Last Lecture of ECS 20

On Being a ~~Computer Scientist~~ Human Being in the Time of Collapse

Phillip Rogaway
Department of Computer Science
University of California, Davis, USA

22 June 2022

This essay is an edited transcript of my final lecture in ECS 20, a lower-division class in discrete mathematics. The 80-minute lecture was delivered 10 March 2022. It evolved from seminars given on 17 Sept 2020 (Reykjavik University, Iceland and Gran Sasso Science Institute, Italy) and on 5 Nov 2020 (Pomona College, USA).

Abstract

We stand today at the brink of civilizational and environmental collapse. What role have we computer scientists been playing in alleviating—or exacerbating—this threat? What role *should* we be playing? These are questions that many of us reflexively shun. At the same time, proceeding with business-as-usual, as most of us do, seems like madness.

Slide 1 I struggled for what to do with this last lecture. I had thought about spending another lecture on graph theory, or doing a review session, or giving you a taste of cryptography, which is my own area and a topic we haven't gotten to touch on. In the end I decided not to do any of those things.

I promised at the beginning of the term that every now and then we would pause to think about where we are—not just as individuals, but as a society, as computer scientists, and as students in a technical discipline. This lecture represents such reflection.

Slide 2 But first I should congratulate you on surviving my class. I know that many of you didn't want to be here—you wanted into the other section. But it was full; what could you do? In the end, maybe you benefited in some way, and perhaps came to realize that there is reason to my madness: if my homeworks were super hard or my lectures quite challenging, it was only because I wanted to make you better at thinking in the manner of a mathematician or theoretical computer scientist. Which is a different aim—one that requires more of you—from getting you to learn and recapitulate some basic material.

Still, even if I was successful at this goal, you really should ask: *was it worth it?* I fear that many

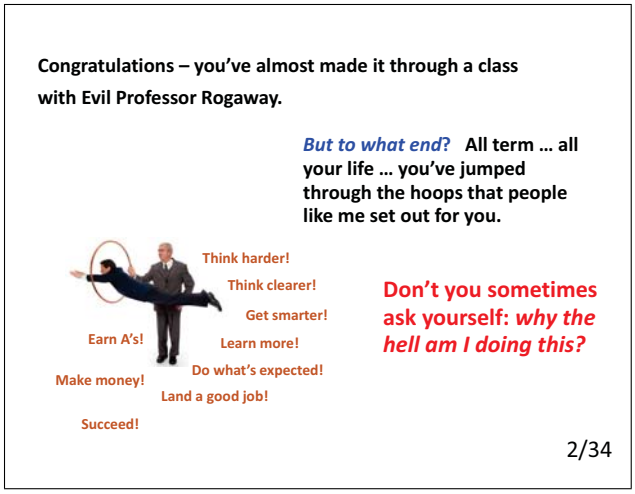

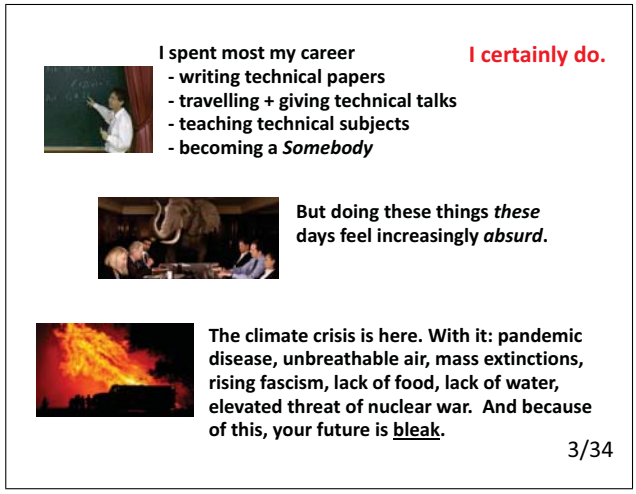

students don't ask questions like this. They just do as they're told and move on.

You've been pushed since you were very young to learn *this*, do *that*, and mostly you did it, appeasing parents, teachers, and coaches. And with all the pressure, both societal pressure and the more particularized pressure from the individuals close to you, and with all of this practiced capitulation, it can be hard to stop and be genuinely introspective.

But I think that if you're not adequately introspective, the value of life is radically diminished. Carrying on one's life in such a way that all you really do is to conform to the expectations of people around you is a stupid way to live.

If you've fallen into that trap, get out of it. Stop. Breathe. Think deeply about whether what you are doing is what you *should* be doing. Does it feel *right*?

Slide 3 I certainly have been re-examining my own choices. I've had that tendency since childhood. There's a personality test you can take, the Myers-Briggs personality test, that includes a quality P, for *perceiving*, which looks at your tendency to re-examine choices, postpone decisions, adapt to new conditions, change course. Being extreme in this trait is not necessarily convenient; for example, it can make it hard to finish things. For better or worse, I score a perfect 100% for it.

So what have I been doing—perhaps wrongly—for the last 30-odd years? Well, I've spent a lot of time time teaching classes like ours (as you may have noticed, I'm pretty obsessive about my teaching). Still, I've spent more of my life thinking about and writing papers in cryptography, and going around and giving talks on them. And it's been kind of fun. I've gotten to travel or live in about 65 countries. I saw myself as an itinerant intellectual.

I would like to think that my main motivation has been a pure desire to advance human knowledge—to be a scientist who lived the Mertonian ideals. That vision is what hooked me as a child. Yet I know very well that I have also harbored a desire to be a "big name" in my field—to be a *Somebody*. I find that motivation arrogant and elitist. Regardless, this entire way of spending one's life has come to feel increasingly selfish and outmoded. There is an elephant in the room that we seem to studiously avoid attending to.

The elephant in the room is the climate crisis. It is here. And with it comes all the attendant disasters: pandemic disease; unbreathable air; mass extinction (we have precipitated our planet's sixth mass-extinction event); the rise of fascism (not only in the US, but around the world); food shortages; water shortages; and an elevated risk for nuclear conflagration.

And because of all this, your future is bleak. It is bleak in the simplest sense: that many of you, I

UC Davis
Spring 2020

Our world already feels radically diminished.

4/34

The view outside my window, 9/2020

5/34

expect, will die from problems on that list.

The likelihood of civilizational and environmental collapse is something that should be profoundly affecting the psyche of us all. But I don't think it does. Probably because we're so good at not seeing that which is inconvenient to see.

Slide 4 At least to me, the world already feels radically diminished. The perception of decline became most acute in Spring term of 2020, when the university closed down for COVID-19. I remember walking here one day, around noon, in the middle of the week. There was nobody. Not in the buildings; not outside of them. I was Will Smith, walking the streets of New York in I am Legend (2007). It was all just—gone.

To be sure, life did go on. The birds were doing well; you could hear them in the trees. The squirrels seemed particularly raucous. At least some non-human life had survived our apocalypse.

Slide 5 The pandemic was only one factor that made it seem as though, by 2020, we were well into the collapse. The other was the wildfires, which by then had now become annual.

I was first forced to flee Davis in 2018. I have asthma, you see, and the smoke isn't just an annoyance for me—it's dangerous. That November smoke from the Camp Fire descended on us as I imagine the Great Smog of London did in 1952.

By the time the 2020 smoke arrived I had learned what to do. Having already planned to take a sabbatical in Portland, Oregon, I hurried it up, left town, and rented an apartment 600 miles to our north.

It didn't help. The fires and smoke pursued me.

This is a picture from outside my window in September of 2020. When the air outside is unbreathable, you feel yourself a target. The smoke is an enemy that wants to invade your home and destroy you from within. It is an implacable enemy, and a patient one.

Slide 6 Here's what it looked like inside my apartment. I put this sticky sort of Saran wrap along the edges of the door to try and keep the bad air out.

This is the bathroom vent, sealed with the same material. I didn't know it before, but the bathroom vents that are supposed to let air out can also let some in. Without sealing the vent, the smoke would worm its way inside.

I developed a strange relationship with the one air purifier I was able to buy before they were gone from all the stores. It wasn't nearly powerful enough to handle the whole apartment, so we hunkered inside a single bedroom of our flat. I entrusted my life to this device, so of course I felt grateful. At the same time, I saw in it a symbol of my incarceration.

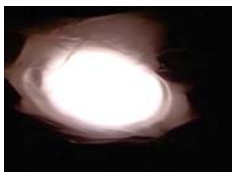

Inside my apartment

6/34

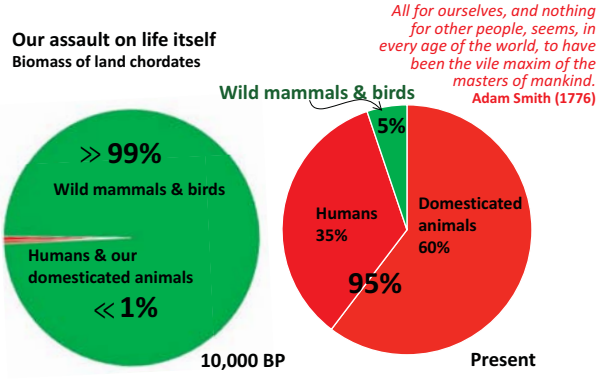

8/34

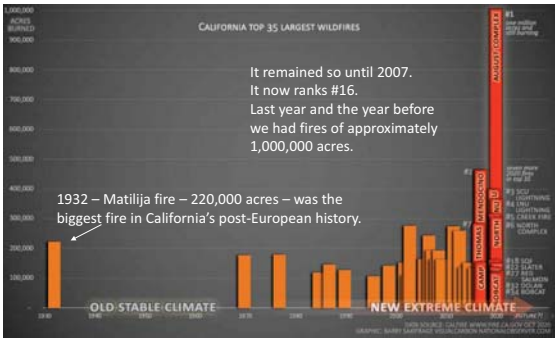

7/34

Slide 7 The massive fires are now an intrinsic part of living in California. This isn't just my anecdotal sense of things; it's what the data shows.

The largest fire since records were kept in California had been this 1932 blaze, the Matilija Fire, which burned 220,000 acres. This wasn't eclipsed until 2007. But since 2007, we've had *fifteen* fires larger than that 1932 blaze—some of which have been four or five *times* as large. For the last several summers—and now year-round—there'd be periods of dangerous air when large swaths of California's forests were going up in smoke.

Those forests used to do an admirable job of cleaning our air. That is how many people speak of it—as though the forests' purpose was to provide an air-cleaning service for *Homo sapiens*. Regardless, they don't do this "job" anymore. Since 2015, California forests have been so degraded that they're now net *polluters*. That which once seemed to be a key part of our potential redemption, the forest, has become yet another liability pushing us towards collapse.

Slide 8 When we look in other directions we see other indicators of the collapse of the natural world. Here's a chart showing the percentage of bird and mammalian life that is us. Humans and our livestock now comprise 95% of the bird and mammalian biomass. Basically, we allow other animals to live only if we plan to eat them. In a geologic instant we went from less than 1% of the bird-and-mammal biomass to 95% of it.

Our eradication of land vertebrates is a stunning expansion of Adam Smith's "vile maxim": *All for ourselves, and nothing for other people, seems, in every age of the world, to have been the vile maxim of the masters of mankind.* The quote speaks to greedy individuals taking all for themselves, leaving scraps for the poor. Now the target of our acquisitiveness has been enlarged, growing to no less than the entire biosphere. Now we must enlarge Smith's admonition: *All for ourselves, and nothing for other species, has*

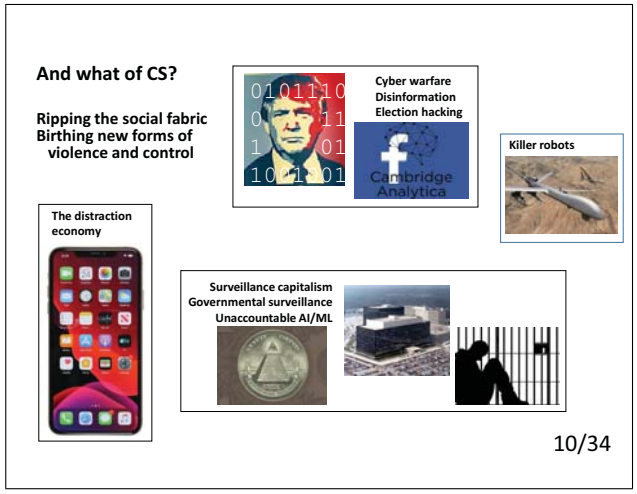

become the vile maxim of mankind.

Slide 9 At this point it seems reasonable to maintain believe that the earth needs to be saved from us humans (and that most humans, in turn, need to saved from the most rapacious). Yet sometimes I wonder if we are worth saving. It hits home each day I wander around campus and observe all the students who seem to me like a fundamentally new kind of being. The students—like most people—have become a cybernetic union: a biological entity attached to a smartphone. It goes far beyond an addiction; it is better conceptualized to say that we have become cyborgs, a union of our former biological selves and our technology. It may not have been realized in 2007—or even today—but Apple's *real* product—the iBorg, we could call it—is *us*.

The slide I have up—I don't know exactly what sort of event this was. It was a little before COVID. You'd see scenes like this. The students—the biological parts—are doing some sort of activity that their phone-part instructs. A treasure hunt or something. I'm not sure I want to know.

Slide 10 If you think of our world as having degraded to the point that people are no longer well conceptualized as autonomous biological entities, and if you have spent your professional life doing computer science, then it is natural to ask what role you and your colleagues have played in bringing about this extraordinary change. And the answer that I invariably come to is that we've played a central and overwhelmingly negative role.

I think you can put at our feet, at the feet of us computer scientists, many of the worst changes in our world. To begin, there's what we've just been talking about: the distraction economy (if you're optimistic) or the melding of humans and phones (if you're not). Intertwined with this is the rise of the surveillance state and surveillance capitalism. Nowadays you are tracked and analyzed continuously: through the analysis of your calls, messages, and emails, and through cameras that pervade our physical environment. You are monitored and manipulated by corporations that want to get your money or create brand associations within your brain, and by governments that want to detect potential threats and head off potential instability. None of this would be possible without CS. The changes are not an attack on privacy; they are the end of autonomy.

Then, relatedly, there is the rise of authoritarianism and the increasingly effective means for influencing elections and minimizing democracy. While propaganda and advertising don't fit squarely into any one academic discipline, the changes that have made them more pervasive and effective largely derive from us. Surveillance and behavioral intervention is the

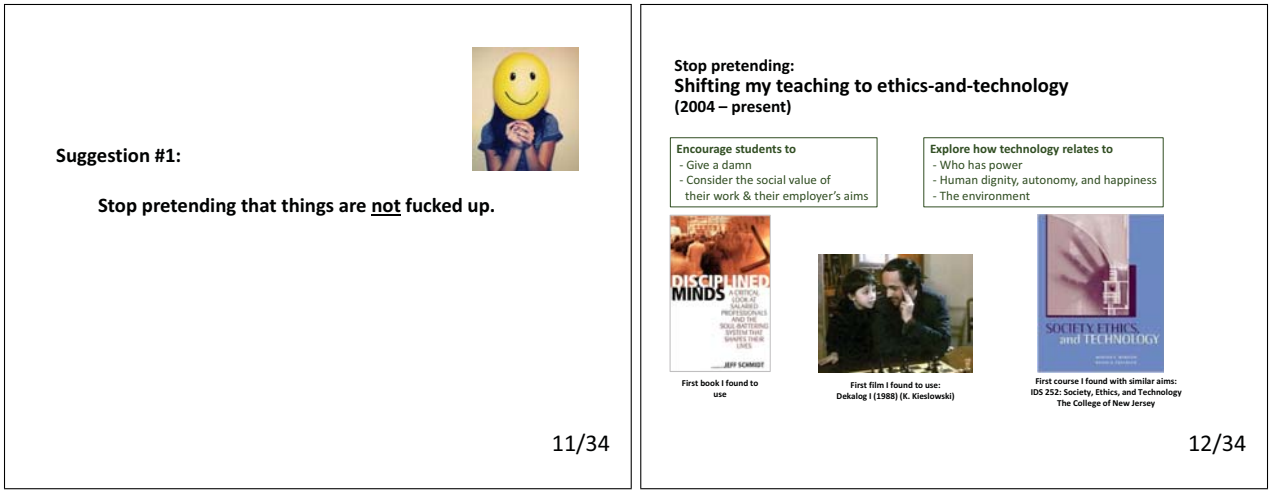

primary purpose to which AI/ML is put.

And let's not forget the increasingly effective ways of killing people from a distance, which again depends on CS. The business-end of the system might be a bomb on a helicopter or drone, but the communications, targeting, and control is mostly on us.

Because we are so much in the thick of it, you might expect that many computer scientists would be deeply concerned about these problems, working towards solutions, and serving as strong voices of warning. This is not remotely the case. Indeed we are among the biggest cheerleaders for change, however ill-conceived. We'll discuss why in just a bit.

Slide 11 I have spoken at some length without yet telling you anything concrete that anyone might *do*. I do want to give some concrete suggestions in this talk. Critiques devoid of suggestions may feel disempowering.

My first suggestion is simply this:

[Suggestion #1] *To stop pretending that things are not seriously fucked up.*

I think we've all been doing an extraordinary job maintaining that pretense. The fiction is embedded into your university experience, your familial experience, and the social experience that we craft for one another. We should question the routineness of all of these interactions in light of the fact that the climate is in crisis, democracy and human autonomy are in tatters, and your future is correspondingly grim.

For many years I have been trying to follow my own Suggestion #1. It isn't easy. There is strong social pressure not to do this—and strong psychological resistance, too. This talk embodies one way of discarding the pretense that all is okay.

I would like to emphasize that, as computer scientists, we absolutely *do* pretend that things are going great. We live this pretense in virtually every professional interaction. When I attend department meetings, read papers, assess grant proposals, or discuss work with colleagues, everyone speaks and acts as though everything is going well. It is the vibe we help instill in nearly every classroom, too. It feels as though we are playing out a scene from Titanic (1997). The musicians are stalwartly performing, playing their violins until moments before their icy end. They emblematize us scientists and technologists.

Then again, maybe this is not the right metaphor. The musicians didn't design the iceberg, profit from its deployment, or tout its myriad virtues.

Slide 12 Another way I've carried out Suggestion #1 is to shift the bulk of my teaching to ECS 188, a class called *Ethics in an Age of Technology*. I take a broad view in that class, addressing questions like: how does modern technology impact who has power

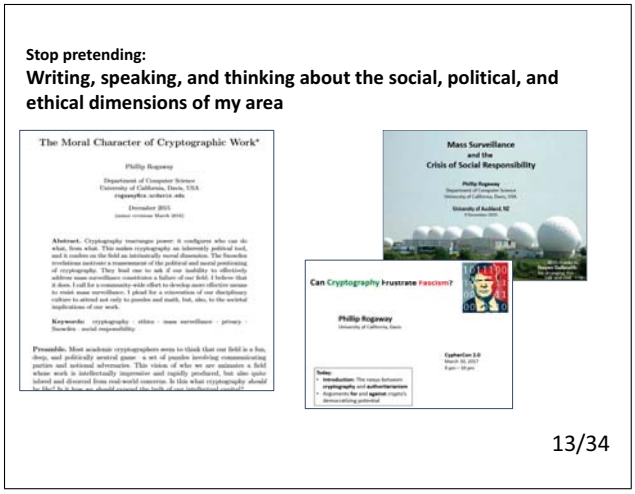

in society? How does it impact the environment? Democracy? Autonomy? War? Teaching CS from a lens bereft of such questions—what we do in all *other* CS classes—helps to create amoral students who casually contribute to our collective grief.

The course's main goal is to push students to give a damn. I want them to care what it is their prospective employer is doing in the world. The professional work you do is the main thing that determines whether you have a positive or a negative impact in this world. You can do all of the charity work you like on the side, but if your day-to-day work is helping Facebook create their shit, then you are exerting a negative influence on our world.

Slide 13 Another thing I've done is to go around to give *non*-technical talks. Talks like today's. Often I would try to communicate the simple message that what we are collectively doing as cryptographers, computer scientists, or technologists has profound ethical and social implications. And that when we do our work without seriously attending to those consequences, we are doing something gravely wrong. Such talks can reach a few people—especially younger folks, who have less invested. But most people, I know, will carry on exactly as before.

Slide 14 Finally, I actualize Suggestion #1 by occasionally engaging in some direct political activity. For example, I authored the IACR statement repudiating mass surveillance. The IACR is the professional organization for cryptographers. The statement just says that *The membership of the IACR repudiates mass surveillance and the undermining of cryptographic solutions and standards. Population-wide surveillance threatens democracy and human dignity. We call for expediting research and develop development of effective techniques to protect personal privacy against governmental and corporate overreach.*

The statement might not sound like much, but it's hard to get a professional organization to say or do *anything* in the political sphere. The primary mission of such organizations is to protect the interests of their members, and "getting political" is seen as antithetical to the interests of many of those members.

Professional organizations often have language in their charters or canons of ethics that speaks to social responsibility. But that doesn't mean that those aims have been internalized or made operationally significant either by members or the organization's leadership.

Slide 15 By now you may be feeling that I have expressed an overly dark view of where we are at: a world already in collapse, with folks like me smugly

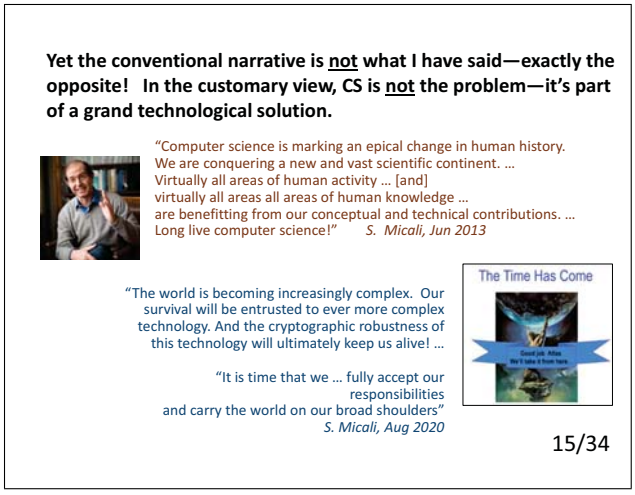

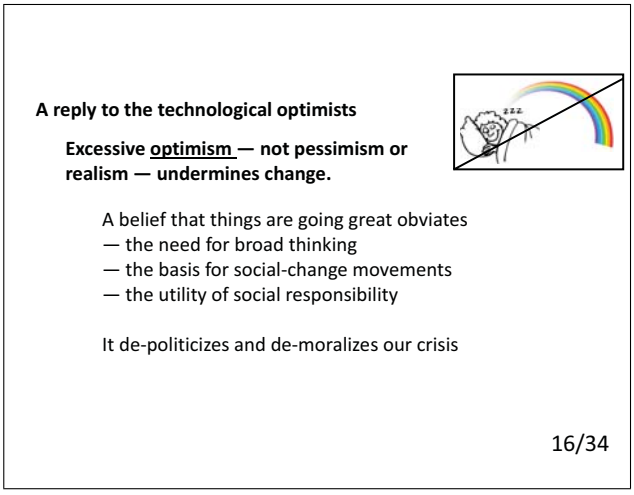

contributing to that collapse. A minute *past* midnight (using the metaphor of the Doomsday Clock). Perhaps you will take comfort from knowing that mine is a minority view.

My own advisor, Silvio Micali, nicely expresses the antithetical view. In Silvio's Turing Award acceptance speech he says that *Computer Science is marking an epical change in human history. We are conquering a new and vast scientific continent. Virtually all areas of human activity and virtually all areas of human knowledge are benefiting from our conceptual and technical contributions. Long lives computer science!* You can't get much more optimistic than that.

In a more recent talk, Silvio says that *The world is becoming increasingly complex. Our survival will be entrusted to ever more complex technology. And the cryptographic robustness of this technology will ultimately keep us alive. It is time to fully accept our responsibilities and carry the world on our broad shoulders.* The accompanying graphic shows the computer scientists taking over the work from Atlas.

I appreciate how optimistic Silvio is, and can smile with the tongue-in-cheek ebullience. But as someone who has worked in cryptography my entire working life, I am keenly aware of just how fragile are the things that we create. You change the slightest thing in a cryptographic protocol and it totally breaks. You change a single word or symbol in a cryptographic definition and it has an utterly different—and typically inadequate—meaning. Things as banal as power usage or timing-characteristics due to cache interactions will cause catastrophic failures. Cryptographic mechanisms are mathematical snowflakes. That is why I love them. The idea that we are going to entrust ever *more* to these snowflakes, not less, seems to me unwise to an extreme.

Slide 16 The technological optimists would have you think that technological pessimism is a bad attitude, in need of correction. A barrier to our coming utopia. I would answer that it is technological optimism, not pessimism, that has become so debilitating. Unbridled optimism eviscerates social change movements. It enfeebles the environmental movement.

Suppose you believe that everything is rosy, and that technology will get us out of whatever problems that technology has engendered. Then, as a technologist, there's really no need for broad thinking or political engagement. Your social contributions will arise through your technical contributions. *That* is how we will build a better world.

In this way of thinking, social change movements are in error; we are already on a desirable path. And whether or not you care about any of this stuff is ultimately irrelevant, because technological advance will inevitably see us through. In this way we depoliticize our current crises. We look towards the scientists and

17/34

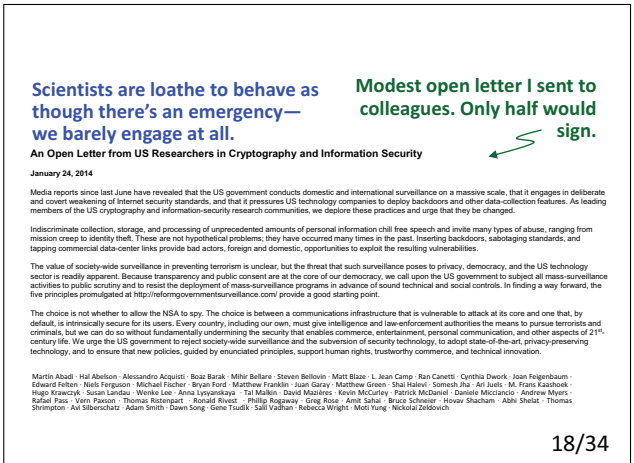

18/34

engineers and say, *Go forth! Invent us fusion power. Make us healthier. End the tedium of work.* And we wait for the *Deus ex machina* that we will learn to build.

Slide 17 Personally, I'm tired of all this optimism. It rings shallow. It doesn't arise from some reasoned assessment; it's just a religion. One whose effect is to prop-up the *status quo*. Gretta Thunberg expresses it nicely (she expresses many things nicely) when she says: *I don't want your hope. I don't want you to be hopeful. I want you to panic and act as if the house was on fire. Because it is.*

Slide 18 Few computer scientists seem worried, let alone in a panic, about our burning home. Most of us are optimistic and apolitical actors.

After the Snowden revelations I co-organized an open letter. I wanted to get prominent cryptographers and computer security experts to jointly express our outrage at secret programs for mass surveillance. These largely mooted whatever claims we might make that our work was adding meaningful protection to real-world systems.

The first version of the letter didn't seem especially strong to me, but virtually no one was willing to sign it. After more than 900 emails it got watered down to the point that half of the the people approached were willing to sign.

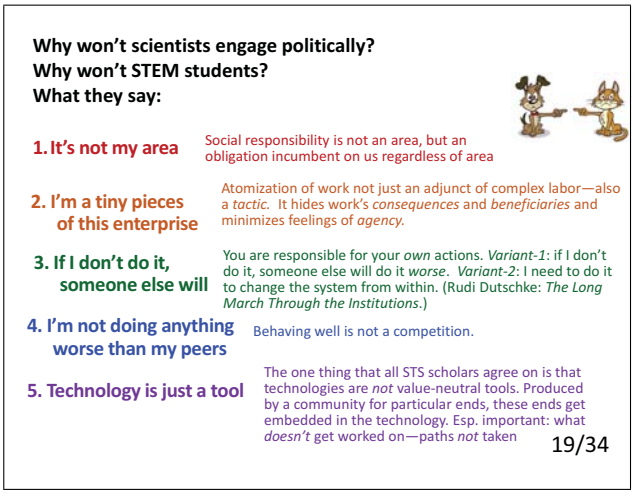

When I would ask people about their unwillingness to sign, I kept hearing the same answers. They were also the same answers I would hear from my ECS 188 students about why *they* were unwilling to engage in anything political. Let's go through the popular reasons voiced.

Slide 19 (1) *It's not my area.* Many computer scientists don't want to engage in social or environmental questions pertaining to technology because they feel it's outside their area of expertise.

It's easy to understand where this comes from. Computer science, like most academic disciplines, has become extremely specialized. When people come to give seminars in areas of CS outside my own, I don't understand much after the introduction. Even in my own area, I'm unlikely to fully follow a talk in any area I haven't worked on. So how could I, without any specialized training, pretend to speak knowledgeably about social, political, or ethical matters, none of which are close to my research? Someone else, presumably, works on such stuff.

It is a stupid excuse. Social responsibility is incumbent on *every* computer scientist.

That isn't some fringe opinion—it is, officially, a well-settled manner. The very first imperative in the ACM Code of Ethics, for example, is to *contribute to society and to human well-being.* The second imperative is to *do no harm.*

Then again, norms that aren't routinely followed or felt aren't exactly norms.

(2) *I'm just a tiny piece of this enterprise.* Complex work is routinely partitioned into tiny pieces, with individual workers lacking a clear view of the activity as a whole. We have come to accept this as though it were a necessary organizational principle. But atomization of work is not a necessary adjunct to complex labor so much as a tactic used to organize that labor. One purpose of this tactic is to hide the likely consequence of one's labor, obscuring the identities of the work's beneficiaries and victims. And atomization of work lessens the perception of agency among workers, who are less likely to organize when they're unclear about what it is that they're really doing.

(3) *If I don't do it, someone else will.* This is one of the most frequent and annoying excuses. It ignores one of the most basic ethical tenets: that you are responsible for your own actions in a way that is fundamentally different from your responsibility for the actions of others.

A variant is: *if I don't do it, someone else will do it* ***worse***. I always love how, in our mental model of the world, everyone else ends up less ethical than oneself.

Another variant is that I *should* walk into the belly of the beast because, when the time is right, I will impart my *own* values on the institution. This vision gained popularity in the student-protest movements of the 1960s. It is particularly associated to the German activist Rudi Dutschke, who spoke of the *long march through the institutions.* Idealists would join evil institutions and slowly rise to positions of power within them. They'd be sleeper agents. When the time was right they would transform the companies from within.

I think we have enough history by now to know that this approach does not work. Almost invariably, the institution changes you, not the other way around. Your chance of making it from idealistic newhire to top-management with your former values in tact is close to zero. If you're smart, you'll get out of the organization before the institution has changed you into something that your former self would have found abhorrent. More likely, you won't even notice how much you've changed, or you'll come to view your former self as naïve.

I regularly advise students who have ethical concerns about an employer to drop any fantasy of changing them. It is better is to find an organization with values compatible with your own. For principled individuals, this can severely limit potential employment.

(4) *I'm not doing anything worse than my peers.* People will explain that what they're doing is no different or worse than what lots of other people do. The rationale extends beyond employment. For example: surely it is okay to eat beef, as most people around me do.

It should go without saying that behaving well is not a competition. Still, I see indignant resistance when I suggest to someone that some routine behavior is unethical. Aldo Leopold explains that an *ethic, ecologically, is a limitation on freedom of action in the struggle for existence.* So understood, it feels *unfair* to people that they should live a life less free than others. They think: it would be unfair for my ethics to limit my options.

The answer to this complaint is a theme of Leopold's essay: that as we evolve, our moral universe expands. What was once seen as permissible is no longer. The time will come, I hope, when many

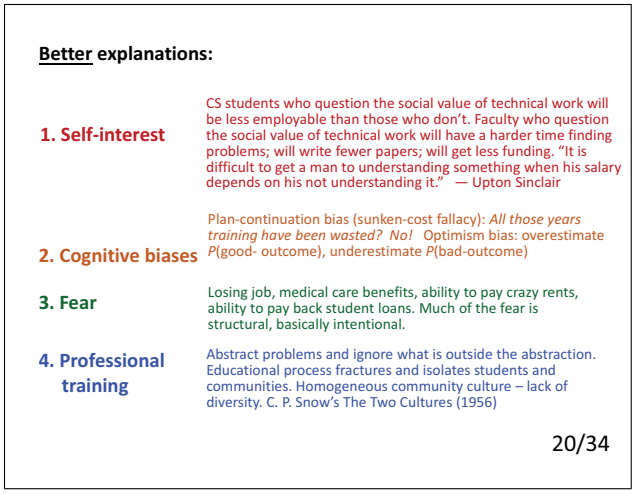

behaviors that are seem as ethically permissible today will be seen as far beyond the pale.

(5) *Technology is just a tool.* Even quite smart people will sometimes make the claim that technology is just a tool. A knife in the hand of an assassin can kill, but in the hand of a surgeon can cure. That sort of thing.

I think that the popularity of this drivel comes from a narrow education among STEM-types. If there's one thing that every STS scholar agrees on, it is that technology is *not* some sort of value-neutral tool. Technologies are produced by a community of people that have a specific history and specific values and goals. All of this gets embedded into the design of whatever they work on.

It is more than the fact that low-level design choices embed designer and company intent. The most important decision is what does and doesn't get worked on in the first place. What doesn't get worked on leaves no historical trace. For this reason, the extent to which our technology is socially constructed is vastly underestimated. There are an infinite number of paths we did not take, and we have no perspective from which to observe or describe those paths.

What determines the technological paths we see and take? Mostly, we see and take paths where someone thinks they will soon find money, power, or prestige. And that's quite a narrow filter.

Slide 20 Let me give some alternative explanations as to why technical people rarely engage with the social, ethical, and political repercussions of their work.

(1) First, it is not in your interest to do so. "It is difficult to get a man to understanding something when his salary depends on his not understanding it" said Upton Sinclair. CS students who question the social value of work will be less employable than their peers: most employers don't want such as employees, even if they claim otherwise. Such people are seen by employers as potential trouble-makers. On the flip side, ethically-concerned students probably won't want to work for many employers, either.

As for faculty members, it's not any different. Those who make a habit of questioning the social value of their own work will have a harder time finding research problems they believe in, will write fewer papers, and will get less funding. Those who question the social value of their colleagues' work, or the appropriateness of their colleagues' funding, will find themselves ostracized. If you think that professors take kindly to having their values questioned by their peers, you are very much mistaken.

(2) Secondly, we share cognitive biases, and are subject to fallacies of reasoning, that make us disinclined to recognize or address the kind of problems we've been discussing. To begin, the *sunken-cost fallacy* reflects a reluctance to accept ideas that would make one's past efforts feel wasted. When I teach ethics to seniors in CS, they strongly resist accepting the possibility that they selected the wrong major. Because it is too late to fix this.

I try to ameliorate the sunken-cost feeling not only by emphasizing that it *is* a fallacy but also by telling students that they are still young enough to change paths, and that if their training in CS was good, then they have been broadly prepared to think, work effectively, and express themselves well. Maybe their time *was* well spent.

Then again: maybe not.

There are other biases that come into play. An *optimism bias* causes most of us to overestimate the probability of desirable outcomes and to underestimate the probability of negative outcomes. This bias is particularly strong for matters technological, where

the culture barely tolerates pessimistic projections.

The *bandwagon effect* leads people to copy the behavior and beliefs of those around us, nudging us towards conforming attitudes and acts. In tech, this works to minimize ethical concerns.

There are technological optimists like Steven Pinker who claim that our cognitive biases run in the opposite direction—for example, that the *availability bias*, when coupled with a media landscape that favors spectacle, pushes us towards pessimism and fear. I think this is ridiculous. The messaging from "mother culture," as Daniel Quinn calls it, is the real winner of the availability bias. And it tells us that technology is the engine of man's ascendance.

(3) Outright fear keeps us disengaged. I remember telling a former student who took a job at Google that maybe he should organize a lunchtime reading group to discuss Shoshana Zuboff's *Surveillance Capitalism*. He laughed. He said it that sounded like a good way to get fired—or at least to come under the microscope.

While companies may try to create the impression that you're on a long leash, there always *is* a leash. You are forever being observed and evaluated. Anyone genuinely opposed to the aims or methods of the company are not welcome in it.

The fear is rational. You took that job in Silicon Valley, where the least expensive homes were extraordinarily expensive. You have a mountain of student debt. By now you're paying quite a sum just to grab a decent dinner. Your spouse likes nice things, too. Losing your job could mean not being able to cover your mortgage, your student debt, or live the life you've gotten used to. So you come to feel that *can't* push back, nor take the kind of lower-paying job you might otherwise prefer.

(4) Finally, your professional training works to keep you on a narrow path. It starts with the basic design of the curriculum, which separates technical aspects of work from an understanding of how the work is situated. In classes, the latter is virtually ignored, signaling that it *should* be ignored.

But it is more than that: there are aspects specific to computer science that make its practitioners especially disengaged.

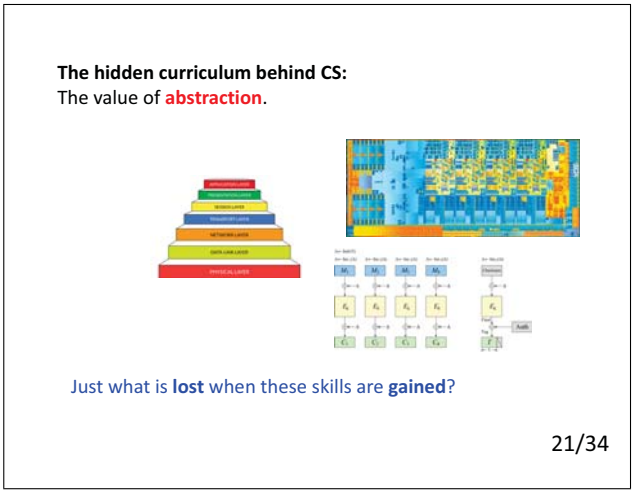

Slide 21 What is the implicit message you get from your CS classes—from classes like the one you just had with me?

When people speak of the *hidden curriculum* in schooling, there's a couple of levels at which this phrase can be understood. One is quite broad. You go to some class and hear the teacher expound on some subject. The subject doesn't matter. The "real" message, arguably, is that the teacher is in front, in a position of authority; you are behind, in a position of acquiescence to that authority; you have an assigned role, that of student; and those with this role are expected to learn what the teacher has already mastered. You'll be drilled on it, assessed on it, then sorted on it. The hidden curriculum, one might say, is to acculturate you to a life of hierarchy, authority, compliance, boredom, mass production, evaluation, and sorting.

The above applies to most every class. But I claim there's also a more *particular* hidden curriculum within different disciplines. Things that are, once again, little seen or spoken of because they're so fundamental to a discipline.

The discipline-specific message for CS is, I claim, *the value of abstraction*. Crafting and respecting abstraction boundaries is the secret-sauce that undergirds almost all CS. Computer science training—at least "good" computer-science training—is about advancing students' skills to employ, respect, and invent

good abstraction boundaries.

To be more concrete: how exactly are you able to do something now—solve some interesting counting problem, let us say—that you couldn't do before our class began? Probably because you are better able to create and use the sort of abstraction boundaries that have proven useful for this domain. Once you internalize what the permutation function $P(n, k)$ and the combinations function $C(n, k)$ count, and once you internalize some basic ideas for using these things, you have a new conceptual universe at your disposal. Just *naming* these functions is already powerful.

Effective and pervasive abstraction boundaries: that is how we've managed to create modern computers and get them to do complex things. At the upper-right of this slide is a picture of a relatively recent Intel chip. There are so many transistors on these things that you can't hope to understand the chip by understanding the role of each. But you don't have to. There are multiple levels at which to see the object. The high-level architecture might entail an arithmetic logic unit (ALU), a bank of registers, a level-2 cache, and so on. It doesn't matter where they are on the chip—it doesn't even matter that they have a physical embodiment. Or you can shift to a different abstraction, looking at the instruction set architecture. Then another, and another. Someone expert in device physics may have a detailed understanding of how each transistor functions, yet no understanding of the chip's architecture. With the right abstraction boundaries we can create worlds of breathtaking complexity.

Being able to pick up or craft new abstraction boundaries is a valuable skill. Yet there is an attendant cost. It can easily become your primary way of thinking—and not just in technical realms. Like all ways of thinking, it can drive out other ways of thinking. If abstraction-partitioning is the wrong way to go in some domain then your training in CS isn't going to help you. It might hinder you instead.

Slide 22 There are, of course, computer scientists who sincerely *want* to do good. This desire is most significantly reflected in the research areas where people choose to work and in the specific topics they choose within those areas.

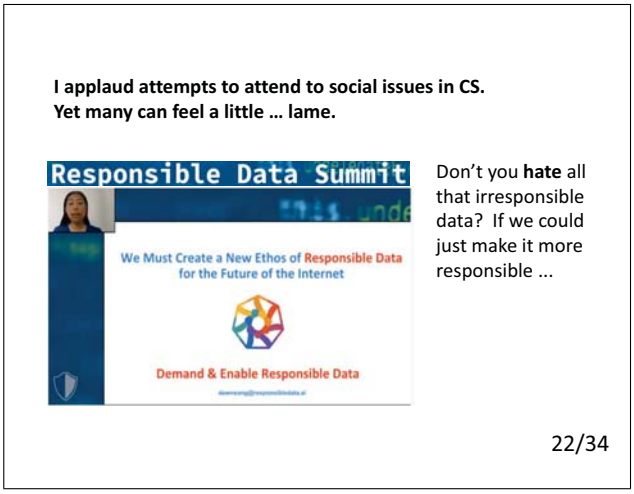

I appreciate almost any effort to inject personally held ethical beliefs into technologists' work. It is far too rare. And if the most ethically engaged colleagues in CS are not my allies, then who?! Just the same, I want to gently express three concerns in this connection.

(1) First, as just suggested, computer scientists often seem to approach social and ethical issues in computing with the same mindset, and even the same techniques, that they bring to "conventional" problems in CS. It's not a good fit. Conventional problems in CS are like puzzles: self-contained, abstract, or mathematical. Problems of social and ethical importance rarely have any of these characteristics.

The slide I've put up opened the 2020 "Responsible Data Summit." Even the name seems potentially revelatory, no? How I hate all that *irresponsible data*! You know, data that just won't conform to our social or ethical norms. We ought to teach it a thing or two.

I don't think it's just careless language. Rather, the language reflects a way of thinking about sociotechnical problems that puts the emphasis on the *things*, tangible or abstract, that we create—instead of placing the emphasis on the people and companies that conceive, design, and use these things.

You could see this mindset in almost every talk. It was as though we could change the world without

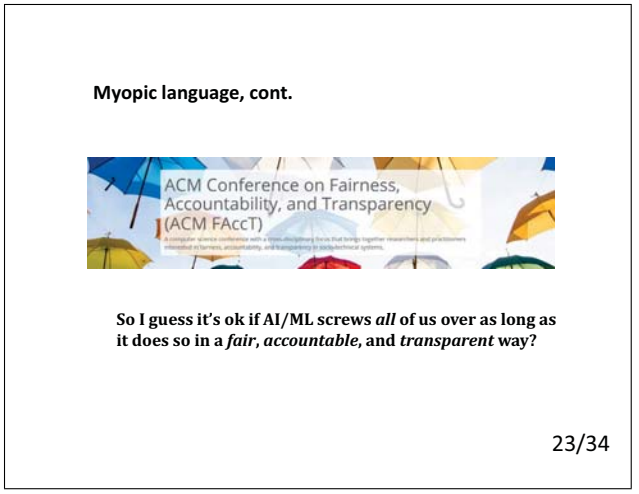

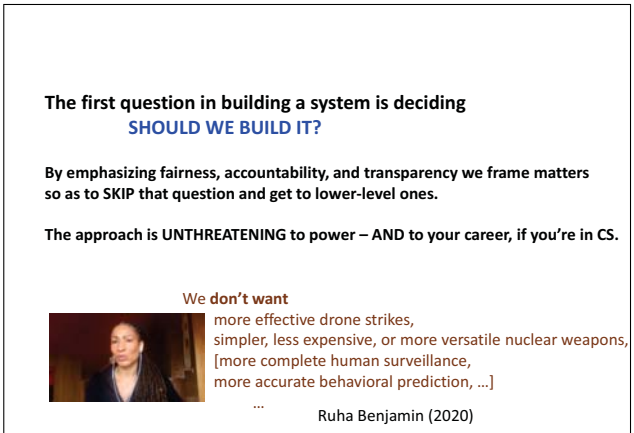

having to change anything significant about our work or ourselves. Convenient.

Slide 23 (2) Second, I want to warn you away from working to improve things that we're better off without.

Here's the webpage for an ACM journal called the *ACM Conference on Fairness, Accountability, and Transparency* (FAccT). The journal is about ethical concerns in ML. But the title would seem to signal a topical breadth of addressing three particular complaints. Is this triumvirate our primary concern with ML? If it were, it would presumably be fine to create systems that screw us *all* over, as long as they do so in a way that is fair, accountable, and transparent. Just what we need.

Slide 24 By framing ethical problems as a quest for under-achieved desiderata we bypass the first and most important question: *whether or not to build the damn thing in the first place.* I believe that the answer to this question is so often *No* or *I don't know* that researchers, developers, companies, and funding agencies really want you to skip this question—hopefully without even noticing you have. That the thing *should* and *will* be built ought to come across as inevitable—almost as a law of nature. Only details of execution are in doubt.

Such a framing is unthreatening both to scientists' careers and to the companies where many of them work. If you're an academic, adding a consideration like *fairness* will give you yet more topics to write about, yet more systems to build. And Google, say, has nothing to fear. It's not as though you'll be asserting that Google needs to stop the massive over-collection of personal data. It will be more like: Google can collect and use data in ways that achieve this property or that—and when we get it all worked out, there will be even *less* for ordinary folks to worry about. Attending to fairness, accountability, and transparency will hinder nothing that company wants to do, and can even ground a narrative that helps bolster their over-collection.

Before one worries about concerns like fairness, accountability, or transparency, make sure to ask if the thing, in any incantation, deserves a place in our world. I like how Ruha Benjamin expresses it when she says that we don't want more effective drone strikes, or simpler, less expensive, or more versatile nuclear weapons. Nor, I would add, better systems for human surveillance, nor better tools nudging people to do what someone wants them to do. Technological improvements to improve a fundamentally rotten enterprise don't make life better. They make it worse.

Going back to FAccT, I should mention that, in browsing recent papers, it seems that many partici-

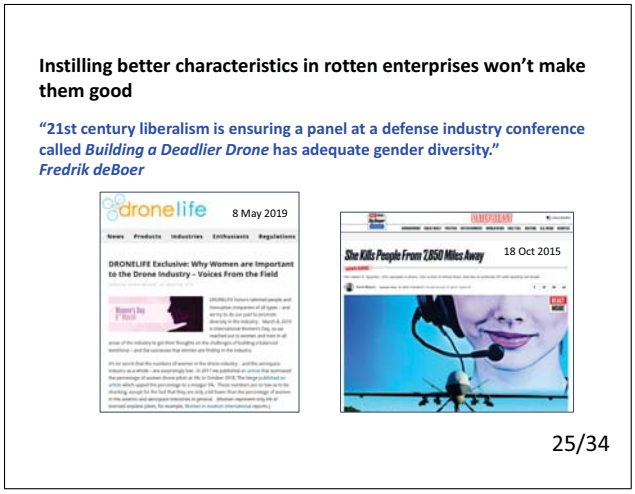

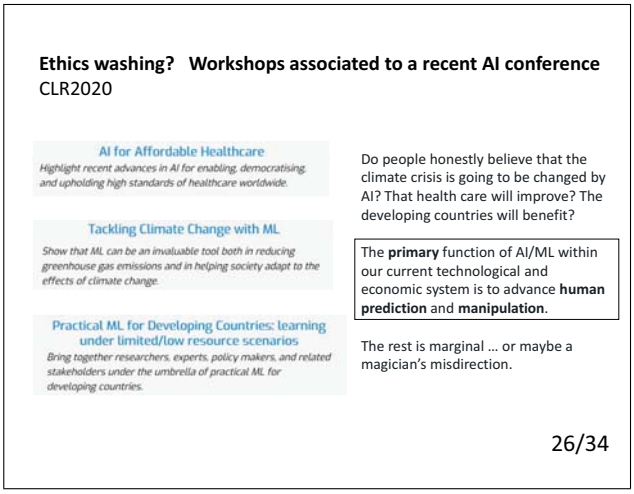

pants *are* willing to ask not just the how-to-build-it-better question, but the should-we-build-it question. This line distinguishes discourse that power will have no problem with and discourse that will be painted as radical, regressive, and inappropriate.

Slide 25 Our teaching assistant Zane Rubaii told me this lovely quote from Fredrik deBoer, that *21st century liberalism is ensuring a panel at a defense industry conference called "Building a deadlier drone" has adequate gender diversity.*

After Zane told me this quote I did a web search and, sure enough, found an article complaining about the lack of gender diversity among the pilots of killer robots. And another about some badass woman who seems to enjoy killing Arabs from her base in Nevada. The articles were serious—not jokes inspired by deBoer's quip.

DeBoer's quip sheds light on the CS approach to diversity. Faculty in my department happily attend to gender, racial, or cultural diversity of faculty applicants, but won't go anywhere near trying to assess if a person's work is desirable or wretched. (Tacitly, all CS research is assumed to be desirable.) Research contributions are a vector $\vec{v}$ that we refuse to asses as anything but a scalar $|\vec{v}|$.

Slide 26 (3) My third and final warning is that *ethics washing* is alive and well. At workshops on ML, for example, you will see sessions like *AI for Affordable Health Care*, and *Tackling Climate Change with ML*, and *Practical ML for Developing Countries*. I wonder if the people involved have much of an understanding why health care is unaffordable in the U.S., why the climate is collapsing, or what plagues the "Developing" world. If they did, would they honestly believe that AI/ML would be of significant help?

When an organization like Google, whose basic business model is founded on human surveillance and behavioral interventions, assembles a research team (then an advisory board, then no advisory board, …) for ethical AI, this feels borderline absurd.

Be skeptical. Attend to what organizations do, not what they say. Many companies are unethical by design. It's in their DNA. Don't be a party to someone else's ethics washing.

Slide 27 Gathering up some suggestion implicit in what I've been saying, I'd like to recommend that we

[Suggestion #2] *Stop touting technological solutions to social problems.*

For scientists and engineers, our technological solutions to societal problems are routinely self-serving, and they almost never actually work. Next, and relatedly, we need to

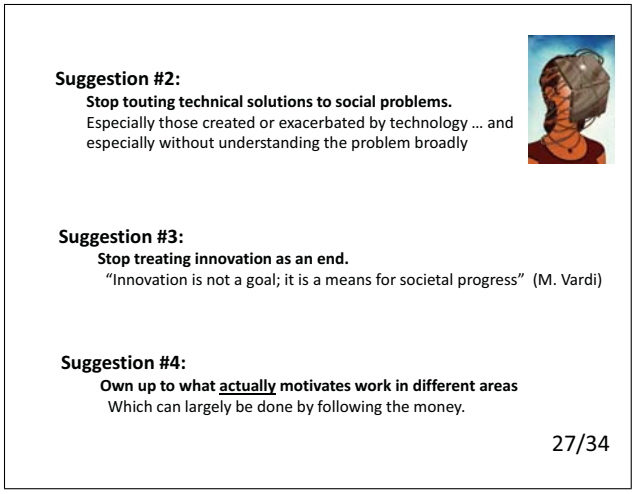

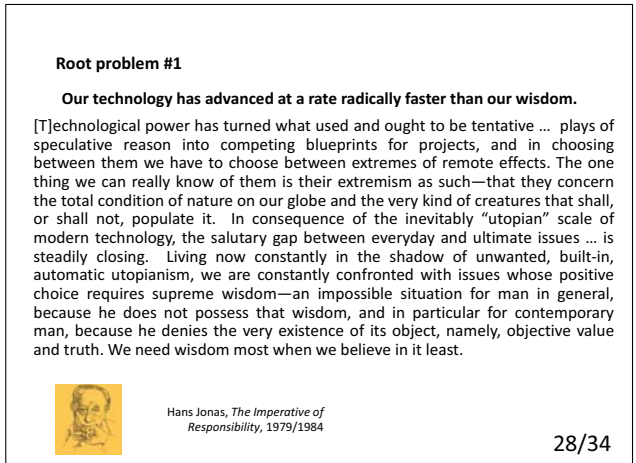

[Suggestion #3] *Stop treating innovation as an end in itself.*

Innovation is something we do for some *other* purpose. It is not an end. Once again, personal interests seem to cloud our judgment. Finally, I think that we should

[Suggestion #4] *Own up to what actually motivates our collective work.*

Almost by definition, we can understand the operative aims of an area by following the money. People who work in AI/ML might talk about how well an ML program can read an x-ray, say, but that's *not* what is pushing the development of the field. The forces that have driven the area's extraordinary growth spring from a desire to more effectively surveil people and influence what they do. If you're a student who gets a job in Silicon Valley using the AI/ML that you study here at UCD, that's probably what you'll be doing.

Slide 28 Let me try identify a few of the basic problems that have put us in the predicament we're in. One is that we are unwise creatures. We are an animal that suddenly and recently acquired the ability to decimate its environment and to reinvent the basic experience of life. We did not yet acquire the wisdom necessary to select well among possible futures. While individuals may see the threats and want to change course, we have not created institutions that are up to the task.

Hans Jonas speaks eloquently of our changed conditions. He writes: *technological power has turned what used and ought to be tentative … plays of speculative reason into competing blueprints for projects, and in choosing between them we have to choose between extremes of remote effects. The one thing we can really know of them is their extremism as such—that they concern the total condition of nature on our globe and the very kind of creatures that shall, or shall not, populate it. In consequence of the inevitably "utopian" scale of modern technology, the salutary gap between everyday and ultimate issues … is steadily closing.*

Slide 29 Another root problem is capitalism—at least capitalism as it is currently exercised. We have embedded our technological march within a particular variety of growth-oriented corporate capitalism. And I don't think it is possible to fix the problems we face within those confines.

It is a big topic, and one that I feel poorly equipped to speak about. But you should ask yourself questions like whether or not humanity can exist in a world in which Chevron does. I don't believe that it can.

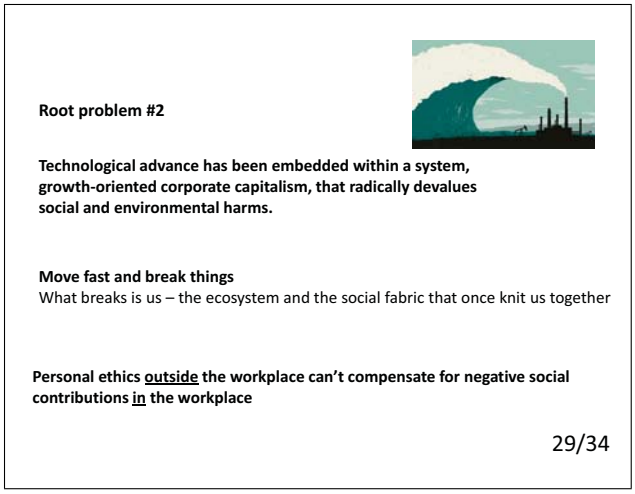

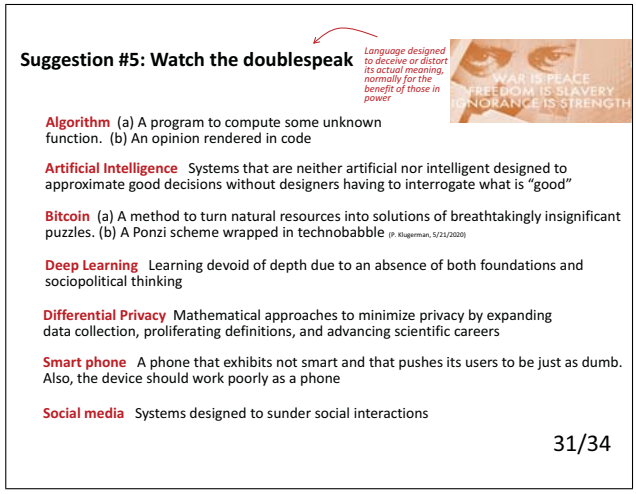

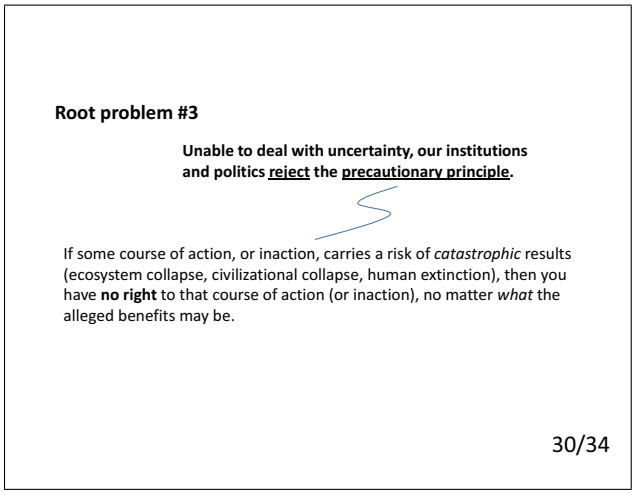

Slide 30 A third problem I would point to, and related to the other two, is that our institutions and politics do not realize the *precautionary principle*.

The precautionary principle says that if there's some activity that you anticipate doing, but the activity carries with it some risk of catastrophic results—say our species will go extinct, human civilizations will collapse, or you'll cause another mass extinction event—then you may not do that activity regardless of its alleged benefits.

The precautionary principle also speaks to inactivity. If carrying on the *status quo* carries with it some catastrophic risk, then you are obligated to change course, even if that change is costly. You must try to eliminate the risk of a catastrophic outcome.

The precautionary principle flies in the face of cost-benefit analyses. Those who would adhere to the principle would insist that cost-benefit analyses are inappropriate in any domain where a possible outcome is catastrophic.

Slide 31 Returning to concrete suggestions, my next one is to

[Suggestion #5] *Stop the Orwellian doublespeak.*

In the last few years, I've come to think of computer science as a veritable treasure trove of doublespeak. A few examples.

When I was young an **algorithm** was an unambiguous method to compute a known, well-defined function. Nowadays it is a program that computes *something*, but nobody knows just *what*. It's somebody's opinion rendered in code. You'll never meet that somebody.

How about a **smartphone**? This is a piece of techno-kitsch that is not smart and that barely works as a phone. It attaches to a biological host, which it drains of intelligence. Thus: a stupid device designed to make you stupid. They are also terrible phones: my Nokia 3310 cell phone, which came out in 2000, was a far better phone, in every significant way, than

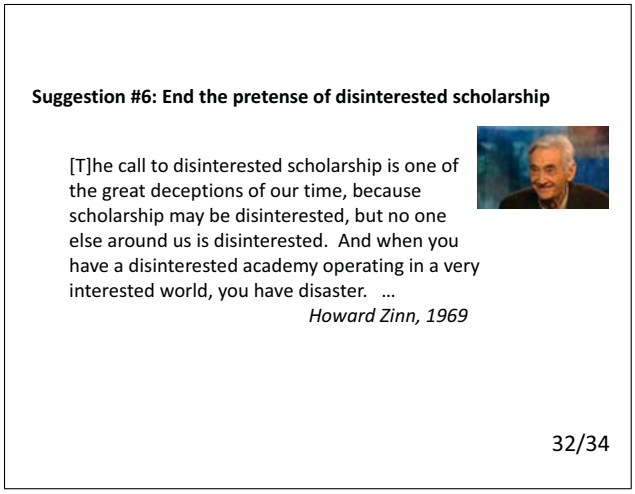

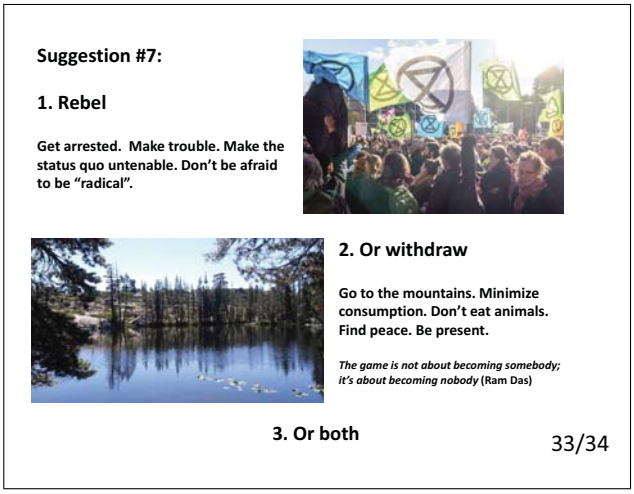

anything on the market today. Smaller, lighter, sturdier, cheaper, better sound, a longer battery life, less distracting, and worse at surveilling you.

One could go on and on. **Artificial intelligence**. **Bitcoin**. **Cloud computing. Deep learning**. **Differential privacy**. **Social media.** When you hear CS lingo, try flipping the meaning of words and see if you get closer to the truth.

Slide 32 My next suggestion is to

[Suggestion #6] *End the pretense of disinterested scholarship.*

Nowadays, few outsiders to an academic discipline would imagine it as a realm of disinterested scholarship. Perhaps we owe this improved understanding to awareness of the tobacco industry's work to get academics to do and say what the tobacco industry wanted. All the same, the pretense of disinterested scholarship remains intense *within* academic departments. For example, if you ask a CS professor who is being funded by DARPA if that funding impacts the work that he does, he will swear up and down that it does not. He'll also tell you that the work does not benefit the military. Beliefs that are as self-serving as they are implausible.

"The call to disinterested scholarship," Howard Zinn said in a talk in 1969, "is one of the great deceptions of our time. Because scholarship may be disinterested, but no one else around us is disinterested. And when you have a disinterested Academy operating in a very interested world, you have disaster."

Slide 33 My final suggestion is here. That we

[Suggestion #7]. *Rebel. Or withdraw. Or both.*

I wish our students were more politically active. Make some trouble. Get arrested. Fight the *status quo*. Because the *status quo* is going to lead to your death or the death of those around you. It's going to lead to the extermination of most animal species on this planet. It will lead to perpetual surveillance, fascism, and war.

It is hard to feel rebellious in the U.S., where power is equated with money; or at a university like ours, where the student body seems so passive. But things can change.

The banners shown in this slide are from an Extinction Rebellion rally. They're a group of activists mostly in the United Kingdom. I wish we saw in the U.S. the same sort of verve. In comparison, the Sunrise Movement comes across as a social club, as students more interested forming friendships than changing the world.

If you can't quite rebel, I understand and sympathize. Perhaps it already feels too late. Or perhaps it seems dispositionally infeasible for you. An

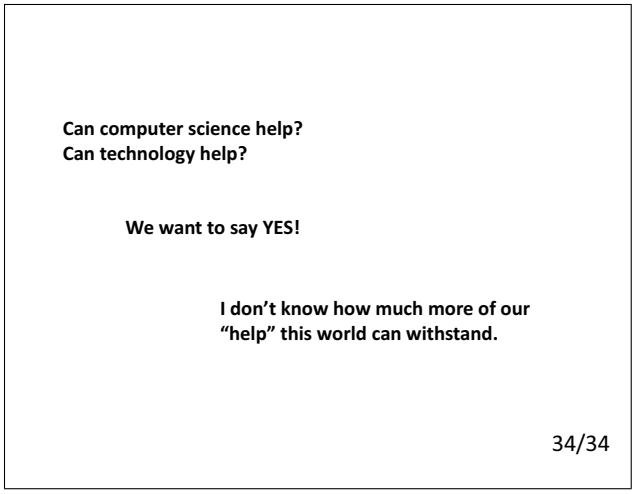

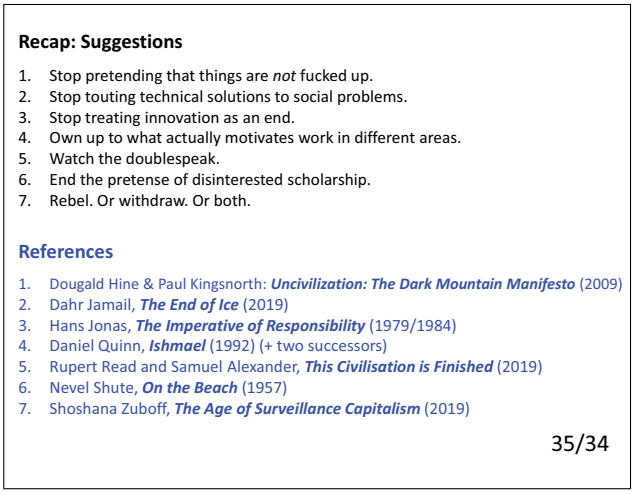

alternative is to withdraw. To try to minimize harm. Withdrawing is okay at a personal level, and maybe necessary to protect yourself in a world gone mad.

In Dahr Jamail's poignant book *The End of Ice*, this is where he came down personally, at least at the time he wrote the final chapter. That he couldn't do much to stop the ice from vanishing, the seas from rising, or the ecology from collapsing. But at least he could bear witness to it. He likens it to sitting vigil at the bed of a dying friend. And it is not valueless to sit at a friend's deathbed. There is dignity in it.

And I am not saying you can't do both of these things—to rebel on weekends and withdraw on weekdays, or whatever. We are complicated beings, people are, limited in our capacities. Be as active as you can. But don't beat yourself up that you can't do more. We are all struggling for what to do.

Slide 34 I want to end by returning to the pragmatic question: can the kind of stuff you learned in our class positively contribute to the existential problems that we face? Can computer science have a net positive contribution? Can science and technology? Not in an imaginary world in which everything is different, but in our current, broken world.

We want the answer to be *yes*. There is among all but the most hardened of us a plaintive yearning to be a part of the solution, not part of the problem. But, as academics, we must be intellectually honest. And I have no reason to think that the kind of thinking that I have helped you develop is going to see us through. And it is worse than that: the reductionist, build-an-abstraction way of thinking where CS shines is entwined with the core of our problems. We may desperately want to help. But I don't know how much more of our "help" this world can endure.

Acknowledgments Thanks to Chanathip Namprempre for helpful comments, particularly on slides 22–26.