

Impact Summary: Cool/Scary AI Sh!t Brooklyn Summit 2025

Hello,

Thank you for being part of our Brooklyn gathering of Cool/Scary AI Sh!t—or for cheering us on from nearby. We're sending you the session summaries and a short synthesis to make the ideas easy to share across your team. Our goal is simple: help mission-driven organizations turn curiosity into responsible, measurable action.

Across the day, four themes stood out. First, culture beats tools. The organizations making real progress pair leadership sponsorship with enterprise literacy, clear policies, and psychologically safe spaces to experiment. Second, data foundations and guardrails matter. Start small, but make your integrations explicit and your ethics, security, and privacy visible. Third, practical automation wins. Unsexy use cases—email triage, onboarding flows, meeting-to-action handoffs, and lightweight internal apps—free capacity you can reinvest in higher-value work. Fourth, experience drives ROI. Borrow proven e-commerce playbooks to remove friction, then use AI plus classic data science to personalize thoughtfully and improve conversion without compromising values.

We also surfaced useful tensions: shipping fast vs. aligning the enterprise, “good enough” data vs. prerequisites, agent hype vs. durable systems, and near-term lift vs. long-term trust. Naming these contradictions helps teams choose deliberately rather than by default.

How you can use this packet: share the 50–60-word summaries with colleagues, pick one pilot you can launch in the next two weeks, and adopt a lightweight “decision log” to document what you tried, what changed, and why. If you want support—facilitation, governance frameworks, or pilot design—reply and we'll connect.

Thanks again for showing up, sharing candidly, and building this community with us. We look forward to seeing what you ship next.

Onward,

Ben Childers and Partners

Cool/Scary AI Sh!t Organizing Team



Summary

Across eight sessions, speakers converged on a pragmatic, culture-first approach to AI in the social-impact sector. They showed quick wins from “unsexy” automation (email triage, SMS onboarding, meeting→PM actions, topic-page blurbs) and rapid prototyping (“vibe-coded” apps), while emphasizing foundations: governance, ethics, data security, and staff readiness. Performance gains come from reducing donor friction and pairing LLMs with classic data science to find real content drivers. Tension points surfaced—speed vs. alignment, imperfect data vs. prerequisites, flashy agents vs. durable systems, short-term CRO vs. mission-aligned virtues—pushing organizations to blend experimentation with guardrails.

Treat AI as an organizational capability, not a tool. Start with shared literacy and governance, define clear guardrails (ethics/data/security/privacy), and run small, well-scoped pilots that you measure and iterate. Connect systems early, publish what changes and why, and prioritize supporter experience—mobile-first, modern payments, personalization—guided by adaptive ethics. Blend LLMs with data science for insights and automated QA. Move fast on prototypes, but earn scale through trust, alignment, and documented learning.

Main Takeaway



Contradictions:

Ship fast vs. go slow

Peskey/Snyder “vibe-code” live apps in minutes and Sacdalan automates 90-day SMS flows—prioritizing momentum and quick wins. Trzop and Kimber/Pogoff argue for deliberate pacing: governance first, policy as a “living garden,” and cautious, enterprise alignment. Perez shows the reality: crises force speed, but approvals stretch months. The tension: prototypes and automations can be spun up instantly, yet sustainable adoption demands slower cross-functional buy-in.

Good enough” data vs. data prerequisites

Kimber/Pogoff and Trzop caution against perfectionism—start small, don’t wait for pristine hygiene. Alani insists the upside appears only after platform integration, consistent classification, and modeling; Peskey warns productionizing vibe-coded tools without secure data practices is risky. Contradiction: pilot now with imperfect data vs. pause to integrate and harden the stack so insights and security are trustworthy.

Agents disappoint vs. no-code apps impress

Alani says LLM-only agents underperform without classic data science; results are mixed. Peskey/Snyder wow with rapid, no-code application building and feature-rich demos. Developers may later reject AI-generated code as unmaintainable—another friction point. The clash: “agents aren’t ready” versus “AI just built our working tool,” raising questions about longevity, handoff, and technical debt.

Maximize conversion vs. uphold mission virtues

Salpietro pushes e-commerce-style personalization (abandonment prompts, fee coverage tuning, aggressive CRO) to lift revenue. Childers’ adaptive-ethics model would, for example, reject suppressing low-income ZIP codes despite better ROAS. Alani adds that emotions like anger may mobilize actions but harm donor giving or brand. The contradiction: optimize ruthlessly for short-term lift vs. constrain tactics to protect equity, dignity, and long-term trust.



Patterns:

Culture-first adoption

(leaders, literacy, governance)

Across talks, culture beats tools. Successful orgs pair active executive sponsorship with cross-functional AI governance, enterprise “AI 101” literacy, and psychologically safe play spaces that include skeptics. Teams use a shared E/D/S/P lens (ethics, data, security, privacy), treat policy as a living “garden,” and document decisions.

The goal isn’t faster copy; it’s trust, alignment, and clarity that scales change. (Pam Trzop; Rachel Kimber & Samin Pogoff; Ben Childers; echoed by Jamie Perez’s culture wins.)

Data foundations + guardrails before scale

Everyone stressed: connect systems, don’t chase perfect hygiene, and wrap AI in explicit guardrails. Adaptive ethics (virtues + review) informs boundaries; security/privacy drive choices like stateless prototypes and avoiding sensitive data in vibe-coded apps. On the performance side, classify unstructured content, add QA loops, and monitor drift. The shared stance: integrate first, pilot safely, and publish changes as you learn. (Ben Childers; Joshua Peskey/Kim Snyder; Pam Trzop; Tareq Alani; Maan Sacdalan.)

Practical automation > flashy demos

The biggest wins came from unsexy work: email triage and routing, 90-day SMS onboarding, meeting-to-PM actions, topic-page “toppers,” and lightweight agenda apps. Speakers advocated small pilots with feedback loops, decision logs, and monitoring checkpoints to earn the next investment. Analytics become triggers, not reports; publish what changed and why. Teams accept timeline friction and iterate. Result: compound time savings without over-automation. (Maan Sacdalan; Jamie Perez; Joshua Peskey/Kim Snyder; reinforced by Rachel & Samin’s pilot-and-measure approach.)

Experience and ROI through personalization + data science

A consistent refrain: optimize supporter experience to drive measurable lift. Borrow e-commerce playbooks—address autocomplete, modern payments, abandoned-donation reminders—and layer AI-driven personalization (including fee-coverage tuning). Move beyond LLM-only “agents” by pairing LLMs with classic modeling to find true content drivers, then test where gains could be 10–20%+ rather than 0.5%. Prioritize mobile, reduce friction, and privilege strategy over tactical tweaks. (Sal Salpietro on CX/CRO; Tareq Alani on LLM + data science; echoed by Maan Sacdalan.)

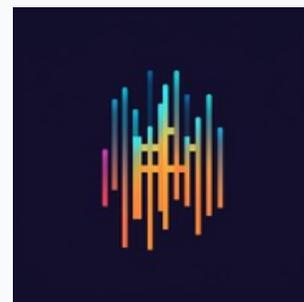


Session Highlights

The sessions at Cool AI Sh*t Brooklyn offered a wide-ranging view into how AI is reshaping work across the nonprofit, creative, and advocacy sectors.

From deeply technical insights to community-driven storytelling, each speaker illuminated a different facet of what it means to engage with AI responsibly and creatively. What follows are brief overviews of each session, spotlighting the practical tools, tensions, and takeaways that emerged. These highlights are meant to serve as a starting point for further exploration, reflection, and cross-sector collaboration.

Tareq Alani explains why LLM-only “agents” underperform and shows how pairing LLMs with classic data science unlocks real marketing gains. He outlines a pipeline—integrate platforms, classify unstructured content via vectors, model drivers (regressions/forests), and continuously optimize with automated QA and a ‘coaching bank.’ Early pilots show 10–20% engagement lifts, \$100–\$1,000 per 1k emails, and 50% content-production time saved. [View Slides](#)



Jamie Perez (Center for American Progress) shares a pragmatic GenAI project: auto-writing “topper” blurbs for 177 website topic pages amid nonstop crisis work. GenAI fakes ML to learn what works, drafts 115; editors deem 40 viable, 14 published—break-even so far, results pending. The deeper win is culture: treating analytics as triggers, building trust in AI-assisted text, and iterating recursively. [View Slides](#)



Session Highlights

Pam Trzop (Firefly Partners) argues AI isn't just a tool; it's organization-wide process augmentation. Adoption succeeds when leadership actively sponsors change, cross-functional workgroups set policy (e.g., transcripts/summaries), data ethics and integration are clear, and teams are trained with trust and governance.

ROI = alignment, clarity, and better supporter experiences—not just faster content. Start small; avoid perfectionism. [View Slides](#)



Joshua Peskey (with Kim Snyder) demos “vibe coding”— using Claude Code and IDEs to spin up a board-meeting agenda app from audience prompts (“Goats of Anarchy”). He outlines best uses (prototypes, internal dashboards, stateless tools) and risks (security, sensitive data, production handoff). A prebuilt Kitty Crochet demo shows guardrails and features: branding, drag-and-drop, PDF export, GitHub repo, Netlify hosting. [View Slides](#)

MEET the MOMENT
RESILIENCE FOR NONPROFITS



Session Highlights

Ben Childers urges a shift from bureaucratic checklists to adaptive, virtue-based ethics for digital transformation.

His loop: anticipate context, set guardrails, pilot small experiments, learn and publish changes, articulate virtues. Example: reject suppressing low-income ZIPs to boost ROAS; add an equity guardrail.

Start small—choose virtues, keep a decision log, set two guardrails, add monitoring—to build trust and agility. [View Slides](#)



Maan Sacdalan shows how AI-driven workflows streamline nonprofit ops: email triage and sentiment routing, 90-day SMS onboarding, a 40-day welcome/fully automated year-end campaign, content repurposing (Notebook LM), and meeting-to-PM actions.

She stresses QA, avoiding over-automation, and change management. Results: a 41% open rate on year-end emails; time saved but new expectations—keep humans in the loop. [View Slides](#)



Session Highlights



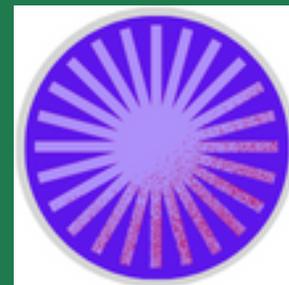
Sal Salpietro argues intentions aren't enough—experience wins. Using a Sicily trip as metaphor, he shows how nonprofits' dated, friction-heavy donor journeys (no receipts, no address autocomplete, limited payment options) depress conversion. Borrow e-commerce playbooks—CRO, abandoned-donation reminders, AI-driven personalization and fee-coverage tuning—to remove friction, lift average gifts and conversion, and meet mobile expectations. The scary future is not adapting. [View Slides](#)



Fundraise Up

Rachel Kimber and **Samin Pogoff** argue that AI "101" is essential enterprise-wide: start with governance (AI + data), shared language, and psychologically safe play, then move quickly to policy and change management.

Center E/D/S/P (ethics, data, security, privacy), align leaders and doers, pilot one use case, measure and iterate. Treat policy as a living "garden"; include skeptics; culture beats tools.



Working Group for a Progressive AI Future initial members



Help shape the future we actually want. Join the Working Group for a Progressive AI Future. Email ben@stratovation.digital to join!



Thank you to our sponsors!



Fundraise Up



Chorus AI

