

Written Testimony

of

Dr. Beth Simone Noveck

Chief AI Strategist, The State of New Jersey

before

The United States

“The Future of Constituent Engagement with Congress.”

United States House of Representatives

Longworth House Office Building

Statement before the House Subcommittee on Modernization and Innovation

Introduction

Chairwoman Bice, Ranking Member Torres, and Members of the Subcommittee, thank you for the opportunity to testify today.

My name is Beth Simone Noveck. I am the Chief AI Strategist for the State of New Jersey, having previously served as the State's Chief Innovation Officer.ⁱ I founded [InnovateUS](#), a nonpartisan, nonprofit that has trained over 125,000 public servants in AI, innovation and engagement skills over the last year. Through our partnership with 40+ states and cities, we expect that number to increase dramatically. I also direct [The Governance Lab](#), a research center that works with governments around the world to design more participatory and effective democratic institutions. In my work, I have spent the last twenty years helping governments—from small municipalities to national parliaments and presidents—use technology to strengthen and modernize how they engage with their publics.

I am grateful for your leadership in holding this hearing and for the Subcommittee's commitment to modernization. You have correctly recognized that Congress requires new tools and better approaches to tapping the collective intelligence and expertise of the American people to enhance our lawmaking processes.

Today, I want to offer both a sense of what is possible and concrete steps that the Subcommittee can take in the coming year to make constituent engagement in lawmaking—what scholars and activists dub “crowdlaw”—more effective and more efficient.ⁱⁱ

At precisely the moment when Congress is most overworked addressing myriad complex and interdependent challenges, we can use the right artificial intelligence tools to enable Congress to listen better, learn faster, and govern more effectively.

Two recommendations from this Subcommittee (Recommendation 174 to study and present options for developing a public-facing interactive platform for constituents to offer their opinions and feedback on pending legislation and Recommendation 176 to study and present options for developing a platform for committees that want to solicit public comment and evidence) provide the perfect place to begin that transformation.

I. Engagement Must Evolve

Shrinking Staff, Increasing Volume of Communications

Each year, Americans send Congress more than 81 million communications—emails, calls, demands for information, tour requests, feedback on legislation, and pleas for help.ⁱⁱⁱ Yet Congress is trying to respond with a staff that is dramatically less than it once was. Committee staff are roughly 40 percent smaller than they were in 1980. Congress’s in-house expertise—at CRS, GAO, and CBO—has shrunk by 45 percent since the mid-1970s. Funding for the legislative branch has grown only 8 percent over the last decade, while 27 percent of the budget now goes to the Capitol Police and to building renovation.^{iv}

Meanwhile, the Supreme Court’s 2024 *Loper Bright* decision curtailing agency deference has shifted responsibility for regulating some of the most complex areas of our economy from career specialists in federal agencies to congressional offices that are already stretched thin. Finally, Americans spend far more on the elections selecting our members of Congress than taxpayers spend supporting the work of making laws.^v Congress simply does not have the staff, time, or analytical capacity to adequately fulfill its constitutional role.

Against that backdrop, “public engagement” can sound like an unreasonable burden. When you can barely staff your offices or answer the phone, how can you invite thousands of people to participate in the legislative process?

On the Web, More Talking Means Less Listening

Indeed, past efforts have shown why such caution has been justified. My colleagues and I who served on the presidential transition team in 2007 invited Americans to propose ideas for the first hundred days of the incoming presidential administration. We received over 125,000 participants and 44,000 ideas. Buried somewhere in that flood were brilliant proposals—but we had no way to find them, let alone act on them.

With the explosion of Internet technologies and social media, the more people participated, the harder it became for institutions to listen. This is why so much public engagement has been in the form of pilots, tried once and rarely repeated.

Other governments have experienced the same bottleneck that depresses public engagement.

- The UK Cabinet Office estimates that a consultation receiving 30,000 responses required 25 analysts working three months to process.^{vi}
- Before the German city of Hamburg modernized its digital participation system, just one public engagement could consume five full-time employees for more than a week.^{vii}
- Singapore Together, launched with great promise in 2019, held a handful of deliberative, 50-person mini-publics before abandoning even these small exercises.^{viii}

Processes Have Been Designed to Increase Inputs Rather than Outcomes

Historically, meaningful participation has been too expensive and too time-consuming. But poor design also impedes institutionalization. Too many processes lack structures that make public input usable for decision-makers. *Decide Madrid* illustrates this challenge. More than 460,000 residents created accounts and submitted over 28,000 proposals in the first seven years of the city’s engagement platform. The design of the process, however, included the requirement to secure support from 1% of the population, and over a year of waiting from submission to final consideration. As a result, only one proposal has ever been enacted, leading to wasted time and squandered trust.^{ix}

Taking Public Expertise Seriously

Compounding these design flaws is a long-standing skepticism about whether ordinary people have meaningful expertise to contribute. As a result, too many engagement exercises are built around conversation without consequences—inviting people to talk, deliberate, or “have a voice” without creating any pathway for their knowledge to shape decisions. These processes may generate goodwill or headlines, but they rarely generate usable intelligence for institutions. The implicit assumption is that public input is either frivolous, polarized, or uninformed, and therefore must be kept at a safe distance from actual policymaking.

Yet every successful example shows the opposite: when institutions ask clear, purposeful questions and create structured ways for people to share their expertise or lived experience, the public produces high-quality insight that decision-makers can use.

I learned this lesson firsthand when we built one of the first expert-intake systems in the federal government for the Patent Office: Peer-to-Patent. Patent examiners needed targeted technical information — not general opinions. By structuring public input around a specific question (“Is there relevant prior art?”), the public contributed high-quality, actionable know-how. After the US, we rolled out the program in the UK, Japan, South Korea and Australia.^x The success across

multiple countries demonstrated that ordinary people can meaningfully contribute to complex decision processes — when the engagement is designed well.

When institutions ask, “What do you think?”, they receive broad sentiment. When they ask, “What evidence should we consider?” or “Which provision affects your community and how?”, people supply knowledge that improves decision-making.

Now AI makes it possible to process large volumes of input, but designing the process to elicit citizen know-how is what makes that input useful. The problem has never been a lack of public expertise; it has been a lack of processes designed to take that expertise seriously.

II. AI Makes It Possible to Listen at Scale—and to Learn from What We Hear

AI Can Make Engagement More Efficient and Effective

Because these next-generation word processors have ingested trillions of words, they can spot patterns in language, making it possible to turn even unstructured input into actionable knowledge.

AI tools make it possible to:

- Synthesize thousands of comments in minutes
- Eliminate identical submissions
- Cluster and categorize input
- Sort and organize spoken as well as written words
- Detect points of consensus and disagreement and other patterns
- Extract and distinguish ideas from opinions

I should note here that the New Jersey State Office of Innovation just won one of only eight grants (out of more than 400) from the Center for Civic Futures to expand our development of AI tools, one of which specifically focuses on summarizing feedback.^{xi}

Around the world, we are seeing governments use AI to tackle information overload at a scale and speed that were unimaginable even five years ago.

- The City of Hamburg’s open-source engagement platform uses AI to cluster comments, extract core ideas, and link input to geospatial maps.^{xii}

- Brazil’s Ulysses legislative AI platform helps staff analyze bills, match them to relevant subject matter experts and prior analyses, flag conflicting amendments, and generate clear briefings for a legislature larger than the U.S. Congress.^{xiii}
- Chile’s CAMINAR legislative system automatically digitizes, compares, and analyzes legislative proposals for constitutional compliance.^{xiv}
- The European Parliament’s Archibot, built on modern large-language models, enables citizens, students, and staff to explore legislative history with unprecedented ease, increasing both transparency and comprehension.^{xv}

III. Implementing Recommendation 176: Soliciting Public Input for Committee Hearings

One of the most achievable, low-risk, and high-impact places for Congress to begin using AI to enhance public engagement is Recommendation 176, which asks the House to study ways for committees to solicit public comment and evidence on matters before them. Congress does not need to invent the process or platforms from scratch. There are well-tested models and open-source, AI-enabled tools that make this kind of engagement both manageable and genuinely useful for Members and staff.

Brazil’s e-Cidadania: A Model for Public Participation in Committee Hearings

Brazil’s Federal Senate offers the strongest real-world model for how congressional committees can meaningfully solicit public questions for hearings. Since 2015, the Senate has invited citizens across the country to submit questions for witnesses through its e-Cidadania portal. For each committee hearing, staff open an “Interactive Event” on the Senate website where citizens submit questions both in advance and during the livestream. A toll-free phone line enables participation from those without internet access. Staff review submissions, categorize them by topic, and relay the most relevant questions to the committee chair, who retains complete discretion over which questions to use and how to integrate them into the hearing.

To improve the quality and diversity of input, staff engage in targeted outreach. They notify individuals who previously participated on related topics—teachers, scientists, farmers, small-business owners, civil society groups—inviting them to submit questions for hearings aligned with their expertise or lived experience. This targeted notification strategy increases participation three- to four-fold and dramatically improves the relevance of questions received. It ensures that committees hear not only from the most mobilized constituencies but also from people with direct knowledge of the issues under consideration.

Brazil's Improvements with AI

Processing thousands of questions per hearing requires curation and judgment. Brazilian Senate staff must filter inappropriate content, avoid redundancy, ensure that selected questions reflect diverse viewpoints, and surface those most useful to the committee's oversight goals.

To support this work, the Senate has begun integrating AI tools that help staff triage submissions more quickly and effectively. AI software summarize and cluster similar questions, highlight novel or underrepresented perspectives, and flag submissions that directly address the central issues under discussion. These tools do not automate decision-making—they accelerate sense-making. Staff retain full control, but AI reduces the manual burden of sorting, enabling them to identify the most constructive questions in time for them to shape a hearing's flow.

Participation has been substantial—46,000 questions across 546 hearings in 2023, and 69,000 questions across 440 hearings in 2024—and remarkably, senators responded to roughly 90 percent of them. Public input is now so routine that, as the head of the e-Citizenship office notes, “If we are late, Committee chairs ask, ‘Where’s the list of questions?’”^{xvi}

Looking ahead, the Senate is developing additional AI-enabled features: chatbots to guide first-time participants through the submission process, multilingual support, and the ability to process voice and video submissions. Together, these improvements lower barriers to participation while giving staff better tools to manage input at scale.

US Precedents

Prior to generative AI, Congress experimented with public engagement in committee proceedings. The Federation of American Scientists (FAS) and its [Congressional Science Policy Initiative](#) (CSPI) invited subject-matter experts including scientists, public health officials, technologists and epidemiologists to submit questions for congressional hearings. CSPI recruited more than 600 volunteer scientists from across the United States who, collectively, contributed information and questions to more than 40 Congressional hearings on topics ranging from clean energy technology to Facebook's digital currency to testing for COVID-19.^{xvii}

In March 2020, the FAS, in partnership with the New Jersey State Office of Innovation and The GovLab, also called upon its community of scientists as part of Ask A Scientist. Ask A Scientist aimed to counter public misinformation about COVID-19 by offering scientist-led advice to the public about the disease and seeks to ease the burden on government agencies fielding large volumes of generic public enquiries. Members of the public could ask questions via the Ask a Scientist website and Alexa app, such as “How can I prevent myself from getting the virus?” or “If I get infected when do symptoms appear?” and the tool would serve them answers prepared

and reviewed by FAS scientists, the network of volunteers from the National Science Policy Network and then edited and translated by the NJ Office of Innovation. Where the tool could not serve a pre-prepared answer to a question, we shared the questions with the network who researched, wrote and returned an answer a day later. Ask A Scientist provided evidence-based responses to more than 1,000 questions during the early days of the pandemic.^{xviii}

What Congress Can Do Now

This subcommittee can begin testing the use of AI to implement Recommendation 176 through a simple, low-risk pilot conducted over the course of a few hearings. The goal would be to develop a manageable workflow that allows Members to benefit from public insight without adding burdens to staff.

Prior to each hearing, the Subcommittee would open a short engagement window via its website inviting the public to propose witness questions.

The website should explain the hearing topic and the goals, specify what makes a constructive question and invite the public to write a question. The form can ask three questions: 1) the proposed witness question, 2) an explanation of the question's importance, and 3) any evidence to support and inform the question.

Until recently, free-text public submissions were nearly impossible to process at scale. Today, AI makes it straightforward to de-duplicate repeated questions, remove off-topic or inappropriate submissions, cluster questions into major themes, and flag those from individuals with relevant expertise or lived experience. Communities across the United States are already doing this:


- **Bowling Green, Kentucky** used a free, open-source AI tool to cluster and analyze more than 8,000 comments during an online engagement about the city's future growth. AI grouped and labeled ideas, identified areas of agreement and disagreement, and produced a concise summary that local leaders said helped them "more accurately represent the community."^{xix}
- After the **Los Angeles wildfires**, the State of California used AI-assisted analysis to process thousands of narrative submissions from affected residents as part of its Engaged California project. AI identified major themes and tensions—such as rebuilding quickly versus rebuilding resiliently. Organizers used AI to rewrite responses into consistently-formatted and actionable policy recommendations.^{xx}

Using existing, free, open-source, fine-tuned AI tested in the similar contexts, subcommittee staff can cluster, organize and sort questions to present to Chair, Ranking Member, and other Members.

In addition to selecting individual questions, staff can use AI to prepare a short briefing memo summarizing themes across all submissions: the major concerns raised by the public, a thematic breakdown of questions, expert or technically informed submissions, and outlier or innovative perspectives.

Importantly, committees retain full discretion over which public questions—if any—to use. Nothing in this process obligates Members to ask a public-submitted question. The pilot simply expands the pool of high-quality, evidence-informed questions available for oversight and examination, while giving staff a manageable, structured workflow for handling input.

With clear directions, time-boxed engagement, AI-supported clustering and summarization, and some help getting the word out, this or another subcommittee can test what works.



Support from AI for Impact

The “[AI for Impact](#)” students I teach from Northeastern University and the New Jersey Council of Community Colleges – who spend six months embedded with government partners like the State of New Jersey, State of Massachusetts and the City of Boston – recently built a free, open-source tool called Open Feedback for the City of Boston to sort, organize, and label citizen feedback submitted through Boston.gov.^{xxi}

Developed this fall, the tool has been deployed publicly by the City. Residents can type in their feedback in natural language. The AI assistant responds with clarifying questions to improve the submission.

On the “back end,” the tool allows staff to sort and organize the comments, extract and analyze sentiment and topic. Boston customized the tool for use in connection with feedback on permitting and it can distinguish between operational, regulatory, communications and other issues and route the comment to the right place. Past comments can be uploaded in bulk for later sorting. The tool also allows staff to create visualizations of all submissions to see patterns.

1. Resident Feedback Intake

Making it Easy for Residents to Share Their Feedback!

How was your experience with permitting?

Need help? Try answering:

- What specific feature or service would you like to see improved?
- Can you describe a problem you encountered?
- What frustrated you about your recent interaction?

Getting the zoning feedback for my business was expensive

Can we contact you with questions about your feedback? *

☒ Yes ☐ No

Submit

YOUR PERSONAL FEEDBACK ASSISTANT

We need more information about your feedback!

Question 2/5 40%

☒ SYSTEM

Could you please provide more specific details about the costs associated with the zoning feedback? For example, how much did it cost and what specific fees or processes were involved?

It was \$200!

☒ SYSTEM

Could you clarify what specific zoning?

Additional details needed

Type your response...

Skip questions and submit feedback

Step 1: Residents speak naturally in their own words (optionally provide contact info)

Enter Feedback Enhancement Extraction Classification Visualization Interact

10

2. Resident Feedback Enhancement

Making it Easy for Residents to Share Their Feedback!

Step 1: Residents speak naturally in their own words (optionally provide contact info)

Step 2: Chatbot improves feedback by clarifying details from user

A system is only as good as the data it gets

Enter Feedback
Enhancement
Extraction
Classification
Visualization
Interact

11

The next AI for Impact cohort starts in January and we would be happy to adapt and fine-tune such a free AI tool for you to use this spring in connection with piloting the use of AI to implement Recommendation 176.

This pilot can be launched within weeks and at little to no cost, enabling the House to evaluate its usefulness and refine the process before scaling it across committees.

IV. Implementing Recommendation 174: A Public-Facing Platform for Feedback on Legislation

Recommendation 174 directs the House to explore ways to gather public input on draft legislation after introduction but before passage. The goal is not mass polling or an open comment box. It is to design a structured, targeted, and time-limited process that allows Members and staff to benefit from public expertise—both credentialed and lived experience—without creating unmanageable workloads.

Structured Participation: The Peer-to-Patent Lesson

When we created the Peer-to-Patent project to help overloaded patent examiners we did not ask the public whether they liked a patent; we asked whether there was prior art that the examiner should consider, and we anchored contributions to specific claims. That request for evidence

produced clear, actionable input with minimal spam that improved the quality of government decision-making. We repeated this process when I joined the White House (and later in New Jersey) and we crafted our Open Government (and other) policies with large-scale public input.

Congress can apply the same principle to legislative drafts: Ask clear, specific questions tied to particular sections of a bill, and provide an easy way for people to supply relevant information. When the contribution pathway is clear, the public will provide exactly the kind of expertise that staff need.

AI Enables Sorting of Comments on Drafts

Several countries have already shown how to gather public expertise during the drafting process efficiently.

Germany has experimented with inviting experts and civil society organizations to annotate draft laws directly, tying each suggestion to a specific clause.^{xxii} Because contributions are anchored to the text, staff can quickly review where language may be unclear, where unintended consequences may arise, and what evidence supports proposed changes. The structured nature of the contribution—not the number of participants—is what makes the input usable.

Iceland’s Better Reykjavik platform offers another instructive model. Residents propose and refine policy ideas, improving one another’s suggestions and articulating the evidence behind them. City officials receive a set of workable, well-developed proposals, rather than thousands of disjointed comments.^{xxiii} The value lies in the refinement and structure of contributions, which is exactly what legislative offices need when reviewing bill language.

Around the world, parliaments, legislatures as well as executive branch agencies are turning to AI-enabled platforms to guide contributors through structured prompts (“Which section are you commenting on?” “What evidence supports your point?”) and to synthesize public input into formats staff can use quickly. The AI helps to provide clarity, structure, and a defined pathway for how public input will be used.

What Congress Can Do Now

Congress can test Recommendation 174 through a simple, low-cost pilot built around a single bill.

A Committee or a Member can begin by selecting a draft bill—or even a single subsection—and inviting the public to provide input through an online form by text or even by voice. Rather than asking, “What do you think?”, the form would pose targeted questions such as: “Where is this

language unclear? What impacts should the committee anticipate? What evidence supports your suggestion? What alternative wording would improve clarity or accuracy?” These prompts elicit expertise rather than general sentiment.

AI can then assist staff by organizing and synthesizing the submissions. Modern tools can turn speech into text, de-duplicate repeated comments, cluster feedback around themes or specific provisions, identify submissions supported by evidence or experience, and flag potential conflicts or unintended consequences. The result is a concise briefing that staff can review in minutes rather than hours.

To close the loop, the committee can publish a short “What We Heard” summary explaining which suggestions were incorporated, which were not, and why. This transparency builds trust and demonstrates that participation is meaningful, not performative.

A pilot of this kind could be launched within weeks using open-source or lightly customized tools. If successful, it would create a scalable model for structured legislative feedback that strengthens drafting, broadens the range of expertise informing Congress, and reduces—not increases—staff workload.

V. Asking Americans What Problems Matter Most—Before Drafting

Long before a draft bill is written, congressional committees must decide which problems merit attention. That upstream moment — the agenda-setting phase — is where Congress currently has the least structured way to hear directly from the public. Paradoxically, even though Congress receives millions of communications each year, it lacks a practical method for converting those interactions into a clear picture of national priorities.



Other governments are beginning to fill this gap by inviting residents to help define priorities before policy is drafted, turning to AI to make this kind of large-scale listening feasible.

Precedents: How Others Are Doing Early-Stage Engagement at Scale

Across contexts — local, state, national, and international — early-stage engagement succeeds when it asks clear questions, provides structure, and uses AI to turn narrative input into usable intelligence.

St. Louis (United States): From ideas → priorities → a bill

When St. Louis needed to decide how to allocate a \$250 million settlement from the Rams football team, the Board of Aldermen turned to residents before drafting any legislation or making policy (much as Bowling Green is doing). More than 16,000 residents participated in a process run using the open source Go Vocal platform. Go Vocal is used by more than 500 governments worldwide. Residents submitted over 1,000 ideas, identified the city's most pressing challenges, and ranked their priorities.^{xxiv}



To make sense of that volume, the city used an AI-supported sensemaking tool that clustered ideas, surfaced common themes, and highlighted where different demographic groups diverged. Each AI-generated insight linked back to original resident comments for verification.

Departments then refined the proposals; residents reviewed them again; and the final bill — previously contentious — passed with broad legitimacy. St. Louis has now institutionalized this approach as an ongoing pipeline for bottom-up policymaking.

California (United States): Post-wildfire agenda-setting and prioritization

California used a similar two-phase structure following the Los Angeles wildfires. The public and especially survivors were invited to identify which recovery challenges mattered most — the issues they believed should shape state and local action. Participants first prioritized among potential focus areas, then reviewed and ranked concrete policy options developed from their earlier input.

They used a tool called Othello to engage more than 2,500 narrative comments, mapped proposals to major themes, and identified both consensus and key tensions — such as speed of rebuilding versus long-term resilience. The resulting community-driven action plan elevated five priority needs, from undergrounding power lines to strengthening water systems and streamlining permitting. Agencies have already begun implementing many of these recommendations.^{xxv} The structured process enabled residents to set the agenda and see their input translated into concrete policy recommendations.

Brazil (National strategic planning): Millions engaged with AI-enabled synthesis

Brazil has taken early-stage engagement to the national scale. As part of its constitutionally mandated four-year Pluriannual Plan, the federal government in Brazil invited the public to help identify national priorities. The response was extraordinary: 1.4 million participants, over 1.5 million votes, and more than 8,200 proposals.

Because no institution can manually analyze input at that scale, Brazil built an AI-enabled process—now part of its permanent *Participatory Brazil* platform—to support structured feedback on draft plans and laws. Public comments are tied to specific provisions, and the AI pipeline automatically clusters submissions, extracts evidence, identifies major themes, and generates clear reports for staff.^{xxvi} Crucially, the system also shows contributors how their input was incorporated, closing the feedback loop and reinforcing trust. The power comes from structure; AI simply makes the process faster, more scalable, and more transparent.

New Jersey (United States): Pairwise priority ranking

New Jersey has repeatedly used the open-source tool All Our Ideas, developed at Princeton and hosted by the nonprofit Citizens Foundation, to run large-scale engagements to identify the most urgent challenges residents believe the state should address. We used it to hear from more than 4,000 people to help us set priorities for our Future of Work Task Force and from over 2,200 workers in connection with our AI Task Force.^{xxvii}

Instead of long free-text comments, residents compare pairs of ideas and “vote” between them, resulting in a ranked list of priorities. AI-supported translation produced multilingual, plain-language content, enabling broad participation. The resulting priority maps guided subsequent policy work.

[About this project](#)[Vote](#)[Results](#)

Over the next five years, what will be the biggest effect from generative AI on the workforce in New Jersey?

New Jersey's Artificial Intelligence Task Force needs YOUR help to understand how generative AI will impact workers in the Garden State.

Start Voting

About this project

Vote

Results

Over the next five years, what will be the biggest effect from generative AI on the workforce in New Jersey?

AI tools designed to address discrimination will create more opportunities for workers from underrepresented groups.

or

AI will change job roles and responsibilities, leading more workers to be classified as independent contractors rather than employees.

Skip

Add your own answer

0 votes of 10 Target (Level 1)

What Congress Can Do Now

Congress could begin with a simple, low-risk pilot that helps a single committee understand what Americans view as the most urgent problems in a particular policy area — before drafting legislation, holding hearings or launching formal oversight.

A pilot might work as follows:

1. Pose one clear, well-defined question.

A committee selects a policy area — such as rural health care, wildfire resilience, cybersecurity, veterans’ reintegration or public engagement in lawmaking — and invites the public to identify the most urgent challenges they experience.

2. Invite short narrative submissions.

Residents briefly describe the problem as they see it. These are not comments on a bill, but accounts of real-world barriers, harms, or needs. Alternatively, provide pre-written statements as we did in New Jersey and invite people to rank the issues.

17

3. Use AI for synthesis, not decision-making.

AI clusters similar submissions, identifies recurring themes and lesser-known concerns, highlights demographic or regional variation, and generates a concise, human-verifiable summary. All Our Ideas produces a rank-ordered list of issues.

4. Publish a short “What We Learned” summary.

A public memo — even one page — signals that input was taken seriously and shows which themes will shape the committee’s work.

5. Repeat only if useful.

If effective, the model can be reused for other topics, committees, or bipartisan initiatives.

This pilot does not require new offices, new rules, or major procurement. It can be launched in weeks using open-source or lightly customized tools. What it offers is something Congress does not currently have: a clear, structured, and scalable way to understand what problems Americans believe deserve attention before drafting legislation.

VI. Risks and How to Mitigate Them

As Congress experiments with AI-enabled public engagement, it should proceed with clear eyes about the risks. None of these risks argue against using AI; they simply underscore that design and oversight matter as much as the technology itself.

1. The Risk of Hallucination

AI models can generate plausible but incorrect statements — the risk commonly known as hallucination. Fortunately, this risk is easiest to manage in the context of public engagement. When AI is used only to summarize and cluster text that the public actually submitted, rather than to generate new content, hallucination risk drops dramatically.

Modern tools allow Congress to:

- **Constrain the model to a defined corpus** (e.g., “summarize only the words provided in these submissions”).
- **Require citations to the underlying text**, ensuring summaries link back to original comments.

- **Keep humans in the loop**, with staff reviewing and validating outputs before anything is used.

In other words: AI becomes a synthesis tool, not a content creator.

2. The Risk of Ineffective or Performative Engagement

The larger risk is not that engagement goes wrong — but that it goes nowhere.

Across the world, we have seen well-intentioned participatory exercises yield little value for either the public or policymakers. People invest their time, contribute ideas, and hear nothing back. Staff receive floods of input with no structure. Both sides walk away frustrated.

People want to participate in ways that are relevant, respectful of their time, and actually used. Policymakers need input that is structured, actionable, and tied to real decisions.

Poorly designed processes — however exciting in theory — create cynicism and undermine trust.

That is why every example in this testimony was designed around the needs of both the institution and the participants: Each one:

- asked clear and specific questions,
- designed structured pathways to contribute, and
- provided visible outcomes, including “what we heard” summaries.

These guardrails ensure engagement helps rather than distracts from the work of governing.

3. The Risk of “One-Platform Thinking”

Lawmaking and oversight are not single events; they are multi-stage processes — agenda setting, drafting, hearings, amendments, implementation. Each stage benefits from different forms of public expertise.

There is no single platform that can serve all of these functions, and adopting any one tool or single process as a silver bullet would be a mistake.

Instead, Congress should view this work as building a **portfolio** of engagement mechanisms:

- brief question solicitation for hearings,
- structured feedback on draft text,
- early-stage prioritization exercises,
- targeted outreach to experts or affected communities, after-action review once laws are implemented.

AI makes this *combinatorial* approach possible by reducing the burden of analyzing input from each stage. The goal is not one platform, but a set of processes that together help Congress collect expertise and intelligence from both representative and non-representative groups — always with human judgment guiding every decision.

These are far from the only risks, which also include the danger of spam, security breaches, AI-talking-to-AI but, with thoughtful design and outreach, these risks are manageable and outweighed by the opportunity to improve how Congress listens and learns.

VII. Beyond the Pilots: What Congress Must Build Next

The pilots outlined in Recommendations 174 and 176 are essential first steps.

We stand at the ready through our AI for Impact program to help with building free, open source, tested platforms for running these engagement processes; with advice and design from the New Jersey Office of Innovation based on our longstanding experience doing engagement at scale and with free training via the InnovateUS initiative, where [we teach how to use AI for public engagement](#).^{xxviii}

Going forward, if Congress wants to keep pace with the rising complexity of governing, it must treat these pilots as the beginning of a systematic effort to forge a more participatory institution that leverages new technology to elicit collective intelligence from the American people.

Other countries have already begun this shift. Brazil built an in-house innovation lab for public engagement in lawmaking. The US Congress needs the same: a dedicated, ongoing capability for innovation in participation to complement the AI and digital work in the Library of Congress and the CAO.

Establish a Participation Innovation Lab for Congress

Congress should create a small, permanent, interdisciplinary team modeled on the innovation labs in Brazil and Chile. This unit would:

- Run and refine ongoing public engagement processes for committees and Member offices.
- Work with outside experts on uses of AI for public engagement.
- Ensure processes are designed around decision-maker needs, not symbolic participation.
- Help committees choose the right engagement method at the right stage — from early-stage priority-setting to structured input on draft text to hearing-question collection.
- Develop templates, workflows, moderation standards, and evaluation metrics so engagement becomes easier, not harder, for staff.
- Coordinate with the CAO Coaches Program, House Digital Service, CRS, LOC and GAO to ensure engagement outputs feed into analysis and drafting work.

In New Jersey, the Office of Innovation was explicitly tasked with developing tools and processes for public engagement so that we could help agencies increase their uptake of public engagement practices. Congress needs the analog: an in-house team devoted to designing, running, and evaluating participatory processes that serve real legislative needs.

Build Congress.ai — A Publicly Governed Legislative AI Model

Congress is behind other legislatures in developing AI designed for the public interest. Without its own capabilities, Congress risks being outmatched by well-funded special interests using advanced AI to shape narratives, craft amendments, and lobby at unprecedented scale. The choice is no longer whether AI will be used in lawmaking but whose interests the technology will serve.

Congress needs a publicly governed, open, transparent congress.ai model trained on legislative data to use for building AI tools optimized for lawmaking.

- The 22 million words of the U.S. Code, nearly 4,000 pages of tax code, and the vast corpus of debates, transcripts, CRS reports, GAO analyses, amendments, committee prints, and historical legislative text form a world-class training corpus for a specialized model.

- Such a model could support drafting, help staff understand complex language, flag inconsistencies, forecast impacts across communities, and detect patterns in lobbying and amendment influence — strengthening both transparency and accountability.
- With AI fine-tuned on congressional data, it will become easier to build both public engagement and other legislative innovation tools.^{xxix}

Invest in People: Systematic AI Training for Congress

Congress currently has no institution-wide AI training for staff or members; many rely on free, insecure tools because approved ones are unavailable or limited. In New Jersey, we have trained and provisioned almost 15K public professionals with AI and the knowledge to understand how to use AI to serve the public.

Congress needs:

- Mandatory AI literacy and safety training for all staff.
- Specialized drafting, analysis, and participation modules for committee counsels, CRS analysts, and legislative directors.

Congress is stretched thinner today than at any moment in modern history. Yet we now have, for the first time, technologies capable of restoring Congress’s ability to listen at scale, learn from the expertise of the American people, and make laws and conduct oversight both more efficiently and effectively.

By implementing Recommendations 174 and 176 through small, thoughtful pilots, this Subcommittee can help Congress engage the public in ways that strengthen lawmaking, increase transparency, and rebuild trust right away.

Thank you for your leadership and for the opportunity to testify. I look forward to your questions.

ⁱ For a complete bio, please see Beth Simone Noveck, <https://thegovlab.org/beth-simone-noveck.html>, accessed December 12, 2025.

ⁱⁱ CrowdLaw for Congress, <https://congress.crowd.law/>. See also Beth Simone Noveck, “Crowdlaw: Collective Intelligence and Lawmaking,” *Analyse & Kritik* 40, no. 2 (2018): 359–380, https://www.analyse-und-kritik.net/Dateien/5be9b083bc696_noveck.pdf.

ⁱⁱⁱ In the State of New Jersey, we are using AI to streamline how the government handles calls.. The team I founded at the Office of Innovation in New Jersey has modernized 13 call centers, enabling them to efficiently handle over

10 million calls in 2024. These upgraded systems increased the number of callers who successfully reach a live agent by 6 to 15 percent, while cutting operational costs roughly in half. **Average wait times have fallen from over 40 minutes to under two minutes across several agencies.** AI now analyzes call transcripts so menus reflect what people actually need, supports self-service tools in English and Spanish, and enables “warm callbacks” so residents aren’t stuck on hold and do not need to call at all. Callers can receive helpful links and information through text message, and managers now receive real-time insights into emerging issues so they can proactively adjust menu options. This is what AI can do when paired with good process design: it transforms millions of unstructured, high-stakes interactions into a manageable, efficient workflow. We are happy to share our playbook to help Member and committee offices manage constituent communications.

^{iv} Congressional Research Service, *House of Representatives Staff Levels, 1977-2023* (Nov. 28, 2023), https://www.congress.gov/crs_external_products/R/PDF/R43947/R43947.9.pdf; LegBranch.org, “Legislative Branch Support Staffing Down 45 Percent Since 1975,” May 30, 2017, <https://www.legbranch.org/2017-5-25-legislative-branch-staffing-down-45-percent-since-1975/>; and Congressional Research Service, *Legislative Branch Appropriations: Frequently Asked Questions* (Report R43397, updated August 8, 2024), <https://crsreports.congress.gov/product/pdf/R/R43397> (see Tables 1-2 on appropriation shares).

^v OpenSecrets, “Total Cost of Election (1990–2024),” accessed December 10, 2025, <https://www.opensecrets.org/elections-overview/cost-of-election>; Congressional Research Service, *Legislative Branch Appropriations: Frequently Asked Questions*, R43397, updated August 8, 2024, <https://crsreports.congress.gov/product/pdf/R/R43397>; and OpenSecrets, “Trends in Spending, 2008–2024,” federal lobbying data, accessed December 10, 2025, <https://www.opensecrets.org/federal-lobbying/trends-in-spending>; Zeller, “Congress Is Good at Shrinking One Part of Government.” *Roll Call*, January 20, 2015, <https://rollcall.com/2015/01/20/congress-is-good-at-shrinking-one-part-of-government/>

^{vi} UK Cabinet Office, Incubator for Artificial Intelligence, “Consult,” accessed December 10, 2025, <https://ai.gov.uk/projects/consult/>.

^{vii} Beth Simone Noveck, “How Hamburg is Turning Resident Comments into Actionable Insight,” Reboot Democracy, October 22, 2025, <https://rebootdemocracy.ai/blog/how-hamburg-is-turning-resident-comments-into-actionable-insight>.

^{viii} Government of Singapore, About Singapore Together, accessed December 12, 2025, <https://www.sg/about>.

^{ix} Beth Simone Noveck, Dane Gambrell, Valeria Gomez Palacios, and Anna Ibru, *Expanding and Strengthening Engagement on Decide Madrid: A Data-Driven Evaluation*, The GovLab, March 2022, <https://files.thegovlab.org/madrid-report-english.pdf>.

^x Beth Simone Noveck, Wiki Government (Brookings Press 2009). See also “Peer-to-Parent: Better Information for Better Patents,” White House Open Government Initiative, <https://obamawhitehouse.archives.gov/open/innovations/Peer-to-Patent>.

^{xi} Center for Civic Futures, “Center for Civic Futures and partners commit \$8.5M for AI solutions that improve safety net program delivery December 9, 2025,” <https://www.centerforcivicfutures.org/resources/center-for-civic-futures-and-partners-commit-8-5m-for-ai-solutions-that-improve-safety-net-program-delivery>.

^{xii} Free and Hanseatic City of Hamburg, “DIPAS - Digital Participation System for Hamburg,” accessed December 10, 2025, <https://www.dipas.org/en>.

^{xiii} Inter-Parliamentary Union, “Brazil: A Digitally Mature Parliament,” June 1, 2022, <https://www.ipu.org/news/case-studies/2022-06/brazil-digitally-mature-parliament>.

^{xiv} Luís Kimaid, “From Paper to Tokens: Transforming Legislative Services in the Chamber of Deputies of Chile,” Bússola Tech, June 11, 2025, <https://library.bussola-tech.co/p/caminar-camara-diputadas-diputados-chile>.

^{xv} Pascale Davies, “An AI Tool is Making the European Parliament’s History Easier to Understand in All EU Languages,” *Euronews Next*, October 22, 2024, <https://www.euronews.com/next/2024/10/22/an-ai-tool-is-making-the-european-parliaments-history-easier-to-understand-in-all-eu-langu>.

^{xvi} Beth Simone Noveck et al., *From Citizen to Senator: Artificial Intelligence and the Reinvention of Citizen Lawmaking in Brazil*, The GovLab, April 2025, <https://files.thegovlab.org/from-citizen-to-senator.pdf>.

^{xvii} For example, in shaping the conversation around foreign researchers during the Senate Finance Committee’s June 2019 hearing on foreign threats to taxpayer-funded research, the CSPI provided committee members with evidence of the importance of global collaboration and the free and open exchange of information for academic research. This helped to balance out a perspective that foreign born students and researchers were possible threats to national security - a view that had seeped into discussion in political circles. In his opening statement, Senator Ron Wyden read from the testimony of a foreign researcher, provided by a CSPI contributor, and pointed to the many contributions to scientific research made by foreign researchers, taking a strong stance in favor of academic freedom and setting the tone for the rest of the hearing.

^{xviii} Dane Gambrell, “Congressional Science Policy Initiative,” The GovLab, October 2020, <https://collective-intelligence.thegovlab.org/case/congressional-science-policy-initiative>.

^{xix} Warren County, Kentucky, and Jigsaw, “What Could BG Be?,” April 2025, <https://www.whatcouldbgbe.com/>.

^{xx} State of California, “Engaged California,” accessed December 10, 2025, <https://engaged.ca.gov/>.

^{xxi} For more about AI for Impact, visit <https://burnes.northeastern.edu/ai-for-impact-coop> and Open Feedback, <https://ai4impact.ai/projects>, accessed Dec 12, 2025.

^{xxii} With the help of the Berkman Center at Harvard Law School and the Governance Lab, the German government crowdsourced input into its AI policy “Key points for a Federal Government Strategy on Artificial Intelligence,” Federal Government of Germany, July 7, 2018, https://www.bmas.de/SharedDocs/Downloads/EN/Topics/Labour-Market/key-points-ai-strategy.pdf?__blob=publicationFile&v=1. See also Beth S. Noveck, Rose Harvey, Anirudh Dinesh, *The Open Policymaking Playbook*, The GovLab, April 2019, <https://thegovlab.org/static/files/publications/openpolicymaking-april29.pdf>.

^{xxiii} Róbert Bjarnason, Gunnar Grimsson, and Gina Joerger, *Better Reykjavik: Municipal Open Innovation*, The GovLab, April 2019, <https://thegovlab.org/static/files/better-reykjavik.pdf>

^{xxiv} Sarah Horton, “St. Louis Involves 16,000+ Residents in the Allocation of Settlement Fund,” GoVocal, January 22, 2024, <https://www.govocal.com/case-studies/st-louis-collects-ideas-from-7000-residents-online>

^{xxv} Micah Weinberg, “The Promise and Potential of Engaged California: Lessons from the Los Angeles Wildfires,” Carnegie California, November 24, 2025. <https://carnegieendowment.org/posts/2025/09/engaged-california-la-wildfires-lessons?lang=en>.

^{xxvi} Christiana Freitas and Ricardo Poppi, “Global AI Watch: Brazil’s Experiment in AI-Powered Participation,” Reboot Democracy, November 5, 2025. <https://rebootdemocracy.ai/blog/global-ai-watch-brazils-experiment-in-ai-powered-participation>.

^{xxvii} Dane Gambrell, “Your Future of Work: What the New Jersey Future of Work Task Force learned from 4,000 workers in the Garden State,” New Jersey Office of Innovation, September 2020, <https://fowtf.innovation.nj.gov/downloads/resources/YourFutureOfWork.pdf>; and Dane Gambrell, “AI for E-governance: Combining Artificial Intelligence and Collective Intelligence to Develop Evidence-Based AI Policy,” in *2025 Eleventh International Conference on eDemocracy & eGovernment (ICEDEG)*, 317-324 (IEEE, 2025), July 2025, <https://ieeexplore.ieee.org/document/11081634>.

^{xxviii} See Reboot Democracy: Designing Democratic Engagement for the AI Era, InnovateUS, <https://innovate-us.org/workshop-series/democratic-engagement>, accessed Dec 12, 2025. We are turning the live workshops into an asynchronous course, which will be ready in spring 2026.

^{xxix} For more on Congress and the use of AI to enhance lawmaking, see Beth Simone Noveck. *Reboot: The Race to Save Democracy with AI* (Yale Press 2026) (forthcoming).