Impact Summary: Cool/Scary AI Impact Forum

San Francisco • November 4, 2025

Hello,

Thank you for joining us at the Cool/Scary AI Impact Forum in San Francisco—or for following along from a distance. This report distills the day's conversations so the ideas are easy to share with your team, your board, and your community.

San Francisco sits at the geographic center of the AI industry, but the questions we explored were not the ones typically asked in this city. Instead of "what can AI build?" our speakers asked: Who gets left behind? Whose bodies bear the cost? Whose voice is missing from the design table? And what would it look like to reclaim agency—as communities, as organizations, as a sector—over how these tools shape our lives?

Across the day, five themes emerged. First, AI's most urgent harms are deeply human—from deepfakes weaponized against women in conflict zones to data labeling pipelines that encode racial bias at the infrastructure level. Second, the digital divide has evolved into something more fundamental: a question of whether people can make real choices in a society being restructured around AI without their input. Third, the social sector has unique structural leverage through coalition models, shared AI workspaces, and participatory design. Fourth, burnout is the operating context for any honest conversation about AI adoption—and a generational opportunity to change how we work, not just what we produce. Fifth, human dialogue is irreplaceable; LLMs exacerbate misinformation, flatten pluralism, and simulate care without delivering it.

How to use this packet: share the session summaries with colleagues, identify one conversation to continue in your organization, and consider what structural change—not just tool adoption—you want to champion this year.

Thanks for showing up, asking hard questions, and building this community with us.

Onward,

Eric Rubin and Ben Childers

Cool/Scary AI Impact Forum Organizing Team

Main Takeaway

AI's most consequential impacts are not technical—they are human, structural, and political. The social sector is uniquely positioned to respond, but only if organizations move beyond tool adoption to address the systemic conditions that shape who benefits from AI and who is harmed by it. Start with the people: the burned-out staff, the digitally invisible communities, the women silenced by deepfakes, the data labelers encoding bias at scale. Build coalitions, embed participatory methods, preserve human dialogue, and use the time AI frees up to rest and create—not just to produce more.

Session Highlights

Eric Rubin (Co-Founder, Stratovation Partners) — AI in Civic-Sector Benefits Enrollment

Eric opened with a case study on See the Stars, a client using direct-response science—typically reserved for fundraising—to close the gap between eligible seniors and the social services they qualify for. Working with a major elder services nonprofit, they apply lead generation, triggered messaging, and predictive analytics to benefits enrollment. The AI opportunity: a single intake site that reviews state-specific eligibility rules, multilingual chatbots that handle the first layer of navigation, and predictive models that identify who is at risk of losing coverage before they churn. Eric framed Medicaid churn as a deliberate policy choice, not a systems failure, and argued that the same data infrastructure nonprofits use for donor acquisition can be repurposed to keep people connected to the safety net. Key tension: trust. People need to believe the platform handling their data is safe, which is why partnerships with trusted intermediaries like AARP Foundation matter.

Jessica Olney (Advisor, PEMA) — AI, Deepfakes, and Democratic Integrity

Jess brought the international human rights lens, drawing on her experience in Myanmar, Sudan, and Yemen. She surveyed the state of AI policymaking in the U.S.—thousands of state-level bills since 2023, most focused on decision-making regulation and synthetic media—and argued that California's transparency bills (SB 942 and AB 853) may become seatbelt-level norms. The scarier material: 95–99% of deepfakes target women, and some tools only work on female bodies. In conflict settings, deepfakes are being used to intimidate women out of public life, with cascading effects on peace processes that depend on whole-of-society participation. Jess challenged the room to stop transposing Western assumptions about media literacy onto other contexts and to recognize that the AI-generated harms happening in places like Sudan—where internet blackouts create information vacuums that disinformation fills—look nothing like the catastrophic risk scenarios discussed in Silicon Valley.

Ben Childers (Founder, Cool/Scary AI) — Digitally Invisible

Drawing on Dr. Nicole Turner Lee's book Digitally Invisible and Amartya Sen's capabilities framework, Ben argued for a paradigm shift from the "digital divide"—which is input-oriented (do people have hardware, software, connectivity?)—to the "digitally invisible," which is output-oriented (can people make real choices and fully participate in society?). He grounded this in a story of a day laborer who cannot work for the last ten days of every month because his prepaid minutes run out. With billions still lacking reliable internet access, the risk is that AI's societal restructuring—driven by a handful of companies—creates a permanent underclass with no agency over the systems being imposed on them. The call: think about structural solutions that expand real freedom, not just access metrics.

Aram Fisher (Co-Founder & CEO, ChangeAgent) — Coalition AI

Aram made the case that Big Tech wants users atomized—isolated individuals generating unit economics—while the social sector thrives on collaboration. His platform, ChangeAgent, builds shared AI workspaces for coalitions and networks: private instances with custom-tuned models built on a coalition's own narrative framework, shared knowledge bases, internal communication channels, and tools like natural-language SQL querying that let non-technical staff access reporting databases without bottlenecking data teams. The New Mexico state table example showed how a progressive advocacy coalition could share talking points, legislative intelligence, and communications materials through an opt-in system that respects organizational autonomy while amplifying collective power. Key insight: when a model is tailored to your mission, you don't need perfect prompts to get useful output—lowering the adoption barrier dramatically.

Amanda Bittaker (Chief of Staff, Product & Technology, Wikimedia Foundation) — Thoughtful AI Integration

Amanda spoke from the unique vantage of the Wikimedia Foundation, where 200,000–300,000 monthly editors curate the world's largest open knowledge base. She laid out the scary-then-cool arc: AI hallucinates (fabricated ISBNs increased roughly 15% after ChatGPT launched), it can be broken by adversarial prompting, and it makes disinformation cheaper to produce than to catch. The cool: AI dramatically reduces editorial toil (the Add-a-Link feature cut a multi-minute task to seconds), and a game jam in Brooklyn showed that 50 non-coders could vibe-code 20 playable games in a single day, generating strategic insights about how young people want to interact with knowledge. Her core framework: keep a human in the loop at every stage, tag all AI-assisted content, and use AI to reduce toil and increase participation—not to replace judgment.

Cassie Grudstein (Plus Impact Partners) — AI in the Age of Burnout

Cassie reframed the AI conversation through the lens of chronic occupational burnout—95% of nonprofit leaders report concern about staff burnout, and 60–85% of knowledge workers identify as burned out at any given moment. She shared a client example: a nonprofit executive used AI to produce a board report in 20 minutes instead of 90 and a $75M pitch deck in 15 minutes instead of two hours. The critical move: don't backfill the saved time with more work. Cassie traced the Protestant work ethic's hold on productivity culture, noted that both John Maynard Keynes (1930) and Bill Gates (today) promised technology would reduce work hours—and neither delivered—and argued that AI presents a generational opportunity to actually change how we work. Her prescription was somatic and structural: close stress cycles through embodied practice, change organizational policies (sabbaticals, four-day weeks), and use the sector's legacy of labor advocacy to push for systemic change.

Emlyn Metz (Lawrence Hall of Science, UC Berkeley) — Preserving Communication and Critical Thinking

Emlyn brought the cognitive science perspective on why LLMs are dangerous for epistemology. Their motives are opaque (shaped by training data biases, developer choices, and user prompts interacting unpredictably), they pretend to care (driving some users to prefer AI companionship over human connection), and they don't learn from conversation to take elsewhere—undermining the progressive education model where students co-construct understanding. She demonstrated how LLMs exacerbate every vector misinformation exploits: confirmation bias (AI helps you dismiss counter-evidence), content overload, polarization, and overconfidence without uncertainty signaling. Her takeaways: check AI output for bias and accuracy, clearly label AI-generated content, distinguish facts from values, practice actively open-minded thinking, and above all—keep talking to other humans.

Annie Brown (CEO, Reliable; AI Bias Researcher, UC San Diego) — Bias in the Data Labeling Pipeline

Annie delivered the most technically grounded talk of the day, revealing how bias enters AI systems at the data annotation stage—before models are even trained. When images are labeled by underpaid overseas workers following instructions from data scientists, both parties' cultural biases fuse into the training data. Result: Midjourney, DALL-E, and Stable Diffusion each produce radically different—and uniformly biased—representations of beauty, with dark skin tones represented at just 18% at best. Microsoft's image classifier rated a woman doing yoga as "racy" with 98% confidence while a man in a similar position scored 3%. Pregnant Black women were significantly more likely to be flagged as "adult" content. Annie's solution:

participatory AI—embedding people with lived experience into the labeling process—which in her work improved model accuracy from 67% to 85% F1 scores while simultaneously reducing bias. The business case and the equity case are the same.

Cross-Cutting Themes

1. The harms are human, not hypothetical

Across speakers, a consistent thread: the most urgent AI risks are not speculative existential scenarios but concrete, present-tense damage to real people. Deepfakes silencing women in conflict zones. Biased classifiers flagging Black women's bodies as explicit. Day laborers locked out of work by connectivity gaps. Medicaid recipients churned off coverage by deliberate policy complexity. The sector's role is not to imagine catastrophe but to respond to what is already happening.

2. From access to agency

Multiple speakers converged on a shift from input metrics (do people have the tools?) to output metrics (can people exercise real choices?). Ben's digitally invisible framework, Aram's coalition model prioritizing opt-in autonomy, and Cassie's insistence on using AI-freed time for rest and creativity rather than more production all point in the same direction: the goal is not adoption but self-determination.

3. Participatory methods are a structural advantage

Annie's participatory data labeling, Aram's coalition AI workspace, Amanda's human-in-the-loop editorial model, and the community game jam all demonstrated the same principle: involving the people affected by AI systems in their design improves both accuracy and trust. This is not a nice-to-have. It is a competitive and ethical advantage that the social sector is uniquely positioned to leverage.

4. Burnout is the operating context

Any AI strategy that ignores the exhaustion of the people who would implement it is incomplete. Cassie's data on nonprofit burnout rates—and her historical argument that every generation's technological promise to reduce work has been absorbed by productivity culture—framed the entire event. The opportunity: use AI to change how we work (reduce toil, create space for creativity and rest), not just to scale existing patterns of overwork.

5. Human dialogue is the foundation

Emlyn's cognitive science framework and the event's own participatory format reinforced that LLMs cannot substitute for human conversation, critical thinking, or the co-construction of understanding. AI collapses pluralism, simulates objectivity, and rewards convenience over reflection. The antidote is not less technology but more deliberate human engagement—which is exactly what events like this one exist to provide.

Productive Tensions Surfaced

Convenience vs. critical thinking: AI makes it easier to produce content, but harder to evaluate it. The friction of human writing and dialogue is a feature, not a bug.

Trust vs. speed: Benefits enrollment, coalition data sharing, and AI-assisted editorial all require user trust—which is slow to build and easy to destroy.

Regulation as seatbelt vs. regulation as patchwork: State-level experimentation is generating useful models, but the absence of federal frameworks risks compliance chaos.

Atomized users vs. networked power: Big Tech's business model profits from isolated individuals. The social sector's coalition DNA is a direct counter—but requires intentional infrastructure to sustain.

Community Feedback

Likelihood to attend again: 4.3 out of 5

Two-thirds of respondents gave the maximum score (5/5), signaling strong intent to return.

What attendees found most valuable

"All the presentations were top-notch and I really loved that the event was small enough to get to interact with everyone."

"Getting a sense of how the nonprofit sector is thinking about and working with AI."

"Meeting other attendees" and "Networking" were recurring themes—the peer connections were as valuable as the stage content.

Sessions that resonated most

Aram Fisher's Coalition AI session, Cassie Grudstein's Age of Burnout talk, and Emlyn Metz's epistemology presentation were each singled out by respondents. One attendee praised the "diversity of approach" across the full program.

What we heard about improving

Attendees asked for more actionable and technical content alongside the structural and societal framing—practical guidance like "what a good AI policy looks like for your nonprofit" and "use this tool in this case, but not that one." Other requests: structured small-group breakouts with opt-in topics, recording or livestreaming sessions, earlier planning timelines, a shared Linktree for speaker follow-up, and fresh catering options.

Topics requested for future events

Practical AI policy templates for nonprofits, continued diversity of speakers and approaches, and deeper dives into specific use-case guidance. One speaker offered to return on AI and burnout—a signal that the content is generating ongoing commitment from the community.

Working group interest

One-third of respondents said "Yes" to joining the working group for a progressive AI future, and another half said "Maybe — send me more info." Strong signal of engagement beyond a single event.

Resources & Next Steps

Books and frameworks referenced by speakers:

- Digitally Invisible by Dr. Nicole Turner Lee (Brookings)
- Development as Freedom by Amartya Sen
- Burnout: The Secret to Unlocking the Stress Cycle by Emily and Amelia Nagoski
- Rest Is Resistance by Tricia Hersey
- The Happy Healthy Nonprofit by Beth Kanter and Aliza Sherman
- The Starfish and the Spider by Ori Brafman and Rod Beckstrom
- What It Takes to Heal by Prentice Hemphill

Policy and research references:

- California AI Transparency Act (SB 942) and extension (AB 853)
- NCSL AI Legislation Tracker (state-level AI bill database)
- Transparency Coalition (model AI legislation)
- Wikimedia Foundation AI principles (QR code shared at event)

Tools and platforms mentioned:

- ChangeAgent (changeagent.ai) — social-sector AI platform with coalition features
- Reliable (reliable.ai) — participatory AI data labeling platform
- Unite Us — coordinated care and closed-loop referral infrastructure

Help shape the future we actually want.

Email ben@coolslashscary.ai to get involved.

Thank you to our sponsors!

Corner Shop Creative • Advocacy AI • Beekeeper Group

Hosted at Google Community Space, San Francisco