

MAY 2024



 higher ground labs

AI EDITION:

HGL Political Tech Landscape Report

Table of Contents

3	INTRODUCTION
5	EXECUTIVE SUMMARY
6	HOW TO READ THIS REPORT
8	GENERATIVE AI POLITICAL TECH LANDSCAPE MAP
9	MEDIA & MESSAGING
15	DATA ANALYTICS & MODELING
18	RESEARCH
21	VOLUNTEER & ACTIVIST MOBILIZATION AND VOTER ENGAGEMENT
26	FUNDRAISING
29	MOVEMENT-WIDE & ORGANIZATIONAL INFRASTRUCTURE
36	RESOURCES
41	ACKNOWLEDGMENTS


Introduction

Generative AI has the potential to be one of the most transformative technologies of our lifetimes. Every industry is grappling with its impacts and how to prepare. Politics is no different.

This is why we are publishing this special edition of our annual Political Tech Landscape Report, focused entirely on the emerging use cases, needs, gaps, and opportunities presented by the introduction of generative AI in politics.

If you looked at the headlines alone, you wouldn't think there was much cause for optimism. Stories about bad actors using deepfakes or spreading misinformation are rightfully generating concern. But for campaigners, these headlines can also obscure AI's potential to revolutionize how we approach every aspect of our work.

AI is moving quickly and changing every day. For organizational leaders, weighing the technology's opportunities and risks and incorporating AI into their workflows is overwhelming—yet we have no choice. Since its launch in late 2022, ChatGPT has experienced faster growth than TikTok, Spotify, and YouTube combined. We know Republicans are already using this technology both defensively and offensively. Choosing not to engage with this technology while the opposition, other industries, and voters adopt AI is to risk lagging behind.



This is a generational opportunity for Democrats to get ahead.

This is a generational opportunity for Democrats to get ahead.

In the research process, we wanted to approach AI like any other tool to judge what makes it effective or dangerous. How is it being used? Who is using it? What are the norms of what's acceptable? Where are opportunities for innovation? This report—the first of its kind—aims to explore these questions and outline how the AI political tech landscape is beginning to take shape.

Our research confirms that we're still in the experimentation stage with AI. The technology itself is still quite nascent, yet generative AI has the potential to help your team run smoother, faster, and more efficiently today.

Using a mix of tools – free-to-use, enterprise solutions, and tools from the emerging market of political AI tech – practitioners have been able to:



Brainstorm campaign ideas and outline the beginnings of a strategy memo;



Increase content production and personalize it for specific demographics and psychographics;



Digest immense amounts of publicly available content and summarize it into actionable insights;



Save time on running fundraising programs by drafting first passes of content for mass email and SMS programs;



Engage volunteers, voters, and other high-priority segments through automated chatbots;

How we experiment with AI now will shape how we invest in and develop this technology for the next decade.

These developments are exciting, but that doesn't mean our community doesn't have concerns to address. We need training on the responsible utilization of AI and should ensure that human review is integral to any process involving AI.

Leaders must also address AI head-on internally with their staff, setting clear expectations on how and when it's appropriate to use AI tools. While many leaders may feel reticent or unprepared, embracing this challenge is paramount for the effective, ethical, safe, and sustainable use of AI.

It's natural to be wary of experimenting, especially in politics—where the stakes are high and time and resources are short. We hope you find this report informative and that it generates ideas on how you might leverage AI in your own work.

How we experiment with AI now will shape how we invest in and develop this technology for the next decade. Together, we can take the first step towards unlocking new potential for engaging volunteers, reaching new voters, and getting insights that were previously out of reach—keeping our campaigns on the cutting edge of how we win.

Executive Summary

Campaigns are leveraging AI for efficiency, but human oversight remains crucial


Campaigns and organizations are leveraging AI to improve the efficiency and speed of content creation, data analysis, and communication. However, human oversight remains critical to ensure quality, accuracy, and ethical use.

Rising AI use in political communications

AI-generated content is becoming most prevalent in political communications, particularly in content generation across images, video, audio, and text. Thus far, AI-generated content has generally performed on par with content created entirely by humans, matching performance as often as not. Testing and maintaining human oversight remain crucial when utilizing this emerging technology.

Political tech practitioners and technologists grapple with AI tool selection and development

Political tech practitioners and technologists must face a host of new factors when selecting and building AI tools. Accessibility, customization, privacy, and security considerations all play a role in organizations choosing between free, general-purpose AI tools or specialized, politics-focused offerings. Politics-centric AI applications are also subject to the discretion of major AI providers, who may modify their Terms of Service to limit the usage of their tools in political contexts.



Ongoing challenges countering AI-driven disinformation and deepfakes

Concerns around AI-generated misinformation, disinformation, and deepfakes are prompting regulatory actions from governments and tech companies. The effectiveness of these policies at curbing abuse by bad actors remains unclear, underscoring the need for increased vigilance and education around the responsible use of AI in political contexts.

Organizations start to prepare themselves for the responsible use of AI

Internal training, policies, and frameworks are critical to guiding the safe, ethical, and responsible use of AI, especially given the accessibility of free generative AI tools. While some organizations are experimenting with internal policies that balance innovation with considerations around data privacy, security, and workers' rights, such practices remain uncommon.

How to Read this Report

Generative AI vs. “traditional” AI

AI is an expansive topic. For the bulk of this report, we chose to focus on the newest and most consequential aspect of AI technology, generative AI—this is what underlies tools like ChatGPT and allows users to both input and output content like language, data, and images.

That said, some of the opportunities, needs, and gaps we identified also call for a “traditional” AI solution, as we didn’t want a strict focus on generative AI to hinder a discussion on the many ways AI might transform our work.

On the term “AI-generated”

Some writing on AI draws a distinction between content that’s generated by AI and content that’s assisted by AI. It implies that AI-generated content is entirely the product of AI without human intervention, whereas AI-assisted content starts with AI but goes through human review. In this report, we avoid that distinction, generally favoring the term “AI-generated” content since all successful experiments with AI-generated content still involve human review.

On the tools featured in this report

In this report, we highlight a mix of purpose-built, politics-specific tools and off-the-shelf, commercially available tools. This is intentional—we think both are essential in providing campaigns and organizations with a robust yet accessible AI toolkit. Getting started with AI does not require a lot of resources, and we also offer critical factors to consider when selecting between the available tools.

How this report differs from previous Landscape Reports

Generative AI is still a nascent technology, and political practitioners’ experiments with this technology are newer still. This report differs from past HGL Political Tech Landscape Reports: Rather than reporting back on all of the ways tech was used over the last year, we’re examining the patchwork of campaigns and organizations experimenting with AI and projecting forward trends for this cycle. While it means many of the trends highlighted here are far newer—and in many cases, smaller-scale—than previous reports, we hope this can provide helpful context and guides for those curious to experiment with AI in 2024. We intend to resume our regular landscape report next year.

AI Glossary

Here are the key words to understand as you read this report—starting from the broadest foundational concepts and ending with the most specific terms.



1. Machine learning

Machine learning algorithms are models that enable computers to learn and improve from experience without explicit programming for each task. Instead of specific instructions, these algorithms are trained on data to identify patterns and relationships, allowing them to make predictions or decisions.



2. Generative AI

Generative AI learns from existing data to produce new content. It goes beyond classification to create original outputs, such as artwork, music, and text, mirroring its training data's style or structure.



3. “Discriminative” AI

Part of “traditional” AI that aims to discriminate between different inputs and classify them into predefined categories (e.g. type of voter).



4. Training data

Each foundational model is “trained” on very large swaths of data (i.e. a large subset of the entire public internet). It consists of examples or instances with known inputs and outputs, which the model uses to learn patterns and relationships in the data.



5. Large Language Model (“LLM”)

A model trained on vast amounts of data, enabling it to understand and generate human-like outputs. These models are trained on massive datasets, allowing them to learn the complexities of language and generate coherent content across various tasks like language translation, text summarization, and conversation.



6. Fine-tuning

Taking a foundational model that already exists and doing a bit of extra training with data from a specific domain to customize it.



7. Hallucination

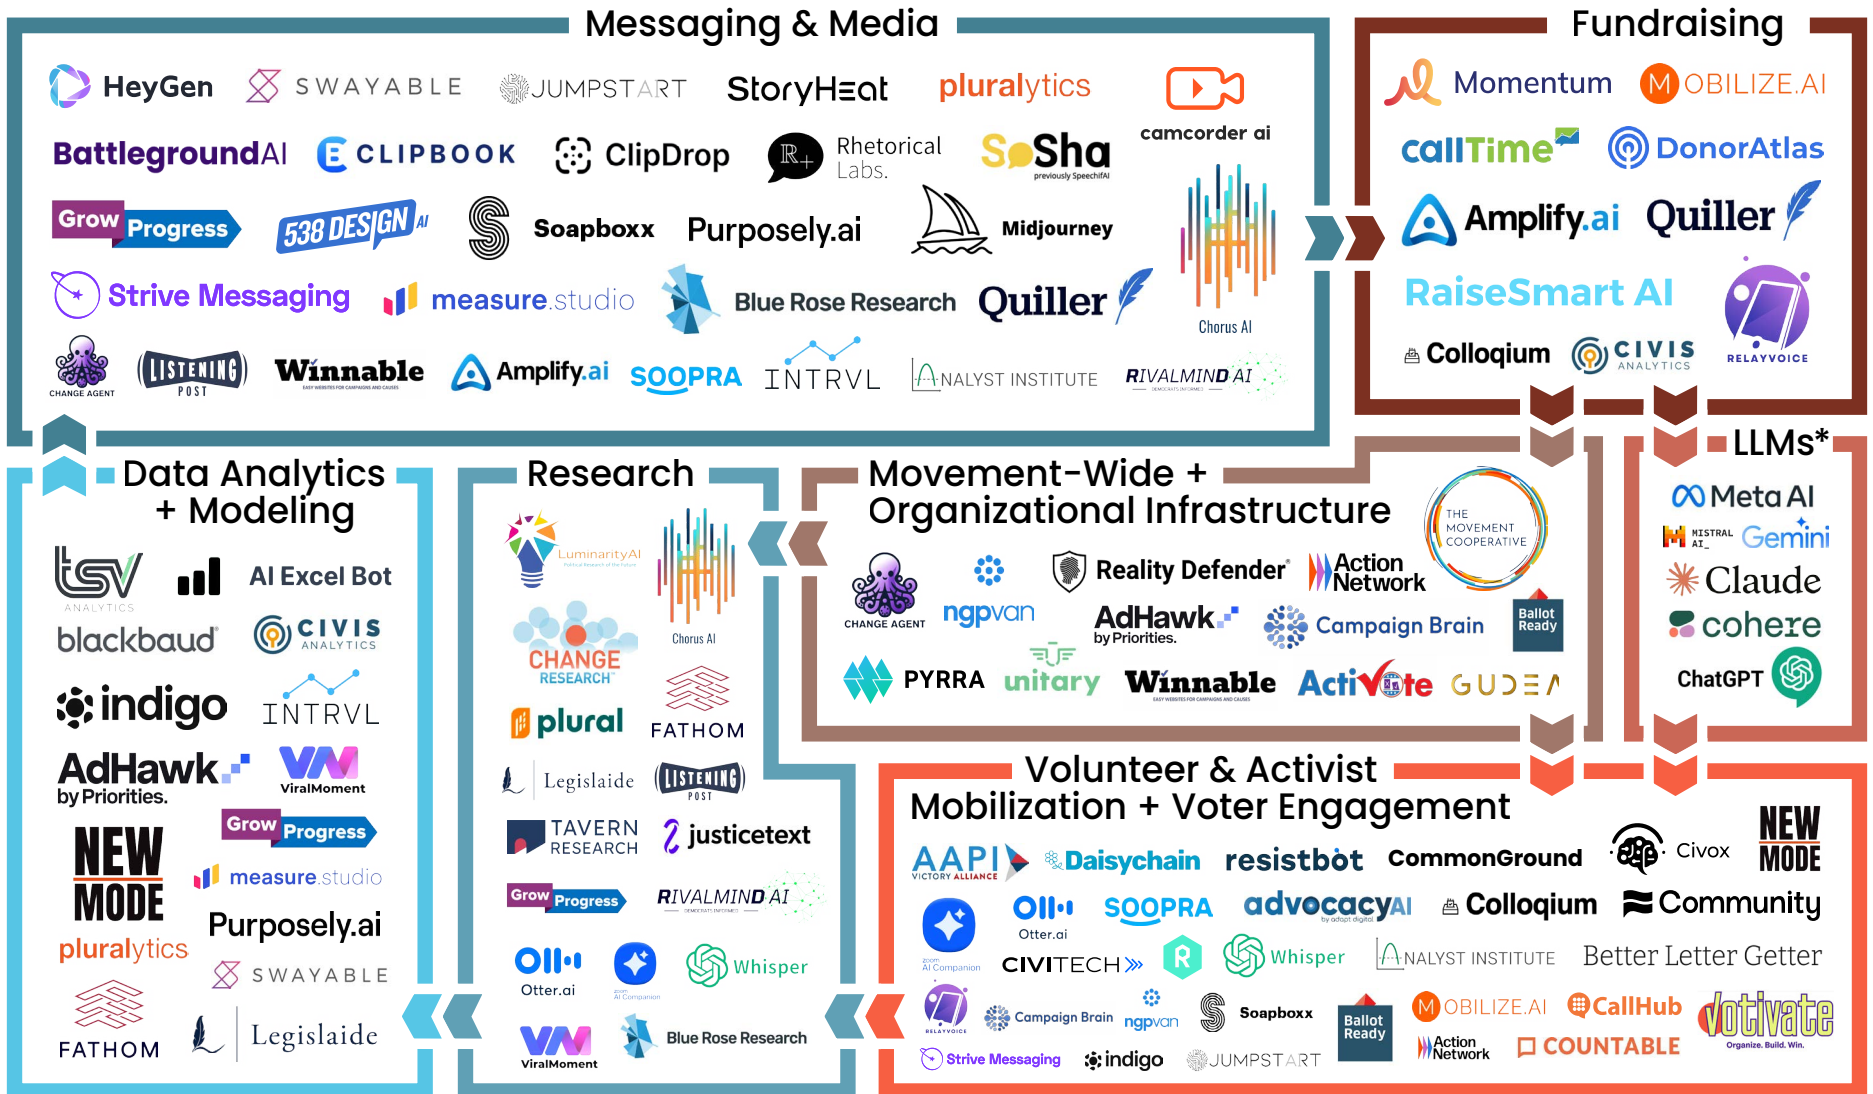
An AI-generated response that is factually inaccurate or nonsensical but presents itself as fact.



8. Deepfake

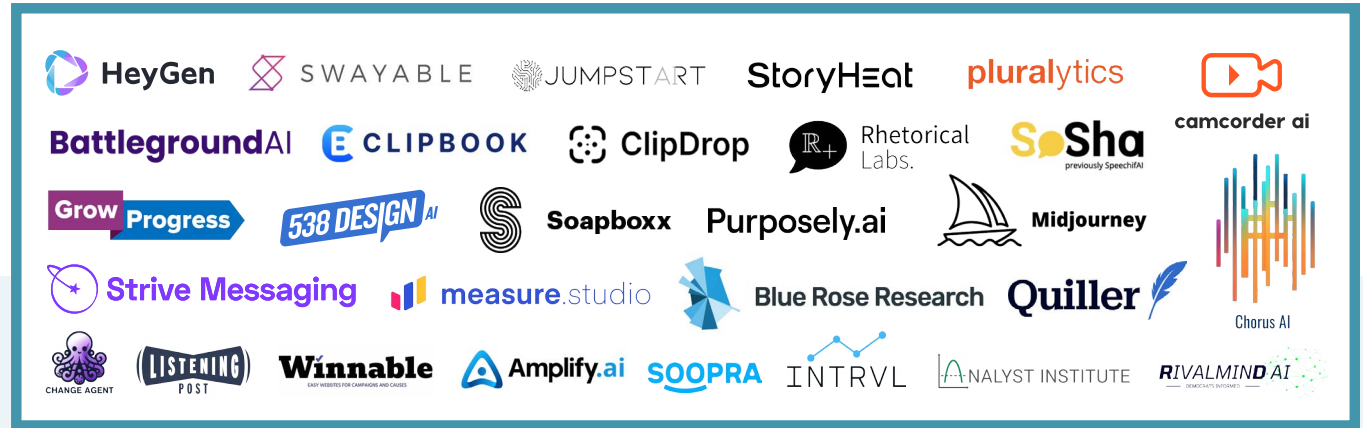
Highly convincing synthetic content generated by AI, including images, videos, and voice. These techniques allow for manipulating visual and audio content to make it appear as if something happened that did not actually occur.

GENERATIVE AI POLITICAL TECHNOLOGY LANDSCAPE MAP



*LLMs that are most commonly used in the ecosystem.

Media & Messaging



Notable Products

538Design AI creates an AI-powered design tool to help campaigns generate design assets in minutes.

BattlegroundAI launches to help create AI-optimized, high-performing political ads for your campaign, candidate, or cause.

Chorus AI creates an AI-assisted media monitoring and content generation tool to help organizations increase their news-driven content production.

Clipbook provides actionable, automated insights on your press clips.

Grow Progress utilizes AI to better understand what drives audience persuasion.

Jumpstart AI develops an AI product to help campaigns quickly create video content.

Quiller launches an AI-powered co-pilot to help campaigns generate first drafts of supporter emails, op-eds, and more.

SoSha pioneers an AI-powered social sharing tool that helps organizations ramp up their social sharing efforts.

Swayable leverages AI to analyze creative testing across messaging, creative development, and media planning.

Emerging Uses

Drafting content more efficiently

Organizations like CASA and Tech for Campaigns are using generative AI tools as a copilot to draft content faster than ever. From advocacy letters to fundraising emails, these organizations use AI to get a first draft of content that's about 70% finished—then a human reviews, edits, and pushes through internal review as they would any deliverable.

These groups—alongside an emerging market of politics-specific AI tools for content generation like Quiller and Chorus AI—frame AI's value-add in terms of efficiency, emphasizing the role of human review in ensuring the final product maintains quality and accuracy.

The impact is significant: Tech for Campaigns' [experiment](#) concluded that using AI to assist with drafting fundraising emails grew dollars raised per work hour by 350-440%. A small-scale test from Winnable found that staffers using the AI assistant in its website creation platform took half as much time to draft content without impacting quality, with the amount of time saved increasing and content quality decreasing the less experienced the staffer.

Aiding message development

AAPI Victory Alliance and Rural Ground Game are among the organizations that have experimented with generative AI to make their message development processes more efficient. AI acts as a thought partner in this application during the message drafting process.

Rural Ground Game experimented with generative AI to refine messaging geared explicitly toward rural voters. They learned that using AI to tailor messaging toward a specific group was an iterative process involving human feedback to produce high-quality content. Rural Ground Game accomplished this by including volunteers in this message development process, which ensured that the messaging generated was neither generic nor stereotypical.

AI's capacity to help with message development isn't limited to chatbots. [Blue Rose Research](#) created an AI-powered predictive scoring tool, Blue Research Messaging Library, that models the persuasiveness of your message based on thousands of message tests they've run in the past. The tool isn't meant to replace message testing but can help campaigners gut-check talking points, experiment with different message framings, or identify less impactful copy in scripts.

Assisting with media monitoring

AI has also proven useful in making media monitoring and response more efficient. Tools like Clipbook and Chorus AI leverage AI to create a more precise news monitoring tool that can understand campaign themes and the search intent of its user—users provide feedback and train the tool to be better at monitoring over time. In one example, Chorus AI partnered with the Texas Border Abuse Campaign to curate news about Governor Abbott, specifically related to immigration and the border. Monitoring efforts without AI produced news clips unrelated to immigration and included articles about the wrong Greg Abbott, showing how AI can make campaigns' monitoring and response efforts more efficient than they would be otherwise.

No effectiveness tradeoff for AI-generated content

Early indicators show that campaigners experimenting with AI-generated content aren't doing so at the expense of performance.

Stanford University researchers found AI-generated messages were as persuasive as messages crafted by lay humans. Messages generated by AI are persuasive across several policy issues.


BattlegroundAI asked 1,000 voters if they could accurately determine if certain political ad copy had been written by a human or a LLM. They found that most respondents could not clearly distinguish between AI- and human-generated content for all messages tested. This provided support to the theory that content written by AI is unlikely to decrease performance.

The DNC also found in its tests that AI-generated content did not perform better or worse than copy generated entirely by humans, often performing as well or better. Trilogy Interactive and Grow Progress reached a similar conclusion when testing AI-generated voiceovers in persuasion ads. The results showed that AI-generated voiceovers didn't necessarily over- or under-perform the human-recorded counterparts. Instead, they offered another creative element they could use to optimize the program.



Testing Human and AI Voiceovers in Persuasion Ads

Trilogy Interactive & Grow Progress

If the 2024 election for Congress were held today in your state, would you vote for the Democratic candidate, Republican candidate, or another candidate?



If the 2024 election for Congress were held today in your state, would you vote for the Democratic candidate, Republican candidate, or another candidate?



AI-generated multimedia explodes

AI-generated images and audio have started regularly appearing in political communications. The Right has done the bulk of public experimentation so far. The RNC, Donald Trump, and Ron DeSantis campaigns all shared AI-generated images in ads or on social media; only the RNC included a disclaimer the image was generated by AI. The DeSantis campaign went a step further, sharing a post on Twitter with AI-generated images of Trump hugging Anthony Fauci, commingling them with real photos of the two together.



Beyond images, Ron DeSantis' Super PAC, Never Back Down, released an ad featuring AI-generated audio of Trump reading one of his tweets attacking Iowa Governor Kim Reynolds. Trump trolled DeSantis' glitchy campaign announcement with a fake Twitter Spaces event that included AI voice clones of Dick Cheney, Satan, and Hitler.

There's room for the Left to experiment with AI-generated multimedia, and we've already seen some applications. In CA-16, Democrat Peter Dixon used AI to depict his life story in a campaign launch video, incorporating an AI-generated creation to bring his past to life. In March 2024, the DNC released an AI-generated song trolling the RNC and Co-chair Lara Trump. The following month, TheShotline.org, a gun reform campaign by March for Our Lives and Change the REF, ran a campaign asking supporters to send AI-generated voice memos to members of Congress from victims of gun violence using their actual voices with consent from their families.



Transparency around how and when campaigns disclose using AI-generated multimedia must remain a top concern as practitioners continue experimenting and developing best practices. Many of the examples above lack a disclaimer indicating that the multimedia was AI-generated. This decision could alienate voters, diminish its intended impact, or heighten the risk of AI-generated multimedia contributing to the spread of disinformation.

Deepfakes dominate headlines

A spate of bad actors using deepfakes stoked fears around AI’s potential to spread misinformation. The week before the New Hampshire primary, an estimated 5,000-25,000 Democratic voters received a robocall with audio of President Biden urging them not to vote in the primary. The consultant who commissioned the call later claimed he did it to spotlight the dangers of AI in politics.

In Harlem, deepfake audio of Manhattan Democratic Party Leader Keith Wright talking trash about a New York assembly member generated waves in local political circles. The 10-second clip was shared with media outlets claiming the Leader was “caught on hot mic,” but the clip was AI-generated.

These incidents have prompted swift action from government officials. In February 2024, the FCC voted to limit the use of unsolicited AI-generated voices in robocalls and announced new action targeting AI impersonation. State legislatures were introducing around 50 AI-related bills per week as of February 2024, with nearly half of those bills addressing deepfakes specifically. At the federal level, Congress is looking less and less likely to act on legislation to address deepfakes and other AI-created misinformation, though Sam Altman from OpenAI has said he’d support regulation to prevent AI election misinformation. In February, several tech companies, including OpenAI, Google, Microsoft, TikTok, Meta, and others, signed an “accord” pledging to work together to combat the deceptive use of AI in the 2024 elections.

No matter how the regulatory landscape around AI-generated robocalls evolves, these rules are unlikely to stop bad actors who operate in the shadows or foreign interference from using AI to their ends.

In Harlem, deepfake audio of Manhattan Democratic Party Leader Keith Wright was shared with media outlets claiming the Leader was “caught on hot mic,” but the clip was AI-generated.



Areas of Opportunity, Need, and Innovation

Keeping human involvement front and center

Practitioners who've experimented with AI, tech developers, and experts are aligned: Human interaction is key to leveraging AI. No AI can create a perfect output. Ensuring every deliverable goes through human review before being shared is critical for overcoming campaigners' adoption skepticism, ensuring AI-generated content remains accurate, and ensuring AI tools are being used ethically and safely.

There is an urgent need for education on AI usage in politics, but managers cannot wait. It is incumbent on leadership to set clear expectations with staff regarding appropriate AI use in the workplace—not just around ensuring humans review any AI-generated product, but to minimize data privacy and security risks, too.

For more on setting an internal AI policy, see ["Movement-Wide & Organizational Infrastructure"](#) and ["Resources"](#).

Education on editing AI-generated content

"Editing AI-generated content does require a mindset shift," said Hillary Lehr, CEO of Quiller, an AI-powered content generation tool. Practitioners who are used to editing copy written by humans may balk when they see common AI errors like inconsistent subject-verb agreement or an incorrect date. Still, Lehr cautions against "throwing the baby out with the bathwater." "It's important to remember that even with the required copyediting, AI generates a solid foundational draft that's still saving you much more time than if you had started with a blank page." Expectation-setting and education around how to edit AI-generated content may help practitioners feel more comfortable experimenting with AI for this purpose.

Tying together messaging and performance data

In the last year, new LLMs have greatly improved the quality of AI-generated content. However, the tools that practitioners use to generate content often live outside performance data gathered by platforms, meaning that the content generated by AI isn't necessarily tied to the performance data. Campaign tech can innovate by finding ways to automatically and closely tie AI-assisted performance data with what's most effective at meeting organizations' KPIs.

Practitioners who've experimented with AI, tech developers, and experts are aligned: Human interaction is key to leveraging AI.

Data Analytics & Modeling

Notable Products

Blackbaud debuts **Impact Edge**, an AI-powered, social impact reporting and storytelling solution.

Change Research launches **Magnify AI** to make custom-modeled targets accessible with smaller sample sizes.

Civis debuts **AI Core** to enable users to analyze their data by asking plain-text questions.

INTRVL leverages AI to rapidly create custom audiences and surveys.



Emerging Uses

Creating more accurate targeting

In the last year, several tools have emerged leveraging AI that allows campaigns to create better audience targeting. Change Research released [Magnify AI Targeting](#), which uses AI to build custom-modeled audiences. In 2023, Boise Mayor Lauren McLean leveraged Magnify AI Targeting to build custom-modeled audiences using a pool of fewer than 600 respondents. Initial analysis of the model's predictive capability showed that it outperformed standard partisanship models, identifying supporters who otherwise wouldn't have been revealed and vice versa.

INTRVL also integrated generative AI and machine learning into its targeting capabilities to build predictive targeting scores that allow campaigns to target users most likely to be impacted by their media. Conservation Voters of Pennsylvania employed this tactic during the 2023 state supreme court race and observed an increase of 2.3-4.2 percentage points, varying by ad type. Additionally, tiers based on AI-assisted scores for mobilization ads found that they were predictive of lift.


Saving time on data analysis

Tools like ChatGPT and AI Excel Bot can make writing first drafts of code and formulas much more efficient. However, it's unclear how widespread the use of generative AI tools by campaigns and organizations for this purpose is. There is much room for practitioners to experiment with this tactic in 2024.

In the meantime, expect political tech to integrate generative AI into their platforms to assist with data-related tasks like querying. Daisychain integrated natural language processing into its targeting features, allowing users to enter their filter criteria in plain English and receive a JMESPath expression in return.

Mobilization Lift Correlated With Predicted Moveability

Tiers built on INTRVL's Max Lift Scores were predictive of lift.
Tier 1 moved ~35% more than Tier 5



Areas of Opportunity, Need, and Innovation

Increasing accessibility for non-analysts

Most current applications of generative AI for data analysis still require technical familiarity. One-off data questions from other departments can be time-intensive for analysts, and AI opens the door for non-analysts to ask plain text questions of an existing dataset and get an answer. Tableau and Civiis added natural language query support to their platforms, while Blackbaud announced Impact Edge, an AI-powered reporting and storytelling tool. Whether these features empower non-analysts on campaigns to take the lead on analysis questions themselves will be something to examine after this cycle. Until this technology matures, it's important any AI-generated analytics code goes through human review by someone with a technical background.

Improving analysis of large datasets

Traditional AI can be used to code and clean data, but generative AI is still not well-suited for analyzing data independently. Tools like ChatGPT can give a broad general data report, but they struggle to choose which data is essential and cannot interpret it in the context managers and leaders desire. Coupled with AI's tendency to hallucinate, there may be a ways to go before campaigners can rely on generative AI to undertake data analysis without significant oversight. A small step to closing this gap is to ask your AI tool to cite how it arrived at certain conclusions.

Instead, as AI's analytical capabilities evolve, practitioners might find more success in using AI as a companion to their ongoing data analysis efforts. Generative AI can be useful for summarizing your existing large report, creating summaries, or generating content from existing reports for other audiences.

AI-powered data visualization

There's room for generative AI tools to lower the lift on data visualization by taking the first pass at data visualization and other dashboards. Campaigns should experiment this cycle with tools that could help analysts create these products more efficiently, freeing up time for other tasks.

Education and best practices around risks

Data privacy and security risks remain the primary concern of practitioners looking to use AI for data analysis, not just in the political world but across industries. Clear internal and vendor policies that appropriately balance data privacy and security concerns when considering AI applications will be key to helping more campaigns grow comfortable exposing generative AI tools to sensitive and proprietary data.

There's also a need for more education about the sources of risk present in training AI on existing datasets. Risks like non-representative data, poor data quality, poisoned data, and more can lead to inaccurate results and exacerbate biases that already exist in data. These risks shouldn't stop practitioners from experimenting with AI for data analysis, but understanding these risks coupled with strong human oversight can ensure they're doing so safely, accurately, and ethically.

Research

Notable Products

Fathom saves analysis time on open-ended qualitative data.

Hill and State Strategies develops **RivalMind AI**, an AI-powered opposition research tool.

Luminarity AI launches an AI solution for opposition research.

Plural launches an AI-powered legislative tracking tool to summarize and analyze legislation nationwide.

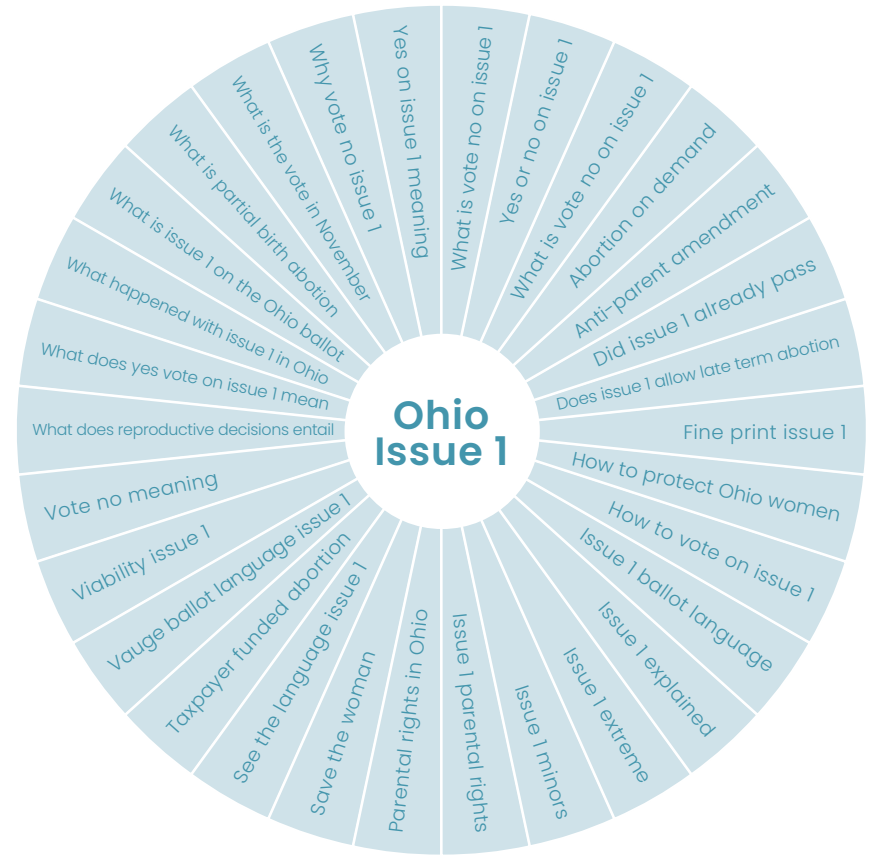


Emerging Uses


Mining data for actionable insights

One of AI's strengths is reading through and pulling insights from large amounts of text quickly, and campaigns are using it to analyze trends in data to get actionable insights. For the Ohio Issue 1 campaign, Red Wine and Blue complemented their relational organizing with AI tools (ChatGPT 4 and Eightify) to data-mine the opposition's YouTube video content and comments around the ballot initiative. This allowed them to discover patterns in messaging, understand how viewers were contextualizing the issue, and predict trends in search behavior from voters who were looking for information on the initiative.

This analysis allowed Red Wine and Blue to identify blindspots in their messaging and opposition targets, which they leveraged to run search ads driving people to pro-amendment messaging that generated a clickthrough rate at 6x the industry average. Their research also helped them uncover and report misinformation on the Ohio Senate website. Taken together, this analysis was a critical tool in helping combat misinformation on the campaign.

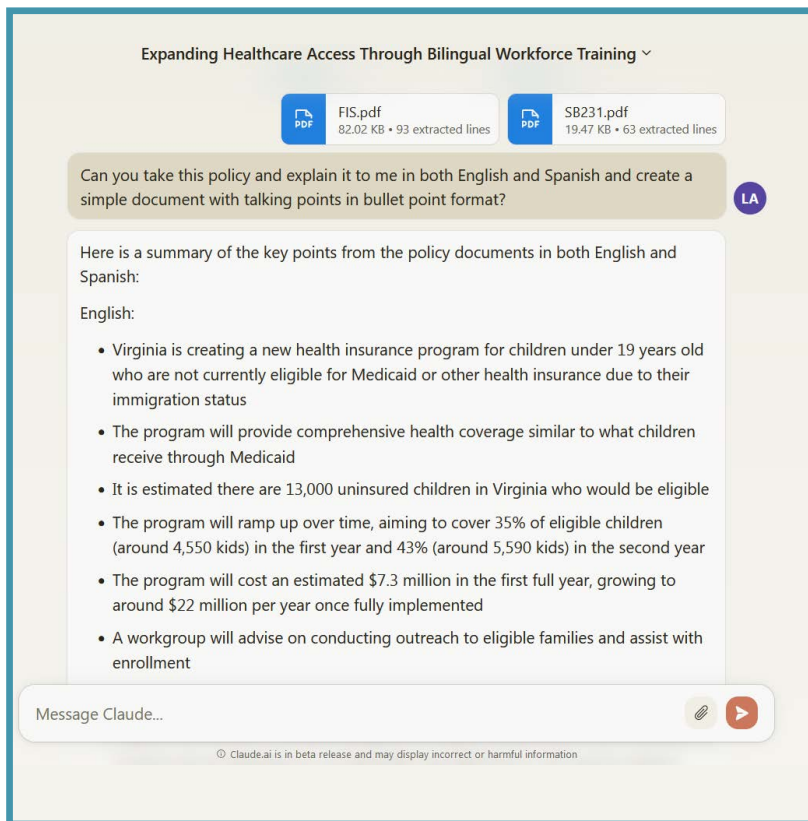


RWB Paid Search Clickthrough Rate vs. Industry Average



Digesting complex policy faster

CASA used AI to quickly summarize and analyze legislation in their Cover All Kids campaign in Virginia. Virginia's legislative sessions are just weeks long, yet lawmakers introduce thousands of bills in that time, meaning it's imperative to act quickly. Using Claude and Plural, CASA was able to quickly summarize new legislation in both English and Spanish, brainstorm calls to action to mobilize their members, and generate first drafts of letters to targets.



Areas of Opportunity, Need, and Innovation

Scaling monitoring across different data streams

AI-powered data mining and research still require human intervention. Results cannot be easily analyzed across multiple platforms and are largely not automated. There's an opportunity for developers to build tools that can passively monitor and analyze multiple data streams, identifying patterns and surfacing alerts in real time. Tools like [Listening Post](#) are starting to take this on by monitoring limited data streams like podcasts, state legislative feeds, and online ad archives. The 2024 cycle is an opportunity to foster innovation and move this capability forward for the entire space.

Helping campaigns with vetting

Vetting significantly consumes the time of research and operations staff on campaigns. AI's capability to pull and clean data from various sources makes it highly efficient for gathering background, social media, and other candidate information, thus saving time.

Focus grouping through AI personas

There's been some interest in using AI to build "personas" for strategists to ask questions as a lower-cost, faster way to perform focus groups. Initial attempts to implement this have fallen flat; [a test of a non-political version called SyntheticUsers](#) showed that the AI was prone to hallucinating and presenting even uncertain information confidently. Should developers succeed in creating a product like this that presents information in a way that users can trust and verify, it could make focus groups more accessible for down-ballot or under-resourced campaigns.

Volunteer & Activist Mobilization and Voter Engagement



Notable Products

Better Letter Getter creates an AI-powered letter-to-target tool to assist supporters with writing letters to decision-makers and for public comments.

Colloquium launches an AI-powered voter contact suite that can help campaigns communicate more efficiently with voters, volunteers, and donors.

Daisychain integrates generative AI throughout its toolset, from response drafting to natural language filter creation.

Emerging Uses

Gaining insights from canvassing

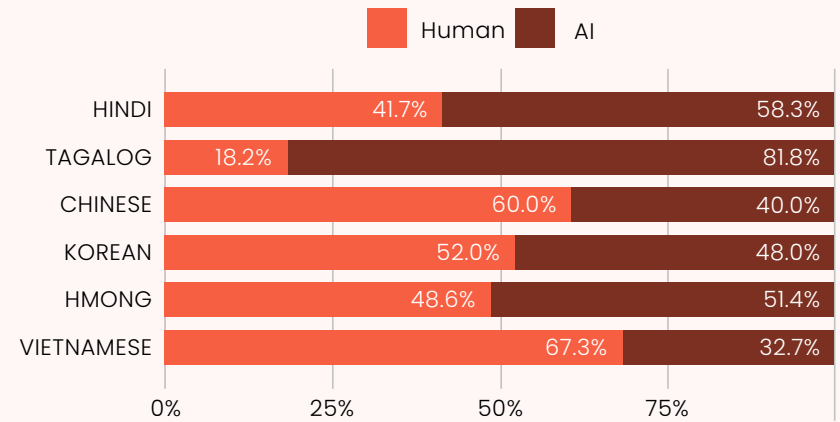
Fair Count used generative AI to capture and synthesize insights from door-to-door canvassing conversations as a part of their work on the 2023 Mississippi gubernatorial election. They trained canvassers to record voice memos summarizing conversations after each door knock, transcribing the audio clips and feeding the text data into the AI model Claude. Claude was able to translate the transcripts and provide county-level sentiment analysis, turnout predictions, and even flag a county where stated voting intentions seemed inflated compared to actual turnout.

Translating copy


Various groups and officials have experimented with using AI to translate copy, reach new audiences, and broaden coalitions. This tactic grabbed headlines in New York after Mayor Eric Adams pushed out an AI-generated robocall of his voice speaking Mandarin, Urdu, and Yiddish—languages the Mayor doesn't speak. Adams said the phone calls promoted hiring for the high number of vacant city jobs.

Using AI to assist with translation can be especially impactful for campaigns when coalitions include a large array of language speakers. The Asian American and Pacific Islander (AAPI) community in the U.S. is extremely diverse, with people originating from over 20 countries and speaking over 100 languages. To study AI's effectiveness in translating and optimizing messaging for this demographic, AAPI Victory Alliance conducted surveys across six major AAPI languages: Simplified Chinese, Korean, Hindi, Tagalog, Hmong, and Vietnamese. They tested human-translated messages against messages that were translated and then optimized by AI on key issues like gun violence, climate, and women's rights.

Human - AI preference by Language



Translation Preference by Translator Type and Issue Area



The results showed an overall 48:52 split preference for human vs. AI messaging with significant variation across languages. Some languages, like Tagalog and Vietnamese, showed a strong preference for human translation, while Hindi speakers marginally preferred AI. They found human involvement in the process was critical to mitigating distracting errors in translation made by AI, but overall, the tactic showed promise especially given that up to 30% of the AAPI community has limited English proficiency. Campaigns may be able to save time and resources on translation efforts by using AI to translate more efficiently and relying on competent cultural language reviewers instead of full translators

Engaging stakeholders with chatbots

To varying success, Dean Phillips, Asa Hutchinson, and Francis Suarez's presidential campaigns all experimented with chatbots to engage voters. Phillips' bot was taken down after OpenAI suspended the creators' account for violating its policy against using its technology for political campaigns. Hutchinson's chatbot drew mockery as users online asked absurd questions unrelated to politics.

Practitioners saw more success leveraging AI chatbots for specific use cases. Progressive Change Campaign Committee used a chatbot to improve the quality of story collection: Users would input their initial story directly, and then an AI chatbot would interview them, incorporating information from the story they had submitted. From there, the AI Chatbot would propose a story submission and provide options for revisions. The AI Chatbot was built by adding new workflows to Better Letter Getter and powered by an OpenAI GPT4 Custom Assistant.

In December 2023, the company Civox made waves with "Ashley," an AI-powered phone agent. Some were excited by the potential to vastly scale up conversations with voters, while others worried about the technology's potential for hallucination and impacts on the important role of volunteers.

Aiding supporter mobilization

Some existing tools added AI components to reduce friction when engaging voters. Resistbot rolled out an integration with OpenAI allowing users to generate letters to targets more easily by interacting with an AI copilot. Daisychain, a peer-to-peer texting platform, also debuted a feature that suggests AI-generated replies when conversing with supporters.



Practitioners saw more success leveraging AI chatbots for specific use cases.

Areas of Opportunity, Need, and Innovation

Providing accurate election information

Concerned about the potential for election misinformation, the major AI chatbots seem inclined to push users seeking election information outside the tool. Google announced that Gemini would simply refuse to answer questions about the election, directing users instead to Google Search. OpenAI announced plans to direct users with voting questions to CanIVote.org.

For now, that may be the safest option for voters: A [study](#) by AlgorithmWatch found that Microsoft's Copilot provided false and misleading information about European elections 31% of the time, while a [study by The AI Democracy Projects](#) tested five popular LLMs and found they provided inaccurate information about elections half the time.

In the long term, this may suggest a need for a purpose-built AI tool that prioritizes providing accurate voting information, especially as users grow to rely on AI chatbots as legitimate information sources.

Countering disinformation across coalitions

Disinformation in private messaging groups is a [growing concern](#) for campaigns on platforms like WhatsApp, which in the U.S. is [much more popular](#) among Hispanics, Asians, and people who identify with two or more racial/ethnic identities. In their experiment testing AI-translated messaging, the AAPI Victory Alliance observed the potential for AI to assist in countering disinformation spread on various platforms. This potential lies in its ability to translate and identify false content across multiple languages swiftly.

Best practices for translation

ChatGPT may be less effective at handling less common languages. In such instances, the users' approach to prompting the AI could significantly influence its accuracy. Hungarian social movement organization noÁr [got poor results](#) from ChatGPT when asking it to generate content in Hungarian. They generated stronger results by making ChatGPT do the majority of its work in English—submitting the original prompt in Hungarian, having ChatGPT translate it to English, running the prompt in English, and then translating the results back to Hungarian. This cycle, there's room for campaigns to experiment and develop best practices around prompting to give practitioners the tools they need to use AI for translation effectively.

Assisting volunteer management

Volunteer management is a top priority for organizing staff. While there already exists a suite of voter-focused organizing AI tools that can be used for volunteer recruitment, developers have an opportunity to expand this technology to assist with training. For example, making slide decks for volunteer training is extremely time-consuming for leadership and staff. It would be worthwhile for staff to experiment with existing non-political AI slide creation tools to assess their utility for their campaign.

Volunteer message testing and optimization

We can leverage AI to optimize every stage of the volunteer recruitment process. This begins with volunteer targeting, where AI has the potential to make building recruitment lists more strategic and efficient. AI could also optimize and scale message testing for volunteer mobilization, enabling the crafting of highly customized pitches to persuade volunteers to get involved. This includes volunteer recruitment via phone, text, email, and social media.

Best practices around transparency in AI use

As AI use becomes more common, practitioners will have to coalesce around standards for when to disclose AI being used in voter-facing materials. Transparency is important for use cases like texting, where AI autonomously interacts with voters in ways that would typically require a volunteer or staff member. However, it might be less necessary in communications like an AI-assisted email drafted from a candidate, where there is less expectation that the communication was personally written and sent by the candidate.

Transparency is important for use cases like texting, in which AI is interacting autonomously with a voter in ways that would otherwise require a volunteer or staff member.

Ensuring constituent-specific content is accurate and appropriate

AI is good at personalizing content for different audiences but still requires human oversight. Although AI can scale re-voicing and message tailoring, it does not eliminate the need to involve the communities you are trying to engage. Practitioners should ensure that any AI-generated content for engaging specific demographics or audiences undergoes the same checks as any other content.

Lowering the lift on voter protection efforts

Voter protection is a critical, resource-intensive function of campaigns and committees that could be made more efficient through generative AI tools. Generative AI tools and chatbots could streamline the intaking and triaging of reports of voter issues, as well as manage and train voter protection volunteers.

Making data entry more efficient

Campaigns could use AI to streamline data entry efforts on campaigns, from transcribing written data from volunteer sign-ups or walk packets or transferring data between unsynced data sources (e.g. from Google Forms to VAN). This shortens data processing time and frees up volunteers for other tasks. While the aforementioned tasks leverage other forms of machine learning, practitioners can also use generative AI to assist with ease of transcription, data cleanup, and schema creation.

Fundraising

Notable Products

DonorAtlas develops a CRM that leverages AI to research potential donors and create donor outreach lists.

Momentum creates an AI-powered fundraising platform that can automatically identify high-priority donors to reach out to and draft copy for that outreach.

RaiseSmart AI creates a platform that automates compiling donor giving history and outreach lists.

RelayVoice launches a tool used to create and send personalized, AI-generated voicemails at scale to help campaigns communicate more personally with donors.



Emerging Uses

Generating first drafts of mass fundraising emails

Several organizations, such as the DNC, Tech for Campaigns, and the Democratic Governor's Association have used or been reported to experiment with AI to enhance their fundraising content. Patti Minter's mayoral campaign in Bowling Green, Kentucky, used Quiller, an AI-powered tool for generating first drafts of social impact content, to reduce email drafting time by over 60% and increase money raised per work hour spent drafting by 267%.

Patti Minter's mayoral campaign... used Quiller, an AI-powered tool for generating first drafts of social impact content... increasing money raised per work hour spent drafting by 267%.

Personalizing and streamlining donor engagement

AI is helping campaigns communicate with donors more efficiently. In 2023, Paul Takac, a candidate for Pennsylvania's state legislature, used RelayVoice to send over 450 AI-generated, personalized voicemails in the candidate's voice to drum up attendance for his annual fundraiser. The campaign drove the same number of attendees compared to the previous year's fundraiser while reducing call time by 20%. 43% of the attendees had received voicemails, and 36% of the event's donations came from voicemail recipients.

These strategies build on scaled, AI-powered donor personalization that campaigns have employed over the last couple years. In 2022, the Fetterman campaign for Senate in Pennsylvania leveraged AI chatbots to improve their digital voter engagement. They partnered with Middle Seat and Amplify.ai to manage over 725,000 comments on Fetterman's Facebook posts in real time. Amplify's software analyzed the comments for sentiment and intent and, when appropriate, automatically started personalized conversations with voters over Facebook Messenger with the goal of driving grassroots donations. Their strategy generated over 330,000 automated personal reply conversations initiated and a 50%+ response rate from voters. This particular tactic did not employ generative AI, underscoring that not all problems require generative AI and that other forms of AI remain useful and effective.

Assisting with donor research

Donor research is one of the most time-intensive tasks for a finance staffer. When DonorAtlas interviewed 35 Congressional staffers about their call time practices, 27 reported that despite having access to useful tools like NGP and Numero, they still did extensive manual research on individual donors through external sources like Google, LinkedIn, and news stories. Emerging tools like DonorAtlas and RaiseSmart AI leverage generative AI to streamline the donor research process by efficiently aggregating publicly available data on potential donors. In early, small-scale testing with RaiseSmart, groups saved 10-15 hours per week with their AI-powered donor research tool and raised 20% more from their pledges by shifting their extra time into follow-ups and event planning.


Areas of Opportunity, Need, and Innovation

Using AI to take notes during call time

Many commercial AI notetakers have come onto the market that can assist campaigns with taking notes during donor calls such as [Circleback](#) and Zoom's AI companion. Politics-specific tools are also working to integrate AI-assisted note-taking with other functions of a donor CRM like giving history, donor communications, and more. Anyone using these tools should check their privacy and data policies against their organization's own AI policy to ensure they're using them responsibly.

Integrating performance data with generative AI

AI tools for generating fundraising content currently rely on either general or user-selected training data for their models. Even if users choose their top-performing data for training, this data must be manually updated with the current best-performing content. As these tools evolve, developers have the opportunity to connect content generation with performance data, allowing generative AI's outputs to adapt to supporters' preferences and best practices.



Movement-Wide Organizational Infrastructure



Notable Products

BallotReady acquires PioneerGov, enabling them to use AI to collect local government data more quickly and accurately.

Votivate AI evolves as a new feature of Votivate, through an AI-assisted campaign manager to advise on campaign strategy, phone bank, and generate personalized media, built by the Working Families Party.

Winnable integrates AI into its tool to help campaigns build websites.

Emerging Uses

Varied excitement around use of AI

Many practitioners are truly excited about AI's potential. There are champions in every part of our ecosystem testing new systems and integrating AI into their workflows. However, many others are understandably unsure how to best approach this new technology. There are several factors that have made staffers apprehensive about experimenting with a new, under-proven technology: a general lack of knowledge on how to use AI, concerns about AI's labor implications, and fear of negative media coverage given the dominant narrative around the technology. Through our research, we've observed that the organizations that have successfully experimented with AI often had a champion in leadership driving that work.

For most major political tech companies, they are maintaining consistency with their core products that drive the progressive ecosystem. There is no rush to fully incorporate AI into their platforms, citing low customer demand, unproven ROI, and a desire to fully address data and privacy concerns. Rather than striving for an AI breakthrough, the primary AI focus in 2024 for these platforms is on launching integrations with new AI tools, ensuring their systems and organizations are AI-compatible, and conducting small experimentation.

Addressing AI's labor implications

Concerns about the potential impact of AI on labor have been at the forefront of conversations on AI since the technology exploded onto the market. A Forbes Advisor [survey](#) found that 77% of workers were concerned that AI would cause job loss, and disagreements about the role of AI were at the center of 2023's Writers Guild of America strike. Among political professionals on the Left, concern about how to adopt AI while respecting workers' rights has been especially salient, driving hesitation and even outright opposition to experimenting with AI in their work.

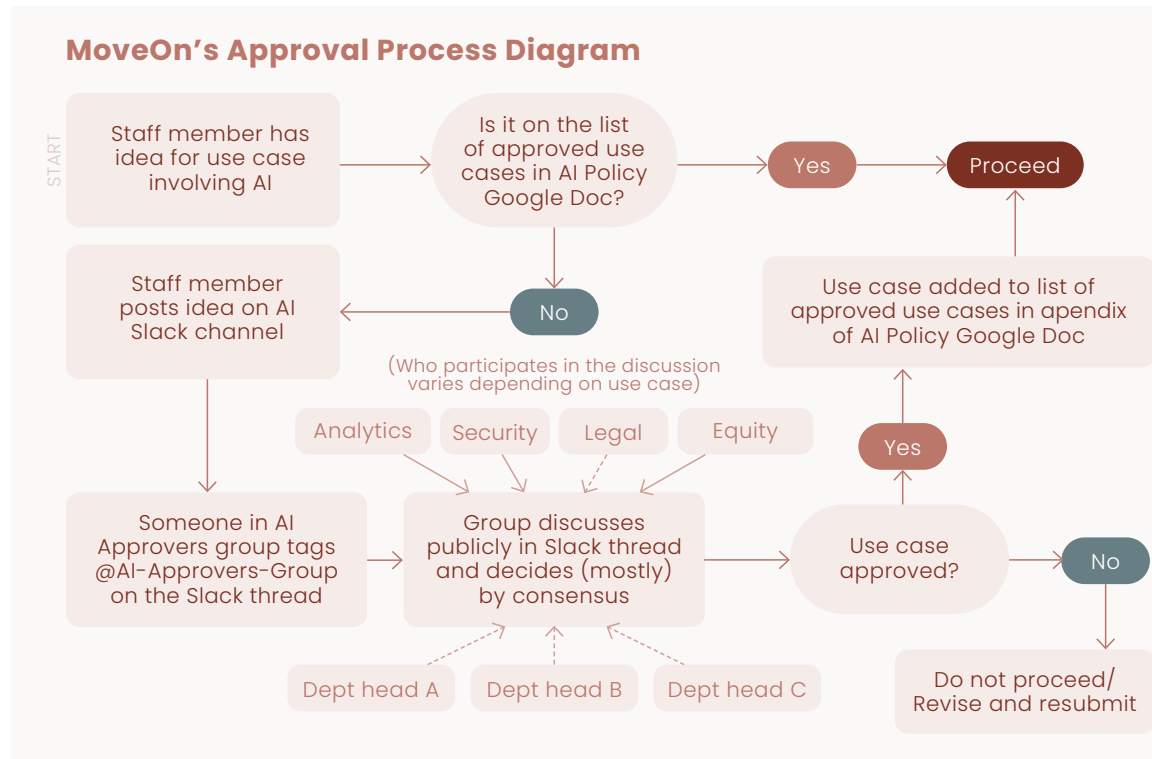
The AFL-CIO's Technology Institute has been taking on the relationship between workers and AI head-on. Rather than advocating to stop the development of AI technologies altogether, they're [working to ensure worker voices are a part of the "co-creation" process](#) for AI technology. This entails giving workers a voice in how AI is developed, regulated, and deployed in workplaces and pushing for stakeholders to prioritize worker rights as the technology evolves. They advocate for collective bargaining as a key mechanism for workers to negotiate over the adoption and impact of AI technologies in their workplaces.

One organization that has taken this approach is Authentic Campaigns, which reached an agreement between the Authentic Union and management around AI usage in the workplace. In a statement, the Union said that the agreement included a commitment that AI technology will only be used to assist workers' responsibilities, not eliminate or replace current or future human workers. Authentic also has a standing agenda item at the firm's Labor-Management Committee meetings for the Union and management to discuss how AI is being used at Authentic.

Internal AI policies to foster innovation

Beyond formal protections for workers, political campaigns and organizations have begun formalizing internal policies that set expectations and guardrails for how staff can experiment safely and ethically with AI tools.

Authentic Campaigns created an “AI Acceptable Use Policy” to compliment the protections for workers in their union contract, which reads, “We want to allow for innovation and efficiency, while ensuring we remain compliant with client contract, confidentiality, and data privacy requirements—and ethical in our use.” Authentic’s policy sets restrictions around using client data in generative AI platforms, producing deepfakes, or using AI to create or distribute any misleading information. The policy also states that staff should “ensure that AI-generated content is properly reviewed and approved by a human before it is published or used in decision-making.”



As a member of Cooperative Impact Lab’s AI Cohort, MoveOn worked to create its own AI use policy that builds in protections for data security, privacy, and human review but is also designed to empower staff to experiment with new ways to use AI. Potential AI use cases are flagged in an internal Slack channel for members of MoveOn’s technical, analytics, legal, and equity advising teams to evaluate alongside department heads. These stakeholders discuss publicly and attempt to come to a consensus, and if the tactic is approved, it gets added to MoveOn’s list of approved tactics.

For organizations looking to create their own internal AI policies, The Movement Cooperative has created a guide of considerations, risks, mitigations, and more to weigh as they navigate the process.

These AI frameworks aren't just for practitioners: When integrated organizing platform Daisychain began adding AI features to its product, it created an [internal framework](#) to ensure it approached that work responsibly.

AI usage policies like these are still not common, but organizations and campaigns must start adopting them for this cycle. Even if the organization has no official plans to experiment with AI in 2024, its staff will likely still use tools like ChatGPT informally. It's important that leaders set clear expectations around how and when staff can use these tools safely and ethically.

Based on their work with MoveOn to develop internal AI policies, Cooperative Impact Lab formulated a set of considerations for organizations interested in establishing their own AI usage policies.

[For more resources on crafting your own organizational AI policy, visit the Resources section of this report.](#)

Cooperative Impact Lab's considerations when setting up an AI approvals process for your organization:



Create a policy document to document high-level guidelines for using AI in the organization.



Keep a record of approved use cases as a reference for staff. Use cases similar to approved ones can move forward without requiring the entire approval process.



Balance curiosity to explore use cases with careful consideration of when it is or isn't valuable to use AI tools.



Maintain process transparency: Conduct discussion in public when possible.



Be inclusive: Different people have different views on AI. Make sure voices from different parts of your org are heard.

Continuing tension between politics and Big Tech

Major technology companies have been quick to restrict how their tools can be leveraged in politics. OpenAI announced they would not let ChatGPT be used for political work, while Meta and Google require disclosure of AI usage in political ads. Early indications suggest that these companies may be unprepared to follow through on their policies: Dean Phillips' PAC was able to launch a chatbot that was taken down only after it was spotlighted in the press and FWIW News found Meta was failing to enforce its own AI policies.

Fortunately, it's unlikely that these policies will stop campaigners from experimenting with AI this election cycle. Open-source and second-tier models are good enough for most expressly political applications this cycle, while many critical campaign uses of AI, like creating training materials and drafting social media posts, are apolitical enough not to be blocked.

A growing choice of tools

The proliferation of free, general-purpose AI tools and more specialized, politics-focused AI offerings have expanded the options available to campaigns and political organizations. While free tools like ChatGPT offer broad accessibility and ease of use, they may lack the nuance and customization needed for political applications.

In contrast, the growing market of politics-specific AI tools can provide outputs more closely aligned with campaign needs and values and stronger privacy safeguards for sensitive information. However, these tools may require more investment and onboarding.

As campaigns experiment with AI this cycle, they'll likely find the ideal mix involves leveraging free tools for more general tasks, while reserving politics-specific AI platforms for mission-critical applications that require tighter integration with campaign data, messaging, and workflows. Careful consideration of privacy, security, and ethical AI use will be key factors in selecting the appropriate tools and configurations, while clear internal policies to guide AI adoption can help organizations strike the right balance between innovation and risk mitigation. As they work to incorporate AI into their workflows, practitioners must also navigate the risk of major AI providers adjusting their Terms of Service to restrict use in the political space.

As campaigns experiment with AI this cycle, they'll likely find the ideal mix involves leveraging free tools for more general tasks, while reserving politics-specific AI platforms for mission-critical applications that require tighter integration with campaign data, messaging, and workflow.

Areas of Opportunity, Need, and Innovation


Training, onboarding, and internal functions

Most campaign-related applications of AI thus far have focused on output, but there is potential for generative AI to enhance internal operations, such as institutional knowledge management, training, and onboarding processes. Generative AI can also help simplify technical concepts for volunteers and staff (e.g. “Explain express advocacy to me like I’m five.”)

Chorus AI has started innovating in this realm, having developed an early product that uses AI to improve how campaigns organize their documents and files through semantic search as well as create new documents. Rather than just typing in a keyword in Google Drive to find a document, their semantic search tool allows the user to ask a question to generate the most relevant resources, such as, “Do we have policy language for child tax credits? I’m trying to draft some legislation for Wisconsin.” Their tool then allows you to create new documents using existing ones as inputs.

AI as a co-pilot campaign strategist

Candidates and campaigns, especially at the down-ballot level, have long expressed a desire for tools that make running for office more simple and affordable, AI could be key in making this happen. Votivate, developed by the Working Families Party, is working to launch Votivate AI, which includes features like a campaign strategy assistant, a low-latency AI voice calling tool, and an AI tool to create high-quality, individualized voter-targeted media. New/Mode is working to launch its Campaign Strategist to provide AI-powered strategic insights to help improve the impact of organizations’ advocacy campaigns. Tools like these are designed to offer smart campaign insights to smaller campaigns that might otherwise lack the resources to access them, and we expect to see growing activity in this space.



Exciting potential for down-ballot

The sophistication and accessibility of generative AI tools have the potential to help under-resourced organizations and down-ballot campaigns the most. From efficiency gains in content production to AI-powered audience modeling from smaller sample sizes, AI has the potential to ameliorate some of the challenges posed by a lack of resources at the bottom of the ballot. Education and training will be key to ensuring practitioners can leverage AI ethically, safely, and strategically. It will be interesting to observe AI's potential impact on down-ballot races this cycle, especially since leveraging these tools still requires keeping a human in the loop to ensure accuracy.

Vigilance around bias

AI models can reflect the biases of the data it is trained on, often mirroring the viewpoints of its predominantly straight, white, male, and college-educated developers. These biases can appear in many ways, like perpetuating stereotypes in content or influencing the language AI uses. To combat this, it is important to critically assess what AI generates and involve a diverse group of people in how AI tools are used. Using multiple AI tools and varying how requests are phrased can also help uncover and correct these biases, ensuring that our communities' diversity is considered.

Increased vigilance around disinformation

The rise of deepfakes marks a new frontier in campaigns' fight against election-related disinformation. Campaigns should be aware not only of deepfakes' potential to deceive but also of their use as a shield by actors who wish to blur the lines between reality and falsehood, as Donald Trump did when he falsely claimed that the Lincoln Project had used AI-generated footage to make him look feeble in an ad. Experts call this the "liar's dividend," or the idea that increased public awareness of the power of AI may incentivize politicians to lie about authentic content.

While deepfakes are obviously a problematic use of AI that would fall outside our own standards and guidelines for responsible use, we know it will be a tool used to proliferate misinformation this cycle. Campaigns and organizations both want and need to identify where deepfakes are used and respond effectively.

The response playbook does not look much different from our strategy around misinformation in the past. New companies are stepping up to fill the detection void. Platforms like Alethea, Memetica, Pyrra, and Blackbird.AI are emerging to help campaigns track and tackle disinformation more efficiently. This election cycle will be a critical learning period to see how campaigns tackle the potential scale of AI-generated disinformation.

Resources

Guardrails When Dealing with AI

From Zinc Labs' [Quick-Start AI Guidelines](#) for Democratic Campaigns

GUARDRAIL 1

Don't put too much trust in AI.

WHAT THIS MIGHT LOOK LIKE

- Failing to fact-check external-facing materials produced by generative AI.
- Failing to quality-check onboarding materials drafted by generative AI.
- Overreliance on generative AI for key functions that have not yet been pressure-tested in a campaign setting, like crafting media plans or accessing data.

HOW TO AVOID IT

Keep a human in the loop, always.

An expert on your team should check and approve every citation, social media post, and code snippet, and asserted fact produced with generative AI for your campaign. And there should always be a human prepared to take over when the AI fails.

QUICK START AI GUIDELINES FOR DEMOCRATIC CAMPAIGNS



In partnership with



Visit AICampaignGuide.com
to learn more and download.

GUARDRAIL 2**Don't forget AI can be biased, too.****WHAT THIS MIGHT LOOK LIKE**

- Perpetuating stereotypes in campaign content, like a social media post about scientific pioneers that only includes men.
- Uploading a job candidate's resume to a chatbot and asking whether they would be a good fit for the role.
- Relying on AI-generated language that skews straight, white, male, and collegeeducated, like the people who predominantly created and trained the generative AI models.

HOW TO AVOID IT**Evaluate AI outputs critically.**

AI is prone to many of the same biases people have, because it's trained on language that was produced by people. But unlike people, AI doesn't have independent reasoning — it doesn't see systemic inequalities or stereotypes, just useful patterns. To manage this:

- Bring diverse stakeholders to the table when deciding how an AI tool and its outputs should be used.
- Try using several generative AI tools for a task to see how their recommendations differ.
- Make a request in several different ways, and have the system explicitly consider factors like racial diversity, gender diversity, and educational attainment.

GUARDRAIL 3**Don't misrepresent AI-generated materials as authentic, and vice versa.****WHAT THIS MIGHT LOOK LIKE**

- Manufacturing a photorealistic image of your candidate speaking to a stadium crowd.
- Using a chatbot to invent an anecdote for a speech.
- Generating convincing video or audio that mimics the likeness of an opponent.
- Claiming without evidence that an authentic video shared by your opponent was created using AI.

HOW TO AVOID IT**Don't misrepresent the source or authenticity of content. Period.**

Public trust in institutions is at an all-time low, and the rapid development of generative AI has created even more confusion about whom and what to trust. Using new technologies to mislead the public about the veracity of content isn't just against democratic principles — it risks destroying your credibility with your constituents. This includes failing to inform voters when they are interacting with an AI, when they could plausibly believe they were engaging with a human. In general, it's safest to be transparent by default, and clearly label when you are presenting AI-generated materials.

GUARDRAIL 4

Don't breach the privacy or security of people who give you their data.

WHAT THIS MIGHT LOOK LIKE

- Uploading donors' personal information to a generative AI tool that might use it for training.
- Using a voter's likeness or story for AI-generated materials without their consent.
- Failing to ask an AI vendor about their tool's security procedures.

HOW TO AVOID IT

As before, keep a human in the loop and have a human review all AI outputs critically.

This is another place where an AI data and usage policy, with clear boundaries for who can share what data to what tools for what purpose, protects your campaign, your staff, and your constituents.

[The Biden White House's Blueprint for an AI Bill of Rights](#) is a comprehensive and evergreen resource for these and other AI risks — it's worth studying carefully and incorporating into your work. You may also wish to consult laws and regulations in your jurisdiction to ensure your organization is in full compliance.



[The Biden White House's Blueprint for an AI Bill of Rights](#) is a comprehensive and evergreen resource for these and other AI risks — it's worth studying carefully and incorporating into your work. You may also wish to consult laws and regulations in your jurisdiction to ensure your organization is in full compliance.

GUARDRAIL 5

Prep for deepfake risks.

WHAT THIS MIGHT LOOK LIKE

- An audio deepfake of your candidate saying something offensive.
- A deepfaked photo of your candidate doing something illegal.
- Deepfake video of a supporter stuffing ballots in your district.

HOW TO AVOID IT

Unfortunately, there's not much you can do to prevent deepfakes from being created, besides advocating for social media platforms to better protect their users against false and misleading information. But there are several steps you can take to reduce or prevent damage if your campaign falls victim to a deepfake:

- First, don't panic. Most disinformation, including deepfakes, can be managed with a calm approach, employing strategies akin to those used in typical comms and rapid-response strategies.
- Create a crisis response plan ahead of time. Deepfakes are an emerging technology. While they may increase the amount of AI-generated misinformation you have to deal with, they follow the same rules as other PR emergencies. Have your communications team prepare a playbook with clear response plans, roles, and timelines. As part of this plan, consider when and how to share the incident with the appropriate Democratic committee or partners, and the best channels to reach the social networks and request a take-down.
- Don't respond to or publicize the deepfake without careful consideration. Just because a deepfake is on your campaign's radar doesn't mean it has reached your voters. A knee-jerk response may just draw more attention.
- Report the deepfake to your organization's senior leadership and legal counsel as soon as possible. They may have resources to help limit its spread.
- The best defense is a good offense. The greatest danger posed by deepfakes isn't that they're easy to create — it's that there's an audience of people ready to believe them. Communicate early, prolifically, and authentically with the audiences most susceptible to deepfakes about your campaign, build trust, and "pre-bunk" likely topics of disinformation, so that if they are exposed to a deepfake, they're instinctively skeptical.

Links

[Quick-Start AI Guidelines for Democratic Campaigns](#) produced by Zinc Labs in partnership with Higher Ground Labs, DigiDems, and Cooperative Impact Lab

[What to consider in developing an organizational AI usage policy and staff guidance](#) (The Movement Cooperative)

[Civic AI Handbook](#)

[Center for Democracy & Technology AI usage policy](#)

[MoveOn's AI Internal Use Policy](#)

Lists of AI tools

[List of AI Workflows and Tools](#) (Cooperative Impact Lab)

[AI Tools in the 2024 Cycle](#) (Leadership Now Project)

Acknowledgments

Report Team

Betsy Hoover | Founder & Managing Partner, HGL
Leah Bae | VP, Programs & Partnerships, HGL
Jose Cornejo | Report Lead, HGL
Lorenza Ramirez | MBA Intern, HGL
Sarah Bloom | Designer, Sarah Bloom Studio

Contributors

Aaron Myran, Tareq Alani, & Sam Landenwitsch, Chorus AI
Andy Richards & Michelle Penson, AFL-CIO
Allie Cashel & Craig Koester, Tech for Campaigns
Damola Ogundipe, Plural
Hillary Lehr, Quiller
Ilona Brand, MoveOn
Isaac Troncoso, CallSmart AI
Jon Warnow & Nathan Woodhull, Daisychain
Kate Gage & Oluwakemi Oso, Cooperative Impact Lab
Keira Stearns, Analyst Institute
Lou Levine, NGP VAN
Loren Merchan, Authentic Campaigns
Luis Angel Aguilar, CASA
Matt Hodges & Ben Resnik, Zinc Labs

Matt Martin, Grow Progress
Maya Hutchinson & Harry Cyranka, BattlegroundAI
Michael Frias, Catalist
Mike Boland, Working Families Party
Raffi Krikorian, Emerson Collective
Ross Morales Rocketto, Run for Something
Sonya Reynolds, The Movement Cooperative
Steven McAlpine & James Owens, INTRVL
Andrew Brown, TargetSmart
Taren Stinebrickner-Kauffman, AI Impact Lab
Nicole Bare & Tom Swartz, Change Research
Tudor Mihailescu, SoSha
Varun Nikore, AAPI Victory Alliance
Will Schrepferman, DonorAtlas