

---

# QUICK START AI GUIDELINES FOR DEMOCRATIC CAMPAIGNS

*Prepared by Matt Hodges & Ben Resnik*



*In partnership with*



# QUICK START AI GUIDELINES FOR DEMOCRATIC CAMPAIGNS

## *About Zinc Labs*

Zinc Labs is a center for tech and talent innovation within Zinc Collective that incubates vital infrastructure across the Democratic Party. Our investments in early stage pilot programs and electoral tech strategy help secure persistent Democratic majorities.

## *About the authors*

**Matt Hodges** is the Executive Director of Zinc Labs and sits on the Advisory Boards of political tech accelerator Higher Ground Labs and of political AI co-pilot Quiller. Previously, Matt was the Director of Engineering for President Biden's 2020 campaign. Matt worked in the tech industry as a Software Engineer and Infrastructure Engineer at companies like Datadog, Lisnr, and Cardinal Solutions.

**Ben Resnik** is the Deputy Director of Tech Strategy at Zinc Labs, where he has helped develop AI strategy across the Democratic tech ecosystem. Ben is also a Political Partner at the Truman National Security Project and an advisor to the Higher Ground Labs Progressive AI Lab. Previously, Ben was Senior Director of Insights & Innovation at Main Street One and Product Manager at Swing Left.



# THESE GUIDELINES ARE PROUDLY SUPPORTED BY



*Zinc Labs is incredibly grateful to the organizations and experts who supported the creation of these Quick Start AI Guidelines for Democratic Campaigns.*

# TABLE OF CONTENTS

EXECUTIVE SUMMARY .....	5
GUARDRAILS WHEN DEALING WITH AI.....	7
Guardrail 1: Don't Put Too Much Trust in AI Tools .....	7
Guardrail 2: Don't Forget AI Can Be Biased, Too .....	8
Guardrail 3: Don't Misrepresent AI-Generated Materials as Authentic, or Vice Versa.....	9
Guardrail 4: Don't Breach the Privacy or Security of People Who Give You Their Data .....	10
Guardrail 5: Prep for Deepfake Risks.....	11
INTRODUCING AI TO YOUR TEAM .....	12
Looking Inward: Write an AI Usage Policy— <i>with</i> Your Team, Not <i>for</i> Them.....	12
Looking Outward: Evaluate Vendors with Your Values in Mind .....	14
GETTING STARTED: PROJECTS TO TRY .....	16
The Elements of a Good Prompt.....	16
Project 1: Drafting Social Media Content .....	17
Project 2: Summarizing and Interpreting Information .....	20
Project 3: Drafting a Call Time Follow-Up Email.....	23
CLOSING THOUGHTS .....	27
Staying Flexible and Thoughtful on AI .....	27
Next Steps and Further Reading .....	28



# EXECUTIVE SUMMARY

Generative Artificial Intelligence (AI) is a technology that will draft a plausible response to requests based on the information it has available and the context you give it.

New developments in generative AI have been arriving too quickly to track, and from the perspective of candidates and campaigns, generative AI may well seem like a superpowered threat to democracy itself. When used thoughtfully, however, it can be the opposite. After all the hype, like any machine, **generative AI is just a tool**. In campaigns, its impact is determined not by the tech, but by people, who decide how it's used, who uses it, and the norms of what's acceptable.

As the people running campaigns, political professionals in positions of leadership have more of a say over AI usage norms than they may think. We are in the moment when the conventional wisdom around political AI is being established. **The time to engage is now. The biggest question is where to start.**

The goal of these Quick Start Guidelines is to be that starting point. This document is designed to kick off Democratic campaign and pro-democracy workers' exploration of AI as a tool, with suggestions and explanations as free as possible from jargon and hype.

(For clarity, we often refer to "AI" in this document, instead of the longer term "generative AI." Though generative AI is just one kind of artificial intelligence, a lot of the principles we'll discuss apply to AI broadly. We'll note when we're talking specifically about generative AI when the distinction is important.)

We explore three main ideas to set your organization off on the right foot:

**First, there are clear, simple guardrails that can help your campaign get the most out of AI while mitigating the risks.** These are covered in more depth in the Guardrails section below, but in short:

**Don't put too much trust in AI tools.** This technology should not be deployed unattended, so always keep a human in the loop. A person should check and approve every citation, social media post, code snippet, or other output produced with generative AI.

**Don't forget AI can be biased, too.** AI is prone to the same assumptions and stereotypes that people are. Bring diverse stakeholders to the table when evaluating AI inputs and outputs, and remember that a recommendation isn't objective just because it comes from a machine.

**Don't misrepresent AI-generated materials as authentic, or vice versa.** Public trust and good faith in your campaign is too important, and too fragile, to risk.

**Don't breach the privacy and security of the people who give you their data.** Know exactly how the generative AI tools you use can and cannot use the data you provide, and keep personally identifiable information out of the mix by default.

**Prepare for deepfakes.** Deepfakes are an evolving technology, but they present similar risks as other PR emergencies. Draft a response plan and be ready to follow it.

**Second, you can prepare your organization to use AI for the long-haul by building staff feedback into the decision-making process, and by evaluating vendors critically.** It's the right moment to create an AI usage policy for your organization — with your team, not for them — and to create a clear vision for the role of AI in your movement-building work.

**Finally, there are low-investment, high-value uses of generative AI on campaigns that are worth exploring today.** For example:

Generative AI can help you quickly draft quality **social media content** — as long as you keep a professional in the loop to sense-check and finalize the draft.

Generative AI can help you get **valuable insights out of dense information** — as long as you're careful with how the model uses your private data, and always verify the conclusions it draws.

Generative AI can help you **reach more supporters, more quickly**, with messages that are consistent with your campaign — as long as you remember that people are still necessary every step of the way.

---

If there's one takeaway from these guidelines, it's that AI has a role to play in campaigns and organizing. It isn't the silver bullet that will solve every challenge, but it isn't a pure harm, either. Once again, the future of AI in politics, like the future of politics itself, will come down to the vision of the people doing the work. **As with all movement-building, the best way to start is to dive in, and the best time to start is today.**

*Matt*



**Matt Hodges**  
Executive Director  
Zinc Labs


*Ben*



**Ben Resnik**  
Deputy Director,  
Tech Strategy  
Zinc Labs

# GUARDRAILS

## WHEN DEALING WITH AI



*Generative AI is a technology that's worth exploring — but not one you should adopt blindly. As with any tool, AI creates opportunities for harm, originating both inside and outside your office, and coming from both poor usage and malicious intent. Below are some of the most common risks campaigns may encounter related to generative AI, and how to avoid or mitigate their threat.*

### 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, 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.

## GUARDRAIL 2

# Don't Forget AI Can Be Biased, Too

### WHAT THIS MIGHT LOOK LIKE

**Perpetuating stereotypes in campaign content**, like proposing a social media post about scientific pioneers that only includes men

**Uploading a job candidate's resume** to a chatbot and asking it questions to inform a hiring decision

**Relying on AI-generated language that skews straight, white, male, and college educated**, like the people who have predominantly created and trained 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, or 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 a violation of 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 are 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 around 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.*



## 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 situations.

**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 through which 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.

# INTRODUCING AI TO YOUR TEAM



*To bring AI into your organization for the long-haul, your team needs more than to understand and trust the tools. They need to be active partners in defining how AI can and should be used. That work never ends — but there are a few ways you can help build a healthy, sustainable relationship between your organization and AI from day one.*

## LOOKING INWARD

### Write an AI Usage Policy—with Your Team, Not for Them

An AI usage policy cements your organization's relationship with AI and helps ensure that its uses align with your values. When setting your AI policies, it's important to involve your staff from the very start: They are best positioned to say whether a policy will empower or undercut them, and they have the direct, lived experience necessary to know which challenges AI is needed to solve most urgently in their work. Here's a brief workflow with steps to consider when designing an AI usage policy:

**Pick a champion, and get informed.** Whether it's you or another member of your staff, empower someone to track the latest AI tools and developments, and make sure they have the resources to do so. At Zinc Labs, we regularly update a reading list of useful resources, trainings, newsletter recommendations, and expert perspectives.

**Spread the knowledge.** Give your entire organization an overview of generative AI, what it's good at, and what it's not, so they can be thoughtful participants and savvy consumers.

**Do some discovery.** Borrowing a best practice from the world of product management, conduct “customer discovery” interviews with your team to identify the areas of greatest need. The two biggest questions to ask: “What tasks do you spend the most time on in your role?” And, “If that were taken off your plate, what would you like to spend your time on?”

**Find the pressure points.** Reflect on the pain points you’ve uncovered and identify which generative AI is best suited to address. If you find something critical/urgent that doesn’t involve AI, prioritize that instead, and come back to AI later! Technology should never be a solution in search of a problem.

**Recommend and refine.** Come back to your teams with concrete AI program recommendations that have a concrete value and are based in both your sense of the organization’s needs and the themes from the discovery interviews. Bake in several rounds of back-and-forth with your team; it’s the best way to make sure you’ve gotten everyone’s perspective.

**Set the rules of engagement.** Identify the guidance your team needs to be successful in executing the programs you’ve identified. For example: What kinds of data can be uploaded to generative AI tools? Which content can be initially drafted by generative AI, and when can generative AI only be used as a thought partner? At what point do humans quality check AI outputs, and who is responsible for ensuring that happens?

**Get data, keep updating.** Generative AI is advancing too quickly for guidelines to be set in stone. Plan to refresh your AI usage policy every 6–12 months (or even more frequently if there are major new developments), and don’t be afraid to let staff solve emergent challenges with generative AI as long as they stick to the guidelines. Make sure you’re tracking what’s going well and what’s causing confusion to improve your guidance with time and keep your team actively involved.

**Always remember, empower, don’t replace.** Generative AI is a tool to empower staff, but it can’t and shouldn’t take on a person’s full role.

*A process like the one above might produce a document like [this one from the Center for Democracy and Technology](#). Notice that this policy is short and accessible, tailored to the needs of the organization, and free of overly specific rules that could quickly go out of date.*

## LOOKING OUTWARD

# Evaluate Vendors with Your Values in Mind

The inboxes of most campaigns are already full of political technology vendors offering their services, and the dawn of generative AI has made it even harder to discern which tools belong in your technology stack. After you've thought through your organization's AI usage policy, there are a few rules of thumb that may help you identify the most worthwhile conversations with potential vendors.

### WHEN SPEAKING WITH VENDORS, LOOK FOR THESE GREEN FLAGS

**There is a clearly defined, tailored use-case that would support a particular role or department on your team.** Generative AI tools are best at solving specific, narrowly tailored problems. Many high quality political AI tools will specialize in a specific task or asset.

**There are clear data privacy, security, and ownership policies.** There should be no questions regarding how vendors are allowed to use your data, and they should be able to share a formal written policy when asked.

**The tool is intrinsically human-in-the-loop.** In other words, review and feedback of the generative AI's outputs is baked into the core user experience.

**There is clear AI expertise on the team or advisory board.** Generative AI has only been in the headlines for the past few years, but AI is an area of expertise that goes back decades. It should be clear who on the team has the experience to answer the thorny data, privacy, cybersecurity, and infrastructural challenges that go into making this technology work.

**BE VERY SKEPTICAL IF YOU CATCH THESE RED FLAGS**

**The vendor frames their tool as a replacement for staff, or underplays the role of human oversight.** There's a good chance they haven't fully thought through the problem you're trying to solve or your needs as an organization.

**They make too-good-to-be-true promises on efficiency gains or ROI.** Generative AI may be able to help your team, but it probably won't, for example, reduce work time on a project by 99 percent or quadruple your fundraising overnight. As with any tool, if something seems too good to be true, it probably is.

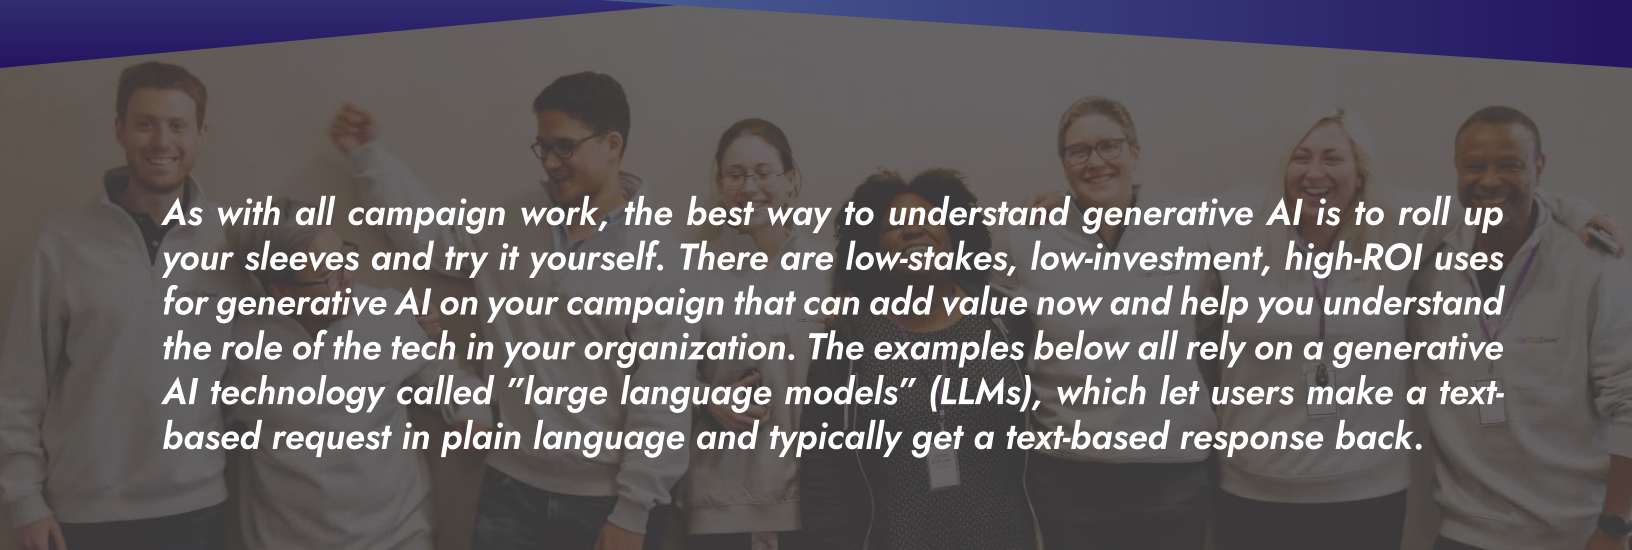
**They have unclear data ownership policies.** Again, it should be clear and explicit how vendors may and may not use your data. A vendor that can't produce a clear, written answer to this question is at a higher risk of exposing your sensitive or proprietary material.

**There is no credible source of AI expertise.** If no one affiliated with the vendor has worked on AI projects before, they likely have blind spots in their product or process, which could have security, quality, or reputational implications for your organization down the line.

*For further reading, take a look at [this case study](#) on how MoveOn has built an AI application approval process into its workflow. You'll notice that it's built on a culture of transparency, collaboration, and experimentation.*



# GETTING STARTED: PROJECTS TO TRY



*As with all campaign work, the best way to understand generative AI is to roll up your sleeves and try it yourself. There are low-stakes, low-investment, high-ROI uses for generative AI on your campaign that can add value now and help you understand the role of the tech in your organization. The examples below all rely on a generative AI technology called “large language models” (LLMs), which let users make a text-based request in plain language and typically get a text-based response back.*

## The Elements of a Good Prompt

When it comes to generative AI, one of the most valuable tools in your kit is the “prompt”: the specific request you make of the service. Prompting is as much an art as a science, and the latest guidance changes often, but there are a few characteristics that effective prompts have in common:

**Clear Instructions:** What exactly you want the generative AI system to do, and how. Consider this the definition of success for the tool. The more direct and prescriptive, the better.

**Details:** Additional context that will make the response specific, not generic. (Avoid including personally identifiable information when using commercial products without data controls.)

**Format:** How the response should be structured. Consider this another element of the definition of success.

**Examples:** Providing a few samples of what a good output looks like can improve the quality of a generative AI tool’s output, and get it a step closer to writing in your campaign’s voice.

*OpenAI’s prompting guide is a great resource to go deeper and keep up with best practices as they evolve.*

FURTHER  
READING



## PROJECT 1

# Drafting Social Media Content

### WHY TRY IT

It's hard to keep up a steady stream of quality online content that connects with your constituents. Generative AI models can jump-start your thinking and get you to the editing phase faster.

### WHAT TO TRY

Give a prompt like this to a chatbot such as [ChatGPT](#), Google's [Gemini](#), Microsoft's [Copilot](#), or Anthropic's [Claude](#):

Write a one-paragraph Facebook post for a State Senate candidate in Pennsylvania celebrating the Fourth of July and tying together the themes of the holiday with the theme of "an economy that works for everyone." The tone should be casual but passionate, and the target audience is suburban women.



Here's a real result:

*Happy Fourth of July, friends! As we celebrate our nation's independence today with backyard barbecues, parades and fireworks, let us also reflect on the principles of freedom, equality and opportunity that led to the American Revolution. Our founders envisioned a nation where every person has the chance to life, liberty, and the pursuit of happiness. But too many Pennsylvanians today feel left out of our economy, struggling with stagnant wages, rising costs, and fewer benefits. In the State Senate, I'll fight for economic policies that live up to our highest ideals - quality jobs that pay family-sustaining wages, affordable healthcare, excellent schools, and a secure retirement. The Fourth of July reminds us that when we pull together and lift each other up, we can build that more perfect union. Our independence and prosperity as a state depends on an economy that works for everyone!*



This post certainly isn't perfect — it's a little long, for starters — but it matches the format, content, and structure that we were going for. And because we specified the audience as well as the tone and policy point we wanted to make, it's able to take on a persona (notice how it mentions "when we pull together") and transition smoothly into an argument that feels powerful and relevant to your intended audience. This draft is a good starting point for you to clean up, refine, and make your own.

### WHY IT WORKS

The request above succeeds because it contains three of the **elements of a good prompt**.

**Instructions:** Platform, points to make, tone

**Details:** Speaker, audience, event

**Format:** Length, structure (one paragraph)

Here's an example of a less effective prompt for this same exercise, which lacks these elements:

Write me something nice for Fourth of July.



And here's the response:

*Happy Fourth of July!*

*As we gather under the summer sky, may the crackle of fireworks echo the joy in our hearts. Today, we celebrate not just the birth of our nation, but the spirit of freedom and unity that binds us all. May the stars and stripes remind us of our shared history and the bright future we build together.*

*Let the barbecues sizzle, the laughter ring out, and the flags wave high. May this day be a tribute to the courage of our forefathers, the dedication of our heroes, and the endless opportunities that lie ahead. As the fireworks light up the night, let's hold our loved ones close and remember the true meaning of this day.*

*Wishing you a Fourth of July filled with pride, joy, and an abundance of happiness. Let's celebrate the land of the free and the home of the brave!*

*Happy Independence Day!*



This post is far too long because we didn't specify the length or structure, and the lack of information on audience and tone results in something that reads as generic and forced. There's also no mention of a fair economy, because we didn't give that guidance. It will take substantial edits to make this response useful — maybe even more time than it would have taken to write the post from scratch.

### WHAT NOT TO DO

**Don't use content without human review and feedback.** Generative AI in many ways resembles an advanced version of auto-complete: It gives the answer that seems likeliest to fit in with the words around it, but it doesn't know what you're trying to say — or even if its answers are true or false. As a result, generative AI may make things up in its responses, and deliver them with the same confidence as it would delivering the truth. That means generative AI can and will get things wrong, from basic facts to the nuanced ideas and tone of your campaign, and even *explanations of its own behavior or logic, which you can't take at its word*. It's a helpful writing aid to workshop ideas and get to the editing stage faster, but if you completely outsource your online communications with your voters to generative AI, you waste those relationships and risk losing your audience's trust.

### WHAT TO DO NEXT

**Iterate.** Thinking about the levers you have to pull (Instructions, Details, Format, and Examples), give the tool feedback on the draft, such as "Make this more casual" or "Include the word 'justice.'" Generative AI tools often don't get it right the first time, but their outputs improve with guidance, and you'll learn what to include in the prompt to get a better draft. Most generative AI models can refer back to what you've talked about over the course of a session (especially if it's a shorter session), so you can give edits and feedback.

**Edit.** Editing AI-drafted content requires a shift in focus compared to traditional copyediting. While AI excels in areas like grammar, it is more prone to errors with proper nouns, contextual details, and numbers. Editing should prioritize verifying factual accuracy and the appropriate use of terms. This adjustment doesn't diminish AI's usefulness but is a shift from how people usually edit content.

**Switch platforms.** Once you have a draft you're happy with, have the AI tool convert it into a post for a different social media platform, or a script for a direct-to-camera video. Generative AI tools are effective at changing up formats once you've defined success for them.

**Quickly build a content backlog.** Draft a list of state, federal, and local holidays and have the AI tool draft social media posts for all of them. If you describe what ideas and themes to include, it will produce content that feels more specific to your campaign.

## PROJECT 2

# Summarizing and Interpreting Information

### WHY TRY IT

Political campaigns often rely on easily accessible metrics like doors knocked and donations raised. Generative AI can unlock richer insights to inform campaign strategy and communications in new ways.

### WHAT TO TRY

**Make dense information more accessible.** Have a generative AI tool summarize a bill you support in plain language, and make it relevant to a broader audience.

Here is an example prompt:

Summarize the attached piece of legislation at a 10th grade reading level, and suggest 3 or 4 examples of how it may make the lives of specific groups of Americans better.



Once again, this prompt contains several of the components of a good prompt:

**Instructions:** Clear tasks (summarize, then suggest), number of examples

**Details:** Document to analyze, nature of the document (legislation), exact reading level of output, request for specific examples tied to specific groups

**Format:** Summary, then examples

**Pull out strategic insights.** Paste in your campaign's policy platform or "Issues" page, as well as the mission statement or "About" section of an organization that is considering endorsing your campaign. Have it identify the points or keywords of strongest alignment between your candidate's positions and those of the organization. Consider what you find as you fill out their endorsement questionnaire or prep talking points for your candidate ahead of a meeting.

**Process information faster.** Download the transcript of a YouTube video (there are several free online tools that will do this for you) and have a generative AI tool call out any statements that touch on a certain subject or may be controversial. Once again, verify the output is faithful to the source material before doing anything else!

### WHY IT WORKS

**This experiment plays to one of generative AI's greatest strengths: Rephrasing and remixing existing information.** Generative AI is effective at finding new ways to say the same thing, which is great for rephrasing an idea, simplifying content, or making materials accessible to audiences you may not have had the time to write for previously.

### WHAT NOT TO DO

**Never give private or proprietary information to a generative AI tool before you know how that data may be used.** Generative AI tools work by taking in a huge amount of data, finding patterns in that data, and creating answers based on what has appeared before in that pattern. As a result, these tools will sometimes save users' input, add them to their huge dataset, and may use that data to respond to other users in the future. That means, under unpredictable circumstances, a generative AI tool could emit private or confidential information if it's the best response to someone else. If you've given AI a draft strategy document or voters' contact information as input, for example, you risk exposing that information to the public at random.

To protect against this, always check how the tool or service is allowed to use your information — look for explicit statements like “We will not use your data to train our models” in places like their terms of service, as well as any language that carves out exceptions.

Using data for training shouldn't always be a dealbreaker. For example, allowing a generative AI tool that only serves Democrats to train their model on your data will mean better outputs for you and the movement at large in the long term. But you should always go in informed and make a conscious choice about data privacy.

## WHAT TO DO NEXT

**Write a clear AI data policy.** Set clear boundaries for who in your organization can share what data to what tools for what purpose. Having simple, understandable redlines about what your staff can't do with AI will help avoid mistakes — and laying out what your staff can do with AI will encourage them to explore and find creative solutions to challenges. We have shared more guidance on this in the previous section.

**Improve your internal materials.** Once you are comfortable with a generative AI tool's usage policies, give it the selected or non-sensitive documents you use to explain new information to people on your campaign — training materials for volunteers, for example, or policy talking points for voters — and ask for ways to tighten or clarify them.

**Turn canvassing conversations into intelligence.** Ask your volunteers to write or dictate brief notes after each voter contact while canvassing. Have an AI tool review the content and find emerging themes and common points of friction. (Note: This is one of the instances where a generative AI model is most likely to make up data, also known as “hallucinating.” To mitigate this, request exact quotes from the reports that justify its analysis, and then, crucially, search for those quotes to confirm the data is real.) If you're confident in the results, consider integrating what you learn into your walk scripts or other campaign programs.



## PROJECT 3

# Drafting a Call Time Follow-up Email

### WHY TRY IT

Building a donor network is the lifeblood of most campaigns. But there's often a tradeoff between quality and quantity — crafting personal outreach is time-consuming, meaning you may touch fewer prospects or fail to get committed money in the door before an important deadline. It is a general fundraising rule of thumb that the quicker you execute your follow up, the more likely you are to receive a contribution. Generative AI can help thread that needle, letting you craft messages that are more tailored than a stock email response but quicker than a note written from scratch.

### WHAT TO TRY

Prepare three sources of information: 1) A follow-up email template (as a Word document or other text file) with pre-approved language to send from the candidate after a call time conversation; 2) A bullet point list of notes from the conversation, included as plain text or an attachment; 3) Any relevant links (call time contribution link, relevant articles, candidate website). **(Note: Only do this after you have carefully reviewed the platform's data usage and privacy policies, as discussed in the last project!)**

Share this information with the generative AI tool. Then have the tool draft a follow-up email based on the call details and email template, along with references to your campaign's talking points as appropriate. Be sure to specify a request or call to action you'd like to include in the follow-up.

**Note:** Make sure your template email includes the full URL of any links but do not include the actual contact information for yourself, your candidate, or the donor — you can add that yourself later.

**Note:** Some AI models can embed links directly into the text with direction, but you should still have a human double check them all before the email is sent.

Here's an example prompt:

Please review the below bullet points and use the attached email template to draft a follow-up email to a prospective donor, Elizabeth, from the perspective of a candidate, Kyra Johnson, based on the notes. Stick exactly to the Email Template document, only customizing the elements [inside the brackets]. Extra info should be no more than 1-2 sentences. Reference the ideas in the sample talking points document as appropriate. Include an ask of a \$3,300 contribution to the campaign with a link to donate (<https://example.com/donateKyraforCongress>) and embed all links directly in the email.

Here are the notes:

- Elizabeth is traveling to the Bay Area next week
- Elizabeth and her husband Todd have a wedding anniversary next week.
- Elizabeth cares deeply about electing more women to Congress
- Elizabeth grew up in Texas



And here's an example response:

*Dear Elizabeth,*

*I'm writing to follow up on our conversation just now. Thank you again for taking the time to connect. I enjoyed hearing about your plans to travel to the Bay Area next week and your upcoming wedding anniversary with Todd. Your passion for electing more women to Congress and your roots in Texas inspire me and align closely with our campaign's goals.*

*As I mentioned on the phone, I am running for Congress in GA-14 to defeat Marjorie Taylor Greene. This week, the Cook Political Report updated our race from Likely Republican to Leans Republican and cited me as MTG's toughest opponent yet. You can read the article [here](#). I've also attached the campaign's prospectus below with more information on the race.*

*And thank you for considering a contribution of \$3,300 to my campaign before our upcoming deadline on March 31. You can do so at [this link](#).*

*Please feel free to call me back at XXX-XXX-XXXX or email Hannah on my team at the email above if you have questions. You can also learn more about our campaign on [our website](#).*

*Thanks again. Look forward to staying in touch.*

*Best,  
Kyra*



## WHY IT WORKS

**This experiment again follows the three components of a good prompt we're already familiar with:**

**Instructions:** Clear task (read, then draft), reference particular materials at particular points in drafting process

**Details:** Author, reference materials, ask, tone, style, any links as full URLs

**Format:** Follow-up email template

**It adds the fourth pillar of a good prompt: Examples.** The email template gives clear guidance on structure, word choice, and the relationships between ideas that generative AI tools are optimized to find, and discourages the model from improvising language that may not be right for your campaign. This further clarifies the definition of success for the generative AI model, saving you time while keeping your message consistent.

**It keeps a human in the loop.** Throughout this process, human oversight is inherently necessary to take notes, select appropriate talking points and personal details to call out, decide on the ask, edit and personalize, catch flaws or deviations from the template, and ultimately make the draft effective. Without a human at the controls, this entire process simply would not work.



## WHAT NOT TO DO

**Do not try to replace staff with AI.** Misguided organizations might interpret generative AI's operational efficiencies as a rationale for workforce reductions, potentially substituting junior fundraising positions with automated solutions. This is a short-sighted approach that contradicts our core values as Democrats. Our dedication to advocating for the dignity of work and fostering economic opportunity is fundamental to our work, and replacing staff with AI would not only diverge from these principles but also introduce substantial risks. Such a move could undermine the stability, integrity, and trust we strive to maintain with our supporters and the broader public. Moreover, from a strategic standpoint, relying on unsupervised generative AI creates a risk of embarrassing missteps that could erode trust in a campaign. In an era where authenticity and connection are paramount, preserving human insight, oversight, and judgment in our operations is not only a reflection of our values, but a strategic imperative. Generative AI tools should upskill and empower our teams, not replace them.

**Again, be extremely careful sharing any proprietary or personal information with a generative AI tool.** As we discussed in the project above, you should never share anything private unless you know exactly how the AI service provider is allowed to use your data. By default, scrub contact information — both yours and the donor's — such as email and cell phone numbers from your documents before uploading them.

## WHAT TO DO NEXT

**Create a writing sample document.** If you want to be less prescriptive, you can share a few writing samples instead of a template, and the generative AI model will attempt to replicate the writing style you have shown it. The more examples of good content in the voice of your campaign/candidate the tool has to work from, the higher the quality its drafts will be.

**Run a test.** Measure the average turnaround time between your candidate finishing a conversation during call time and sending the follow-up email before and after this process, or use another metric your team finds valuable. Hard evidence will help you decide whether an approach like this should be part of your campaign long term.

**Land and expand.** A process like this can translate elsewhere within a department (for example, to small-dollar donor outreach) or across departments (for example, to grassroots relationship management). Who else in your campaign faces similar challenges to what you're tackling here?

# CLOSING THOUGHTS

## Staying Flexible and Thoughtful on AI

In September 2023, tech investor Rex Woodbury [observed](#) that when the iPhone debuted in June 2007, just 6 in 100 Americans owned a smartphone. A year later, driven in part by the iPhone's success, smartphone adoption had nearly doubled — to 11 percent. But the ideas that put smartphones in the pockets of nearly every adult were still years away: WhatsApp wasn't released until 2009; Uber and Instagram until 2010; Snapchat until 2011; and TikTok until 2016. Tech observers in 2007 knew something big was coming, but nobody could have guessed what.

Generative AI is at a similar moment. The 2024 elections are taking place only two years after generative AI first hit national headlines, but the full effects of this emerging technology on campaigns and organizing — for better and for worse — may not be clear for many more years to come.

For now, one of the best things you can do is to keep an eye on what's developing. Below are a few areas to watch especially carefully:

**Voter data:** A host of organizations are exploring how generative AI can improve staffers' relationship with data. It's possible that creating custom audiences by writing super detailed Boolean queries could soon be a thing of the past, replaced with plain-language descriptions of the people you want to reach.

**Fundraising:** One of the very first uses of generative AI in politics was drafting content for mass fundraising emails. The tech is now being applied to other elements of fundraising, from drafting custom donor follow-ups to quickly and cheaply creating videos for digital ads.

**Staff support:** Political AI is also positioned to take a lot of tedious work off of staff and volunteers' plates. Tools are already available, for example, to pre-write responses for volunteers doing peer-to-peer texting. (The volunteers have the chance to review and edit before hitting send.) And one of the most popular uses of generative AI in the commercial world has been the creation of internal knowledge bases, which organize files and let staff answer questions with plain text prompts that previously required distracting back-and-forth in email or Slack.

Beyond that, the best thing to do is to consider what should come next. Once again, AI is a tool: It's in your and your team's power to decide if, when, and how to pick it up. It may be impossible to predict the future of generative AI — but there is power in defining, and working toward, a future worth striving for.

## Next Steps and Further Reading

If you have found these Guidelines helpful, please share them with your network.

This document will be updated periodically. Check back from time to time for the latest resources and guidance. For those updates, as well as additional resources and ways to get in touch with us about AI for Democratic campaigns, visit [AICampaignGuide.com](https://AICampaignGuide.com).

Finally, take a look at these resources below and throughout these Guidelines if you want to go deeper.

*Cooperative Impact Lab has a series of case studies on [how political organizations and nonprofits are using AI today](#).*

*The City of Boston was early to release [sensible, high quality guidelines](#) for city employees. They're worth reviewing as you design your own AI policy.*

*The Movement Cooperative has also published [a thoughtful guide](#) on this subject.*

*"[Preparing for Generative AI in the 2024 Election](#)" contains strong recommendations for both offense and defense around generative AI.*

*If you want a good look under the hood, "[What We Know About LLMs](#)" is an excellent primer on core concepts in generative AI.*

*[Higher Ground Labs' AI Landscape Report](#) has a detailed overview of major AI vendors in politics today, and a discussion of where the technology is going.*



LEARN MORE AND  
DOWNLOAD THE GUIDE AT

**AlCampaignGuide.com**



POWERED BY  **ZINC**  
COLLECTIVE  
[zinccollective.com](http://zinccollective.com)

