9

Programming

> Interface evolves toward transparency. The one you have to devote the least conscious effort to survives, prospers … The real-deal cyborg will be deeper and more subtle and exist increasingly at the particle level, in a humanity where unaugmented reality will eventually be a hypothetical construct, something we can only try, with great difficulty, to imagine—as we might try, today, to imagine a world without electronic media.
>
> —William Gibson

Coral reefs smell of rotting flesh as they bleach. The riot of colors—yellow, violet, cerulean—becomes ghostly white as their flesh turns translucent and falls off, leaving the skeletons underneath fuzzy with cobweb-like algae.

Corals exist in relationship. The coral animals live in symbiosis with a type of algae. During the day, the algae photosynthesize. During the night, the corals open their mouths and catch passing food. Just one degree Celsius of ocean warming can break down this coral–algae relationship, for when it is placed under stress, the algae leaves. Without the algae, the corals "bleach." After repeated or prolonged bleaches, corals starve and become diseased. Eventually, they reach an unrecoverable dead state.

Australia's Great Barrier Reef, actually a 2,300-kilometer system made up of nearly 3,000 separate reefs, has suffered severe bleaching in the past few years. Daniel Harrison, an Australian oceanographer, is looking at what might be done to buy more time for the

Great Barrier Reef. "When the reefs got bleached back to back so badly, two years in a row, we kind of just formed a little informal working group, and were like, 'You know, if we can send people to the moon, and to Mars maybe soon, then surely we can stop the reef from bleaching, you know? From just overheating.'" The situation is getting dire. "There might be as little as 25 percent of coral cover left from pre-anthropogenic times," he tells me. "We don't really know, because nobody started surveying before 1985 ... It's incredible, isn't it? I mean, you've got less than 1 percent of the ocean in coral reefs, and 25 percent all marine life. And it's not just the Great Barrier Reef that's in dire trouble, obviously. You know, we're looking at losing all of that really quite quickly, in evolutionary terms. Quite quickly, in human lifetime terms."

The Australian working group formed teams to look at different ideas that could help the reef stay alive. Their investigations showed that most of their exploratory, out-of-the-box ideas wouldn't scale too well. For example, they wondered: Since the ocean is full of cooler water at deeper depths, could we just pump up some of that in order to cool of the reef? Harrison explains: "You might be able to do something to protect small areas of the reef. Or to maybe protect important coral larvae source reef, and that sort of thing. But none of the other ideas—I mean, it's just infeasible to move up enough cool water to kind of cool the whole reef." After considering different options, the researchers honed in on the idea of marine cloud brightening—a form of solar geoengineering—as something worthy of further study. Brighter, more reflective clouds could cool the area. If small salt particles were sprayed into the air, tiny water droplets could condense around them, and these micro-droplets would make the clouds brighter. Harrison is doing modeling research to better understand the feasibility of this idea. The first stage of the modeling his team has done indicates that it might be possible to cool the water by 0.5° to 1°C.

Another research effort, the Marine Cloud Brightening Project, thinks that this could be a scalable approach with some promise for reefs. I talked with Kelly Wanser, the executive director of the nonprofit organization SilverLining and senior advisor to the MCBP, which is led by atmospheric scientist Robert Wood and colleagues at the University of Washington. Wanser describes even more ways scientists are thinking of sustaining corals: they could be genetically modified, or otherwise bred to withstand warmer waters. Robust corals could also be moved into new areas and replanted. But doing these on ecosystem-wide scales would entail a tremendous undertaking just to restore tiny parts of the whole. "The Great Barrier Reef, that's like reinforcing the Rocky Mountains. It's massive."

"Essentially, it's heat stress that's killing corals," explains Wanser. "They're affected by other stressors. Heat compounds the other stressors. Heat makes acidity worse. It's a compound stressor. Heat is the mother of all stressors on corals, and it's just like there's a certain delta, and they start to go." She recalls a recent Ocean Studies Board meeting, where a scientist offered a grim prognosis: there's maybe twenty years left for 95 percent of the world's corals. Year by year, the maps went red. "If I didn't work on solar geoengineering, I would have had to leave, because it was very emotional because to see that. It's stunning."

What would brightening marine clouds actually look like? Essentially, it consists of engineering devices to spray seawater. "There's certainly some technical challenges to be overcome, but the basic process of just taking sea water and filtering it and then spraying it out, at submicron size, is not that difficult a technical challenge," Harrison says. His modeling results suggest that there would probably need to be some stations far offshore, beyond the edge of the continental shelf, which would require floating platforms or ships. This gets pricey. The maintenance costs would probably be the larger part, and the whole project could cost around $300 million. Expensive, but then again, the reef brings in an estimated $6 billion to the Australian economy. In Harrison's conception, you wouldn't want to brighten the clouds all the time, or even every summer. Rather, it would be done when the coral was at risk of bleaching, which would require about two weeks of forewarning in order to cool the water down to the maximum extent.

"But, I mean, there's some real unknowns here, right?" Harrison says. "Because no one's ever done any field work on this. So, it's quite unknown. You know, there's a general belief here that you can only target low-lying marine stratocumulus clouds, if they're already occurring. But then there's also quite a large body of evidence that shows that production on the reef influences the local climate, and influences cloudiness in an actual analog to what we're sort of thinking about." Essentially, coral can produce a chemical that makes clouds form—though this research on how reefs modify their own climate is in its early stages. His research team is interested, and worried, about how that will change as the reef bleaches, and whether there might be a positive feedback loop, wherein less clouds mean even more bleaching. "So, to some degree, we might be putting the system back towards where it was with aerosol production on the reef. But we really don't know, so I don't want to over emphasize that. And it's probably impossible for us to know, because we started monitoring the reef too late in human history."

Indeed, marine cloud brightening comes with a lot of unknowns—in part because cloud–aerosol interactions are not well understood in climate science, more generally. For a big-picture look at what climate models can and can't tell us, I spoke with Ben Kravitz, an atmospheric science professor at Indiana University who coordinates a project that compares geoengineering model simulations. He explains: "The climate system is inordinately complex. It's one of the more complicated systems that we know how to deal with. A great example of this is clouds. If you look out the window on an airplane, you can see clouds with all sorts of different structures. They're moving, some of them are a couple meters across, some of them are tens of kilometers across. Some of them are organized, some of them are not. Basically, you can't model all of that behavior in any single model, because we don't have the computational power. If there were a way to understand how clouds behave, in such a way that we could parameterize those behaviors and put them in models that we could actually run, that would solve some of the largest uncertainties in climate science."

Newer climate models are better at controlling for clouds' varying sensitivities to aerosols, so perhaps there will soon be better tools to get information about the effectiveness of marine cloud brightening.

But for that to happen, there has to be funding. Kelly Wanser of the Marine Cloud Brightening Project says that applied cloud-brightening research could actually help us understand some of these basic unknowns. However, potential funding organizations may see controlled outdoor field experiments into cloud–aerosol interactions as geoengineering related, because they have geoengineering applications. For the US-based project, their next step is to actually test the nozzle with seawater, which they would like to do on the California coast. The association with geoengineering has made it difficult to raise funding to actually build and test these nozzles. "I think we talked to all of the relevant government agencies who could support this, and essentially there's no one willing to say, 'We'll just do it as the cloud–aerosol basic science.' They're like, 'No, the cat's out of the bag, this is geoengineering. We would have to get approval.'" So, on one hand, there's a potential technique that could have global applications, as well as regional or local ones for particular marine ecosystems—but we don't know how well it would work, or what it would take to do it. On the other hand, it's been difficult to fund the research needed to get those answers because of the stigma of geoengineering.

Another aspect of marine cloud brightening that lies at the edges of scientific understanding is the teleconnections in the system—for example, how clouds in one place are connected to weather in another place. Anthony Jones, a climate modeler, has simulated regional solar geoengineering using stratospheric aerosols. His work has examined what happens when only certain parts of the system are modified. He tells me, "I think it scares me, the thought of doing marine cloud brightening." Because of all these weird teleconnections that we don't understand, I ask? "Yeah, the teleconnections. I've been looking at that a bit in some of our [stratospheric aerosol] simulations recently. So if you cool the North Pacific, you can actually shift the position of the jet stream ... You get cold temperature on the western half of America, and warmer temperature on the eastern half of America," Jones explains. "The teleconnections are almost unavoidable, and if you can cool a certain area significantly, you are going to change the climate and the weather response." For this reason, Daniel Harrison thinks any attempt to use marine cloud brightening on a global scale would bring up major questions around governance: "If you want to cool the whole planet by doing marine cloud brightening, you know, some places are going to cool more than others. You're certainly going to alter, to some degree, global weather patterns. Maybe not that much, in the scheme of things, but it might not have to be very much to disadvantage some group of people living in some certain place, while advantaging everybody on the average." On the other hand, the concern about shifting weather patterns in remote places is less severe when it comes to brightening marine clouds over a reef, versus trying to modify temperatures globally with this technique. Those are two different goals. It would be better to consider something like area-specific marine cloud brightening for reefs to be a form of radical adaptation, rather than geoengineering.

So who cares about coral reefs, besides enamored children watching movies about clownfish? Coral reefs are not just a backdrop for colorful fish and exotic species. Reefs protect coasts from storms; without them, waves reaching some Pacific islands would be twice as tall. Over 500 million people depend on reef ecosystems for food and livelihoods.[1] Therefore, keeping these ecosystems functioning is a climate justice issue. Again, over 99 percent of corals would be wiped out at a two-degree-Celsius temperature rise, and perhaps 70 to 90 percent would be lost at 1.5 degrees.[2] And even if temperatures eventually stabilize at 1.5 degrees of warming a century or two from now, it's not known how well coral reef ecosystems would survive a temporary overshoot to higher temperatures.

Are we basically agreeing to give up on coral, and all the other animals and plants and unique forms of life in reefs, and the human economies and communities they support? On a societal level, it seems so. Many coral scientists, however, aren't willing to give up, though they oscillate between hope and despair, as ethnographer Irus Braverman finds in her research. In her book *Coral Whisperers*, Braverman describes how this polarity maps on to the rift between those conservation scientists who believe that it is possible to use traditional conservation methods (like withdrawal of human impacts), and others who take a more interventionist approach (on the basis that the natural systems are already fundamentally altered). Interestingly, she notes that "female scientists, many of them young and with diverse backgrounds, have taken the lead in promoting narratives of hope and models for assisted evolution."[3] But the restoration efforts in which many coral scientists are engaged stop short of intervention in the climate; some scientists in Braverman's book describe these efforts as holding together a patchy safety net, or as reef gardening—according tiny spaces of management that would hopefully survive an overshoot, like an outdoor aquarium in which to keep them until global warming is managed.

Is the survival of these life-forms and lifeways important enough to warrant research and discussion of geoengineering, or to justify coming up with a different conceptual category and language around geoengineering that would include more targeted interventions? It seems not. Nonhuman life is relatively absent from the anthropocentric geoengineering discussion, even though, as the saying goes, extinction is forever.

"The corals are a little bit like the canary in the gold mine," Harrison says. "They're very, very temperature sensitive. I really do think it's just a harbinger of things to come. You know, the coral ecosystem might collapse first, but I think there might be quite a few more ecosystems that'll follow it. I think that life is very resilient, but ecosystems as we know them aren't." Other ecosystems are also at high risk from even small changes in global mean temperature: Arctic ecosystems, mountain glaciers, and the Redwood forests in California, for instance. So are species that can't move quickly and find another suitable ecosystem. "It's the things that already live at the kind of extreme ends of the scale, and that can't move, right? So coral reefs, you know, they're stuck in already some of the warmest waters. If it gets too hot for them there, then (a) they can't move, and (b), they've got nowhere to go anyway. And the same with the extremely cold ecosystems. And the same with the Redwood forests. I guess. Trees can't up and move quickly enough to keep up with climate change."

Once you delve into the temporalities of the climate change problem—and especially the permanence of some of these changes, like extinction—it's easy to see how the idea of solar geoengineering makes its entrance. You're not reading this book in 1990, when carbon dioxide concentrations were still in the neighborhood of 350 parts per million. At this point, most people would agree that there's at least a chance we don't decarbonize before we lock in dangerous change—and for sensitive species and ecosystems like coral, the danger threshold has already been passed. The question, then, emerges: Can one use solar geoengineering to keep ecosystems on life support and forestall climate tipping points, while also decarbonizing?

Enter the "peak shaving" scenario, which uses solar geoengineering to "shave the peak" off of warming while carbon dioxide levels are being brought down. While we've discussed marine cloud brightening above, "solar geoengineering" in this context usually implies stratospheric aerosol injection, which would be a global-scale program.

In a nutshell, the most basic version of this peak-shaving scenario (depicted in Figure 2) means using specially designed high-altitude aircraft—perhaps a small airline's worth—to constantly fly aerosol precursors into the stratosphere. These aerosol precursors would cause the formation of particles made from sulfur, calcite, or some yet-to-be-determined substance. Why, if the world is trying to reduce particulate air pollution, would we put more particles up there? The particles are injected into the stratosphere—a layer of the atmosphere above where clouds form, and higher than planes usually fly. This means they would not fall back to earth in only a few days, as pollution from trucks and factories tends to. Rather, they would circulate around the whole planet, staying aloft for a year or so. Nevertheless, such an undertaking would not be without human health impacts. One study estimated an additional 26,000 deaths per year with enough sulfur-based geoengineering to offset one degree Celsius of warming, due to air quality and ultraviolet exposure; for comparison, 4 million people currently die each year from degraded air quality.[4] Indeed, the idea is to create a blanket

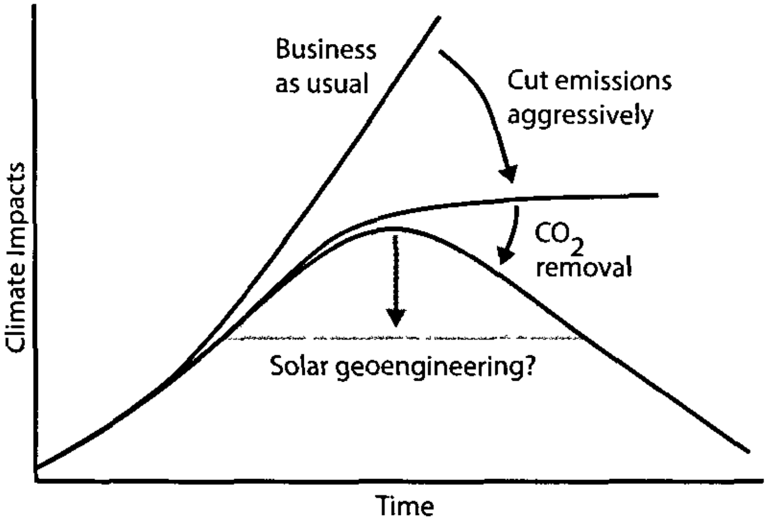

Figure 2. *Conceptual diagram of using solar geoengineering to "shave the peak" off a temperature overshoot. Sometimes called the "napkin diagram," based on a presentation by John Shepherd at Asilomar in 2010. Version source: Doug MacMartin*

of intentional, high-altitude pollution that would reflect something like 1 to 2 percent of incoming sunlight, perhaps less. Depending on one's perspective, this may sound mundane, or it may sound like an alarmingly unprecedented intervention into a poorly understood system.

We've come full circle, then, since the beginning of this book. Is "buying time" for carbon removal a legitimate reason for doing a limited amount of solar geoengineering? Or is it a weak justification for a project that will send the earth system careening down a dangerous road? Once started, how would people make sure the carbon really does get removed? In this chapter, we will look at some best-case and worst-case scenarios for a climate intervention program that includes solar geoengineering.

Kingston, Jamaica, July, 92°F / 33°C

Rain clouds hang in the Blue Mountains but never descend. The sidewalks in Kingston are jammed with people in business suits walking to work and street vendors selling sweets, making their way through fumes and honking and reggae blaring from car stereos. It is a glorious morning. I catch a ride up to the university, where there's a meeting about solar geoengineering research governance—the first such convening here in Jamaica.

The country has recently suffered from a drought, and while people here are accustomed to dry periods, yearlong droughts are something new. "Unfamiliarity is transforming our sensitivity into a vulnerability," explains the first speaker, climate scientist Michael Taylor. Jamaica has steep slopes and narrow coastal plains; with limited water storage, farms are largely dependent on rainfall. Livelihood, well-being, and water access here and elsewhere in the Caribbean are linked to the rains. But now the rain is more variable; the "nature" of rain is changing. Nighttime temperatures are soaring; one has to run the fan straight through the night. The climate will keep changing, Taylor says, with 98 percent of days here being "hot" by the 2090s—he emphasizes that only 2 percent of them will be cool.

It's sweltering in this room, and we're shifting uncomfortably on our wooden seats. There are a few fans, an adaptation to the broken AC, but even the local residents are sweating beads. Someone makes the inevitable joke about our prospects for controlling the global climate when we can't even control the climate in the room—a joke that's been made at approximately one-third of the geoengineering meetings I've ever been at. We press on.

Most people in the room are confronting the idea of geoengineering for the very first time, and their initial thoughts are diverse. A policymaker explains that he's read an article that suggested climate warming could create a whole new economic zone of mining and exploration—which would suit Northern countries, while those in the tropics would lose opportunities. Another person says that he's just installed solar panels to take his house off the grid: Will solar radiation management impact those? One speaker recalls a James

Bond movie he saw, in which someone controls an orbital laser beam. Someone else notes that manipulating the climate is like manipulating genes; it can be done in the wrong direction. An ethicist says that there's always some price to pay for engaging with technology, and we are in this problem because of technology. There's promise and peril; we are flying in the face of God.

But the conversation keeps returning to two themes. One is equity: it's the norm for Jamaica to be on the receiving end of things. Someone asks: "How can that be changed?" Jamaicans have regional alliances, they network, they negotiate as a block of Caribbean or small island developing states. But there's a disparity in terms of population when you're a small country. Discussion turns to this disparity in terms of historical emissions, and inequality also comes up in carbon trading, which can let polluters off the hook. "In an ideal world, we have iron-clad politics *before*" confronting something like this, one speaker asserts. Another asks: Can we afford *not* to look at an issue like this?

A second theme is capacity. A policymaker explains that the research system here was originally set up to train agricultural researchers for plantations, not to teach classical subjects. "Those of us in countries like Jamaica need to develop basic research," he says, because the problems people are trying to solve here may be different than those in other places. But here, they don't have the computing resources to run many computationally intensive climate models.

When it comes to designing a solar geoengineering program, both metaphorically ("program" as in a course of actions) and literally ("program" as in coding on a computer), a small developing country like Jamaica has a limited capacity to write it. The organization that pulled together the meeting, the Solar Radiation Management Governance Initiative, now coordinates a fund for researchers in developing countries, which provided an initial round of $430,000 toward eight projects that look at how solar geoengineering could impact things like droughts in Southern Africa or the spread of cholera in South Asia. This is an important step for the philanthropic sector, and its NGO and academic partners. Yet it's still a drop in the bucket compared to what would be needed for genuine inclusion in research design.

Algorithmic governance

Who *does* get to write the program for geoengineering? The verb "program" is rooted in the Latin *-graph*, connoting a written plan. Geoengineering would be a program to be authored, to be written, with real choices about what goes into the plan. In the end, though, program might not be the best metaphor, because it still has resonances of something "fixed" that you'd receive on a disk (even if these are now obsolete technology, as our software auto-updates). Solar geoengineering requires a more dynamic practice. "Responsive" or "adaptive" governance tries to connote this; but still, "responsiveness" seems like a tacked-on quality, something that modifies geoengineering after the fact, rather than being written into the fabric of what it is.

Let's follow these overlapping meanings of "programming," and also draw in another fuzzy term: "algorithm." "Algorithm," on a basic level, signifies a set of instructions; a recipe of sorts. Today, though, algorithms have taken on a broader meaning. They have become agents that determine aspects of our social reality: helping a billion-plus people get where they're going, assisting us in finding information, driving cars, manufacturing goods, assigning credit, shaping financial markets, and more, as science historian Massimo Mazzotti describes.[5] Wouldn't solar geoengineering inevitably be yet another one of these domains governed by algorithms, aided by an invisible computational hand?

We can be certain that aerosol geoengineering would be implemented in some kind of institutionalized program: with research milestones, perhaps "stage gates," flight operations, monitoring operations, and so on. Geoengineering with aerosols would also involve the use of computer tools to design a literal *program*—a recipe or operation—that would figure out the optimal way to put the particles in the stratosphere in order to achieve a combination of climate goals while minimizing negative impacts. In short, someone, somewhere, would write code for a system that would block a certain amount of sunlight, monitor the effects, and then adjust the system again. Researchers call this a "feedback control algorithm" because it would guide geoengineering using feedback from climate observations.

This is complex enough: now add the complexity of possibly using multiple climate engineering techniques. For example, a recent paper from a collaboration between scientists in China, India, and the United States simulated "cocktail geoengineering," which involved using two different geoengineering strategies—stratospheric aerosols and cirrus cloud thinning—to best restore preindustrial temperatures and precipitation.[6]

Then, add a temporal dimension: stratospheric aerosol geoengineering would likely take place over a time span of 150 years or more (at which point, if enough of our descendants make it through the twenty-first century, they will hopefully have not only found better ways of removing carbon, but also improved upon our rusty old technology for deploying and monitoring the particles).

You can see how heavy computing would be crucial to a problem this complex. The resulting climate could be seen as some kind of human-machine-nature collaboration or dialogue, with constant back-and-forth retuning. In the parlance of the scientists working on it, it's about feedback and adjustment. Humans would input the goals. These could involve changing global temperatures, reducing sea level rise, stopping Arctic sea ice loss, or some other combination of ends that would likely be subject to long negotiations. It's quite possible that a set of decision rules for solar geoengineering could be created in a quasi-democratic matter—likely by the United Nations, where technical-expert delegations from various countries would hammer out the goals and the scheme for monitoring results, and so forth. But you can also very readily see the possibility that it will *not* be done that way—just consider how countries like Jamaica experience international decision-making processes.

Looking at solar geoengineering as an algorithm allows us to draw from the emerging literature on "algorithmic governance," which questions how algorithms are used to make decisions that pattern an increasing number of aspects of our lives. One key issue is about transparency and the black-boxing of algorithms: How can a geoengineering system be designed for openness and "algorithmic accountability"—that is, explainability in real time?

There is also the danger that bias could be coded into the program. This could happen because the underlying climate data is uneven; for example, war-torn countries are going to have gaps in their data. Or bias could be introduced due to variations in how problems are defined. Droughts, for example, can be hydrological, agricultural, or meteorological, each of which would be defined using different thresholds. Theoretically, the system could disadvantage vulnerable peoples without the explicit intention to do so, just because of poor data or poor problem definitions. (Of course, this is on top of the basic bias that determines who gets the education and power to even be in the position of writing computer code and making decisions.)

Given all the advances in computing, we might take for granted the ability to carry out the programming part. However, researchers point to varying constraints on this—both in terms of computing resources and qualified personnel. This isn't just the case for small island states or developing economies. Even the US public sector is stunningly constrained when it comes to running climate models; the computing infrastructure of the commercial cloud is soaring by comparison. Scientists I spoke to in several countries pointed to high-level expertise as another limitation. The labor time from qualified humans analyzing outputs is a constraint on geoengineering research (and on how advanced these algorithms can become), probably more so than computing resources. This is, of course, a matter of both training personnel and funding their research. Labor-time of qualified scientists seems like a worldwide constraint—though in theory, it should be solvable, given that politicians everywhere pay lip service to STEM training.

One challenge is to expand that training globally. In 2017, the geoengineering research program in China held the first geoengineering research course in Beijing for scientists from the developing world.

They provided model results for the entire earth system for students to analyze, because they believed students would best understand how to select the model parameters that were most important in their country. "It should not be some sort of teacher and student thing," clarifies project leader John Moore at Beijing Normal University, speaking about international collaboration more broadly. "It should be an equal relationship." He explains, "The international collaboration, and everybody being a sort of fair and equal partner, is a big priority from the Chinese viewpoint." China has no wish to do geoengineering unilaterally, and wants to avoid being seen as eager to geoengineer, he says. I ask him why he thinks that the Chinese are so interested in international cooperation. "They care about how the country's image is ... I guess it's sort of patriotism in a way, that people have a love of their country, and they don't want to be the international bad guy. They want to be the good guy, the nice guys. Given a choice, that's the natural choice."

International cooperation and collaboration, then, are part of a best-case scenario—one that many people are actively working toward. These cooperative-minded researchers and funders can help with capacity building. Still, they can't fully address the underlying structural inequalities between disparate working contexts. All this is to say that when thinking about the algorithm, we can't forget the material resources needed to make it work—both workers and infrastructure. Beginning an international, interdisciplinary research initiative *now* would create and hold the space for researchers to think more deeply about what transparent, democratic algorithm design might look like.

Keeping humans in the loop

Over the course of the next few decades, as we weigh climate intervention, machine learning and artificial intelligence will no doubt continue to make advances. What does it mean for these two capacities to evolve together? To be clear, scientists thinking about feedback control algorithms for solar geoengineering are *not* interested in mixing in artificial intelligence. Rather, they view these as systems where humans would be very much in the loop.

Ben Kravitz, the atmospheric science professor, points out two important problems with employing machine intelligence. "Number one, you have to actually believe what you're designing is correct." Understanding the underlying physical system, he tells me, is crucial. If you don't, and "you just say, 'Well, I don't really care, let's just wrap a controller around it and be done with it,' then everything could be fine until you get some weird—'weird' is I guess a technical term in this case—some weird behavior that you can't explain but that really messes things up. And Rumsfeld's 'unknown unknowns'. That's always a concern."

"Just because you can automate something doesn't mean it's a good idea to do so," Kravitz cautions. A second problem is that a machine intelligence might see an optimal outcome differently. "That is subjective—and not. ['Optimal'] is a really important word. If you're an economist and you can reduce everything down to dollar amounts, what you call 'optimal' might be different from what say a politician calls 'optimal,' because there are various additional concerns." Kravitz recalls the Three Laws of Robotics, by science fiction writer Isaac Asimov: First, a robot may not injure a human; second, a robot must obey the orders given by humans except where such orders would conflict with the First Law; third, a robot must protect its own existence, as long as that does not conflict with the first two laws. "That's sort of why they were invented, because what a machine calls 'optimal' is not necessarily what a human would call 'optimal.'" Kravitz points to the analogy of the Federal Reserve, a controller for a very complex system: "Depending on whether you call the Federal Reserve 'optimal' ... it's basically a bunch of experts substituting for that machine, doing control theory on a poorly understood system." On the other hand, he notes, there can be problems with human decision processes, too. "Do you want the computer to do everything for you, or do you want people to be involved, even if that means reduced performance?"

To gain a better handle on this, I dropped in on control systems engineer Doug MacMartin at Cornell, who's authored numerous papers in top journals on potential designs of geoengineering systems (he has also collaborated with me on a project about how to incorporate community ideas into geoengineering research). MacMartin offers me some bread that he made, and begins to humor my questions about what it means when geoengineering and artificial intelligence grow up in the same time frame.

MacMartin, like Kravitz, emphasizes that deciding the goal of a climate intervention system is clearly a human activity: "That's a value, those are value judgments. Once you then say 'these are all the things that I care about,' you could essentially imagine that there's an algorithm that determines, given all the information that it has available to it and given the goals—here is the best thing to go do." Conceivably, some complicated deep-learning algorithm could help in this: one that has a much more advanced model of the climate system, and projects the future based on its knowledge of past climates and goals imposed by humans that tell it the performance metric. "'Here's what we care about. This matters this much, we don't want the rainfall here to deviate by more than this. We don't want this to change by more than that amount. And go find the solution in that space that is robust in all of your uncertainties about the current state of the system, and uncertainties about how the things evolve.' It sort of does the best job of balancing in this multidimensional goal space."

This is all hypothetical, of course, and I tell MacMartin that if this program existed, the outcome seems like it would be a human collaboration or negotiation with the program. I begin speculating: "It's so complex that you'd have to say, I care about Arctic ice, and I care about precipitation in this vulnerable region, and I care about XYZ—and then it would come up with something—but then that thing it comes up with would have this other thing that might cause a problem, and then you have to go—"

MacMartin asks: "You've done optimizations before, right?"

"Not really."

"This is the way they always work. Every optimization you ever do is like, 'Here's what I care about.' The computer then comes out and says, 'That's the optimum.' You look it and go, 'That's not what

I wanted.' You realize it's like, 'I only specified these variables, and I didn't specify this one over here, and it found a solution that never even occurred to me where it improves these, but destroys this thing over here.'"

I'm thinking that it sounds like a big mess—but this is what engineers deal with all the time, and a lot of our technological systems *do* actually work much of the time. "What you really want, in some sense," MacMartin explains, "is an iterative process ... And you presumably want a human in collaboration with that process who can then basically say, 'Wait a minute, that might have been what I asked for, but it's not what I wanted.'" MacMartin makes an important point—in the feedback class he teaches, they don't do anything that's optimal, because "optimal" is so tough to pin down. "You tend to do just as well by not optimizing it quite so much, if you know what I mean."

Do we get anywhere by looking at geoengineering as a program, or as software? "I think if we think about it as software, the first thing that comes to my mind is to think back to Star Wars [the Reagan-era missile defense program]. Which, as long as you thought about missile defense as a physics problem, seems solvable. The instant that you think about the missile defense system as a giant piece of software that happens to interface with physics, then you just laugh at it and say there's no way we could ever make this work."

MacMartin, like Kravitz, thinks that letting humans stray too far out of the loop would be risky: "I would say the biggest risk would be engineers being over confident in the ability of computer algorithms and allowing the computer too much leeway to make decisions. I don't personally buy into the idea that the computer eventually becomes sentient to prevent you from allowing you to turn it off. I think you can always turn it off." So, there will be no malicious artificial general intelligence, in MacMartin's view. But he identifies two other risks. The first is that something unexpected happens that is outside the training data that you've used to train a machine learning algorithm. The second is the risk that people become overreliant on the infrastructure and fail to understand the interconnected parts of the system.

In terms of social effects, like unemployment, the risks of these technologies developing together are probably indirect, MacMartin judges: "I think the bigger issue with them maturing at the same time is probably far more, on some level, related to the Trump factor—on steroids ... Just as we look now and we think, 'Wow, George W. Bush. I wish we still had George W. Bush.' It wouldn't surprise me if in thirty years we say, 'I wish we had Trump.' Because if half the country is unemployed and unemployable forever, and there is no foreseeable pathway to get somebody who's forty years old to be employed in any meaningful way, that could have some pretty serious social repercussions—and combining that with something as powerfully upsetting for the human relationship with the universe as taking responsibility for the entire climate, and as inherently globalizing ..." When people are struggling to find employment, he suggests, one reaction is to elect someone like Trump. "That nationalistic tendency is kind of at odds with the global implications of doing geoengineering." He pauses. "I suspect that issues have far less to do with narrowly how AI is being used in conjunction with geoengineering, but broadly in terms of how both AI and geoengineering affect the human relationship with the rest of the world in antagonistic directions. That could lead to really serious problems."

Designing the program

What is the best-case scenario for geoengineering? Scientists often laugh when I ask them that. "I would say the best-case scenario is if we figure out a way not to do it," replied Ben Kravitz, which is actually most people's first answer. "Barring that, I would say the best-case scenario is that we do it in an intelligent way, where we are designing geoengineering so that it will do what the politicians want it to do, minimizing side effects, and that there is an appropriate government structure ... and just people working in a way that I believe they can work. Not the way they usually do." He adds, "We could go out and do geoengineering tomorrow, really poorly. That's what scares me."

Doug MacMartin has a similar view, sketching out a hypothetical scenario where the world manages mitigation that curbs warming to two and a half degrees, and uses solar geoengineering for a century to bring temperatures down to one and a half degrees. "We do that in a way where adjusting aerosol injection at different latitudes balances all sorts of different climate impacts. So that in principle, almost nobody on the planet is actually harmed [by geoengineering itself], and that there is strong international trust in whatever organizations are involved in making decisions, so that everybody on the planet feels that their voice has been heard, and accepts the fact that this limited deployment of solar geoengineering is better in some aggregate sense than not having had any geoengineering, and so it doesn't result in conflict."

"Geoengineers" are often caricatured as being in bed with fossil fuel companies, stranding us in the business-as-usual era; or otherwise, as gritty realists who lack the imagination to see the social transformation that's truly possible. Yet in the best-case articulation of many researchers, solar geoengineering *does* represent a rather-utopian dream. You have to believe that people are really capable of long-term thinking and cooperation to even articulate these scenarios—much less spend your time researching them. MacMartin adds that a best-case scenario, socially speaking, would frame geoengineering as a conscious acceptance of responsibility for the climate, rather than simply an effort to control it. "Accepting responsibility for it, I think, speaks to a maturation, and a way of expanding one's moral sphere further. Expanding one's moral sphere to the rest of the planet to future generations, to nonhumans, and saying we actually have responsibility for the betterment of all them. It's certainly possible that a hard discussion about solar geoengineering could push humanity in that direction."

While climate engineering researchers don't tend to interpret this question of a "best-case" climate scenario through the lens of a comprehensive, long-term solar geoengineering program, they do have a *sense* about what a climate engineering program might look like—contingent on what society is doing with mitigation and carbon removal. Ideally, solar geoengineering would be limited in scope. And it would be limited in time.

"Would a marginal deployment of geoengineering harm anybody?" Peter Irvine, an atmospheric scientist based at Harvard University, is trying to help answer this empirical research question. He explains that modelers often look at a scenario where all warming is offset, and evaluate who wins and loses in this extreme scenario. "Would offsetting 0.1 Celsius with geoengineering dial anything back, or would it amplify certain things? Does a marginal deployment help or hinder? ... How many more increments of cooling can you add before you start running into new issues that come with geoengineering?" Irvine and colleagues' work looks at what happens when just enough stratospheric aerosols are used to offset half of a doubling of CO_2 in the atmosphere for a hundred years. They find that using this smaller amount could avoid many of the previously reported impacts of aerosol geoengineering upon the hydrological cycle, such as extreme or decreased precipitation. With halved warming, everything seems to scale up without major hazards. "Beyond half ... some of the real differences start popping out a bit more." I asked Irvine about his best-case scenario for using solar geoengineering, and like most scientists, he says the best case is strong emissions cuts. "Against that backdrop ... I think some deployment carefully, carefully scaled up, little by little, over the course of a decade or two. ... Gradually halving the warming, possibly halting the rate of warming some decades in, against the backdrop of trying to cut emissions and then bringing them back down."

Even though Irvine is a climate scientist who thinks about the temporal aspects of a solar geoengineering system, he doesn't necessarily take a long-term, programmatic view. "You've got to think kind of decade by decade. Like, what do we do now, in this decade?" Irvine asks: "Who are we to say what the 2100 climate policy should be?" He thinks it is possible to have a climate policy that includes solar geoengineering but not negative emissions, where it would be decided decade by decade whether to ramp back the solar geoengineering or continue it, and points out that the same questions apply to making policy around negative emissions. "How quickly should we get to zero, and how quickly should we go negative? I think these are questions that, it's a bit daft to impose what we think... to basically, to meet some arbitrary 2°C target or 1.5°C target people in 100 countries agreed on in one meeting somewhere, and to assume that it's going to bind people 100 years from now."

Many scientists see solar geoengineering as a temporary measure that could be phased out if carbon removal was succeeding. Oliver Morton, in *The Planet Remade*, calls this temporary geoengineering scenario the "breathing-space approach," in that it allows incremental use of solar geoengineering to create breathing space for decarbonization.[7]

This intuition that solar geoengineering is a temporary intervention develops, I think, not only from an engineering perspective that thinks about resilient systems, but also from a moral sensibility. Kravitz says that solar geoengineering is not a permanent solution: "It's not something that we want to just do forever, or at least I don't, because I think too much can go wrong." Similarly, MacMartin thinks we'd eventually want to restore the climate using carbon removal, in order to have an exit strategy and avoid inflicting an implied commitment to solar geoengineering on future generations. Most stratospheric aerosol scenarios last 200 years, he says, and there's probably no deployment scenario that's less than a hundred years (with the caveat that it's possible that someone will find a really cheap way to pull CO_2 out of the atmosphere during that time period). Kelly McCusker, a climate modeler who has studied what happens when a solar geoengineering program is ended abruptly, also emphasizes that solar geoengineering will need to be done in combination with carbon dioxide removal or mitigation. "That's my feeling, but I don't necessarily have a good basis for that. I mean, in part, it's like you can look at these plots and just know it would be completely immoral to do it on its own."

The scientific consensus around this feeling is reflected in the IPCC's Special Report on 1.5°C, which explicitly assesses solar geoengineering in terms of its potential to limit warming to this amount "in temporary overshoot scenarios as a way to reduce elevated temperatures and associated impacts." It states that "if considered, SRM [solar radiation management] would only be deployed as a supplement measure to large-scale carbon dioxide removal."[8] This report actually reflects quite closely the views MacMartin, Kravitz, and others have already published, and their sense of how the practice might be best used, given that the aim of the report is to summarize the existing literature—and these researchers are the ones who wrote the existing literature. At the same time, this consensus on how it should be used might not be shared by politicians or industry leaders, who are often thinking on much-shorter timescales, and with different considerations in mind.

Why, exactly, is it so important that geoengineering be temporary? One issue is that of "termination shock," mentioned at the beginning of this book—the phenomenon in which if solar geoengineering was suddenly ceased, temperatures would shoot back up to a level commensurate with the greenhouse gas concentration of the atmosphere.

Some scientists argue that the risk of termination shock is not as high as portrayed in the media or academic literature. For a termination shock to happen, the geoengineering intervention would have to be large: as Oliver Morton writes, if solar geoengineering were "a relatively modest affair, the termination shock would be more a termination shudder."[9] Likewise, Andy Parker and Pete Irvine argue that a sudden termination is unlikely, and preventable.[10] Because solar geoengineering is cheap, they argue, it would take a pretty massive catastrophe to halt it and keep it turned off. For instance, 70 percent of GDP of the United States or China could be wiped out, and they could still deploy solar geoengineering for less than 1 percent of their post-catastrophe GDP. When it comes to such external forces, Irvine says: "If we're talking large-scale nuclear war, everything else that we're technologically dependent on is going to kill us first. I mean, the fact that I don't know how to grow food or catch food or hardly even look after myself ... I think that's what's going to do us in, rather than a slightly greater warming off the back of your mushroom cloud–induced nuclear winter." It's true that people could *choose* to stop geoengineering, but Irvine finds this similarly unlikely, because for termination shock to be a big issue, you're already have to be decades into the program. "You're a generation into this. It's normal. It's as normal as, you know, irrigation and rivers. It's everyday ... If you start to run through this story line, and put yourself in the perspective of a world that's thirty, forty years in ... I think it would be as unthinkable as international shipping or aviation ending." Parker and Irvine also add that it would take months for a disruption to a solar geoengineering program to have any effect upon temperatures, because the aerosols would take months to thin out. "This is crucial for analysing the risks of termination shock, as it means that humanity would have a period of several months in which to resume deployment of SRM in the event of a disruption," they write. They point out that even if solar geoengineering was used to offset a large amount of warming, it could be slowly phased out over the course of decades without a shock. They also suggest criteria for a solar geoengineering system that would be robust against termination shock: it would need to be geographically distributed, affordable enough for multiple actors to maintain independent systems or backup hardware, and slow to lead to damages following disruption. "If back-up deployment hardware were maintained and if solar geoengineering were implemented by agreement among just a few powerful countries, then the system should be resilient against all but the most extreme catastrophes," they write.[11]

Other researchers, however, view the risk a bit differently. In one study, ecologist Christopher Trisos and colleagues modeled the consequences of solar geoengineering termination on other species.[12] Their study looked at the "climate velocity" of different species, which is the rate at which plants and animals have to move to keep their climate the same. Basically, if solar geoengineering was started and then stopped, many species would not be able to move quickly enough to keep up, threatening extinction for corals, mangroves, amphibians and land mammals. I caught up with Trisos one spring day in Washington, DC, to hear more about his research. "To the extent that risk is probability times consequence," he explained, "even if that probability is tiny, if the consequence is worse than anything we could conceive of over a similar time period of climate change without geoengineering; then for me, it's big enough of a risk to really think twice about the geoengineering discussion. Essentially, if we are exposing ourselves to that level of planetary risk, should we even think about geoengineering anymore in the first place?" Trisos thinks that the extreme severity of a termination shock, even if unlikely, raises the bar for geoengineering researchers to show that termination "is not a risk, or at least a low, low probability."

Atmospheric engineering on an ocean planet

There's another key reason that solar geoengineering would need to be nested within carbon removal: we live on an ocean planet. Ocean acidification has been called the "other CO_2 problem," and solar geoengineering wouldn't directly help ocean acidification much. (Solar geoengineering would likely affect ocean acidification through a host of other secondary effects—temperature effects on terrestrial biomass, hydrological cycle changes, changes in marine productivity[13] —but it doesn't directly address the issue of increased CO_2.)

Solar geoengineering may help mitigate sea level rise, though there are some limitations. Sea level rise has two drivers: ocean heating (because warmer waters expand), and melting ice. To stop the former, we'd have to address the energy imbalance in the ocean, which would mean returning it to roughly preindustrial conditions. Imagine that temperatures rose, but then we did geoengineering to lower temperatures back down half a degree, suggests Pete Irvine: this would reduce the direct atmospheric-driven, temperature-driven effects, such as heat waves. However, the ocean has layers. The surface ocean can respond to solar geoengineering in decades, but the deep ocean takes hundreds or thousands of years to respond. For example, if we reached two degrees Celsius and did geoengineering to come back to 1.5, the *rate* of heating would slow, but the ocean would not necessarily stop heating. "It's going to take thousands of years to reach a new balance."

By bringing temperatures down, one could stop ice from melting in places like Greenland, Irvine notes, which is otherwise going to experience a slow runaway feedback loop in which it loses mass over thousands of years. But in parts of Antarctica, on the other hand, staving off melting might not be possible, since Antarctica has points that become unstable with only a little bit of warming. "Because of the way the glaciers meet the ocean, when they start to retreat, they have kind of a runaway retreat. Again, very slow, like a couple of centuries. Five centuries. But once it starts, it's not a temperature-driven thing; it's a dynamic-driven thing ... Once the ice shelf is sheared off or melted away, it's not there to hold the ice sheet back and there's this kind of dynamic response." Irvine explains that some ice may have crossed that tipping point already: "Quite big chunks of West Antarctica may have already crossed this threshold, and it might not be possible to dial them back. But there's other parts that may trigger a little later. Or may trigger much later ..." It's not like everything is safe below 1.5 Celsius, and then there's a sudden step change when the world reaches 1.5. "I think there's lots of little mini steps. And the sooner you arrest warming, the fewer of these steps you cross." Some things are not reversible by solar geoengineering. Irvine continues: "You could imagine quite easily if that migrant species' last refuge gets pushed off the top of a mountain because it's getting too warm, well, it can't come back because it's dead. So I think there are examples where you could say, the atmosphere can respond back and we can dial it back. We can dial it back five decades later and get it back to where it was." But some aspects of the ice sheet response can't be dialed back. "And obviously if certain harms and damages have occurred, things have died, glaciers have melted away permanently ... You know, there's certain things you can't take back."

Best-case solar geoengineering: A temporary measure for conservation?

One of the best reasons to consider solar geoengineering, in my view, could be to preserve species during an overshoot—but would this even be possible? I asked Christopher Trisos, the ecologist, about this. He's hesistant to say that geoengineering could save species, he says, because while some ecological impacts could be forecast, they really depend on the social choices that are made around geoengineering, and Trisos's impression is that "the scenario's uncertainty is still so big that ecologically, it is just anyone's game." There are areas that could be investigated, though, such as biome changes or wildfires, as well as disease vectors like mosquitoes, which we know a fair bit about. "What happens to the insects that are vectors for a lot of really nasty diseases like malaria, chikungunya, Zika? Is a geoengineered world better or worse for Zika in the Americas?" Or would it be worse for cholera in Asia? I wonder. It could be the opposite; maybe the conditions it sets in place climatologically are more preventative than promotive. We don't know. "A geoengineered world with potentially more Zika, more cholera, more malaria ... I'd much rather take climate change without geoengineering."

There's so much that's unknown, but to me, it's noteworthy that not very much research has been done on the ecologies of geoengineered worlds. I ask Trisos to speculate on this gap. He hypothesizes that ecologists don't want to lend the idea credibility. "The 30,000-foot view of what geoengineering could mean for ecosystems—my sense is a lot of people are reluctant to do that, or just not interested, because they view it more as the sci-fi fringe with the crazy people at climate conferences, and the real work of ecologists is to try and show how bad climate change can be. To promote greenhouse gas emission reductions. And also, I think in a normative way, geoengineering goes against what a lot of ecologists hold dear about promoting resilience of the planet, and natural recovery of ecosystems, and giving things space and time to have an adaptive capacity. The idea of putting your thumb on the thermostat of the planet is antithetical to that." The study of speculative futures under geoengineering is on the cutting edge of climate change ecology research, Trisos adds, which also makes it a tough sell for ecologists. "If you're going to push that research frontier forward, they would rather focus on a conventional climate change scenario than look at the more unconventional, potentially fringe ones like geoengineering."

Of course, actual conditions in a world with higher CO_2 and lower temperatures are unknown, and no one would have any reason to model it, unless they're thinking about geoengineering. For example, Trisos points to the possibility of a high-CO_2 world where woody vegetation encroaches on grasslands: Could the Serengeti be turned into a shrubby landscape that thickens into a forest under geoengineering? "In a lot of Southern Africa, where I'm from, savannas are really beautiful landscapes. They're important for ecotourism. They're important for a lot of livelihoods around grazing and livestock. They have a lot of endemic species. They're ancient grasslands; fire has been a large part of the maintenance of those grasslands for hundreds of thousands of years. And if we turn down the temperature and have increased atmospheric CO_2, the extent to which that increases the invasion of these grassland and savanna areas by woody shrubs, I think, is important to know. It's potentially really concerning if you lose large swaths of grassland as a result of geoengineering. You could get bald patches. I mean, we haven't run the models yet. We don't know."

I understand the reluctance of ecologists to spend their time on the speculative imaginary of solar geoengineering when there's so much other essential research to do. Yet I worry about the idea of solar geoengineering moving into political discourse without these ecological aspects being investigated—for it is not only risks that could be ignored, but potential uses, too.

For Kelly Wanser of the Marine Cloud Brightening Project, the purpose of solar geoengineering is to keep ecosystems stable and intact, so that they don't start to break down in an unrecoverable way while the world gets its act together on climate change. I ask her about the best-case scenario for this. "A good scenario would be you use just the amount of solar geoengineering you need to keep all the coral reefs from disappearing," she answers. "To keep the ice sheets stable. To keep the methane trapped in the ground and things like that, for as long a period as you need in order to safely keep those systems stable, while you're bringing down the greenhouse gas concentration and you determine that you have a safe situation to proceed … The best-case scenario is, nothing really bad happens.

You use enough solar geoengineering to prevent the really devastating changes in the system, and you take as many measures as you can as quickly as you can to restore the underlying balance of forces in the atmosphere." She returns to the medical metaphor: "If you do it sooner and you keep heat levels from rising higher, you need less, and it would therefore be safer. Like any intervention. If we let temperatures continue to rise, then the counterforcing we're going to need is going to be pretty strong, and at that point you may also have other changes in the system going on."

Indeed, waiting a long time to do solar geoengineering, and then suddenly deciding to do it in the wake of an "emergency" like a massive drought, would not be the best-case program—because during that waiting period, the world is breezing past some irreversible tipping points, like species being senselessly lost forever. On the other hand, it is difficult to understate just how preliminary the idea of solar geoengineering is—basic crucial questions around how ecosystems may respond have scarcely been explored.