

Constructing a New Knowledge Infrastructure

An environmental knowledge commons could support evidence-based policymaking, but it will require long-term coordination across the many communities monitoring pollution and local conditions.

In the northern part of the San Francisco Bay Area, four communities are simultaneously navigating the past and future of energy. At oil refineries in Richmond and Martinez, it's business as usual. Recently, Chevron Richmond and PBF Energy-owned Martinez Refining Company (MRC) have battled state regulators over stricter emissions rules. MRC has had a string of notable accidents and toxic dust releases that have worried and angered nearby residents, who have watchdogged the refineries for decades.

In neighboring Benicia, on the other hand, Valero Energy is in the process of closing a third oil refinery. Residents of the small city worry about the future of the site and the local economy. Meanwhile, since 2020 two former oil refineries—Phillips 66 in Rodeo and a second refinery in Martinez—have converted to renewable fuels facilities, trading crude oil for grease and cooking oil to produce renewable diesel and other fuels. Their operations still involve hazardous materials and the potential for accidents. On top of these facilities, a carbon pipeline that would cut through the area has been proposed. Neighboring communities oppose the pipeline, arguing that it endangers residents and perpetuates use of fossil fuels.

In this context, maintaining healthy communities requires information. To respond effectively to accidents and releases, residents and emergency responders need real-time measurements of hazardous chemicals in the air. Residents who observe soot or dust blanketing their communities

need information about its chemical composition to share with their health providers. Closing the refinery in Benicia requires still more information, so the city can understand the levels of toxins left in the soil and the risks of further exposures from clean-up processes. Faced with the choice between the razing of an oil refinery and its conversion to a renewables facility, communities should be able to compare the status quo with expected emissions and safety risks for multiple future scenarios.

Creating the kind of knowledge base necessary for such consequential decisions would require long-term coordination across the many communities affected by energy infrastructure. Places like Benicia, Martinez, and Rodeo would need a place to store data about pollution before, during, and after major changes at nearby energy facilities. They would need to have a way of sharing their data and analyses with other similarly situated communities if they chose to do so, and they would need to be able to access data and analyses from other communities just as easily. Academic and nonprofit researchers with a bird's eye view of the issues could also enhance knowledge infrastructures if they had access to data shared by communities and a way not only to disseminate their findings, but to share their methodologies for communities to adapt and deploy.

Existing data infrastructures can't support this kind of collective learning about environmental issues. Both the technical and governance aspects of the infrastructure

would need significant upgrades, and the customary models for funding science in the United States don't offer the kinds of investment that would be necessary. Funding is typically structured around short grant cycles and discrete deliverables, making it difficult to support the long-term, shared stewardship that this infrastructure requires. Addressing these hurdles could enable creation of a robust environmental knowledge commons maintained by a plethora of users and contributors. Such a commons could ensure the continued capacity to generate new insights about the impacts of pollution and environmental change, forming a durable basis for evidence-informed public policy, whether or not the federal government continues to support environmental science. An environmental knowledge commons could, moreover, offer a model for ongoing advancement in other fields of science where traditional funding models have become precarious, even as their knowledge remains essential to public well-being.

Overcoming a disjointed information landscape

Thirty years ago, the data necessary to conduct broad analyses of environmental and community health did not exist. But sustained activism by local communities over the past three decades has made the Bay Area an information-rich environment. For example, state law now requires oil refineries to monitor ambient air quality at their fencelines in real time. The Bay Area Air District has taken steps to ensure industry compliance and has set up additional monitoring programs, including one to understand particle fallout. Communities themselves remain active in producing data through a community-run monitoring station in Benicia, a community oversight group for Phillips 66's fenceline monitoring system, and an extensive network of particulate monitors.

Yet even with this wealth of information, communities, regulators, and researchers have limited ability to combine data streams to create a richer picture of what's going on. Environmental monitoring results are scattered across platforms; they are owned by myriad organizations, most of which do not publish licensing information; and they remain cut off from tools that allow residents to report their embodied experiences of pollution as well as from government-generated health statistics.

This fragmentation undercuts Bay Area communities' efforts to make the most of their own data, and it inhibits the ability of communities in similarly industrialized areas to learn from one another. Refinery-adjacent communities in Louisiana and Texas, where little local-level data is available, might be able to refer to analogous data from California refinery fencelines. Conversely, Benicia residents might look to Philadelphia, where a refinery closed in 2019 and has since been largely dismantled, to better understand the kinds of issues that require active community

advocacy. Environmental research, too, is thwarted by the fragmentation and inaccessibility of data: Academic studies that could result in findings relevant to fenceline communities nationwide aren't possible if researchers can't assemble trans-community datasets.

The solution to this disjointed information landscape is easy to envision but hard to achieve. What is needed is a knowledge commons, a shared space where communities and their allies can store, steward, and access multiple types of environmental data, enabling them to conceptualize common problems, tailor new frameworks for interpreting data, and design strategies to collect actionable information. Such a commons should not be owned by anyone, nor should it rely on proprietary technology. Instead, an environmental knowledge commons should include a system of collaborative stewardship: users and contributors participating in establishing norms and policies, seeing that the space evolves to meet emerging needs, and ensuring continuity over the long term. It would draw on and promulgate "open" (i.e., openly licensed and freely reusable) hardware, software, and data repositories with the flexibility to adapt to the changing needs of its various users. Through open infrastructure and collaborative governance, the commons could be designed to promote the advancement of knowledge, especially actionable knowledge, by researchers from frontline communities, advocacy groups, and academic institutions alike.

This is not pie in the sky. Scientific endeavors in which collaborators share resources without sacrificing the flexibility to pursue their own priorities, from foundational physics infrastructure (White Rabbit) to microscopy platforms (OpenFlexure) to agricultural knowledge systems (OpenTEAM and the emerging Agricultural Knowledge Concordance), already exist. However, attempts to establish an environmental knowledge commons have foundered because of systemic issues. First, the long-term attention needed for stewardship tends to be overshadowed by the urgencies of local campaigns, funding cycles, and academic career advancement. Second, funders have tended to prioritize short-term, place-bound environmental research projects over investments in infrastructure that could support many communities far into the future. Finally, while it has seemed appropriate to expect the federal government to play a leading role in creating and maintaining information infrastructures, reacting to unstable government involvement has in fact siphoned energy away from the creation of durable community assets.

Creating a knowledge commons requires a radical rethinking of how researchers, funders, activists, government, and communities work together to generate and steward the information needed to support meaningful solutions to environmental problems. Much as countering environmental injustice calls for a systemic analysis and structural response, we argue that environmental knowledge should be

understood as a collective, infrastructural problem for which communities, institutions, and funders share responsibility. As government support for environmental science and data decreases, committing to a knowledge commons through the collective work of governance—including establishing models for collaboration and shared authority—will enable both shared purpose and individual agency over important data in the long term.

Creating an environmental knowledge commons will undoubtedly require innovation, both in technology and in organizational structure. It does not, however, require starting from scratch. There is much to be learned from decades of experimentation across open-source movements and commons-based initiatives. Further, a successful environmental knowledge commons could inspire and inform efforts in other areas of knowledge production whose sustainability is threatened by federal disinvestment in research.

The power of knowledge commons

The past two decades have given us several models of knowledge commons. Examples include the now-defunct community-built nonprofit Public Lab and the global, decentralized Reagent Collaboration Network, or Reclone project, that provides molecular biology reagents to researchers. These projects are quite different, but they illustrate the power that commons have to advance both generalizable knowledge and local problem-solving. Their experiences also illustrate the challenges that knowledge commons must overcome to achieve the scale that enables them to be effective while also being sustainable.

Public Lab. Public Lab began in 2010 during the massive BP Deepwater Horizon oil spill, which at the time was being minimized by company and government officials. In response, a group of volunteers documented the extent of the damage by devising ways to fly over the spill area to take aerial photographs of the encroaching oil, improvising rigs from balloons and kites, digital cameras, and hacked camera software. They then used open-source software packages to knit images together to create a community atlas of the spill.

Drawing on traditions of open hardware and open-source collaboration, volunteers made their rig design and image-stitching protocols accessible to anyone who wanted to use them. In the aftermath of the spill, hundreds of people from all walks of life took pictures and supplied metadata on weather conditions, interpreting and contextualizing the aerial images through their lived experience of the Gulf Coast. They assembled maps that could inform the public understanding of and response to the spill. These volunteers' efforts made an important contribution at a time when the Federal Aviation Administration had closed the airspace above the spill and traditional media had little access to independent information.

The seven people who organized Public Lab (including coauthor Dosemagen) realized that this idea was bigger than the oil spill and that mapping could be incorporated into local organizing work, enabling communities to benefit from local environmental information being withheld or simply not collected by government officials. To serve those broader needs, they formed Public Lab, a knowledge commons that combined an extensive digital infrastructure, including a wiki, where plans for tools like the aerial imaging rig could be shared, documented, and updated with a diverse network of grassroots groups, volunteers, scientists, students, media, educators, and coders. This community convened virtually and in person to learn from each other, experiment with new tools, and articulate shared values. Guided by those values, they worked in concert to maintain and improve the digital infrastructure and to mentor newcomers.

Community groups and their partners from diverse places—including Brooklyn's Gowanus Canal as well as a community craft market in Uganda—shared their local findings on the wiki. Other participants contributed ideas and plans for new tools tailored to other kinds of environmental problems, including terrariums designed to filter toxins from indoor air and data loggers to monitor water quality parameters such as turbidity and temperature.

Public Lab became a nonprofit organization to provide an administrative backbone for the knowledge commons. With this designation, it was able to raise funds and provide staffing for the infrastructural work of the commons, such as maintaining the collective wiki, organizing convenings, and institutionalizing shared values. The Public Lab nonprofit relieved some of the burden on grassroots organizations to contribute to collective resources, streamlining administrative activities and freeing them up to focus on their local issues. However, having a central administrative body ultimately limited the longevity of the knowledge commons. It shuttered in 2023, when the organization could no longer raise funds for its unique work and no other institution was in a position to maintain the knowledge and resources amassed by the larger community.

Reclone. By contrast, the Reagent Collaboration Network, or Reclone project, follows a decentralized model, which has enabled it both to survive and to scale. In 2020, in response to long-standing supply chain challenges in the Global South that were exacerbated by the COVID-19 pandemic, a collection of academic researchers wanted to make it possible for anyone to produce molecular biology reagents without having to buy them from commercial vendors. In particular, these reagents can be used in the diagnosis of established and emerging diseases. Making reagents more accessible enables people without large grant budgets or extensive scientific infrastructure to conduct research targeted to local problems. Examples of efforts enabled by the Reclone project include tracking gastrointestinal pathogens transmitted

by invasive beavers in Patagonia, developing biosensors and diagnostic tests for human diseases like dengue, and tracking agricultural diseases like the potato virus Y.

Funded through a mix of regional philanthropic grants, UK government funding, and a largely distributed, bootstrapped fundraising model, the Reclone project distributes open-source toolkits, including reagents and backup collections of DNA to react with reagents. Research projects that source materials from Reclone strengthen its capacity and longevity by contributing back new reagents, methodologies, and findings that result from their investigations. As with Public Lab, the Reclone community is larger than its scientific materials. It comprises a global network of collaborating biologists, place-based reagent hubs, and localized technical support.

Unlike Public Lab, however, the Reclone project is fully decentralized, with no core administrative apparatus. While Reclone partners with universities and research institutions that have long time horizons and strong stewardship commitments, those institutions do not own the products of the network or exert control over its directions. Investigators involved in the project choose for themselves what research

The challenge of navigating competing incentives

Both Public Lab and the Reclone project were initiated by people who recognized a need for improved scientific infrastructure and saw the value in working as a collective to address that need. However, in committing to creating a commons, both researchers and communities have to overcome competing incentives that are baked into the way collaborative research is conducted and funded. Mismatches between short-term, place-based, centralized funding structures and the long-term, collective work of building infrastructure put projects like Reclone and Public Lab in a precarious position. Fragmented and inadequate pools of funding often come with demands for deliverables that don't give much back to the knowledge commons itself.

The first set of challenges arises in the design of knowledge commons, which are strengthened by diverse networks of people and organizations participating in the ways that fit them best. These researchers and communities, even as they work together, have different inherent priorities and must navigate separate sets of incentives from funders and supporters.

Creating an environmental knowledge commons will undoubtedly require innovation, both in technology and in organizational structure. It does not, however, require starting from scratch.

is most needed in their particular context and, through their participation in the community, become able to recognize how their innovation, maintenance work, or both can be of value to researchers working in other places.

Without staff devoted solely to maintaining the Reclone project's infrastructure, the commons relies on a handful of core labs with a shared vision and a sustained commitment to contributing actively to the commons. This structure has enabled Reclone to survive and reach meaningful scale without a central administrative body, in part because its early participants were academic scientists with established labs, existing infrastructure, and strong networks who could immediately respond to a clearly felt need. Today, its DNA collections are housed in roughly 500 labs worldwide, with about 30 labs actively contributing and 10 serving as core stewards. But without dedicated infrastructure funding, growth has been steady rather than rapid. For instance, expanding to several thousand labs and building a larger base of 20 sustained core labs would require additional resources, even as overall scale is naturally bounded by the finite number of molecular biology labs globally.

Ideally, knowledge commons are structured to mediate between these different interests and incentives to accommodate everyone involved. Governance agreements and processes must be built to acknowledge contributors' differential capacities for infrastructural work and stewardship, without marginalizing smaller organizations in decisionmaking or long-term planning. Community values need to be broadly shared and to provide the basis for dividing up work and resources in a way that is fair, even if not equal. The Reclone project, for example, has a handful of core labs that contribute actively to building and maintaining the commons, yet hundreds of labs use its DNA and strains. This structure encourages and makes it easy for less well-resourced organizations to offer their data, innovations, and experiences back to the commons.

All of this is difficult, time-consuming work. Within a scientific culture that prioritizes individual accomplishment and clear, private ownership, researchers have few incentives to engage in this collective work beyond a belief in its intrinsic good. Work done to ensure

that a broader community has the capacity to participate in scientific research tends to be considered by promotion and hiring committees to be of lower worth, when it is acknowledged at all.

But the reward systems of academic institutions are not the only ones that discourage participation in knowledge-sharing. Community researchers who need context-specific information are also under pressure from their supporters, especially grantmakers, to delineate *their* outputs and deliverables. Their contributions to collective, infrastructural work can be difficult to characterize in these terms. Community environmental problems, for example, are often chronic but punctuated by events demanding a quick response. When regulators open a 60-day comment period, communities need to quickly marshal data to be included in public deliberations; a fire or chemical spill creates both an urgent need to collect and interpret data and new opportunities to garner media attention for local environmental hazards. Although having a robust knowledge commons enables a quick response to such calls, the need to act urgently can undermine grassroots groups'

idea gained popularity, investments expanded, but often in ways that were segmented. Environmental funders tended to support community-based efforts tied to local outcomes or broad topical campaigns and coalitions. Science and technology funders, on the other hand, invested in research, tools, and platforms. What remained harder to fund was the perennially complicated backbone that connects this work: the ongoing maintenance, coordination, and translation required to weave these efforts into a shared and durable infrastructure.

Because most foundation and government grants provide funding for no more than five years, any investments in infrastructure that a particular grant may allow are necessarily short-lived. These short time frames make it difficult for knowledge commons contributors to think at the decadal scales appropriate to infrastructure. They also make it harder for participants to build shared values and trust. In this sense, today's funding structures not only fail to support knowledge commons; they may even undermine attempts to envision it.

A further barrier to long-term collaborations is grantmakers' practice of typically awarding funding either

In committing to creating a commons, both researchers and communities have to overcome competing incentives that are baked into the way collaborative research is conducted and funded.

capacity to document lessons and translate their experiences into shared knowledge, even when they see the value of doing so. As a result, even organizations that want to share resources with the scientific community may find it more viable to set up projects that they can retain control over, rather than envisioning them as true commons.

Barriers to investing in foundational architectures for knowledge

These cross-wise incentives for academic and community researchers' participation in knowledge commons are strongly influenced and structured by the funding landscape. Securing funding for infrastructural work often entails navigating disconnects between the ways grantmakers define fundable projects and the kinds of work necessary for a knowledge commons to flourish. This constant friction requires that contributors to the commons find ways to talk about and structure their work that conform to grantmakers' expectations.

For example, Public Lab's support for community-led science first garnered funding at a time when the value of nonscientists' contributions to environmental knowledge was just starting to be institutionally recognized. As the

to principal investigators or to lead organizations. Even when collaboration is required, only one entity is designated as the primary grant recipient. Although this structure may serve grantmakers' desire for accountability, it also pushes researchers and nonprofits to structure activities in ways that give them control, rather than enabling the creation of decentralized structures that can be governed and maintained collectively. Our reflections about these barriers are informed in part by collaborating with Guillermina Actis, Katie Hoerberling, and Cathy Richards to create the Digital Toolkit for Collaborative Environmental Research, which involved considering how environmental researchers and open infrastructure practitioners might better align funding structures with long-term digital stewardship.

Moving from government to governance

As the environmental data community has grown over the past two decades, there has been a persistent hope that the federal government, which has played a significant role in maintaining public data infrastructure in the United States—including weather satellites and ambient air monitoring stations—would support the creation and storage of this kind of knowledge.

That hope has been dashed as both Trump administrations have moved in the opposite direction, restricting access to information that was formerly publicly accessible and taking down resources, such as EJScreen, that combined multiple data sources to enable users to get a bigger picture of environmental impact. In response, environmental researchers and advocates have mounted large-scale, collaborative efforts to rescue data likely to be lost and to preserve analytical tools on which communities and local governments rely.

Although these rescue efforts are important, the primary obstacles to adequate environmental knowledge infrastructures are not hostile presidential administrations. In fact, government-sponsored environmental initiatives have never formed a robust basis for a knowledge infrastructure that can foster advances in generalizable knowledge alongside community problem-solving. Environmental data collected by government agencies themselves tends to be fragmented and inaccessible. Monitoring and emissions data reported to the government by industry, for example, is often made available through platforms such as WebFIRE, which are organized by the rules under which data is collected rather than the issues to which it pertains. In addition, not all data is available in machine-readable formats: PDFs are still all too common.

Furthermore, even when the federal government has been supportive of community-based environmental research, it has done little to ensure that place-based research contributes to understanding larger environmental trends. For example, in November 2022, the Environmental Protection Agency (EPA) awarded a total of \$53.4 million in grants to 132 communities to conduct local air quality monitoring. The grants included activities to ensure that monitoring contributed to community empowerment, skills-building, and/or policy action to reduce local air pollution. However, the program did not encourage applicants to allocate funds to making data accessible to researchers beyond the community or to collaborate with other organizations to set up shared infrastructure for data storage and management. Further, the grants came with bureaucratic requirements that presented challenges to small nonprofits unaccustomed to administering federal funding. Projects had to develop Quality Assurance Project Plans detailing their methods for collecting and validating data and have them approved by EPA staff, a time-consuming process. Teams planning to have contractors conduct the actual monitoring had to put the work out for bid, a hurdle that took still more time. These barriers restricted the participation of smaller and less bureaucratically savvy organizations—further widening the gulf between federal initiatives and the diversity of a knowledge commons.

For communities and researchers who want to bring an environmental knowledge commons into being, responding to government initiatives is not time well spent. Even in favorable administrations, opportunities for federal support are stubbornly shaped by the expectation that research will be conducted by entities with significant technical capacity and not collectivized in ways that might serve the goal of creating knowledge that can apply across multiple locations.

The sweeping freeze of federal funding imposed by Trump shortly after taking office in 2025 abruptly halted work on community air-monitoring projects, many of which subsequently received notices of termination. This abrupt reversal underscores a third reason the federal government can't be counted on as the locus of environmental knowledge infrastructure: Its priorities are too changeable. Just as two- and three-year grants do not offer researchers and advocates the time horizons needed to develop a robust knowledge commons, four-year presidential administrations and two-year congressional election cycles almost guarantee ongoing upheavals to government support for environmental research. When communities focus on preserving resources created by relatively supportive administrations while waiting for hostile ones to fall out of power, time—and opportunities to build knowledge more broadly—slips away.

Making progress toward an environmental knowledge commons

Today, many supporters of environmental policies, including those who see community-based knowledge production as playing a significant role in guiding those policies, crave a return to a federal government interested in evidence-based policy and willing to invest in environmental research. Decades of experience, however, have shown federal policy to be too mercurial to provide the foundation for steady advancement in environmental knowledge that cuts across places without losing sight of local issues.

Instead of planning for a return to normalcy, people who care about knowledge with the potential to inform environmental action should combine forces to develop a civil society-based knowledge commons—one that includes community agreements and decisionmaking procedures as well as data management platforms—to which government agencies at all levels *could* contribute, but which does not depend on government support or participation to thrive. The establishment of a scalable, sustainable knowledge commons will not only put more data in the hands of communities made vulnerable by pollution and climate change; it will also help generate insights into the ways traditional knowledge systems have operated to the advantage of the already-powerful

and enable environmental researchers (academic and community-based alike) to generate analyses that shift those power dynamics.

As the examples of Public Lab and Reclone show, much can be accomplished through decentralized sharing of data and resources for scientific research. A knowledge commons built on open infrastructure can foster the advancement of generalizable findings and techniques, while also facilitating place-specific research and problem-solving. The system-level hurdles to the infrastructural work necessary to create and maintain a knowledge commons could be overcome by a wide range of civil society groups working together. In particular, funders have a great deal of discretion to change the parameters and requirements of grants—and, in doing so, to substantially reshape incentives.

Making true progress toward an environmental knowledge commons requires three major shifts in the approach to environmental research. First, building on the environmental justice movement’s recognition that environmental injustice is a structural problem, systemic coordination across communities to address the root causes of environmental degradation should be extended

to incentivize the knowledge commons. Minimally, this support would involve large initial investments to permit collaboration at scale and subsequently structuring grants to enable a portion of the funding and staff time to be devoted to the work of the knowledge commons. But we urge foundations to go further. Using both their financial resources and their convening power, major funders could help contributors to the commons surface shared goals and values, develop governance agreements, and set up administrative structures that permit the commons to remain decentralized. As collaborators in the knowledge commons, well-established foundations could also help ensure the long-term viability of shared resources by prompting periodic stock-taking and supporting updates, revitalization, and course corrections.

Finally, we call on everyone interested in producing environmental knowledge to ensure community health to stop reacting to ever-shifting political winds. Instead, all parties should use the project of the knowledge commons to reconstruct their capacity to act collectively without shutting down local creativity, to work toward

Instead of planning for a return to normalcy, people who care about knowledge with the potential to inform environmental action should combine forces to develop a civil society-based knowledge commons.

into the epistemic space. Frontline communities and their allies should commit to advancing collective knowledge of environmental problems in tandem with building information resources at the local level.

In making such a commitment, grassroots groups, nonprofit supporters, and academic researchers would accept the obligation to contribute actively to the knowledge commons, providing data, labor, and/or funding as they are able. Resisting the urge to own or control the knowledge commons, they would instead deploy their organizing know-how to demonstrate its utility and to create community agreements and oversight boards that enable coordination without a loss of autonomy.

In practical terms, it will be crucial for a knowledge commons to start big. The initial scale should be such that potential participants can’t miss it and thus won’t squander energy on attempting to duplicate it. This need for scale makes a collaborative mindset and good governance agreements all the more important. It also requires vision and commitment from funders.

Financial supporters of environmental research, including government agencies, should enable and

common purposes while respecting individual agency. All should, in other words, engage earnestly in the hard work of governance, paired with sustainable stewardship, to enable the flourishing of knowledge. Creating effective governance structures is undoubtedly the aspect of the knowledge commons that will require the most care, experimentation, patience, and goodwill. Yet the underlying challenge is not unique to environmental research. Across domains, from public health to agriculture to physics, the long-term stewardship of shared research infrastructure is increasingly precarious. Learning to build and sustain a knowledge commons here may help establish a durable model to sustain collective capacity wherever science depends on shared data, tools, and governance.

Shannon Dosemagen is founder of the Open Environmental Data Project and currently serves as senior fellow. Gwen Ottinger is professor of politics and environmental studies at Drexel University and author of The Science of Repair: How People Who Believe in Facts Can Build a Better Future (Oxford University Press, 2026).