethics trainings or consultations.

Ethics committees should be considered as a component of building a data and AI ethics ecosystem within and between organizations.

03

Why Use an Ethics Committee?

A committee-based oversight model has been used effectively in several contexts, such as: protecting human research subjects (Institutional Review Boards or IRBs); providing policy perspectives, clinical guidance, and patient consultations in medical contexts (hospital ethics committees); providing oversight on embryonic stem cell research (Embryonic Stem Cell Research Oversight Committees or ESCROs); and ensuring legal and humane care of research and laboratory animals (Institutional Animal Care and Use Committees or IACUCs). Committee-based oversight models like these have several features that make them well-suited for building organizational data and AI ethics capacity. A well-designed committee will:



Bring together people with the range of expertise — technical, legal, ethical, and organizational — needed to effectively analyze, assess, and respond to complex problems



Be responsive to rapid advances in technological capabilities and to novel applications



Develop standards, cases, precedence, and resources to be used in decision-making processes



Constitute a governance body that can learn, adapt, and be a repository for institutional knowledge

The ethical issues and questions associated with the use of big data, machine learning, and artificial intelligence are emerging and dynamic. There is not an established field of data ethics with well-developed and comprehensive resources that can be drawn from to address them, and a strictly compliance, monitoring, and enforcement model is inadequate. In this context, a collaborative approach to social and ethical analysis and evaluation that incorporates multiple perspectives, areas of expertise, and areas of concern can be effective. Moreover, an ethics committee will contribute to developing substantive guidance and resources not only for managing digital risks and addressing ethical concerns for the organization, but also for helping to shape data collection, sharing, and use norms and practices more broadly.

An ethics oversight committee can be an important component of organizational capacity for promoting responsible data, machine learning, and AI practices and uses. However, building a well-functioning and effective committee requires that it be adequately resourced, sufficiently empowered, and thoughtfully designed.

COMMITTEE-BASED RESEARCH OVERSIGHT IN THE UNITED STATES

Committee-based oversight is an established approach to ensuring compliance with legal, regulatory, and ethical standards governing scientific research involving human and nonhuman subjects. In the United States, committee-based oversight is a legal requirement for nearly all research involving humans and vertebrate animals. The committees that oversee human subjects research are called Institutional Review Boards (IRBs). Review by an IRB is mandated for all human subjects research that is federally regulated or receives federal funding.

Research subjects oversight arose as a direct result of ethical misconduct, such as in the Tuskegee syphilis studies and the Willowbrook children studies.¹⁴ In these cases, researchers directly misled research subjects, chose research subjects from vulnerable populations, and exposed research subjects to unnecessary risks and harms. In response, regulators sought to enshrine standards for ethical research in law. In the United States, protections for human research subjects is governed by three primary values: Respect for Persons, Beneficence, and Justice. These values are embodied in general principles, such as those requiring informed consent, fairness in choosing research subjects, and harm minimization. Given the breadth of the values, the generality of the principles, and the diverse forms of human subjects research, it is not feasible to enshrine specific and exhaustive guidelines in law to ensure that ethical research is accomplished. Instead, it is the role of IRBs — whose composition is subject to federal guidance to ensure competence — to make judgments about whether research projects are consistent with the values and principles. No IRB-governed research may begin prior to IRB approval.¹⁵

Another form of research that is subject to committee-based ethics oversight in the United States is human-derived embryonic stem-cell (hES) research, which is overseen by Embryonic Stem-Cell Research Oversight (ESCRO) committees. The development of ESCRO committees for overseeing hES research is notable in the context of thinking about data and AI ethics committees because they are not federally mandated. In 2001, the United States government placed significant restrictions on the use of federal funding for hES research. Yet no federal regulations nor other forms of federal guidance were developed for addressing the distinctive ethical issues in this domain of research. In response, the National Academies convened experts and stakeholders in order to develop guidelines for hES research, which led to the establishment of ESCROs as the mechanism for oversight.¹⁶

As discussed earlier, there are no established models for building effective data and AI ethics committees. However, there are many examples of the effective use of ethics committees in other domains from which to draw. Some of these are mandated by law, while others are extra-legal. Some require pre-approval and have the power to prevent or halt projects, even as others are only consultative and advisory. Some are guided by highly specified guidelines and protocols, with others operating in a less legalistic context. This report does not recommend a specific form of data and AI ethics committee. Rather, the goal of the report is to draw from knowledge gained in the development and functioning of ethics committees in other domains, in order to highlight questions that must be asked and decisions made in the process of developing effective ethics committees in the data and AI domain.

PART TWO:

BUILDING AN ETHICS COMMITTEE

There are not yet data and AI ethics committees with established records of being effective and well-functioning, so there are no success models to serve as case-studies or best practices for how to design and implement them. However, it is possible to learn from the development, features, and functioning of other types of ethics oversight committees, such as ESCROs, IRBs, hospital ethics committees, and IACUCs. While no one analog is perfect, together they provide a rich picture of how to build well-structured and thoughtfully designed data and AI ethics committees.

A roadmap for building an ethics committee



01 Organizational Function

Why is the committee being created?


A foundational question in the formulation of any ethics committee is what its organizational function is meant to be. That is to ask: Why is the committee being created? A committee that

is intended to create broad policy guidance or perform as a consulting body at a high organizational level would be a very different sort of committee — in constitution, organization, charge, powers, participants, and procedures — than one that is intended to be a decision-making body with respect to complex cases of data use in research, AI systems, or product design. Without a clear conception of the intended institutional function and responsibilities of the committee, it is difficult to design and implement an effective one.

02 The Ethical Content

What should guide the committee's activities and decision-making?

The ethical content is the norms, rules, standards, and goals that guide the committee's work, as well as the values that underlie them. The ethical content is constituted by Foundation Values, Core Principles, and Deliberative Resources.



Foundational Values

What are the basic values that the committee is meant to protect or promote?

Concerns about data and AI ethics are frequently motivated by considerations like privacy, security, transparency, and fairness. To consistently address these concerns in practice, it is necessary to be clear about why they matter. **What are the underlying values that are the basis for concerns about them?**

For example, privacy is not an absolute or basic value. What makes privacy important is the way in which control of information about oneself is related to individual autonomy, protection from harm, and prevention from discrimination. It is when data practices impugn on these that privacy concerns arise. Similarly, transparency and explainability are not absolute or basic considerations. They are required or desired in some contexts because they promote individual autonomy, respect for persons, and procedural justice and fairness.

Being explicit about the underlying values is important to ensure that guiding norms and principles (for example, informed consent, explainability, and anonymity) are defined and operationalized in appropriate ways, and so that decision-making in complex cases is attentive to what really matters ethically.

The values that are the basis for ethical decision-making may not be the same or have the same salience across all organizations. Legal considerations, professional responsibilities, or institutional missions may contribute to differences in values that organizations prioritize. For example, medical organizations might prioritize beneficence and justice; child and family service organizations might prioritize child welfare and parental rights; educational institutions might prioritize student learning and well-being; financial institutions might prioritize fairness, fiduciary obligations, and risk minimization. However, in each instance, an effective ethics committee will need clarity on the basic values the committee is meant to protect and promote.

Core Principles

What are the primary guiding principles in support of the values?

Basic values indicate what the committee is meant to care about or protect and promote. The **core principles** are the norms or rules (or “oughts” or “prescriptions”) that describe in general how to do this. Requiring informed

consent for data collection, use, and sharing is a norm to protect the basic values of user agency, autonomy, and rights. Requiring explainability for financial decisions is a norm to ensure the basic values of procedural and distributive justice as well as respect for consumers.¹⁷

Core principles describe in general what ought to be done ethically or what is legally required in order to protect and promote the basic values. Clearly articulating these core principles is crucial because it will guide the work — the assessments and decision-making processes — of the ethics committee.

For example, if informed consent and maintaining anonymity are core principles or requirements for sharing user data with other firms or organizations, then the role of the ethics committee is to ensure that any proposal to share data — any program, technical design, partnership, or initiative that would do this — meets the standards of informed consent and anonymity. As another example, if avoiding adverse impacts on socially and economically disadvantaged groups is a core principle for using AI or machine learning in decision-making, then the role of the committee is to ensure that design and implementation adequately considers, monitors, audits, and adjusts to ensure this.

The core principles embody and protect the basic values. They are the first component in articulating the norms and standards that guide an ethics committee's decision-making.

Substantive Content and Deliberative Resources

What is required in practice to satisfy the core principles?

The core principles provide general norms. It is the ethics committee's responsibility to assess whether the norms are satisfied in a particular case, for example with respect to a new product, data use, or organizational policy.

Suppose that explainability is a norm adopted by a firm: the reasons or basis for decisions that are made regarding a consumer (or client, patient, or citizen) must be adequately provided to them. What, **in practice**, does this require? Does it require a narrative explanation? Providing a set of factors on which the decision was based? A set of like cases decided in a similar way? Does it require that the consumer can discuss the decision with someone if they do not understand or if they think a mistake has been made? Does the explanation need to be intelligible to anyone, or to people with a certain amount of background knowledge, or a certain language proficiency? Do the standards of explanation differ by context, by what is at stake, or by the sorts of decisions being made? These are the types of questions that ethics oversight committees will need to answer **in practice** regarding how to operationalize the core principle of explainability. For any core principle — fairness, informed consent, explainability, anonymity — substantive content and deliberative resources are needed to operationalize it.

As the example on the next page illustrates, moving from basic principles to the substantive content of what is required to satisfy them will involve significant work by ethics committees — work that both contributes to and is informed by the field of data and AI ethics. Effective ethics committees require a robust field of data ethics, which in turn requires that there be a robust data ethics community of practitioners, researchers, officers, and oversight participants moving the field from the articulation of broad or general principles to **substantive or operational content**. There is tremendous need for developing illustrative case studies, a body of precedents, conceptual resources, a shared knowledge base, critical perspectives, analytical tools, professional expectations, and detailed standards and rules for particular types of contexts and decisions.

No ethics committee can itself develop all these. But each can contribute (along with the research community, practitioners, policy makers, and others) to building the field by drawing on existing expertise, resources, and analogs to best operationalize their core principles, developing new standards and solutions when needed for the types of questions and problems they face, and sharing knowledge gained through experience. Part of what makes a committee-based approach to data and AI ethics so apt is that the field itself is still in development, and the questions and problems that arise as it grows often require a collaborative effort among people with diverse skills and knowledge to identify and address them.



FROM BASIC VALUES TO SUBSTANTIVE GUIDANCE: INFORMED CONSENT IN MEDICINE AND RESEARCH

The use of informed consent in medicine and research to authorize a procedure or ‘enroll a study’ is illustrative of what is involved in moving first from basic values to a core ethical principle, and onward to operationalized or guiding content.

Informed consent is a widely accepted principle for protecting the basic values of patient autonomy and respect for persons. It is typically taken to require that patients (or their proxies) are provided clear, accurate, and relevant information about the situation, that they adequately understand or comprehend the information and situation more generally, and that their decision is made voluntarily rather than being pressured or coerced. But the conditions of information, comprehension, and voluntariness must themselves be explicated.

What information should be provided in order to satisfy the informational requirement? What level of understanding constitutes comprehension? How do practitioners provide professional opinions and information in ways that sufficiently inform and guide patients without nudging them too hard toward a decision? Moreover, how do these standards differ by situation — for example, between emergency medicine and clinical practice? These are the sorts of questions that bioethicists, IRBs, and hospital ethics committees address in the context of thinking about consent agreements, professional standards of practice, and research design that involves human subjects. Working through them often involves going back to the basic values to determine how the principles should be applied. In the case of emergency medicine, respect for a person’s life often justifies life-saving treatment even in the absence of informed consent, unless the patient has provided a legitimate advanced directive to withhold treatment, in which case respect for their autonomy may take precedence. Consideration of foundational values helps to shape how principles are operationalized in deliberative practice, how conflicts between principles are resolved, and what should be done in cases where it is unclear what principles require.

03

Committee Composition

Who should be on a data and AI ethics committee?

To ensure that oversight committees are able to effectively identify ethical concerns and apply core principles in particular cases and contexts, they must have the necessary range of expertise, diversity of perspectives, and be resistant to bias and conflicts of interest. In some cases, it may also be important that the committee be composed in ways that foster input from the public.

What are the types of expertise needed?

Achieving sufficient diversity of expertise will typically require having the following types of committee members:



Technical Experts that have a good understanding of and ability to explain the technical details of the technologies, applications, practices, systems, and research projects that fall within the purview of the committee. A key role of technical experts is to ensure that other committee members have an adequate understanding of the technical aspects of cases or issues that come to the committee.



Ethical Experts that have a good understanding of and ability to explain key value concepts and core principles. Ethical experts will help the committee to work carefully through ethical analysis and evaluation, as well as provide information on analogical cases and potentially useful resources from other areas of ethics (for example, research ethics, public health ethics, and business ethics).



Legal Experts that are well-versed in the legal dimensions relevant to the types of applications, technologies, policies, and projects that fall within the committee's purview. Legal experts are crucial to ensuring legal compliance, as well as for identifying areas where current legal guidance is absent, inadequate or ambiguous.



Subject Matter Experts that have detailed knowledge of the issues, context, and practices within the intended domain of application or intervention. In medical delivery contexts, clinical practitioners and managers may be the appropriate subject matter experts. In social service contexts, case workers and administrators may be. Subject matter experts are crucial to identifying


the ways in which new technologies, systems, policies, and decision-making processes will impact current practices, empower and disempower people, and change clients'/users' experiences.



Citizen Participants that can represent public concerns and perspectives. In some cases, organizations may wish to include committee members that can raise potential or actual community concerns and views, and that can take a civic or community-oriented perspective. Citizen participants can also contribute to avoiding problematic biases and conflicts of interest.

Many organizations will not have the full range of expertise “in house” or will in some cases need to consult domain-specific experts. For these reasons, the committee should be constituted in ways that make such consultation possible. Part of the committee’s responsibilities include identifying when additional expertise beyond the core or standing committee is needed.

Hypothetical committee on data use



This figure illustrates the possible composition of a data ethics committee. It is meant to exemplify the types of expertise that might be included on the committee, as well as the types of specialists that might provide such expertise. For any given ethics committee, members should have appropriate domain knowledge as well. For example, the consumer advocate, social scientist, subject matter expert, information ethicist, and internal counsel appropriate for a committee whose domain is health care would be different from one whose domain is financial services.

How can the committee avoid bias and conflicts of interest?

It's important to recognize that biases and conflicts of interest may be unintentional but nonetheless present. This should always be considered, but in some organizations, for example those that have a social or public mission, it will be especially crucial to compose committees in ways that avoid conflicts of interest or the appearance of them.

Oversight committees in other areas — for example, IACUCs and ESCROs — often have at least one member of the committee that is not a member of the organization. These non-institutional members help to ensure that there is a voice that is not subject to organizational pressures or constrained by organizational goals.

In addition to including a non-institutional member, conflicts of interest and bias can be avoided by ensuring committee members have sufficient organizational protections to feel comfortable raising objections and concerns, even when there might be a perceived cost to the organization. This might be done by establishing policies that prohibit retaliation against committee members for their decisions or by providing some form of anonymity to committee members. For example, it might be that all decisions by the committee are reported by a specific member of the committee, while the votes, decisions, and comments of individual members are anonymized.

Organizations should also think carefully about whether and how to disclose the judgments of the committee. Transparency about the use of ethics consultation and review, its authority within an organization, and the results can further organizational goals and engender trust. However, it is also important that committee deliberations and findings are sufficiently confidential that committee members are protected.

How are committee members selected and how long do they serve?

Organizations should set clear guidelines for how decisions about staffing the committee are made, what roles people will fill on the committee, and how long members will serve. This includes rules for removing, changing, or adding members as necessary, as well as whether there are term limits for serving on the committee. Clear and transparent staffing policies that promote diverse perspectives are necessary to ensure that conflicts of interest are avoided and that the committee has sufficient independence in its work.

How large should a committee be?

The size of the committee will depend on the size of the organization and the full scope of its activities. A very large organization engaged in many different forms of data and AI use might have multiple committees with purview over different areas of

research, product design, service delivery, and policy. Smaller organizations or those with a narrower focus may have a single committee with fewer members. In such cases, committee members might consult more frequently with external experts. It is common in other forms of ethics committees that there is some overlap between roles. For example, the citizen representative might also be an ethics expert if there is no in-house person for that role.

Large organizations involved in a broad portfolio of data- and AI-related activities might consider whether to have a multi-leveled structure. When there are several ethics committees within an organization, having a further committee, perhaps made up of members of the underlying committees, can help with quality control, ensure that principles are being applied consistently, and develop a shared set of procedures and practices. An alternative to a multi-level structure for addressing the issues of quality control and consistency in very large organizations is using a review process (see figure on page 19) and information sharing.

04 Organizational Position and Powers

Where should an ethics committee be situated within an organization and what should be its powers?

In order for a data and AI ethics committee to effectively accomplish its function it must be appropriately positioned within an organization, its responsibilities and purview must be clearly defined, and it must be sufficiently resourced and empowered.

Where is the committee housed within the organization?

Particularly for larger organizations or organizations engaged in a broad range of potentially ethically sensitive activities, part of establishing an ethics committee involves deciding whether it is located within a department or unit or sits independently of them. It must also be decided at what organizational level it is situated, to whom it reports, and who is responsible for overseeing its activities. For smaller organizations, where it is not feasible to staff a committee internally, it may be necessary to engage with external expertise in a consultancy arrangement.

What is the purview of the committee?

Not every activity of an organization involving data need be subject to ethics committee oversight. It is crucial to clearly define and communicate the data-related activities, policy decisions, and applications that the committee is charged with evaluating. For example,

data collection and analysis to help improve supply chain efficiency or to reduce loss and waste internal to the organization might not warrant ethics committee oversight, whereas projects that aim to automate decision-making about users, clients, or consumers using AI systems might warrant review. The purview of a committee could be defined in a number of ways, such as by the sector involved (e.g. health care), the type of technology (e.g. facial recognition and biometrics), or the type of activity (e.g. sharing data outside of the organization or using it for purposes for which it was not originally collected).

In other domains with committee-based oversight, whether an activity is subject to committee review is often a matter of law. In the United States, for example, research on invertebrate animals is not subject to mandatory committee-based oversight, whereas the use of many vertebrates, such as primates, requires strict review. However, since there are at present few legal requirements to guide decisions in the data and AI ethics space, organizations will have to make decisions about which types of projects, technologies, applications, and activities warrant committee-based review.

When should the committee be consulted?

Once the purview of the committee is defined, it is also necessary to determine when the committee should be consulted. For example, review could be done at the start of a new project, prior to implementation, or only when it appears that there are potential issues. Alternatively, committee review might be done collaboratively with research or product teams in an effort to make ethical consideration a part of the design process. An advantage of this approach is helping to institutionalize that ethically aware development is a crucial component of good product, service, and technology design.

Who is required/permitted to consult with the committee?

Organizations must decide which sorts of projects, applications, or data usage require review by an ethics committee, as well as which can access review or consultation. It must also be determined which parties are responsible for initiating and facilitating a review. For example, it may be that specific domains or activities are sufficiently sensitive that all new projects, practices, or systems require at least a consultation or pre-review. This might be the case, for example, with applications that employ facial recognition technologies, collect information about minors, make use of genetic or medical data, or involve AI decision-making for service, benefit, or financial product eligibility. Other data-rich domains or activities, such as supply chain management or gathering feedback from users to improve service or product performance, might have a lower presumption of ethical concern, and therefore might not require ethics review or consultation. However, those working on such projects might have access to request a review if they become concerned about privacy, data-oriented risks and liability, or potential for misuse.

What authority does the committee have?

There is a wide range of authority that might be afforded to an ethics committee. Extant research oversight committees have the power to prevent research from being conducted, the power to monitor research, and the power to halt research for a variety of violations. Analogous authority in the data and AI space might include pre-authorization for the development or deployment of a large-scale data collection and sharing system, or the authority to pre-assess and monitor a significant new AI product launch. This level of authority may be appropriate for only sensitive projects that could have wide-ranging organizational or societal impacts. In other cases, such authority and stringent oversight may be overly burdensome and inefficient. Weaker authority could involve consultation reports that identify ethical challenges or concerns for developers and project managers to consider, or requiring that steps be taken to monitor for possible problems as development or implementation progress.

As these considerations make clear, there is a wide range in the possible powers and positions that a data and AI ethics committee could have within an organization. On one end of the spectrum, engaging with the committee can be strictly voluntary and consultative, and the committee's findings only recommendations. On the other end, engaging with the committee can be required and the committee can have the power to prevent research or halt a product launch until the ethical issues identified are satisfactorily addressed. Committees can also be housed entirely within an organization or they can be external to it and accessed in a consultory fashion.



05 Procedures and Governance

What are the processes and procedures that govern the work of the ethics committee?

A key component of building an effective ethics committee is establishing clear procedures for making use of the committee, as well as efficient processes by which the committee does its work. These issues of governance will depend in significant part on the authority of the committee and its position within an organization. However, choices also need to be made about such questions as how cases are referred to the committee, what information the committee requires or has access to, what roles members of the committee have, the format and structure of meetings, how decisions are reached, and how, when, and to whom decisions are communicated.

What is presented to the committee?

Oversight committees must be presented the cases they are to evaluate or the issues on which they are to advise in an appropriate format. In the case of IACUCs, IRBs, and ESCROs, researchers submit what is known as a protocol, a detailed explanation of a research program that includes information such as the number of research subjects involved, consent procedures, and justification for the experimental procedures in layman's terms.¹⁸ This method of presentation is useful when research or development requires pre-authorization by the committee. An alternative to protocol review is that teams present their projects to the committee. The committee can then ask questions before making a determination, providing feedback, or helping to identify ways the project could better satisfy organizational aims and foundational values. This may be more appropriate if the committee functions primarily as a consulting body.

Whatever option an organization chooses, there should be clear guidelines about what those submitting a case to the committee should provide. For example, in a context where the primary ethical issues involve data collection and use, guidelines might require that teams explicitly address data collection practices, data storage procedures and safeguards, explanations of how sensitive data will be managed, how data will be anonymized if applicable, which particular uses the data will be put to, who precisely will have access to the data, and what is being done to ensure that data is not used for other projects without review.

Clear guidelines not only help the committee to evaluate cases, they promote reflection by those working on the projects or products being evaluated. It requires them to think more carefully through the implications of their data or research practices.

Is there an expedited review process?

One way to increase efficiency in the review process is to have mechanisms for expedited review for cases unlikely to raise significant ethical concerns. This is common in other committee-based oversight contexts. For example, researchers submitting protocols involving animal subjects research can indicate that the research will not involve significant amounts of pain and suffering. Such cases are eligible for expedited review by a subset or single member of the committee who can either confirm that the project does not require a full review or refer it for full review. Similar procedures can be designed for data and AI contexts by identifying which kinds of cases are unlikely to be of concern, as well as by building up a body of precedence to guide those determinations.

What are the procedures and rules for external consultation?

There may be instances where the committee feels that their evaluation would benefit from consultation with experts not on the committee. It is important to establish guidelines regarding who can and who cannot serve as an external consultant, as well as what form external consultation should take. This can be important for protecting IP, avoiding conflicts of interest and bias, and ensuring that the committee has access to the expertise needed for an informed and ethically sensitive evaluation.

What are the standards by which committees make judgments?

Ethics committees perform best when it is clear what standards of judgment they should use when evaluating cases. In other areas, such as in human and animal subjects research, the standards are largely provided by law. In the case of animal subjects research, the overall goal is to minimize the suffering of animal subjects consistent with realizing valuable scientific ends. The guidelines for IACUC members are thus to evaluate whether an experiment is likely to generate useful knowledge, given the details of the experiment, and whether steps have been taken to achieve those results at minimal cost to the animals. As discussed earlier, in the case of data and AI ethics, there are not adequate legally prescribed standards. Therefore, the standards of judgment used by committees should be informed by foundational values and core principles as well.

Possible standards include:

- 
- Whether the project under review**
advances organizational aims and foundational values to an extent that it justifies any organizational and social risks or costs
 - Whether the project is likely to violate**
any hard constraints, such as legal requirements or fundamental organizational commitments/principles
 - Whether an impartial citizen would judge**
that the organization has done due diligence in considering the ethical implications of the project
 - Whether it is possible to secure**
the sought benefits in a way that better aligns with organizational values and commitments and without any significant additional undue burden or costs
 - Whether reputational risks could be significant**
enough to damage the brand value in the concerned market or in other places where the organization operates

These possible standards are meant to be illustrative. What standards a data and AI ethics committee adopts will vary depending upon its organization's mission and value commitments, as well as the types of projects and activities that it reviews. For example, in deciding whether to adopt an AI decision-making system, a governmental organization that provides social services will reasonably adopt standards for satisfying access to service concerns that are different from a privately-owned company that provides financial advising services. Both organizations must consider such things as fairness and privacy. However, the social service organization's standards must be informed by its public mission and accountability to citizens as well, whereas a private company has to consider its fiduciary responsibilities.

How do committee members collectively issue a decision?

There will be cases where not all committee members agree in their judgments. It is essential to have clear guidelines about how committee judgments are aggregated and reported. This is especially important where committee approval is required for research to commence or products to launch, but it is also needed when committees serve in a more consultative fashion.

If committee decisions are made by vote, it is necessary to decide in advance on such things as what constitutes a quorum, the voting procedure (for example, is it open or anonymous), what percentage of the vote is required for a positive outcome, whether any members exercise special authority (for example, some members might be able to issue an overriding negative vote), and whether and how votes will be reported outside the committee. If the committee issues a written statement of its findings, it is necessary to decide whether there will be a single report or whether the organization will allow for dissenting reports in the case of disagreement.

What is the timeline for committee review?

Guidance should be provided on how much notice committees should be given before they are expected to conduct a review, how frequently committees will meet, and how long committees have to conduct a review before making decisions or issuing reports.

How are committees audited and evaluated?

In order to ensure the efficacy of the committee, there need to be procedures to audit or review the committee's work. Guidelines should be put in place for how often committees will be evaluated or what circumstances can trigger an audit, as well as the methods by which they are to be audited and evaluated. Options for auditing include having individuals external to the committee observe committee meetings, having outside experts evaluate the committee's procedures and decisions, or submitting test cases for the committee to review.

Auditing and evaluation should be formative so far as is possible. No ethics oversight committee will be the same on day zero as it is two, five, or ten years into operation. Committee functioning and their value to organizations should improve over time. Data should be collected on a regular and routine basis to assist with improving the committee in all the respects discussed above — for example, refining foundational values, improving how it is constituted and staffed, streamlining evaluation processes, developing substantive guidance, and providing organizational value. Committees will have different sets of success criteria depending upon the types of organizations they are in and the organizational functions that they perform. For some organizations, the primary function may be to mitigate risk and liability associated with ethical lapses. For other organizations, the primary function may be to promote positive social outcomes. Therefore, each organization will need to develop appropriate metrics by which they measure committee success and collect the data needed to assess and improve committee functioning.



CONCLUSION:

GETTING STARTED

Data and AI ethics committees can play a crucial role in developing ethics ecosystems within organizations and society more broadly. However, building an efficient and effective committee is a substantial undertaking. It takes organizational commitment and resources. The aim of this report has been to help identify the types of questions that need to be answered in the building of an effective ethics committee.

Ethics committees can take a wide variety of forms and roles. Crucial to beginning the building process is putting together the right team of people and engaging with organizational stakeholders to begin to think through the key questions about function, values, principles, location, composition, and process. Once the basic outlines of the committee are established and the initial committee is formed, quite a lot of the operational details will be developed in the context of the committee's work. As has been emphasized, there is no existing data and AI ethics committee template and the field of data and AI ethics is still maturing. Creating meaningful and effective ethics committee oversight models not only offers benefits and protection to the organization, it is critical to the broader data and AI ethics development process.

References

- ¹ Accenture, "Technology Vision 2019," February 7, 2019.
- ² Joshua Brustein and Mark Bergen, "The Google AI Ethics Board With Actual Power Is Still Around," Bloomberg, April 6, 2019. Axon, "AI Ethics Board." Tom Simonite, "Tech Firms Move to Put Ethical Guard Rails around AI," Wired, March 16, 2018. Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018. John Basl and Eric Schwitzgebel, "AIs should have the same ethical protections as animals," Aeon, April 26, 2019.
- ³ Steven C. Tiell, "The case for data ethics," Accenture, 2015. Accenture, "Building digital trust: The role of data ethics in the digital age," 2016.
- ⁴ Accenture, "Facilitating ethical decisions throughout the data supply chain," 2016.
- ⁵ Many organizations have adopted or proposed general data and AI principles and value statements. For example, Microsoft, "Our approach to AI." Partnership for Artificial Intelligence to Benefit People and Society, "About." Google, "AI at Google: our principles," June 7, 2018. Future of Life Institute, "Asilomar AI Principles," 2017. Accenture, "Universal principles of data ethics: 12 guidelines for developing ethics codes," 2016.
- ⁶ Partnership for Artificial Intelligence to Benefit People and Society, "Tenets."
- ⁷ Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018. Kara Swisher, "Who Will Teach Silicon Valley to be Ethical?," The New York Times, October 21, 2018.
- ⁸ DeepMind, "DeepMind Ethics & Society Fellows." Jillian D'Onfro, "Google Launches Advisory Council To Help Company Question Assumptions On Ethical AI," Forbes, March 26, 2019.
- ⁹ Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018.
- ¹⁰ Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018. Accenture, "Informed consent and data in motion," 2016.
- ¹¹ Accenture, "Facilitating ethical decisions throughout the data supply chain," 2016. Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018. IBM, "AI Fairness 360 Open Source Toolkit."
- ¹² Microsoft, "FATE: Fairness, Accountability, Transparency, and Ethics in AI." Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018.
- ¹³ Joshua Brustein and Mark Bergen, "The Google AI Ethics Board With Actual Power Is Still Around," Bloomberg, April 6, 2019. Axon, "AI Ethics Board." Tom Simonite, "Tech Firms Move to Put Ethical Guard Rails around AI," Wired, March 16, 2018. Darrell M. West, "The role of corporations in addressing AI's ethical dilemmas," Brookings Institute, September 13, 2018. John Basl and Eric Schwitzgebel, "AIs should have the same ethical protections as animals," Aeon, April 26, 2019.
- ¹⁴ Beecher, Henry K., "Ethics and Clinical Research", New England Journal of Medicine. 1966.
- ¹⁵ National Institutes of Health, "Protecting Human Research Participants," September 26, 2018. Department of Health and Human Services, "45 CFR Subtitle A," October 1, 2016.
- ¹⁶ Institute of Medicine and National Research Council. Guidelines for Human Embryonic Stem Cell Research. Washington, DC: The National Academies Press. 2005.
- ¹⁷ Accenture, "Universal principles of data ethics: 12 guidelines for developing ethics codes," 2016.
- ¹⁸ For example protocols for human subjects research and animals subjects research see: National Institutes of Health Office of Laboratory Animal Welfare, "Animal Study Proposal." Children's Hospital of Philadelphia, "Writing a Protocol."

Author Information

RONALD SANDLER

Professor of Philosophy
Department of Philosophy and Religion
Director, Ethics Institute
Northeastern University
r.sandler@northeastern.edu

JOHN BASL

Associate Professor of Philosophy
Department of Philosophy and Religion
Researcher, Ethics Institute
Northeastern University
j.basl@northeastern.edu

Accenture Contact

STEVEN C. TIELL

Sr. Principal, Responsible Innovation
steven.c.tiell@accenture.com

About Accenture

Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. Combining unmatched experience and specialized skills across more than 40 industries and all business functions—underpinned by the world’s largest delivery network—Accenture works at the intersection of business and technology to help clients improve their performance and create sustainable value for their stakeholders. With 477,000 people serving clients in more than 120 countries, Accenture drives innovation to improve the way the world works and lives.

Visit us at www.accenture.com.

Ethics Institute at Northeastern University


Northeastern University’s Ethics Institute is a multidisciplinary research and teaching institute dedicated to bringing ethical analysis and evaluation to bear on complex social and environmental issues. The Institute supports ethics research and education throughout Northeastern University’s global network of campuses, and is home to a diverse range of research and education programs on information ethics.

www.northeastern.edu/ethics
www.northeastern.edu/infoethics

This document makes reference to trademarks owned by third parties. All such third-party trademarks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such trademarks is intended, expressed or implied.

Copyright © 2019 Accenture
All rights reserved.

Accenture and its logo
are trademarks of Accenture.



Published under Creative Commons Attribution—Non-Commercial-ShareAlike 4.0 International Public License