Open Knowledge

THE TECH WE WANT

Read This Before You Build

A field guide from the Open Knowledge Foundation on **how to build public interest tech *with* communities**

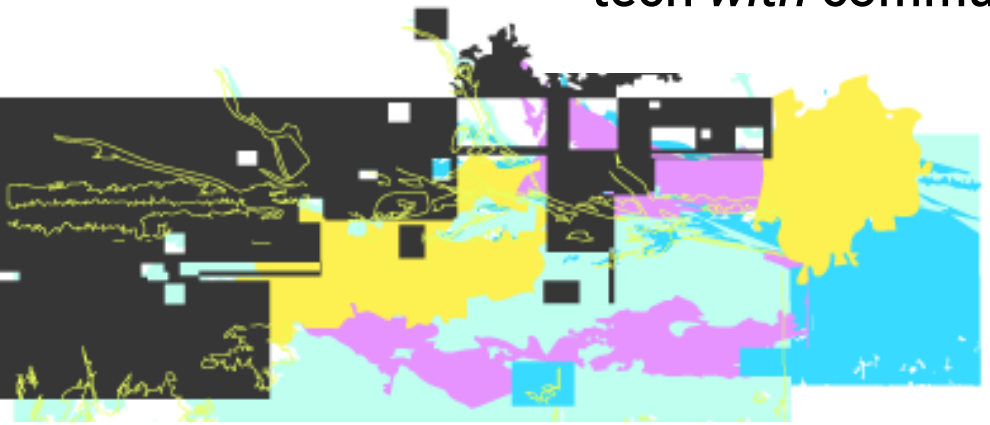

SEPTEMBER 2025

INDEX

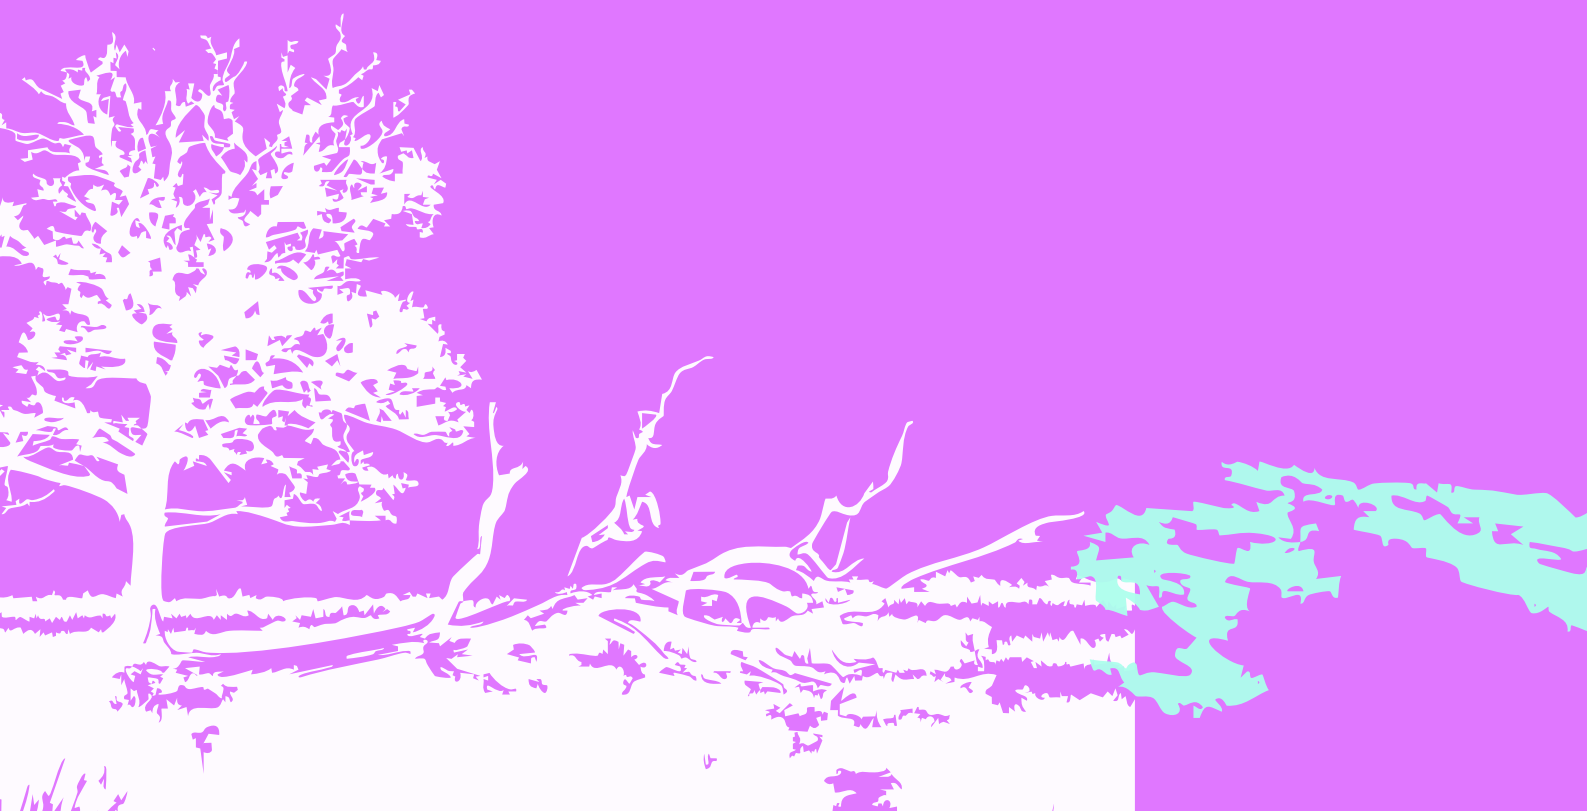

ReadMe

On learning and unlearning

Over the past 20 years, the Open Knowledge Foundation (OKFN) has pushed for openness to solve the most pressing problems of our times. We support people who apply open knowledge and data to their work, including with tech that is adaptable to different realities. For instance, we understand the challenge of working through daily internet outages—or just how much labour goes into a public dataset before it ever sees the light of day.

Lately, we have been reflecting on our approaches and inviting public discussion about *The Tech We Want*. It's about the futures we envision for open knowledge and the digital tools we wish to create, given the unprecedented power concentration there is around the tech we have.

As builders of public interest tech, we are pressed to defend how we add value to a tech ecosystem dominated by the most powerful corporations in the world. We don't have billion dollar infrastructures, and we don't have capacity to update our software as frequently as the tech giants do. It's not a level playing field and this is exactly the problem.

We know far too well that technology and data can be tools of oppression, just as they can be instruments of liberation. It depends on the values, priorities and incentives of their creators. For technology to line up with the values, aspirations and hopes of the people we serve at OKFN, something must change. We need to rethink how we build and govern digital tools that truly work for *us* to ensure that we have options, now and in the future.

In an age of tech companies that 'move fast and break things' we are doing the opposite. How do we build slowly and deliberately in ways that differ from norms dictated by Silicon Valley? Collectively, we have so much to *unlearn* about what makes software successful and worth sustaining.

At the same time, we at OKFN, are questioning our own practices and learning how to do things anew. We have faced our mistakes, decommissioned and redesigned several of our projects. Gone are the days when a lone developer would design software without input from real people. Gone is the presumption, for us, that less technical communities don't belong at the core of a design process.

It's about *unlearning* to desire digital tools that can do all things for all people. Simpler, faster and more focused apps for *specific* people can be more practical for everyone, including their code maintainers. It calls on us to have frank discussions about what makes building *The Tech We Want* genuinely worthwhile—and about what not to build too.

At OKFN, we have turned the tables and distilled what we learned in this short guide for anyone building public interest tech, including apps or any other kind of interoperable software that becomes infrastructure for open knowledge movements.

Whether it be in the commercial, civil society or public sectors, how do we root out practices that are counterproductive for the digital and physical transformations we wish to see worldwide? There is so much we can learn, still, by designing and creating alternatives, together.

8 questions to ask before you build

Are you getting ready to build public interest tech?

We have compiled eight important questions to discuss with your team, accompanied by field notes from the OKFN from our recent experience building the Open Data Editor (ODE).

Today, this is a simple app that makes it easier to detect errors in datasets—but it took a lot of learning (and unlearning) to get there.

Open Data Editor (ODE) is a free desktop app by OKFN that makes discovery of problems in datasets—like inconsistent labels and dates—easier and faster.

In 2023, we launched the beta version of ODE after two years of development. But almost no one understood how to use the app.

We had detoured from best practices for software development, including getting early and iterative feedback from a clearly defined audience.

We had relied mainly on input from technical experts rather than from actual communities working on public interest data projects.

So, we chose to rebuild collaboratively with our community.

We put their needs at the center of our planning. And we chose to build very simply and modestly to increase the life span and sustainability of our project.

Our feedback improved—and the process itself was rewarding due to all that we learned about how to best serve our communities.

Mongolian data enthusiasts during an Open Data Editor training course run by the School of Data/Public Lab Mongolia. Photo: Tsenguun Tumurkhuyag

8 questions to ask before you build

1. Who are we building this for?

What we advise now (the 'after' story)

Identify your target audience and establish a genuine relationship with several people who are part of it. Otherwise, you develop tools that work for no one.

Our learning journey with ODE (the 'before' story)

We knew we wanted people with less-technical backgrounds to use ODE, but we hadn't defined our audience narrowly enough to engage well with them.

2. What problem are we solving?

Ask specific people from your audience how they usually work and exactly what they need and know how to use, making sure their input is core to what you design.

After two years of development, ODE had so many functions, it was hard to explain simply to our audience what the app was for, and all that it could do.

3. What should we NOT build?

To stay focused, set clear limits on what tools shouldn't be able to do. With limited resources you need to prioritise core functions and cancel or postpone others

Our communities also use data tools from major tech companies. It was pointless for us to try to compete in areas where their needs were already being met.

What we advise now (the 'after' story)

Our learning journey with ODE (the 'before' story)

4. Does this work for our audience?

Discuss with communities what they understand technically and what works on their specific devices and operating systems. Not just once, but many times.

We built an overly complex app that few people could use. We relied too much on technical experts in the development process and failed to address that there are often knowledge and access gaps.

5. Is our language clear?

Check whether the language you use is understandable and accessible or is technical jargon. Release updates often to get quick feedback—and make changes!

Sometimes the language we used on buttons and in menus was not intuitive enough for everyone to understand. We didn't realise how much of a barrier this can be.

6. Do we need AI in this?

It's what everyone is talking about, but you should carefully consider whether AI really adds value to your project, and whether you are ready to deploy responsibly.

We introduced a generative AI plug-in to ODE that received so much community pushback over lack of transparency and reliability that we later replaced it.

What we advise now (the 'after' story)

Our learning journey with ODE (the 'before' story)

7. How long lasting will this be?

Public interest tech should ideally be designed for stability and permanence. Will it expire or require updates? Less features can also mean longer lifespan.

After the beta release of ODE, one of our main software dependencies went offline. We didn't have the capacity to fix everything that broke.

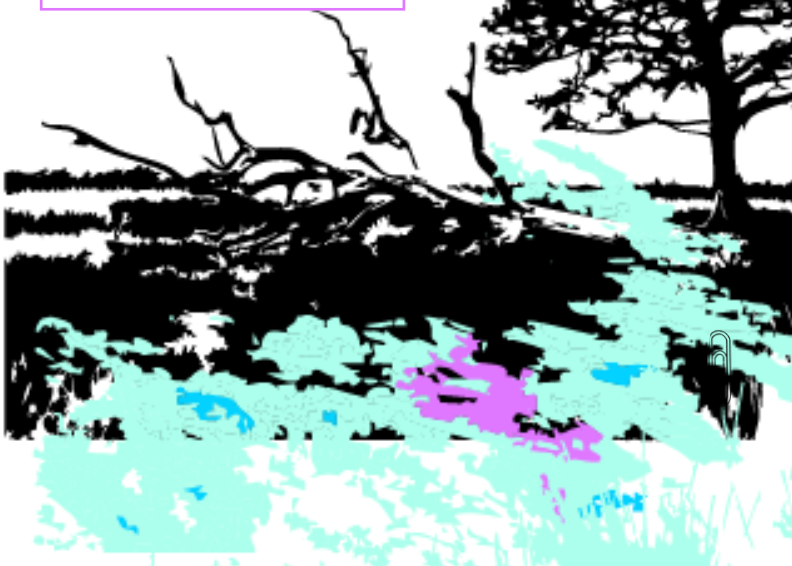

8. Who will maintain this?

Build public interest tech that is easy to maintain so you can have a realistic process for how to govern and enable communities to localise and upkeep on their own.

We had a team of five people working on ODE. Every change and update had a real cost. Only later did we start to consider how to keep updating and funding the app long term.

Students, developers and representatives of civil society learn about ODE in Benin. Photo: Saliou Abdou

Some tips on building with communities

Do user research at the beginning, but keep consulting with your community while building—and after release.

Building with communities is the slow cooked meal of software development. We need to simmer, stir, taste a dish, and add seasoning many times before it's done.

OKFN has been a software building organisation for nearly 20 years. Until recently, we had grown a little overconfident in our ability to guess what people need based on a few conversations.

In the case of Open Data Editor (ODE), we did ask a number of potential users about their challenges working with data. But we never checked in with them again to see if they could use what we developed. This was a mistake.

The point of ODE was to help people save time and effort finding errors in large datasets, but instead we ended up creating new pain points with the app itself.

We tried to anticipate needs by working on a desktop app with locally stored files that would be better for privacy and unreliable internet connections. But in the feedback process, we turned only to our usual open data collaborators. The result—we later realised—was that even installing ODE was too complex for less technically savvy folks.

We also included several features that were a bit buggy and difficult to understand intuitively. For instance, technical terms like "metadata panel" or "execute" can be confusing or intimidating if you have never seen them before.

In the end, we turned things around by working directly with communities, and also by changing how we work and communicate internally ourselves.

Today, ODE has been collaboratively redesigned and is used by researchers, journalists, non-profits, and government workers around the world. We are proud of how simple and accessible it is now.

We'd like to share some tips based on what we did. The secret ingredient is user testing after release and at every stage of development. It's slow and deliberate, and sometimes steers us in a different direction than we originally intended.

To us, the added benefit of the process is that it unlocks new digital literacies. Building together and actively supporting each other, is learning together. We get to understand more about our community's needs and abilities, which helps us in our mission to impart skills in ways that can be scaled. We train people, so they can train others. This emphasis on peer learning and mutually beneficial exchange goes beyond simple user research.

We occasionally refer to our audience as "non-technical" even when it is highly skilled professionals working with large datasets. They are not primarily technical (or coders), which is why ODE can be helpful.

Build tech for humans, not abstract audiences

What we advise	*How we redesigned Open Data Editor*
Identify pilot projects from your target community. Offer a fee for their time and insights. Ensure you include a diversity of backgrounds and experiences—including on your internal team—for better outcomes.	*Our pilot projects included journalists, researchers, and local governments in 10 countries on four continents.*
Find a person who is good at 'translating' between your technical and less-technical contributors, internally and externally. Don't write a single line of code until you know the exact problem you will solve.	*For our user research process we chose to work with someone who has both data and journalism skills.*
Release fast, get feedback, and improve again. Keep iterating together. Have pilot projects write journals of their experience, and commit everyone to a certain number of calls.	*The feedback from pilot projects was meticulously logged, and we brought more transparency to our own decision process.*
Don't use tech jargon. Create an atmosphere where people feel comfortable asking questions. Prepare materials (beyond standard documentation) for people to learn how to use what you build.	*Several pilot projects enrolled in training programs for trainers to help even more people learn how to use ODE.*

StoryData, a Barcelona-based data journalism agency, participated in the first pilot test of the Open Data Editor in 2024. Photo: StoryData

More tips: Ask and listen

- *Start with broad and open questions, like:* "Tell me more about how you work with data"?
- *Be an active listener.* Do not ask questions to prove that your ideas are the right ones.
- *Talk to organisations working on projects similar to yours* to learn from them.

Meet an ODE pilot project:
In Argentina, The Civil Association for Equality and Justice (ACIJ) monitors for housing inequality in massive volumes of data.

"The leap in quality that ODE enabled us to make was fundamental. In a country where access to housing is critical, the ability to analyse data quickly strengthens our impact."

Eduardo Ferreyra, Co-Director of ACIJ

Meet an ODE pilot project:
The Demography Project in Kenya is a youth-led science and civic tech group that collects citizen data about the environment and democracy.

"With ODE, we have a tool that helps us (...) ensure that we have some form of coherence in the tabular data from our air quality monitors, phones and action cameras."

Richard Muraya, Executive Director

What makes a good pilot project?

- A dedicated person for you to be in touch with.
- Support of their organisation for their participation.
- Good ability to deliver constructive feedback.
- Genuine interest in peer learning and collaboration.

What to do about AI?

Lately it feels like all tech conversations lead to artificial intelligence (AI)—and this goes for public interest tech as well. Do we want it? Do we need it? What should we do?

As frequent critics of the extractive business models of major AI companies, we are relatively cautious in our approach to integrating AI in our software.

But even being cautious, we moved forward with a ChatGPT integration in ODE that we later removed and replaced with a more private, open source option.

Through trial and error, we learned the importance of interdisciplinary collaboration and feedback from our community—beyond the technical implementation.

Here is some advice:

- **Create an interdisciplinary team to consider AI:** Don't limit conversation about AI implementation to code. Open up space for multiple voices to be heard—including those who say, "But we don't need AI to do that."
- **Ask many questions:** Don't take anything for granted, ask yourselves questions throughout the process. For example, "Am I certain people understand what data we are sharing when they click the button?"
- **Listen to your community's needs and concerns:** Several journalists and researchers told us they had concerns about using AI because of privacy issues. It was a mistake not to center this input from the start.
- **Don't strive for perfection:** With limited resources, we were still able to try an AI integration quickly that we later improved. Building public interest tech we have to balance many needs, including time and cost.
- **Keep learning as you go:** It's hard to find practical advice about AI models for specific use cases. We struggled to find info on tabular data. OKFN offers a range of AI literacy resources for creators of public interest tech.

Simple. Long-lasting. Ours.

We created this guide and included reflections on our past mistakes—not to overwhelm you about what could go wrong—but to inspire anyone building public interest tech about how relatively small changes in design processes can lead to huge rewards in terms of results.

Choosing to develop tools that are 'simple' in an age where the tech industry is constantly making bigger, bolder and more bullish promises may seem counterintuitive, but we believe it's key to unlocking the full potential value of The Tech We Want.

We often use the examples of a ballpoint pen or a safety pin as basic, lowtech inventions that have remained useful through the ages. They are useful and adaptable to so many different contexts and environments precisely because they are simple. We can draw parallels to the tech needs of many communities working on open knowledge projects around the world, because sometimes it's the very basic, most private, offline tool that works best.

It's admittedly difficult to strike a balance. As technologists, we are easily excited about possibilities and want to push ourselves to build ambitious projects. However, sometimes it is in collaborating directly with communities, we fully realise that the most ambitious path can be to go for something that is also built for longevity and sustainability—especially on budgets that will never compare to those of corporate tech giants and startups.

Let's keep learning together

Did you know the OKFN has many online learning resources and hosts seminars in multiple languages? Contact us if you want to talk to our team.

We believe we share a responsibility in the field of open knowledge to use our resources wisely, and build tools that can be easily learned by all and used for as long as possible.

It requires humility, curiosity, and awareness of the knowledge gaps between technical and non-technical communities, but also a dedication to the mission of building something that we—the global we—can control ourselves, even as free and open technical spaces are actually shrinking around the world.

Our technology systems, whether it be the software, the networks, or pervasive data collection, are vectors of potential control. So it is crucial that we keep building and experimenting with alternative tools that are transparently governed and privacy friendly.

Civil servants in Mexico City after completing an Open Data Editor training course offered by Escuela de Datos. Photo: INFOCDMX

As we showed throughout this guide, good intentions don't isolate us from actually needing to 'do the work' and follow best practices for software development. Even we ourselves, have had to go back to the drawing board to ensure that when we build for specific communities, we develop genuine relationships with people, and hold ourselves to a high standard of proof that what we build works for them. This is something you can do too.

There are knowledge gaps to be bridged in all directions. Everyone involved in a project brings their own personal and professional experiences to the table; and everyone has something to learn (and unlearn) in a process of co-creation. Even as we are self critical, we want to be joyful and optimistic in the process of building something together—no matter if it is imperfect or only a small step in the direction of where we wish to go in the future.

Ready to build?

We prepared a simple self-assessment exercise to spark ideas and discussion about how to build The Tech We Want. Share it with your team and see how your answers line up.

How could you develop more readiness?
Ask us for advice!

Self-assessment

Are you ready to build with a community? Print this page.
Read each statement. Does it apply to you?
Mark (x) on the scale.

1.
Specific people from our audience are ready to collaborate with us.

YES — NO

X

2.
We have a clear goal with our build and it's easy to explain.

YES — NO

3.
We have cut down on features that are not truly necessary.

YES — NO

4.
We have a good feedback process and a person in charge of it.

YES — NO

5.
We check our language with people to weed out tech jargon.

YES — NO

6.
We push back when ideas for AI carry too many risks.

YES — NO

7.
We have a plan for the years and decades after release.

YES — NO

8.
There are folks who can help upkeep our project longterm.

YES — NO

Self-assessment cards

1. Who are we building this for?

Discuss:

Who was the last person you spoke to of your core audience? What will you ask them the next time you speak?

2. What problem are we solving?

Discuss:

If you wrote a message in a bottle about your software and threw it in the ocean, what would it say?

3. What should we NOT build?

Discuss:

You're going to a deserted (tech) island. What's the one key feature of your app you bring for survival?

4. Does this work for our audience?

Discuss:

If your software were clothing, how would it fit your audience? Is it soft or scratchy? Imagine the colors.

5. Is our language clear?

Discuss:

Think of a time you saw a button in an app or on a machine that was hard to understand. Would you rename it?

6. Do we need AI in this?

Discuss:

If your tech project was a house, what would you automate inside it? Laundry, climate control, or...?

7. How long lasting will this be?

Discuss:

How do you envision the 20th birthday of your software? With whom will you celebrate and how?

8. Who will maintain this?

Discuss:

Will you have funds to pay a developer? What governance model could help others make it their own?

Further resources

Ready to keep reading? Here are some resources from the Open Knowledge Foundation — and a few that we enjoyed from others too.

Building Together: Simple Tech for Stronger Communities
OpenKnowledge Foundation

A video from The Tech We Want Summit with speakers from Zooniverse, Global Voices, Mastodon, Platoniq Foundation, and the Restart Project.
archive.org/details/building-together-the-tech-people-want

Don't Build It. A Guide for Practitioners in Civic Tech
MITGov/Lab

This guide aims to help you avoid bad projects, structure the team right, ship and learn quicker, and mature longer. Available in English, Spanish, and Japanese.

mitgovlab.org/resources/dont-build-it-a-guide-for-practitioners-in-civic-tech

Compost engineers and sus saberes lentos: a manifest for regenerative technologies
Coding Rights

An epistemological, historical, political, and creative exercise to expose harmful Western-centered logics and imaginaries that guide tech development.
codingrights.org/docs/compost_engineers.pdf

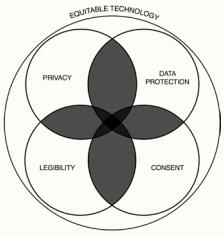

Human Rights Centered Design Methodology
Caroline Sinders and Natalie

Are you an organization building, designing and researching technology for vulnerable or at-risk people? This methodology and curriculum by human rights researchers, activists and designers is for you.
humanrightscentered.design

Design Justice Network

A network of people and projects practicing principles of design justice to support care, healing, liberation, joy, and deep sustainability.
designjustice.org

PJMF Learning Hub

A centralized place for supporting nonprofits in becoming more data and AI-enabled organizations.
learn.mcgovern.org

How to Use AI Without Selling Your Soul (or Data)
Open Knowledge Foundation

A video from The Tech We Want Summit with speakers from Aspen Digital, Correctiv, University of Virginia, and the Digital Policies and Digital Transformation Section of UNESCO.
archive.org/details/how-to-use-ai-without-selling-your-soul-or-data-the-tech-people-want

Open Knowledge
FOR A FAIR, SUSTAINABLE AND OPEN FUTURE

By: Open Knowledge Foundation

Illustrations and design:
Constanza Figueroa Bustos

September 2025

Recommended citation style
Open Knowledge Foundation (2025): Read This Before You Build: A field guide on how to build public interest tech with communities, London.

Open Knowledge Foundation (OKFN) is a not-for-profit organisation incorporated in England & Wales with company number 05133759. okfn.org

We are building a world open by design where all knowledge is accessible to everyone.

Patrick J McGovern
FOUNDATION

We are thankful for the support of the Patrick J. McGovern Foundation. Learn more about its funding programmes here.

Open Knowledge
FOR A FAIR, SUSTAINABLE AND OPEN FUTURE

#TheTechWeWant

A collective conversation about practical ways to build software that is useful, simple, long-lasting and focused on solving people's real problems.

Join us:
okfn.org/the-tech-we-want

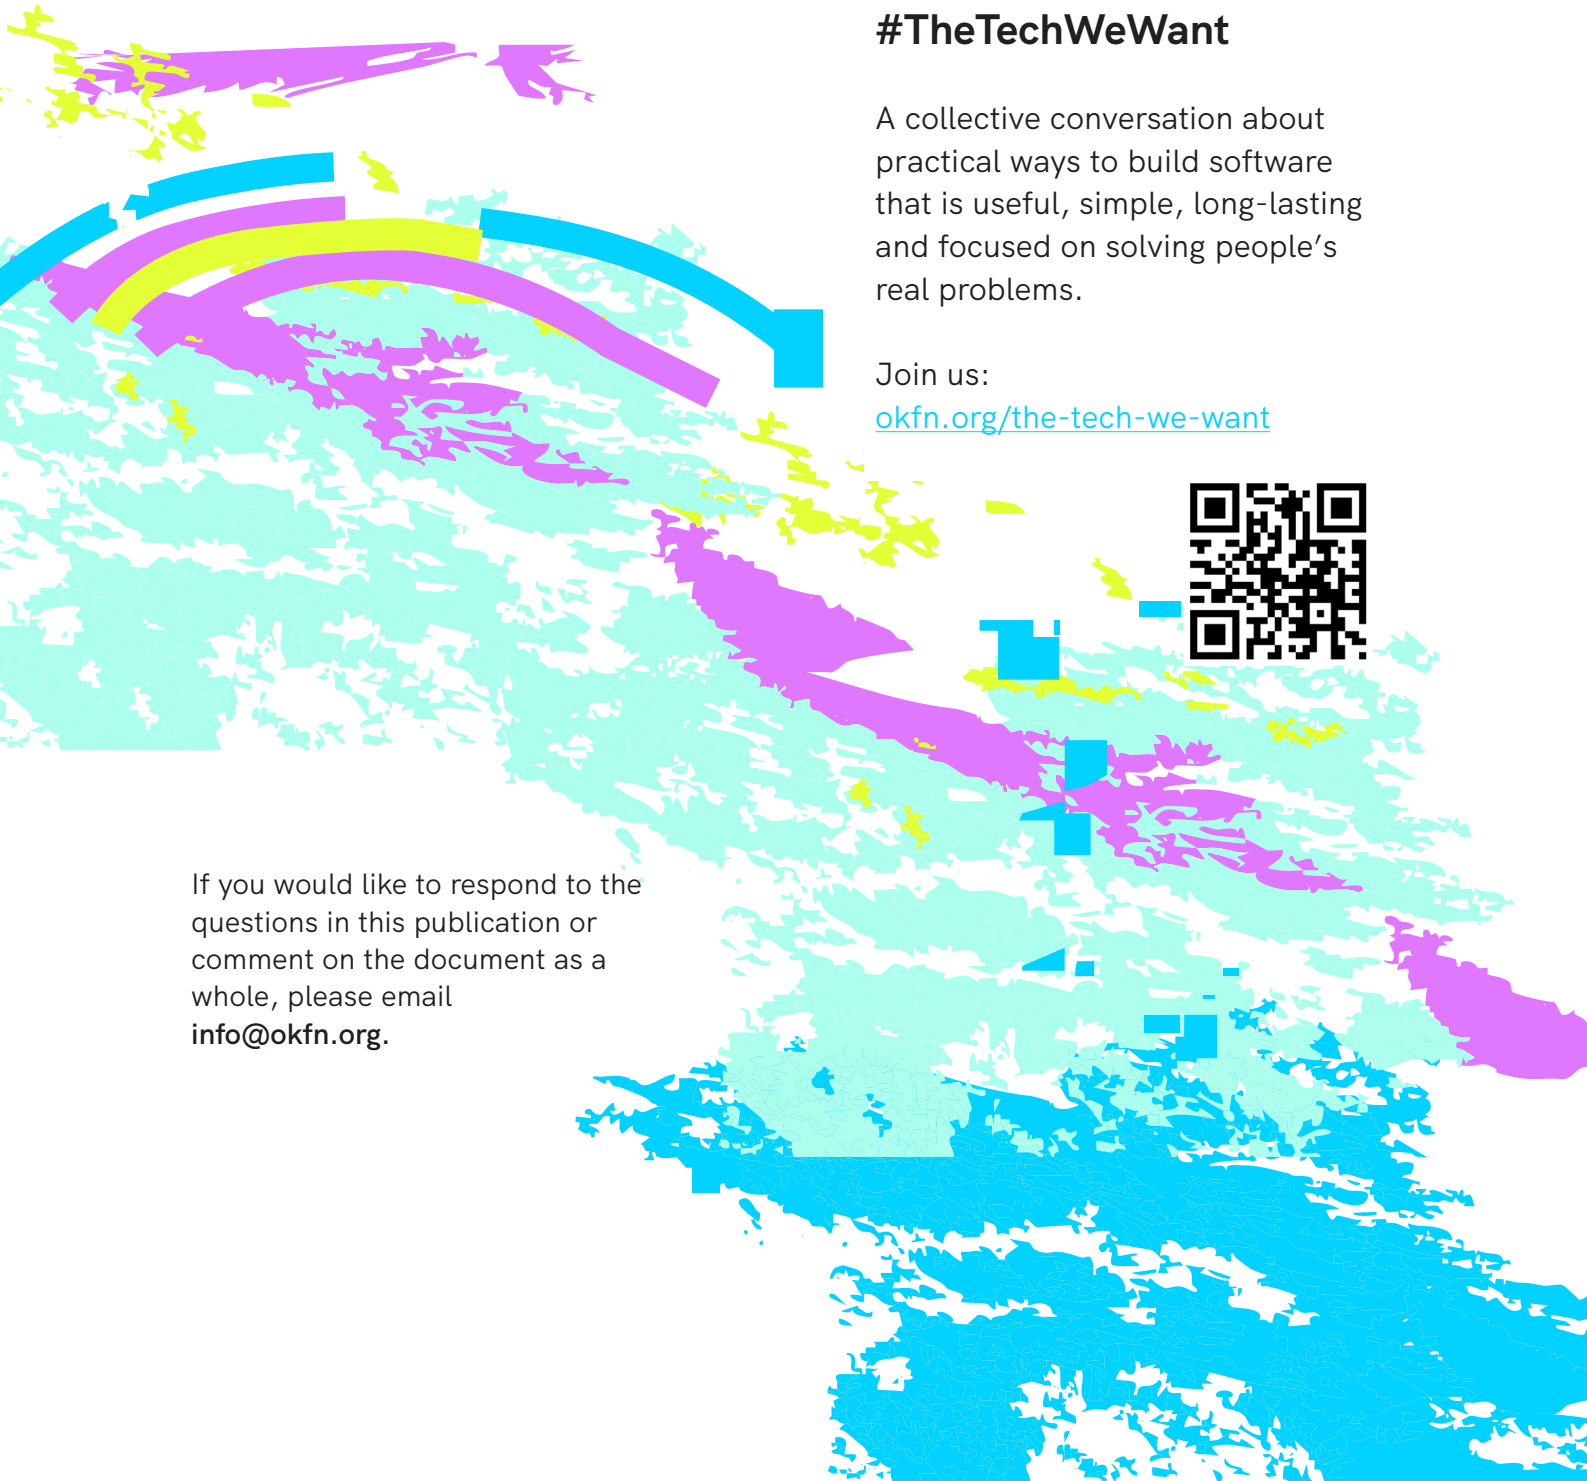

If you would like to respond to the questions in this publication or comment on the document as a whole, please email **info@okfn.org**.