

Technology and Engineering Practice: Ethical Lenses to Look Through

MARKKULA CENTER FOR APPLIED ETHICS

at Santa Clara University



©2020. This document is part of a project, *Ethics in Technology Practice*, made possible by a grant from [Omidyar Network's Tech and Society Solutions Lab](#) and developed by the Markkula Center of Applied Ethics. It is made available under a [Creative Commons license \(CC BY-NC-ND 3.0\)](#) for noncommercial use with attribution and no derivatives. References to this material must include the following citation: **Vallor, Shannon, Irina Raicu, and Brian Green (2020). *Ethics in Technology Practice*. The Markkula Center for Applied Ethics at Santa Clara University. <https://www.scu.edu/ethics/>.**

Technology and Engineering Practice: Ethical Lenses to Look Through

Shannon Vallor, Irina Raicu, Brian Green



Conceptual frameworks drawn from theories that have shaped the study of ethics over centuries can help us recognize and describe ethical issues when we encounter them. In this way, a basic grasp of ethical theory can work like a field guide for practical ethical concerns.

Many ethical theories can frame our thinking; here we focus on several that are widely used by both academic and professional ethicists, including tech ethicists in particular. However, these are ethical theories developed largely in the context of Western philosophical thought; our selection must not be taken as a license to ignore the rich range of ethical frameworks available in other cultural contexts. A brief discussion of such frameworks is included on p.11 of this guide.

Each section below includes a brief overview of one ethical theory, examples that reflect its relevance to technologists, and a list of helpful questions that technologists/design teams can ask to provoke ethical reflection through that lens.

The Rights Perspective

Overview

This ethical lens focuses on *moral rules, rights, principles, and duties*. It tends to be more universalist than other kinds of ethical theories—that is, the rules/principles are usually intended to apply to the majority or all possible cases. This creates a special kind of challenge when we encounter rights, principles, or duties that conflict—where it is impossible to follow one moral rule without breaking another, or to fulfill our moral duties to one stakeholder without failing to fulfill a moral duty to another. In such cases we must determine which rights, principles, and duties carry the most ethical force in that situation, and/or find a solution that minimizes the ethical violation when ‘no-win’ scenarios (often called ‘wicked’ moral situations) occur. Thus the rights perspective still requires careful ethical reflection and judgment to be used well; it is not a simple moral checklist that can be mindlessly implemented.

Rules-based systems of ethics range from the most general, such as the ‘golden rule,’ to theories such as that of W.D. Ross (1877-1971), which offers a list of seven *pro tanto* moral duties (duties that can only be overridden by other, morally weightier duties): fidelity; reparation; gratitude; justice; beneficence; non-injury; and self-improvement.

Even more widely used by ethicists is the *categorical* imperative, a single deontological rule presented by Immanuel Kant (1724-1804) in three formulations. The two most commonly used in practical ethics are the *formula of the universal law of nature* and the *formula of humanity*. The first tells us that we should only act upon principles/maxims that we would be willing to impose upon every rational agent, as if the rule behind our practice were to become a universal law of nature that everyone, everywhere, had to follow. The second, the *formula of humanity*, states that one should always treat other persons as ends in themselves (dignified beings to be respected), never merely as *means* to an end (i.e., never as mere tools/objects to be manipulated for my purposes). So, benefiting from one’s action toward another person is only permissible if that person’s own autonomy and dignity is not violated in the process, and if the person being treated as a means would consent to such treatment as part of their own autonomously chosen ends.

The ethical issues and concerns frequently highlighted by looking through this ethical lens include, but are not limited to:

- *Autonomy* (the extent to which people can freely choose for themselves)
- *Dignity* (the extent to which people are valued in themselves, not as objects with a price)
- *Transparency* (honest, open, and informed conditions of social treatment/distribution)

Examples of Rights-Related Ethical Issues in Tech Practice:

In what way does a virtual banking assistant that is deliberately designed to deceive users (for example by actively representing itself as a human) violate a moral rule or principle such as the Kantian imperative to never treat a person as a mere means to an end? Would people be justified in feeling wronged by the bank upon discovering the deceit, even if they had not been financially harmed by the software? Does a participant in a commercial financial transaction have a moral right not to be lied to, even if a legal loophole means there is no legal right violated here?

Looking through the rights lens means anticipating contexts in which violations of autonomy, dignity, or trust might show up, regardless of whether there was malign intent or whether material harms were done to people’s interests. Many violations of such duties in the tech sector can be avoided by seeking meaningful and *ongoing* consent from those who are likely to be impacted (not necessarily just the end user) and offering thorough transparency about the design, terms, and intentions of the technology.

However, it is important to remember that concerns about rights often need to be balanced with other kinds of concerns. For example, autonomy is not an *unconditional* good (technologists should not empower users do anything they want). When user autonomy poses unacceptable moral risks, this value needs to be balanced with appropriately limited moral paternalism (which

is *also* unethical in excess.) An excellent example of this is the increasingly standard design requirement for strong passwords.

Rights-Related Questions for Technologists that Illuminate the Ethical Landscape:

- What *rights* of others & *duties* to others must we respect in a particular context?
- How might the *dignity & autonomy* of each stakeholder be impacted by this project?
- Does our project treat people in ways that are *transparent* and to which they would *consent*?
- Are our choices/conduct of the sort that I/we could find universally acceptable?
- Does this project involve any conflicting moral duties to others, or conflicting stakeholder rights? If so, how can we prioritize these?
- Which moral rights/duties involved in this project may be justifiably overridden by higher ethical duties, or more fundamental rights?

The Justice/Fairness Perspective

Overview

This widely used ethical lens focuses on giving individuals—or groups—their due. This includes the appropriate distribution of benefits and burdens, taking into consideration ethically relevant distinctions among people—what is known as distributive justice.

Aristotle (384-322 B.C.) argued that equals should be treated equally; however, that leaves open the question of which criteria should be used in determining whether people are “equal.” Aristotle, for example, unjustly *excluded* women and non-Greeks from those entitled to equal treatment. Philosophers have argued that justice demands that we consider, as part of our analysis, elements such as need, contribution, and the broad impacts of social structure on individuals and groups.

The justice lens also encompasses retributive justice. “An eye for an eye” was a call for such justice; current laws that aim to punish criminal wrongdoing also reflect this notion of justice. A third type, compensatory justice, is reflected in efforts to compensate injured people (whether or not the injury that they suffered was committed intentionally), or to return lost or stolen property to its rightful owner.

Justice and fairness also demand impartiality and the avoidance of conflicts of interest. Philosopher John Rawls (1921-2002), for example, argued that a fair and just society should be organized under a “veil of ignorance” about characteristics over which people have no control and which should not determine participants’ role and opportunities in society—characteristics such as age, gender, race, etc. Unable to tip the scales on behalf of their own characteristics, he argued, people would then create a more egalitarian and fair society.

Like the rights perspective, the justice lens is also related to notions of human dignity. It focuses more, however, on the relationships among people—on the conditions and processes required in order to implement societal respect for human dignity.

The ethical issues and concerns frequently highlighted by looking through this ethical lens include, but are not limited to:

- *Equality, Equity and Fairness* (a morally justifiable distribution of benefits, goods, harms, and risks)
- *Diversity and Inclusion* (ensuring that all impacted stakeholders, and particularly those in vulnerable or marginalized groups, are active participants in the process of determining the right distribution)
- *Due Process* (establishing procedures and conditions most likely to lead to just results)
- *Universality/Consistency* (holding all persons and actions to the same moral standards)
- *Power and Opportunity* (recognizing that not all parties are similarly situated, acknowledging historical and systemic injustices as a background condition)

Examples of Justice/Fairness-Related Ethical Issues in Tech Practice:

How does a digital advertising app that allows people to place custom housing or job ads that target only people under 40, or only people in specific zip codes, impact *fairness* and *justice*?

How does design impact the accessibility of various products for people with disabilities, or for people who can only afford older, cheaper technology tools?

How should facial recognition tools (or other algorithmic tools that are typically trained on unrepresentative datasets and are therefore far less accurate for some people and groups than for others) be used, if at all, and in which contexts, if any?

Looking through this lens enables us to see that technologies often distribute benefits and harms unevenly, and frequently exacerbate or perpetuate preexisting unfair societal conditions. It also stresses the need for technologists to consult stakeholders who may be very differently situated from themselves, in order to truly understand (rather than assume) the potential benefits of a product, as well as to be made aware of harms that they might have otherwise missed.

Justice/Fairness-Related Questions for Technologists that Illuminate the Ethical Landscape:

- What are both the benefits and the burdens created by this design/project, and how are they distributed among various stakeholders?
- What are the ethically relevant differences among potential users? How should we adjust for those?
- Have relevant stakeholders been consulted, so that their views inform the project?
- Are those stakeholders most likely to be impacted by the project included as active participants and leaders in the design and development process?

- Have multiple options been considered, to serve individuals and groups with different needs?
- Do the risks of harm from this project fall disproportionately on the least well-off or least powerful in society? Will the benefits of this project go disproportionately to those who already enjoy more than their share of social advantages and privileges?

The Utilitarian Perspective

Overview

Utilitarian ethics, in its most complete formulation by John Stuart Mill (1806-1873), asks us to weigh the overall happiness or welfare that our action is likely to bring about, for *all* those affected and over the long term. Happiness is measured by Mill in terms of aggregate pleasure and the absence of pain. Physical pleasure and pain are not the most significant metrics; although they do count, Mill argues that, at least for human beings, intellectual and psychological happiness are of an even higher moral quality and significance.

Utilitarianism is attractive to many engineers because in *theory* it implies the ability to quantify the ethical analysis and select for the optimal outcome (generating the greatest overall happiness with the least suffering). In practice, however, this is often an intractable or ‘wicked’ calculation, since the effects of a technology tend to spread out indefinitely in time (should we never have invented the gasoline engine, or plastic, given the now devastating consequences of these technologies for the planetary environment and its inhabitants?); and across populations (will the invention of social media platforms turn out to be a net positive or negative for humanity, once we take into account all future generations and all the users around the globe yet to experience its consequences?)

The requirements to consider equally the welfare of *all* affected stakeholders, including those quite distant from us, and to consider both long-term and unintended effects (where foreseeable), make utilitarian ethics a morally demanding standard. In this way, utilitarian ethics does *not* equate to or even closely resemble common forms of cost-benefit analysis in business, where only physical and/or economic benefits are considered, and often only in the short-term, or for a narrow range of stakeholders.

Moral consequences go far beyond economic good and harm. They include not only physical, but psychological, emotional, cognitive, moral, institutional, environmental, and political well-being or injury, or degradation.

The ethical issues and concerns frequently highlighted by looking through this ethical lens include, but are not limited to:

- *Happiness* (in a comprehensive sense, including such factors as physical, mental, and many other forms of well-being)
- *Balancing of stakeholder interests* (who is benefitting and who is being harmed, in what ways and to what degree, and how many)

- *Prediction of consequences* (some consequences can be predicted and others cannot; still, one should account for all reasonably foreseeable effects of this action)

Example of Utilitarian Ethical Issues in Tech Practice:

When technologists designed the first wave of apps to be maximally ‘sticky,’ to keep users coming back to their apps or devices, most did not anticipate doing their users moral harm. By now, most technologists have had to abandon this level of ethical naivete about issues like technology addiction. While rights and fairness issues are also in play here (since addiction compromises people’s cognitive autonomy), even a utilitarian reading tells us we went wrong, since many of the harmful consequences of ‘sticky’ design *were* in fact foreseeable.

Technologists *falsely* assumed that people’s technological choices were reliably correlated with their increasing happiness and welfare. This was not a *reasonable* assumption, because people make themselves unhappy with their choices *all the time*, and we are subject to any number of mental compulsions that drive us to choose actions that promise happiness but will not deliver, or will deliver only a very short-term, shallow pleasure while depriving us of a more lasting, substantive kind. Leading device manufacturers now increasingly admit this by building tools to fight tech addiction.

Related Questions for Technologists that Illuminate the Ethical Landscape:

- Who are all the people who are likely to be directly *and indirectly* affected by this project? In what ways?
- Will the effects in aggregate likely create more good than harm, and what *types* of good and harm? What are we counting as well-being, and what are we counting as harm/suffering?
- What are the most morally significant harms and benefits that this project involves? Is our view of these concepts too narrow, or are we thinking about *all* relevant types of harm/benefit (psychological, political, environmental, moral, cognitive, emotional, institutional, cultural, etc.)?
- How might future generations be affected by this project?
- Have we adequately considered ‘dual-use’ and downstream effects other than those we intend?
- Have we considered the full range of actions/resources/opportunities available to us that might boost this project’s potential benefits and minimize its risks?
- Are we settling too easily for an ethically ‘acceptable’ design or goal (‘do no harm’), or are there missed opportunities to set a higher ethical standard and generate even greater benefits?

The Common Good Perspective

Overview

Another approach to ethical thinking is to focus on the *common good*, which highlights shared social institutions, communities, and relationships (instead of utilitarianism’s concentration on

the aggregate welfare/happiness of individuals). The distinction is subtle but important. Utilitarians consider likely injuries or benefits to discrete individuals, then sum those up to measure aggregate social impact. The common good lens, in contrast, focuses on the impact of a practice on the health and welfare of *communities* or groups of people, up to and including all of humanity, as *functional wholes*.

Welfare as measured here goes beyond personal happiness to include things like political and public health, security, liberty, sustainability, education, or other values deemed critical to flourishing community life. Thus a technology that might *seem* to satisfy a utilitarian (by making most individuals personally happy, say through neurochemical intervention) might fail the common good test if the result was a loss of community life and health (for example, if those people spent their lives detached from others--like addicts drifting in a technologically-induced state of personal euphoria).

Common good ethicists will also look at the impact of a practice on morally significant institutions that are critical to the life of communities—for example on government, the justice system, or education, or on supporting ecosystems and ecologies. Common good frameworks help us avoid notorious *tragedies of the commons*, where rationally seeking to maximize good consequences for every individual leads to damage or destruction of the *system* that those individuals depend upon to thrive. Common good frameworks also share commonalities with cultural perspectives in which promoting social harmony and stable functioning may be seen as more ethically important than maximizing the autonomy and welfare of isolated individuals.

For practical purposes, it can be helpful to view the utilitarian and common good approaches as complementary lenses that provoke us to consider both individual and communal welfare, even when these are in tension. Using the utilitarian and common good lenses while doing tech ethics is like watching birds while using special glasses that help us zoom out from an individual bird to survey a dynamic network (the various members of a moving flock), and then try to project the overall direction of the flock's travel as best we can (is this project overall going to make people's lives better?), while still noticing any particular members that are in special peril (are some people going to suffer greatly for a trivial gain for the rest?).

The ethical issues and concerns frequently highlighted by looking through this ethical lens include, but are not limited to:

- *Communities* (of varying scales, ranging from families to neighborhoods, towns, provinces, nations, and the world)
- *Relationships* (not only among individuals, but also relationships in a more holistic sense of groups, including nonhuman animals and the natural world as well)
- *Institutions of governance* (and the ways in which these networked institutions interact with each other)
- *Economic institutions* (including corporations and corporate cultures, trade organizations, etc.)

- *Other social institutions* (such as religious groups, alumni associations, professional associations, environmental groups, etc.)

Example of Common Good Ethical Issues in Tech Practice:

In our discussion of technology addiction (under the analysis of utilitarianism above), we noted its impact on individual happiness and well-being. However, we also know that technology addiction can harm the *common good*, through damaging family and civic ties and institutions that are essential for healthy communities, for example, democratic institutions.

Technology use, data storage, and energy-intensive training of AI models also impact the environment in ways that implicate the common good.

In addition, with everything from weapons to pacemakers now being connected to the internet, cybersecurity has become one of the conditions required for the common good.

Related Questions for Technologists that Illuminate the Ethical Landscape:

- Does this project benefit many individuals, but only at the expense of the *common good*?
- Does it do the opposite, by sacrificing the welfare or key interests of individuals for the common good? Have we considered these tradeoffs, and determined which are ethically justifiable?
- What might this technology do for or to social institutions such as various levels of government, schools, hospitals, churches, infrastructure, and so on?
- What might this technology do for or to the larger environment beyond human society, such as ecosystems, biodiversity, sustainability, climate change, animal welfare, etc.?

The Virtue Ethics Perspective

Overview

Virtue ethics is more difficult to encapsulate than the other ethical frameworks. Virtue ethics essentially recognizes the necessary incompleteness of any set of moral rules or principles, and the need for people with well-habituated virtues of moral character and well-cultivated, *practically wise* moral judgment to fill the gap.

Aristotle (384-322 B.C.) stated that ethics cannot be approached like mathematics; there is no algorithm for ethics, and moral life is not a well-defined, closed problem for which one could design a single, optimal solution. It is an endless task of skillfully navigating a messy, open-ended, constantly shifting social landscape in which we must find ways to maintain and support human flourishing with others, and in which novel circumstances and contexts are always emerging that call upon us to adapt our existing ethical heuristics, or invent new, bespoke ones on the spot.

Virtue ethics does, however, offer some guidance to structure the ethical landscape. It asks us to identify those specific *virtues*—stable traits of character or dispositions—that morally excellent

persons in our context of action consistently display, and then identify and promote the *habits of action* that produce and strengthen those virtues (and/or suppress or weaken the opposite vices). So, for example, if honesty is a virtue in designers and engineers (and a tendency to falsify data or exaggerate results is a vice), then we should think about what habits of design practice tend to promote honesty, and encourage those. As Aristotle says, ‘we are what we repeatedly do.’ We are not born honest or dishonest, but we become one or the other only by forming virtuous or vicious habits with respect to the truth.

Virtue ethics is also highly *context-sensitive*: each moral response is unique and even our firmest moral habits must be adaptable to individual situations. For example, a soldier who has a highly-developed virtue of courage will in one context run headlong into a field of open fire while others hang back in fear; but there are other contexts in which a soldier who did that would not be courageous, but rash and foolish, endangering the whole unit. The virtuous soldier reliably sees the difference, and acts accordingly—that is, *wisely*, finding the appropriate ‘mean’ between foolish extremes (in this case, the vices of cowardice and rashness), where those are always relative to the context (an act that is rash in one context may be courageous in another).

While other ethical lenses focus our moral attention outward, onto our future technological choices and/or their consequences, virtue ethics reminds us to also reflect inward—on who we are, who we want to become as morally expert technologists, and how we can get there.

It also asks us to describe the model of moral excellence in our field that we are striving to emulate, or even surpass. What are the habits, skills, values, and character traits of an *exemplary engineer*, or an *exemplary designer*, or an *exemplary coder*?

Virtue ethics also incorporates a unique element of moral intelligence, called *practical wisdom*, that unites several faculties: moral perception (*awareness* of salient moral facts and events), moral emotion (*feeling* the appropriate moral responses to the situation), and moral imagination (*envisioning* and skillfully *inventing* appropriate, well-calibrated moral responses to new situations and contexts).

The ethical issues and concerns frequently highlighted by looking through this ethical lens include, but are not limited to:

- *Habits of character / virtues and vices* (the features of a person’s character, both good and bad)
- *Context* (the particular contexts in which decisions about technology developments are made, as well as the contexts in which the technology will be deployed)
- *Expression of already-existing virtues and vices* (technological products will express virtues and vices that already exist in creators and users of technology)
- *Cultivation of new virtues and vices* (technological products may cultivate or reinforce virtues and vices in their users)

Examples of Virtue-Related Ethical Issues in Tech Practice:

Many websites have given people access to news sources that are subsidized only by online

advertising, elevated to mass visibility by popularity and page views, and vulnerable to being gamed by armies of bots, trolls, and foreign adversaries. Has such access helped to make us wiser, more honest, more compassionate, and more responsible citizens? Or did it have very different effects on our intellectual and civic virtues?

There is probably no better conceptual lens than virtue ethics for illuminating the problematic effects of the attention economy and digital media. It helps to explain why we have seen so many pernicious moral effects of this situation even though the individual acts of social media companies appeared morally benign; no individual person was *wronged* by having access to news articles on various social media platforms, and even the individual consequences didn't seem so destructive initially. What has happened, however, is that our habits have been gradually altered by new media practices not designed to sustain the same civic function.

Not all technological changes must degrade our virtues, of course. Consider the ethical prospects of virtual-reality (VR) technology, which are still quite open. As VR environments become commonplace and easy to access, might people develop *stronger* virtues of empathy, civic care, and moral perspective, by experiencing others' circumstances in a more immersive, realistic way? Or will they instead become even more numb and detached, walking through others' lives like players in a video game? Most important is this question: what VR design choices would make the first, ethically desirable outcome more likely than the second, ethically undesirable one?

Related Questions for Technologists that Illuminate the Ethical Landscape:

- What design habits are we regularly embodying, and are they the habits of *excellent* designers?
- Would we want future generations of technologists to use our practice as the example to follow?
- What habits of character will this design/project foster in users and other affected stakeholders?
- Will this design/project weaken or disincentivize any important human habits, skills, or virtues that are central to human excellence (moral, political, or intellectual)? Will it *strengthen* any?
- Will this design/project incentivize any vicious habits or traits in users or other stakeholders?
- Are our choices and practices generally embodying the appropriate '*mean*' of conduct (relative to the context)? Or are they *extreme* (excessive or deficient) in some ways?
- Is there anything unusual about the context of this project that requires us to reconsider or modify the normal 'script' of good design practice? Are we qualified and in a position to safely and ethically make such modifications to normal design practice, and if not, who is?
- What will this design/project say *about us as people* in the eyes of those who receive it? Will we, as individuals and as a team/organization, be *proud* to have our names associated with this project one day?

Global Ethical Perspectives

There is no way to offer an ‘overview’ of the full range of ethical perspectives and frameworks that the human family has developed over the millennia since our species became capable of explicit ethical reflection. What matters is that technologists remain vigilant and humble enough to remember that whatever ethical frameworks may be most familiar or ‘natural’ to them and their colleagues, they only amount to a tiny fraction of the ways of seeing the ethical landscape that their potential users and impacted communities may adopt.

This does not mean that practical ethics is impossible; *on the contrary*, it is a fundamental design responsibility that we cannot make go away. But it is helpful to remember that the moral perspectives in the conference room/lab/board meeting are never exhaustive, and that they are likely to be biased in favor of the moral perspectives most familiar to whoever happens to occupy the dominant cultural majority in the room, company, or society. Yet the technologies we build don’t stay in our room, our company, our community, or our nation.

New technologies seep outward into the world and spread their effects to peoples and groups who, all too often, don’t get a fair say in the moral values that those technologies are designed to reinforce or undermine in their communities. And yet, we *cannot* design in a value-neutral way—that is impossible, and the illusion that we can do so is arguably more dangerous than knowingly imposing values on others without their consent, because it does the same thing—just without the accountability.

Technologists *will* design with ethical values in mind; the only question is whether they will do so in ways that are careful, reflective, explicit, humble, transparent, and responsive to stakeholder feedback, or in ways that are arrogant, opaque, and irresponsible.

While moral and intellectual humility requires us to admit that our ethical perspective is always incomplete and subject to cognitive and cultural blind spots, the processes of ethical feedback and iteration can be calibrated to invite a more diverse/pluralistic range of ethical feedback, as our designs spread to new regions, cultures, and communities.

Examples of Global Ethical Issues in Tech Practice:

Facial-recognition, social media, and other digital technologies are being used in a grand project of social engineering in China to produce a universal ‘social-credit’ system in which the social status and privileges of individuals will be powerfully enhanced or curtailed depending on the ‘score’ that the government system assigns them as a measure of their social virtue—their ability to promote a system of ‘social harmony.’

Many Western ethicists view this system as profoundly dystopic and morally dangerous. In China, however, some embrace it, within a cultural framework that values social harmony as the highest moral good. How should technologists respond to invitations to assist China in this project, or to assist other nations who might want to follow China’s lead? What ethical values should guide them? Should they simply accede to the local value-system (especially in an authoritarian society in which people might be reluctant to express their real values)? Or should

technologists be guided by their own personal values, the values of the nation in which they reside, or the ethical principles set out by their company?

This example illustrates the depth of the ethical challenge presented by global conflicts of ethical vision. But notice that there is no way to *evade* the challenge. A decision must be made, and it will not be ethically neutral no matter how it gets made. A decision to ‘follow the profits’ and ‘put ethics aside’ is not a morally neutral decision, it is one that assigns profit as the highest or sole value. That in itself is a morally laden choice for which one is responsible, especially if it leads to harm.

It may be helpful, where possible, to seek ethical dialogue across cultural boundaries and begin to seek common ground with technologists in other cultural spaces. Such dialogues will not always produce ethical consensus, but they can help give shape to the conversation we must begin to have about the future of global human flourishing in a technological age, one in which technology increasingly links our fortunes together.

Questions for Technologists that Illuminate the Global Ethical Landscape:

- Have we invited and considered the ethical perspectives of users and communities other than our own, including those quite culturally or physically remote from us? Or have we fallen into the trap of “designing for ourselves”?
- How might the impacts and perceptions of this design/project differ for users and communities with very different value-systems and social norms than those local or familiar to us? If we don’t know, how can we learn the answer?
- The vision of the ‘good life’ dominant in tech-centric cultures of the West is far from universal. Have we considered the global reach of technology and the fact that ethical traditions beyond the West often emphasize values such as social harmony and care, hierarchical respect, honor, personal sacrifice, or social ritual far more than we might?
- In what cases should we *refuse*, for compelling ethical reasons, to honor the social norms of another tradition, and in what cases should we incorporate and uphold others’ norms? How will we decide, and by what standard or process?

References and Further Reading

Aristotle [350 B.C.E.] (2014). *Nicomachean Ethics: Revised Edition*. Trans. Roger Crisp. Cambridge: Cambridge University Press.

Benjamin, Ruha (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge, UK: Polity Press.

Broussard, Meredith (2018). *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press.

- D'Ignazio, Catherine, and Lauren F. Klein (2020). *Data Feminism*. Cambridge, MA: MIT Press.
- Ess, Charles (2014). *Digital Media Ethics: Second Edition*. Cambridge: Polity Press.
- Jasanoff, Sheila (2016). *The Ethics of Invention: Technology and the Human Future*. New York: W.W. Norton.
- Kagan, Shelly (1989). *The Limits of Morality*. Oxford: Clarendon Press,
- Kant, Immanuel. [1785] (1997). *Groundwork of the Metaphysics of Morals*. Trans. Mary Gregor. Cambridge: Cambridge University Press.
- Lin, Patrick, Abney, Keith and Bekey, George, eds. (2012). *Robot Ethics*. Cambridge, MA: MIT Press.
- Lin, Patrick, Abney, Keith and Jenkins, Ryan, eds. (2017). *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*. New York: Oxford University Press.
- Macmillan. Shariat, Jonathan and Saucier, Cynthia Savard (2017). *Tragic Design: The True Impact of Bad Design and How to Fix It*. Sebastopol: O'Reilly Media.
- Mill, John Stuart [1863] (2001). *Utilitarianism*. Indianapolis: Hackett.
- Noble, Safia (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. NY: NYU Press.
- Robison, Wade L. (2017). *Ethics Within Engineering: An Introduction*. London: Bloomsbury.
- Ross, William David (1930). *The Right and the Good*. London: Oxford University Press.
- Sandler, Ronald (2014) *Ethics and Emerging Technologies*. New York: Palgrave
- Selinger, Evan and Frischmann, Brett (2017). *Re-Engineering Humanity*. Cambridge: Cambridge University Press.
- Tavani, Herman (2016). *Ethics and Technology: Controversies, Questions, and Strategies for Ethical Computing*, Fifth Edition. Hoboken: Wiley.
- Vallor, Shannon (2016). *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. New York: Oxford University Press.
- Van de Poel, Ibo, and Royakkers, Lamber (2011). *Ethics, Technology, and Engineering: An Introduction*. Hoboken: Wiley-Blackwell.

Online Resources

The Ethics of Innovation (2014 post from Chris Fabian and Robert Fabricant in Stanford Social Innovation Review, includes 9 principles of ethical innovation) https://ssir.org/articles/entry/the_ethics_of_innovation

Ethics for Designers (Toolkit from Delft University researcher) <https://www.ethicsfordesigners.com/>

The Ultimate Guide to Engineering Ethics (Ohio University) <https://onlinemasters.ohio.edu/ultimate-guide-to-engineering-ethics/>

Code of Ethics, National Society of Professional Engineers
<https://www.nspe.org/resources/ethics/code-ethics>

Markkula Center for Applied Ethics, Technology Ethics Teaching Modules (introductions to software engineering ethics, data ethics, cybersecurity ethics, privacy)
<https://www.scu.edu/ethics/focus-areas/internet-ethics/teaching-modules/>

Markkula Center for Applied Ethics, Resources for Ethical Decision-Making
<https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/>