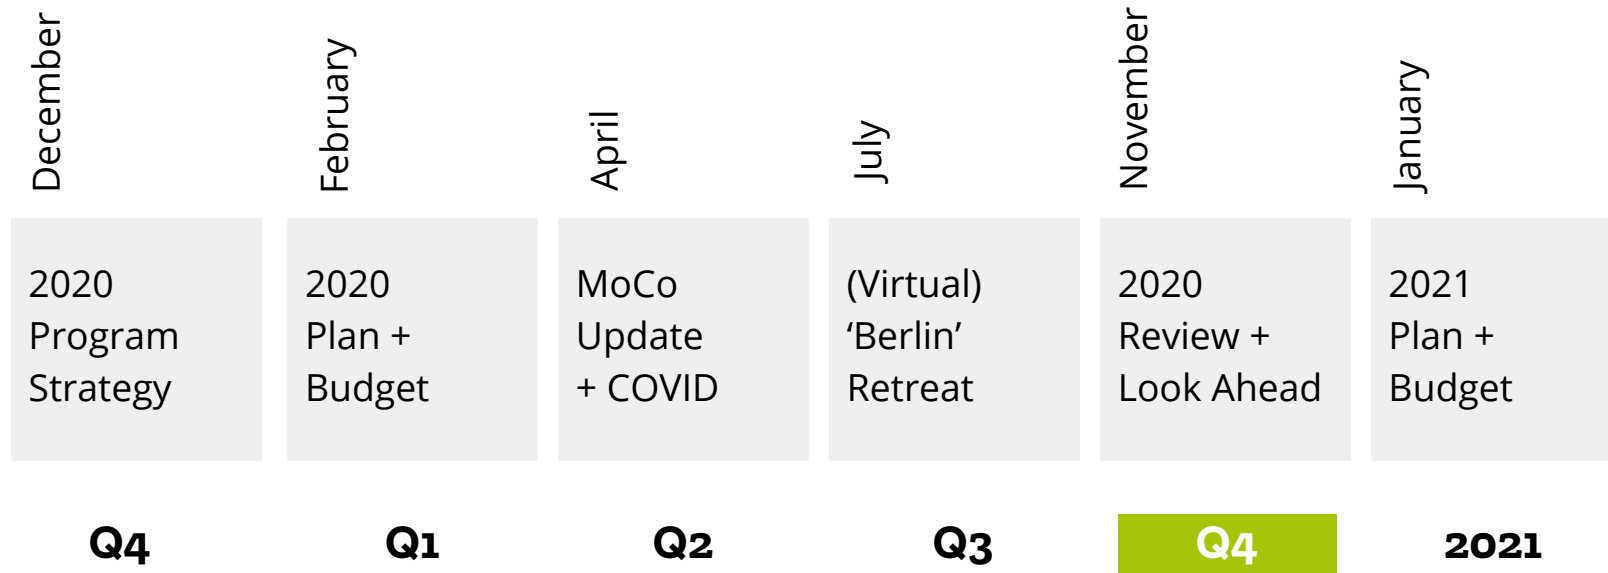


moz://a

Mozilla Foundation Board Meeting

November 18, 2020

2020 board workflow



Agenda

1. Board Business
2. 2020 Financial Update
3. MoFo 2020 Review + Look Ahead
4. Executive session



Board Business

decisions

Board business

1. Approve minutes from July 2020 board meeting.
2. Board recruitment update.
3. Data Futures Lab grant approvals.
4. Thunderbird update.



Thunderbird update

Revenue continued to be strong, with the latest release (78) driving an additional increase in donations. For the first 9 months of the year, Thunderbird contributions were up 51.5% over the same period in 2019.

Main themes for 2020 were: a) **trust, privacy & security** b) **choice & compatibility** and c) **ecosystem health & success**. Looking ahead to 2021 and Thunderbird 90, focus will be on UX and performance improvements, simplified message encryption and unified account setup.

As the MZLA board, our **main concern remains business leadership of the project**. There is a new Council and they have brought on new team members. Priority is to work with them on this issue and ensure long-term sustainability.

update

2020 Finances

Summary

In **July we shared a forecast that reflected changes in response to COVID-19**. At that time, we expected revenue to be down \$0.8M and had adjusted unrestricted spending down by \$1.3M. We were confident in our financial position but there was still a lot of uncertainty in the economy.

In the end, our financial position is better than we projected. Total **revenue is forecast to be 31% more than we forecast in July**. Unrestricted revenue is forecast to be up by 7% due to better than expected MoCo earnings, as well as strong individual giving. We also secured a number of large grants.

Combining this increased revenue and our July plans for reduced spending, **we now forecast an underspend against the approved budget**. Also, we expect to spend only 10% of the \$1M committed spend from capital reserves for work in Africa. We will end the year well above the unrestricted net asset floor in the budget resolution.



MoFo 2020 Program Review + Look Ahead

Our focus *(enduring)*

moz://a

Internet Health Report **2019**

How healthy is the internet?

Our 2019 compilation of research and stories explains what's key to a healthier internet across five issues, from personal experience to global concerns.

MoFo's overall focus remains creating a healthier internet by **supporting a movement of people and allies around the world** who share our cause ...

Our focus *(right now)*

moz://a

Internet Health Report 2019

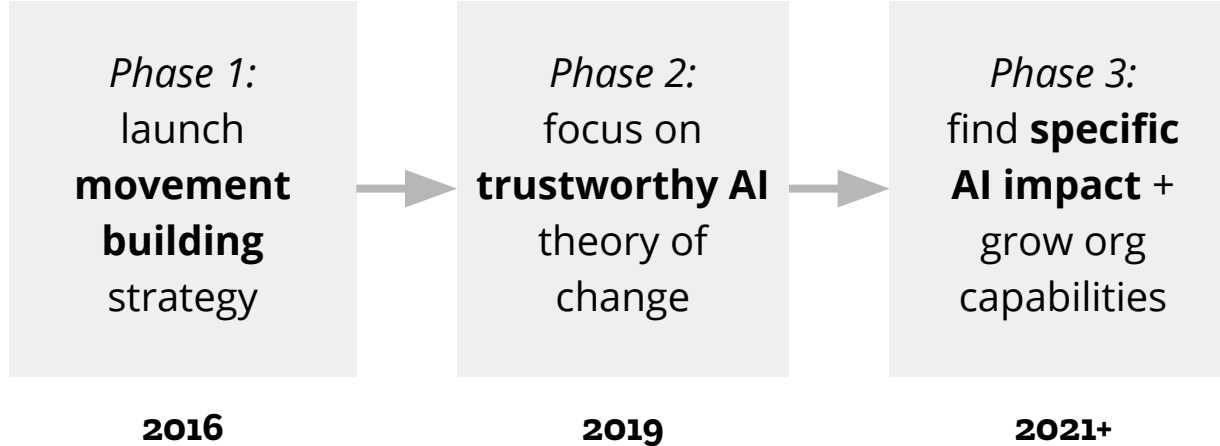
Spotlight: Let's ask more of AI



... and **our trustworthy AI work adds focus and momentum** to our overall internet health agenda.

Where are we?

We are moving into the **next phase of our internet health movement building strategy** ...



AI Theory of Change

Working on in 2020

Could add to our work

Short term outcomes
(1-3 years)

Medium term outcomes
(3-5 years)

Long term outcomes
(5+ years)

Long term impact

Clear "Trustworthy AI" guidelines emerge, leading to new and widely accepted industry norms.

Engineers, product managers, and designers with trustworthy AI training and experience are in high demand across industry.

Diverse stakeholders — including communities and people historically shut out of tech — are involved in the design of AI.

There is increased investment in and procurement of trustworthy AI products, services and technologies.

Shifting industry norms

The people building AI increasingly use trustworthy AI guidelines and technologies in their work.

More foundational trustworthy AI technologies emerge as building blocks for developers.

Transparency is included as a feature in more AI enabled products, services, and technologies.

Entrepreneurs develop — and investors support — alternative business models for consumer tech.

The work of artists and journalists helps people understand, imagine, and critique what trustworthy AI looks like.

Building new tech and products

Trustworthy AI products and services are increasingly embraced by early adopters.

Agency

All AI is designed with personal agency in mind. Privacy, transparency, and human well-being are key considerations.

Trustworthy AI products and services emerge that serve the needs of people and markets previously ignored.

Consumers are increasingly willing and able to choose products critically based on information regarding AI trustworthiness.

Citizens are increasingly willing and able to pressure and hold companies accountable for the trustworthiness of their AI.

A growing number of civil society actors are promoting trustworthy AI as a key part of their work.

Generating demand

Consumers choose trustworthy products when available and demand them when they aren't.

Accountability

Companies are held to account when their AI systems make discriminatory decisions, abuse data, or make people unsafe.

In a world of AI, consumer technology enriches the lives of human beings.

Governments develop the vision, skills, and capacities needed to effectively regulate AI, relying on both new and existing laws.

Progress towards trustworthy AI is made through wider enforcement of existing rules like the GDPR.

Regulators have access to the data they need to scrutinize the trustworthiness of AI in consumer products and services.

Governments develop programs to invest in and incent trustworthy AI.

Creating regulations and incentives

New and existing laws are used to make the AI ecosystem more trustworthy.

2020 in review

It's been **an odd year, to say the least**. A pandemic. A reckoning on racial justice. A deeply worrying US election. There were many moments when it felt like we could go off the rails.

Yet, despite all this -- or maybe because of it -- **our work on internet health and trustworthy AI has felt more relevant than ever**. Work, school, community all moved online. The internet was the terrain of protest and politics. What happened with big tech and AI mattered, a lot.

Impressively, every one of MoFo's teams kept their goals in sight while also adapting their work to meet the moment. With few exceptions, our work maintained or grew momentum this year.

As we go into 2021, we are **teed up to move into the next phase of our trustworthy AI work**. *And, we are also headed into the next phase of the pandemic*. To keep momentum, we need to keep an eye on energy levels and workload, giving people space to go slow when needed.



What we did 2020 (OKRs)

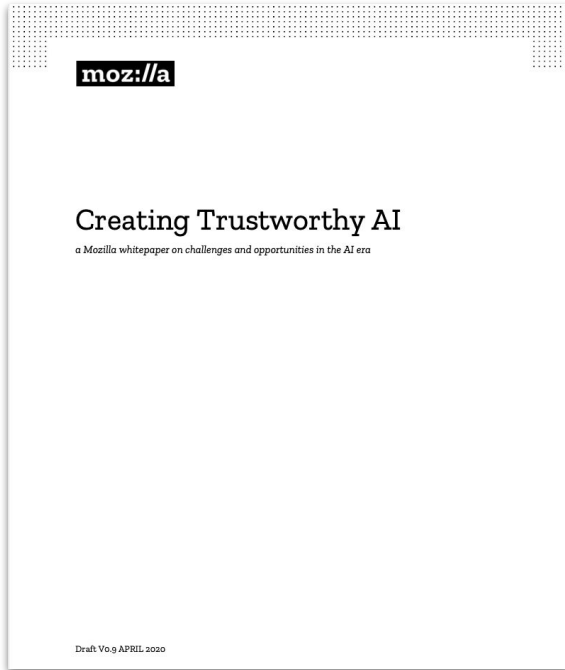
OKR1	Test out our AI theory of change and establish Mozilla as a credible thought leader	Published paper based on theory of change. Honed thinking on where to focus our work.
OKR2	Increase data stewardship innovations that can accelerate the growth of trustworthy AI	Raised ~\$6M for Data Futures Lab. Released initial research and awarded first grants + fellowships.
OKR3	Use pivotal moments to pressure companies to make 'consumer AI' more trustworthy	Used campaigns to make strong connections between AI + pandemic / racial justice / election.
OKR4	Partner with diverse movements at intersection of their primary issues and trustworthy AI	Launched new funding programs connecting our work to orgs in Europe and the global south.
OKR5	Update our organizational capabilities so our ambition can grow over multiple years	Drafted fundraising plan and created financial model for next three years.

** a full list of OKRs including final results is included in Appendix A.*



thought leadership: 2020 review

Big picture goal: clear trustworthy AI guidelines emerge to **drive new and widely accepted industry norms.** *ToC STO 1.1*



Objective: test out our AI theory of change and establish Mozilla as a credible thought leader.

What we did: wrote and honed **Creating Trustworthy AI paper**. Scoped partnerships on shifting industry norms.

Learning: the 'AI guidelines' space is too abstract and broad for Mozilla. The most interesting partners want to work on concrete AI projects. Moving forward, we **will pick an AI theme (transparency) and work with builders, consumers and policymakers on concrete projects** and best practices documentation (not guidelines).

thought leadership: 2020 results

Objective: test out our theory of change in ways that both give momentum to other orgs taking action on trustworthy AI and establish Mozilla as a credible thought leader.

<i>Key result</i>	<i>Notes</i>	<i>Progress</i>
Publish a white paper theory of change	'Creating trustworthy AI' paper's focus on practical action across industry, policy and consumers well received by our collaborators.	Complete
250 people + orgs participate in AI project mapping and offer feedback on white paper	Strong engagement on paper from civil society. Not as much industry engagement as hoped. Mapping focused on projects not people.	147
25 collaborations with partners working on concrete trustworthy AI projects	Many possible collaborators. The best want to work on concrete projects, not guidelines. Will pick narrow, concrete set of projects for 2021+.	35

~~thought leadership~~: 2021 thinking

As per our learning, we will shift from general thought leadership to doing and funding concrete projects on specific AI themes. We will start with 'AI transparency' (plus data stewardship).

Possible objective: work with tech builders, consumers and policymakers to **model what AI transparency looks like in practice.**

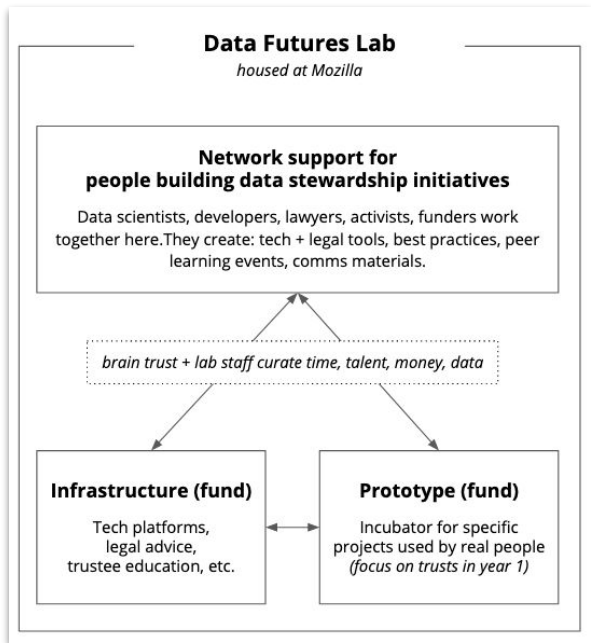
Existing 2021+ commitments: our work on **misinformation + platform transparency** (Regrets Reporter, Ion) will continue with a focus on data informing policy deliberations in the EU (DSA + AI regulation). We have a fellow working on **AI transparency in consumer products.**

Additional 2021+ opportunities: design challenge to demonstrate transparency features in recommendation engines. Collaboration w/ Turing Institute + Consumers International to test consumer-facing transparency features in financial services.

note: notional 3 year arc for this work is included in Appendix A

data stewardship: 2020 review

Big picture goal: more **foundational trustworthy AI technologies emerge as building blocks** for developers. *ToC STO 2.1*



Objective: Increase the number of data stewardship innovations that accelerate the growth of trustworthy AI.

Progress: launched 'proto' version of **Data Futures Lab** with Luminate. Released report on the state of alternative data governance. Reviewed potential for a Mozilla run data trust.

Learning: people are talking about alternative data governance, but there are few examples in practice and none working in concert. It's like open source in the mid-1990s. Mozilla should **invest in practical examples to accelerate the field**. We should also 'walk the talk' in our advocacy and data donation work.

data stewardship: 2020 results

Objective: increase the number of data stewardship innovations that can accelerate the growth of trustworthy AI.

<i>Key result</i>	<i>Notes</i>	<i>Progress</i>
\$3 million raised as an indicator of growing philanthropic support in this area	Data Futures Lab established with core funding from Luminate and Siegel, project funding from Gates plus German and French governments.	~\$6M
10 awards or fellowships for prototypes or other concrete projects re: data stewardship	Fellows acting as initial brain trust for Lab. 'Proto' grants re: data governance in consumer rights, gig economy, activism and voice tech.	8 fellowships 3 grants
4 concentric "networks of practice" utilize Mozilla-housed Data Futures Lab	Landscape analysis helped us understand complexity of the field -- and also helped us build foundations of the Lab network.	4

data stewardship: 2021 thinking

Possible objective: stick with 'increase the number of **data stewardship innovations** that can accelerate the growth of trustworthy AI' as objective. Develop new key results for 2021.

Existing 2021 commitments: we have **set up the Data Futures Lab** to network practitioners plus fund prototypes and infrastructure in this space. An initial cohort of grantees and fellows will do work throughout 2021. We will convene a '**state of the union**' on **data governance** at MozFest. Also, we will bring the Mozilla Common Voice project under the Lab umbrella.

Additional 2021+ opportunities: we could **model good data stewardship at Mozilla**, developing new and better practices to underpin our advocacy and data donation work. We could also build a slate of **policy research** and work related to advancing this topic.

note: notional 3 year arc for this work is included in Appendix A

consumer power: 2020 review

Big picture goal: **consumers are increasingly willing and able to pressure and hold companies accountable** for the trustworthiness of their AI. *ToC STO 3.3*



Objective: use pivotal moments to pressure companies to make 'consumer AI' more trustworthy.

Progress: better than expected engagement and fundraising driven by **campaigns that connected AI to timely topics like the pandemic, elections and racial justice issues.**

Learning: like misinfo and privacy, issues related to **bias and AI** resonate strongly with the public. However, the lack of concrete ways to address bias means there is no real consumer demand. In 2021, we should work with developers and other social movements to fix this.

consumer power: 2020 results

Objective: mobilize an influential consumer audience using pivotal moments to pressure companies to make 'consumer AI' more trustworthy.

Key result

Notes

Progress

3m page views to content on Mozilla channels, a majority of which focuses on trustworthy AI

Consistently created content responding to big events. E.g. Privacy Not Included for video apps as pandemic started. This approach drove traffic.

2.6M

75k new subscribers drawn from sources (partnerships, contextual advertising, etc.) oriented towards people ages 18-35

Used this same content, plus new features like Mozilla News Beat, to reach out to new and non-tech audiences on platforms like Instagram.

65K

25k people share information (stories, browsing data, etc.) in order to gather evidence about how AI works and what changes are needed

Developed the YouTube Regrets Reporter browser extension to get public involved in 'watching' what tech platforms are doing.

22K
(Regrets Reporter)

~~consumer power~~: 2021 thinking

Big 2020 learning: we need to engage the public on AI issues people can relate to. This includes misinfo, privacy -- and AI bias. We should **add an exploratory theme on AI bias in 2021**.

Possible objective: build bridges between AI developers + civil society to **grow our collective ability to tackle bias in AI**.

Existing 2021 commitments: **grow our relationships with the AJL Bias Bounty project** (two fellows already). Deepen our connection to the MoCo Applied ML team + their *PRESC* project. Promote CMAs that focus on Black artistic and technological exploration of AI.

Additional 2021+ opportunities: as we did with data stewardship, we should undertake **landscape research on concrete projects tackling bias**, as well as what areas of bias are of greatest concern for our civil society partners. Also explore whether data donation platforms like Ion could be used to uncover bias cross-platform and at scale.

note: notional 3 year arc for this work is included in Appendix A

movement building: 2020 review

Big picture goal: a **growing number of civil society orgs are promoting trustworthy AI** as a key part of their work. *ToC STO 3.4*



Objective: partner with diverse movements at intersection of their primary issues and trustworthy AI.

Progress: documented extensive movement building work across MoCo. Also, launched **Tech and Society Fellows** (6 countries) and **European AI Fund** (€4.6M for intersectional work on AI) despite lack of travel.

Learning: our partnerships are most impactful -- and *received as partnership* -- the greater the involvement across teams. We are respected but asked to bring greater clarity and focus. **Future = clear, long term collaborations.**

movement building: 2020 results

Objective: partner with diverse movements at intersections between their primary issues and internet health, including trustworthy AI, so that we increase shared purpose.

Key result

Notes

Progress

30% increase in partners with whom we have both created something that includes shared language, methodologies, resources, or events

MoFo is a strong collaborator. This was the first year we document who we work with. We've already uncovered collaborations w/ ~200 orgs.

22.%(IHR + MozFest data still coming)

75% of partners from these diverse movements report deepening intersection between their issues and internet health/AI

We had several collaborations with unexpected partners (e.g NBC + CogX) that led to a deepened and more nuanced exploration of our work.

71%

4 new partners in Global South report deeper intersection between their work and ours

While it's too early to measure impact, we did launch Tech and Society Fellows with Ford plus civil society partners in 6 countries.

0

movement building: 2021 thinking

Possible objective: continue to, “**partner with diverse movements to deepen intersections between their primary issues and ours**” while honing our strategies for doing so.


Existing 2021 commitments: **Tech & Society** cohort has first full year. Convene Ford’s regional program officers @ MozFest to explore expanding into new countries and deepening in current geographies. Complete movement landscape analysis. Racial justice commitments continue. Operationalize MoFo **Africa Mradi** workstream, moving forward fully w/ 2 year / \$2M commitment made in 2020.

Additional 2021+ opportunities: in addition to our racial justice commitments, **identify key movements with which we have most synergy**, per the landscape analysis. Commit to improving our movement building capabilities, with interoperability across teams. We will choose specific movements to partner with, and have points of contact in each of our teams.

note: notional 3 year arc for this work is included in Appendix A

org effectiveness: 2020 review

Big picture goal: ensure we have an org structure and systems that match our ambition to **be a sustainable, collaborative and effective player in growing the internet health movement.**



Objective: update our organizational models and capabilities so that our strategy and people can succeed, and our ambition can grow over multiple years.

Progress: laid important groundwork through projects like the **job description audit, financial scenario planning, CRM review**, and Core Capacity Assessment Tool (CCAT), building on infrastructure put in place over last few years.

Learning: must put focus on better communication efforts internally. Org effectiveness measures must be integrated into program planning and require buy-in across org.

org effectiveness: 2020 results

Objective: update our organizational models and capabilities so that our strategy and people can succeed, and our ambition can grow over multiple years.

Key result

Notes

Progress

Organizational effectiveness baseline established

Ran org assessment process. Strengths = vision, org learning, financial mgmt. Weaknesses = succession planning, adaptability, tech capacity.

Complete

A long-term funding model that matches our ambitions exists

Developed new budget model to support long range planning. Also, developing formal three year fundraising plan.

First draft

100% of job descriptions have been audited and are up-to-date

Undertaking job description audit as step towards understanding the skill sets we have -- and need -- in order to meet our goals.

Phase 1 complete

org effectiveness: 2021 thinking

Possible objective: continue to update our organizational models and capabilities in order to drive a shift toward more data-driven decision-making.

Existing 2021 commitments: **Roll out the fundraising strategy** created in 2020. Create and deploy MoFo-wide data, impact and measurement strategy. Transition MoFo to Salesforce.org and new email marketing platform. Identify and implement measures to **address priority areas identified in CCAT results**. Create career development framework and succession plan.

Additional 2021+ opportunities: **Model better data stewardship internally** through work we are doing on the CRM and other initiatives.

note: notional 3 year arc for this work is included in Appendix A

discussion: possible 2021+ objectives

1. **Increase the number of data stewardship innovations** to accelerate the growth of trustworthy AI.
2. Work with tech builders, consumers and policymakers to **model what AI transparency looks like in practice.**
3. Build bridges between AI developers + civil society to **grow our collective ability to tackle bias and discrimination in AI.**
4. **Partner with diverse movements** to deepen intersections between their primary issues and ours, including trustworthy AI.
5. Continue to update our organizational models and capabilities in order to **drive a shift toward more data-driven decision-making.**

Appendix A

Trustworthy AI Narrative Arcs

3 year arc: making AI transparency the norm

2020

Public + policymakers increasingly concerned about misinfo + lack of platform accountability

Mozilla develops citizen-driven research tools to study the platforms (Regrets Reporter + Ion)

Mozilla disseminates model policies for transparency - e.g. bulk ad disclosure

2021

Mozilla + partners develop and test transparency features in consumer tech

Greatest gaps in platform transparency identified and studied through using Ion along w/ civil society partners

Mozilla develops model policies for transparency of AI-enabled systems; first focus on EU

2022

Consumers pressure tech to integrate proven transparency features via data donations, campaigns, + PNI

Additional transparency features developed + tested

Researchers + civil society collaborations shape policy agenda; policies mandating transparency gain traction

2023

High-use consumer tech tools have robust explainability features for end consumers

Differences (or gaps) in tech transparency features rated in PNI

Professional networks develop transparency resources for AI builders

Platform regulations include AI transparency mandates

STO 2.2: Transparency is included as a feature in more AI enabled products, services, and technologies.

STO 1.1 Best practices emerge in key areas of trustworthy AI, driving changes to industry norms.

STO 4.3: Regulators have access to the data and expertise they need to scrutinize the trustworthiness of AI in consumer products

3 year arc: modeling good data stewardship

2020

Alternative data stewardship analysis scopes work & establishes us w/in networks of practice

Initial cohort of people with data stewardship innovations help us build Data Futures Lab

Fellows and staff identify areas where Mozilla can model “good” data stewardship

2021

Data stewardship prototype projects up and running

Initial infrastructure projects seed work in diverse communities and geographies.

Mozilla/DFL convene networks to share & learn

Early experiments in leveraging CCPA to create data rights collectives

2022

The most successful prototypes have constituency level impact; attract users, press, attention

New infrastructure and building blocks begin to emerge

Policy makers endorse the idea of data rights collectives, Mozilla and others step into this space

2023

The ‘penny drops’ as people start building on infrastructure and prototypes we seeded - including Mozilla, potentially.

Data collectives grow membership, drive companies to improve products and services

Data stewardship innovations fuel growth of trustworthy AI

STO 2.1: more foundational trustworthy AI tools emerge as building blocks for developers.

STO 3.2: Consumers are increasingly willing and able to choose products critically based on information regarding AI trustworthiness.

STO 4.1: governments develop the vision, skills, and capacities needed to effectively regulate AI.

Focus: rooting out -- and fixing -- biased AI

2020

Mozilla tests resonance of AI bias issue through work on Amazon Ring + RegretsReporter

Initial investments in bias, race and AI made via fellowships (AJL), awards (CMA) + MoCo ML team

2021

Mozilla and civil rights groups develop wide ranging platform policy agenda that includes bias as a key issue

New tools emerge to help developers detect bias; focus on data sets

Groundbreaking art explores bias, imagines different kinds of AI

2022

Early bias detection tools begin to be used in industry to assess and correct bias

Consumers increasingly consider bias as factor in choosing AI-enabled products - PNI, art + movement partners drive public agenda

Policies requiring use of debias tools/processes gain traction

2023

Narratives shift, the public is keenly aware and watching for AI bias as a result of art + mov't partnerships

Accountability mechanisms re bias AI a feature of laws and company policies

Developers have easy access to tools to root out and fix bias in AI, they routinely use them

STO 1.3: Diverse stakeholders, including people historically shut out of tech, involved in design of AI.

STO 2.4: artists and journalists help people understand, imagine, and critique trustworthy AI.

STO 3.4. growing number of civil society actors promote trustworthy AI as part of their work.

3 year arc: growing across movements *racial justice example*

2020

Established pan-mozilla racial justice commitments (personal, institutional, programmatic)

Launched mvmt landscape analysis & F&A evals.

CMA's focus on black artistic and technological exploration of AI. (unpack multi pronged approach to walk thru trajectory)

HBCU engagement pan-mozilla: 3 school MoCo investment, Spellman @ MozFest; 4xHBCU @ RCS

2 x FIR = Bias Bug Bounty

IHR spotlight: Race & Software

2021

Mozilla commits to work RJ cross movement thru 2023

Cohort of civil society grantees wrestling w data sovereignty, mis/disinfo, bias in ML, disenfranchisement thru connectivity & democracy mobilization in communities of color, funded.

Highlander center is anchor host org for fellows. Host org cohort chosen. HC host organizer training for Fellows.

CMA showcase @ MozFest D&D focus. Keynote.

2022

Indigenous Data Trust prototyped that transcends U.S. Canadian border.

Fellows, DFL Lead & grantee partners center affected communities & contribute to US, EU and Brazilian encryption policy recommendations.

U.S. grantee partners & host orgs co-launch 2x PNI w concurrent US policy play. Includes tech co. pressure, best practice tech launch w MoCo, Policy landscape analysis launch.

2023

Movement building methodology & commitments socialized & implemented across all OKRs & team strategies.

The strategies, calls to action and constituents of our movement and partner movements overlap in greater collective purpose

Movement building methodology & commitments socialized & implemented across all OKRs & team strategies.

Significant increase in number and make up of core constituency.

moz://a

Mozilla Foundation Board Meeting

November 18, 2020