YOUNG ANDERSON
BARRISTERS & SOLICITORS

AI AND LOCAL GOVERNMENTS

NOVEMBER 20, 2026

Jordan Adam and James Barth

AI AND LOCAL GOVERNMENTS

I. INTRODUCTION

There is no shortage of commentary on artificial intelligence and how it impacts various fields, and it's unlikely that this is the first paper you have read on the subject, but we hope to provide some fresh insight on the topic from a municipal law perspective. First, we provide a very brief and high-level overview of the history of artificial intelligence. Then, we explore some of the issues and challenges that local governments face related to the use of AI tools by the public. Finally, we discuss how local governments can protect themselves from potential pitfalls and liability associated with AI usage by local staff and elected officials.

II. A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

We begin with a very brief overview of the development of artificial intelligence, from early questions about whether machines could think to the tools now used by the public and in the workplace. While far from exhaustive, it will provide some context for the issues discussed in this paper.

1950: Alan Turing publishes "Computing Machinery and Intelligence", which opens with the question "Can machines think?" and imagines the possibility of computers that simulate intelligence and delves into ideas of intelligence testing and machine learning.

1956: John McCarthy christens "Artificial Intelligence" at the Dartmouth Summer Research Project on Artificial Intelligence. The goal of the project was to investigate ways in which machines could be made to simulate aspects of intelligence.

1950s – 1970s: Research into many artificial intelligence-related areas continues, including heuristic search, computer vision, natural language processing, mobile robotics, machine learning, artificial neural networks, and expert systems.[1]

1980–2001: Artificial Intelligence moves from theoretical research toward practical applications. Researchers develop systems capable of performing specialized tasks that ordinarily require human expertise. In 1997, IBM's Deep Blue defeats chess champion Garry Kasparov.

[1] Peter Stone et al, "Appendix I: A Short History of AI", in Artificial Intelligence and Life in 2030: One Hundred Year Study on Artificial Intelligence, Report of the 2015–2016 Study Panel (Stanford University, September 2016), online: https://ai100.stanford.edu/reports/2016-report/appendix-i-short-history-ai

2001–2018: Artificial Intelligence shifts toward machine learning and deep learning. AI systems learn patterns from large datasets, supported by advances in computing power. AI becomes more integrated into search engines and digital assistants like Siri. In 2011, IBM's Watson wins Jeopardy, and Google's AlphaGo beats one of the world's top Go players.

2018 – 2022: In 2018, OpenAI launches GPT-1, which could answer questions in a humanlike way and generate text in response to prompts.[2] In November 2022, OpenAI launches ChatGPT, built on GPT-3.5, which is rapidly adopted by the public.

2023 – Present Day: Large Language Models (LLMs) continue to advance with Google, Meta, and Apple unveiling new models and products; agentic AI gains prominence, and enterprise software (such as Microsoft Co-Pilot) and businesses begin to integrate AI tools into their operations.

III. ARTIFICIAL INTELLIGENCE USE BY THE PUBLIC

In a recent case in BC's Court of Appeal, *1182573 B.C. Ltd. v Gerber*, 2026 BCCA 399 [*Gerber*], the Court directly addressed common pratfalls associated with AI use. The sole shareholder and director of the appellant company was self-represented, and he admitted to using AI tools to prepare his submissions. Noting the rise in self-represented litigants relying on AI to assist with their claims, the Court offered words of advice and caution to litigants using AI to prepare court materials, warning specifically of the following risks:

- Hallucinations;
- Attaching numerous and irrelevant legal authorities; and
- Bad legal advice.

Although the Court's words of warning on these risks were made in the context of formal legal disputes, similar issues also arise when the public interacts with local government staff. In an article for the Journal of Law and Society titled "'Chucking Their Stuff Into ChatGPT and Hitting Send': AI Slop in the Administrative State", Jed Meers of York Law School examines the impact of the use of AI tools by the public on administrative officials. Using a qualitative survey of 512 public-facing civil servants, Meers identifies three components of public AI use that negatively impact the administrative state:

- Its recourse to often lengthy, if also often inaccurate, legal arguments (*volume*);
- A significant tonal shift in the person accessing the service (*voice*); and

[2] Ivan Belcic & Cole Stryker, "What is GPT (Generative Pretrained Transformer)?", IBM, online: https://www.ibm.com/think/topics/gpt

- An indefatigability, where replies were met almost immediately by other lengthy and legalistic responses (*velocity*)[3].

There is significant overlap between the issues with AI tools raised by the Court in *Gerber* in relation to use by self-represented litigants and those identified by Meers in relation to use by the general public when communicating with administrative officials, and we will explore these issues further in this section. While we can't offer much in the way of solutions to these issues, we hope to help local governments better understand the challenges and perhaps begin to think about how systems can be changed to address them.

A. Hallucinations, Mistakes and Irrelevant Material

Generative AI tools can hallucinate and produce outputs that sound plausible on a first read, but that are actually factually wrong, irrelevant, or completely fabricated.[4] An article written for computing giant and AI pioneer IBM indicates that it is unlikely that AI tools will ever reach a point where generated content is error free, noting that generative AI models "will always be vulnerable to inadvertently producing some misinformation" because of their predictive nature, and that the "models are statistical machines that, effectively, are trying to predict some sort of outcome… at the end of the day, [the] models can still predict an incorrect answer or outcome"[5].

In *Gerber*, the Court acknowledges this risk:

> The risk that AI will generate hallucinations, by making up case authorities that do not in fact exist, making up purported case quotes that do not exist, or incorrectly summarizing real cases, is by now fairly well known.[6]

There are many examples of AI hallucinations in the legal world. In the BC Court of Appeal case *Wu v. Murray*, 2025 BCCA 365, a self-represented appellant admitted to using generative AI to prepare submissions which contained two false cases that supported a central point of her application. In *Mazaheri v. Law Society of Ontario*, 2026 ONLSTH 112 [*Mazaheri*], a lawyer admitted to using AI to prepare his factum and supplementary factum and not verifying what was produced, with the Court noting that the materials:

> … referred to Tribunal cases that don't exist. For example, the applicant cited and hyperlinked the case of "*Law Society of Ontario v Mercer*, 2024 ONLSTH 136". But the hyperlink takes you to Law Society of Ontario v Deokaran, 2024

[3]Jed Meers, "'Chucking Their Stuff into ChatGPT and Hitting Send': AI Slop in the Administrative State" (2026), Journal of Law and Society [forthcoming], online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6898138, at page 9.
[4]Chrystal R China, "What Are AI Hallucinations?" (updated 31 July 2026), IBM, online: https://www.ibm.com/think/topics/ai-hallucinations
[5] Alice Gomstyn & Alexandra Jonker, "AI Misinformation: Here's How to Reduce Your Company's Exposure and Risk" (8 January 2025), IBM, online: https://www.ibm.com/think/insights/ai-misinformation#635426478
[6]*Gerbe*r at para 81.

> ONLSTH 136. This was because there is no case titled Law Society of Ontario v Mercer.[7]

The Court noted that the pleadings referred to cases that did not exist, even though the case references included properly formatted and realistic citations and hyperlinks. The pleadings also referred to real cases as well as the Tribunal's Rules of Practice and Procedure, but the references didn't support the legal point for which the lawyer was relying on them. Ultimately, the Court determined that the pleadings were completely useless, noting that "[n]ot a single proposition of law in any of the applicant's materials was supported by a reliable authority. The result was gibberish."[8]

This demonstrates the ability of AI to hallucinate and generate entirely fabricated information, make the hallucinations look very plausible (by formatting proper citations and hyperlinking fake cases, referring to specific paragraphs of real cases, and providing fake legislation that is formatted and written like real legislation) and present the hallucinated material in a very confident way.

Hallucinations are a type of mistake that AI tools can make, but they're not the only type. As IBM notes, AI mistakes can also include providing inaccurate information, such as outdated facts, making biased recommendations, providing overly generic answers, and making subtle logical errors that can lead to poor or incorrect outcomes. In the legal context, AI can make mistakes by misunderstanding and misapplying legal principles. While there are legal principles that are widely applicable across different legal systems and jurisdictions, many legal principles are only applicable to a particular country or province. Because BC local government law is province-specific, and relatively little writing on BC-specific local government law is available on the internet, it is not uncommon for AI tools to fill the gaps by applying irrelevant legal principles. In *Gerber* at paragraph 89, the Court refers to the problem of receiving lengthy, often irrelevant materials:

> Again, from personal observation, it appears that increasingly the courts are receiving pleadings and submissions that are unduly lengthy and verbose, containing "legal" language untethered from established legal principles, and which are incomprehensible. These types of materials do nothing to advance a party's legal case. It seems likely that at least some of these materials are AI generated. It appears that AI can be quite delusional about one's legal rights and how to present legal argument.[9]

Indeed, the writers have received several long and detailed letters from opposing parties that are grammatically perfect and error free, but that refer to legal concepts that are not applicable in the context of municipal law in British Columbia. Additionally, while AI tools can generate

[7] *Mazaheri v. Law Society of Ontario*, 2026 ONLSTH 112, at para 54.
[8] *Ibid* at para 59.
[9] *Gerber* at para 89.

lengthy and authoritative sounding missives, they often fail to address the real issues, especially if the individual providing the prompts doesn't fully grasp the issue at hand.

This phenomenon is not unique to the legal field, and it's likely to arise more often in contexts where localized knowledge is required or specific expertise is applied, including local government environments. As Meers notes:

> Long, complex and multi-faceted correspondence was leading to 'pointless…time consuming and frustrating' text-based conversations where officials felt they had to wade 'through pages of word salads that makes no difference to the outcome of their case'. In some contexts, particularly for other participants dealing with FOI requests, officials are obliged to respond to each of the points raised in correspondence, even if inaccurate, 'which we then have to work through to explain why it's not relevant which wastes time'.[10]

Given the potential for hallucination and mistakes, it is more important than ever to closely review material submitted by the public. While AI generated content is usually very polished and confidently presented, one cannot assume that it is relevant or accurate. Of course, thoroughly reviewing content takes time, which brings us to our next issue: volume and velocity.

B. Volume and Velocity

Not only can AI tools hallucinate or generate inaccurate material, they can also produce a lot of irrelevant content very quickly, and that irrelevant content can distract the recipient from the real issues that the submitter is trying to address. The volume often goes hand in hand with bad advice, because in order to generate content, the AI tool may focus on unrelated, unimportant, and irrelevant issues. Volume also magnifies the negative impact of hallucinations, mistakes, bad advice, and irrelevant information on the receiver, because they are required to wade through large amounts of information to determine which information provided is accurate, relevant, and requires a response. In the legal context, this may mean that a lawyer has to review every authority included in a list to determine if the authorities actually support the point being asserted. It also means reviewing and responding to every point that is raised, even if just to note the irrelevance. Anecdotally, litigators in our firm have seen pleadings and submissions from self-represented litigants increase substantially in length. Where previously such documents were usually one to ten pages, it's not uncommon to see 20–30 pages, and each page and point raised has to be reviewed, because some of the content may be valid and require a response, even though most of it does not.

[10]Meers, *supra* note 3 at page 9.

The same is true in the local government context. As Meers notes, participants in his study reported that often "fairly straightforward interactions come with wide-ranging legal arguments bolted on as a standard".[11] As a result, staff are forced to review the arguments and determine their relevance, and potentially still respond to baseless or incorrect assertions in order to outline how these arguments are irrelevant. Meers sums up the issue:

> Long, complex and multi-faceted correspondence was leading to 'pointless…time consuming and frustrating' text-based conversations where officials felt they had to wade 'through pages of word salads that makes no difference to the outcome of their case'. In some contexts, particularly for other participants dealing with FOI requests, officials are obliged to respond to each of the points raised in correspondence, even if inaccurate, 'which we then have to work through to explain why it's not relevant which wastes time'.[12]

The volume of correspondence often wastes time and can obscure the real issues that litigants or constituents are trying to advance or address. Long, formal, and supposedly legally grounded correspondence can also be intimidating, perhaps causing the recipient to needlessly question their positions, which may result in the recipient involving legal counsel unnecessarily.

Meers also notes that the velocity at which inquiries and responses are sent can be problematic. While administrators may spend a significant amount of time crafting a response to an inquiry, thanks to AI tools, they can now receive a lengthy reply within minutes asking for additional detail or providing new arguments. This can create a "rapid, seemingly inexhaustible cycle of lengthy replies where it tends to go round and round in a loop".[13] The velocity of responses can also be indicative of the fact that the submitter is not taking the time to actually understand (or even read) the response provided by a staff member; they have merely plugged the response into an AI tool and generated a reply. The Meers study supports this, as he notes that many of the administrators participating in the study were concerned not only that members of the public were not reading the correspondence they received from administrators, but that members of the public may not even have been reading their AI-generated replies prior to sending them.

Ultimately, the increased volume of correspondence and the velocity at which it is sent creates more work for local government staff and often makes it harder for recipients to understand the real issues being presented and address them appropriately.

[12]Meers, *supra* note 3 at page 9.
[13]Meers, *supra* note 3 at page 12.

C. Voice and Tone

AI tools tend to have a distinct and recognizable voice and style in written correspondence, characterized by formality and perfect punctuation and grammar that feels very different from the way that real people write, even in formal communication. Meers quotes a government official who describes this well:

> 'their last email was all "pls help my bins ain't been emptied in 3 weeks FFS" and now it's a five-paragraph essay with perfect punctuation'.

While the "AI-style" in itself may not be inherently problematic, it can lead to issues when the correspondence takes on an aggressive or combative tone. As the Court notes in *Gerber*:

> I am also concerned by the increasingly aggressive tone I am seeing in some litigation files, not just in the pleadings and submissions but also in correspondence. It causes me to wonder if AI may be leading individuals to lose their human moral instincts to be polite, measured, honest, and ethical…

This tone isn't necessarily restricted to legal proceedings, as it can arise in any situation when an individual tasks an AI tool with drafting persuasive correspondence. It strikes us that taking an aggressive or overly confident stance can be unhelpful, and more likely to create a divide between the parties, especially in text-based communication, where it's already easy to misunderstand or misinterpret the tone someone is taking.

D. Sycophancy

A feature of AI tools that may cause or exacerbate the issues outlined above, and may help us better understand why these issues occur, is the tendency of AI tools to be sycophantic to their users. In relation to AI models, sycophancy is defined as a phenomenon where a model seeks human approval in unwanted ways.[14] In testing sycophancy, researchers led by Dr. Mrinank Sharma came to the following conclusions:

- AI tools can give biased feedback: tools provide more positive feedback about arguments that the user likes, regardless of the content thereo;[15]
- AI tools can be easily swayed: even when tools provide accurate answers, they often modify their answers when questioned by the user and subsequently provide incorrect information;[16]

[14] Mrinank Sharma et al, "Towards Understanding Sycophancy in Language Models" (Paper delivered at the 12th International Conference on Learning Representations, 2024), online: https://arxiv.org/pdf/2310.13548, at page 2.
[15] Sharma et al, *supra* note 13 at page 3.
[16] *Ibid* at page 3.

- AI tools can give biased answers, tending to provide answers that agree with user beliefs rather than providing accurate information;[17] and

- AI responses can sometimes mimic user mistakes: when a user presents an incorrect claim, AI assistants sometimes do not correct the user and instead respond in ways that cohere with the user's beliefs.[18]

It is plausible that the tendency of AI tools toward sycophancy, as well as other specific behaviours identified in the Sharma study, can contribute to the issues identified with AI use in legal and administrative settings. If an individual is seeking a specific outcome, such as an administrative or legal decision being made in their favour, they may implicitly or explicitly nudge AI tools toward providing content that supports their desired outcome, which may result in the AI tool generating hallucinated, erroneous or flawed content if no actual support is available. The AI tool may also tend toward a combative or authoritative tone, and assert full confidence in content that is irrelevant or incorrect.

Hallucinations and irrelevant arguments, amplified by volume, velocity and an overly confident or combative tone, can obscure the underlying message that the submitter is trying to get across and create more work for local government staff. Sycophancy may make these problems worse by validating a user's position rather than providing constructive feedback or appropriate advice. The challenge for local governments is to look beyond the polished exterior of these communications, identify the substantive issues and respond to those issues without being drawn into an unproductive cycle of AI-generated exchanges.

IV. ARTIFICAL INTELLIGENCE USE WITHIN YOUR ORGANIZATION

Given the rapid rise and adoption of generative AI in the last five years, local governments should ensure that their policies consider, control and account for AI usage within their organization. Whether a local government is choosing to implement AI tools in elements of its workflow, or simply wishes to manage informal individual usage of AI by staff and council, it is essential that its policies address key limits and guidelines around AI. This can be done through a standalone AI policy, or through amendment to existing IT, privacy and other related policies.

We wish to highlight five key considerations when developing an AI policy, or reviewing existing policy:

- Scope of applicability: does your policy cover all circumstances in which AI is being used, and everyone who may use it?

- Usage guidelines and training: if AI is being used in your organization, what training is required, and what guardrails and limits are there on that use?

[17] *Ibid* at page 4.
[18] *Ibid* at pages 4-5.

- Human oversight: if AI is being used with some degree of autonomy, when and how is its output reviewed by a person?

- Privacy and confidentiality: is sensitive, confidential and personal information being properly managed in accordance with your organization's statutory obligations?

- AI use in decision-making: have risks been considered and managed if your organization is considering using AI as part of any decision-making processes?

This is by no means an exhaustive list of the factors to be considered in an AI policy. There are other complex legal issue around AI generally which go beyond the scope of this paper, such as: (1) intellectual property rights and the legality of generative AI tools trained on copyrighted material; and (2) potential assumptions, bias and stereotypes that may be baked into AI models These are both nuanced open legal questions which extend past what this paper can cover. The five factors outlined in this paper do however provide a starting point of essential practical considerations around the use of AI in local governments.

A. Scope of Applicability

An important overarching consideration when developing policies around AI is who and what those policies will apply to. For instance, consider whether your policy adequately covers both staff and elected officials. While there should be applicable policy for both, where that policy is implemented and what it covers should be distinct. For elected officials, policy could potentially be implemented in the local government's code of conduct for its board or council. Elected officials may still be subject to general AI policy, but additional rules and guidelines specific to their role can be developed, for instance addressing where AI can and cannot be used in relation to closed meeting material or other sensitive information.

Likewise, your policies can address specific senior staff roles which may incorporate AI uniquely in their workflow, or work with information that requires stricter guidelines for AI use. This may include additional responsibilities for your IT staff in managing and monitoring AI tools, or guidelines for HR staff in where AI can or cannot be used in relation to personnel files. The extent to which policy delineates between specific roles and positions will of course depend on the size of your organization, and to what extent these relevant job duties are divided between different positions or managed by one staff member wearing multiple hats.

In addition to ensuring the scope of your policy covers elected officials and staff, you should also consider other groups and organizations you may wish your policy to address. When working with outside consultants and other third-parties you may contract with, your AI policy should outline how those third-parties' AI usage is to be guided, which should then be reflected in the language of the contracts entered into with those third-parties. If your organization works with volunteers, consider what those volunteers' duties are, and whether AI usage is needed or likely to be used. There may also be circumstances where local government business is done by a mixture of elected officials, staff and people from outside of the organization (e.g.

committees comprised in part of elected officials and outside residents). Would your existing policies for elected officials and staff be sufficient, or does additional language need to be added to cover groups comprised of a mixture of individuals internal and external to your organization?

B. Usage Guidelines and Training

If your organization uses or intends to adopt AI into elements of your workflow, it is important that corresponding training is required and provided for anyone who will be using AI tools, or working with the output of AI tools. Your policy should provide training for all members of your organization to provide a baseline understanding of any AI tools that are adopted – of what those AI tools are capable of being used for, and the implications of using those tools. This basic training should provide a mixture of practical and technical tips (e.g. the functions and features of a given AI tool, and how to use them most efficiently) as well as instruction on responsible AI usage.

You should ensure that a record is kept of the date and time that members of your organization have taken this basic training, and how frequently people who have taken the training are required to take a refresher course. Given the current rapid pace at which generative AI tools are changing and developing, training modules should be updated with greater frequency than other similar training materials your organization may employ. Likewise, you may wish to require refresher AI training with more regularity than other similar training (potentially on an annual basis).

If a person's position allows or requires them to use AI tools, they should be required to complete the requisite training before making use of those tools in their role. In addition to basic AI training provided to all members of your local government, you should also require use-case-specific AI training depending on the unique requirements of a given role.

Once individuals have been trained and are prepared to use AI tools as part of their local government work, they must also follow all usage guidelines established by your policy. Usage guidelines might include prohibitions on entering sensitive, confidential or personal information into AI tools, and requirements for human review and modification of AI generated material (both discussed in greater detail below). Usage guidelines may also include rules around disclosure and attribution. If generative AI is used to produce a substantial portion of a work product, a disclaimer should be provided in the header or footer of that document, setting out the name of the generative AI tool used, a summary of the prompt or task used to generate it, and the name of the employee who reviewed or edited the content.

Usage guidelines should also outline prohibited uses of AI in the workplace. These might include the use of personal AI accounts for local government business; the use of the local government's AI tools for personal matters; entering copyrighted, proprietary or confidential information into an AI tool; or using AI tools as part of a project or initiative that has yet to have AI use, oversight and review planned and approved.

C. Human Oversight

Anytime an AI tool is being used as part of a local government workflow, your policy must clearly establish that there will be a "Human in the Loop" (HITL). HITL ensures that any AI produced content or guidance will be subject to checks and balances via review by a person exercising critical thinking and judgment before the AI output is relied upon. HITL ensures both that someone is checking the use of AI to ensure it is operating with speed and efficiency, and to intervene if it appears that an AI tool is inadequately or inaccurately addressing the task assigned to it.

There are several key aspects to HITL that should be addressed in your policies:

- What "human oversight" entails: it should be clear to the person assigned to review any AI output what they are looking for (both common issues that may arise with AI such as hallucinated citations, as well as subject-specific errors that may commonly arise in a given task), how they should amend and correct errors, and how they will document recurring issues and the steps taken to resolve those issues;

- Who is responsible and accountable for oversight over a given use of AI: it should be clear what position or department is responsible for specific oversight. Assignments should be based on the knowledge-base and expertise of a given role or department. Consider for instance the kind of errors in AI generated work that would be immediately evident to your planning department, but not necessarily readily apparent to other departments;

- When and how will the veracity of AI generated content be evaluated and verified: it should be clear what reliable external sources and databases the reviewing individual(s) should be cross-checking, to ensure that verification doesn't ultimately rely on cross-checking against another source which may have involved AI along the way;

- What the documentation process is for HITL: wherever possible, a standard checklist should be set out for review processes, and all review of AI output should be thoroughly documented; and

- What types of work products require heightened review: AI output should be subject to heightened review, potentially involving oversight from multiple people or departments if AI is being used in any part of the process to develop material that will be public facing content; is being relied upon in any aspect of a decision-making process; or is part of foundational work on which other subsequent work will build and rely on.

HITL is a fundamental principle that should be considered in any instance where AI is being employed in your organization. Before AI is introduced into a workflow, you should first determine how the HITL protocol around that AI usage will be implemented on a day-to-day basis. Once AI use has been incorporated into a workflow, you should ensure that there is a final review and sign-off process by a qualified individual before a work product involving AI use is published, distributed, or acted upon.

D. Privacy and Confidentiality

Your AI policies should clearly prohibit inputting sensitive, confidential or personal information into AI tools. Local governments are subject to a number of statutory obligations which govern how certain information must be handled. Section 117 of the *Community Charter*, for instance, requires council members and former council members to keep in strict confidence any confidential records and information considered in a closed meeting, until and unless that information is lawfully released to the public. The same rules apply to regional district directors, per section 205(1)(d) of the *Local Government Act.*

Local governments must manage personal information in their custody or control in accordance with *Freedom of Information and Protection of Privacy Act* (FIPPA). Personal information is defined in FIPPA as "recorded information about an identifiable individual other than contact information" (contact information referring to information allowing someone to contact at their place of work). This definition of personal information captures both clear and evident personal information, such as a staff member's credit card number or medical history, as well as less immediately self-evident information, such as an anonymous bylaw complaint that contains enough descriptive information to allow an assiduous reader to correctly discern the complainant's identity. Personal information may only be used or disclosed where authorized under sections 32 and 33 of FIPPA. Local governments are also required under section 30 of FIPPA to protect personal information in their custody or control by making reasonable security arrangements against unauthorized collection, use, disclosure or disposal of that personal information.

Given the statutory obligations incumbent on local governments when dealing with confidential and personal information, clear policies should be established which prohibit such information being used in work involving AI tools. Providing confidential or personal information to an AI tool would be, in the overwhelming majority of circumstances, an impermissible use or disclosure of that information.

E. AI Use in Decision-Making

In discussing how AI may be used in your organization's processes, this paper has referred to where AI might be used to generate specific work products or be incorporated into different workflows. However, it is important to note that the potential ramifications of using AI to generate work products, for instance in assisting with creating initial templates for standard forms or providing first drafts for low-stakes public communications, are vastly different from

the potential ramifications of incorporating AI into any kind of decision-making process. While any and all AI usage should be guided and restricted by usage guidelines and mandatory training, HITL, and privacy and confidentiality principles (as set out above), it is nonetheless inadvisable to incorporate AI usage in core decision-making processes, even with clear policy in place around responsible AI usage.

Major risks that stem from the use of AI in decision-making include impact to your insurance coverage, and the potential that the decision is overturned on judicial review. Circumstances in which a local government would usually be covered by its insurance policy for issues flowing from a decision it made may no longer be covered if that decision is shown to have been made by an agentic AI. Even if HITL practices are being followed, it may not be sufficient to show that a human decision-maker reviewed and approved of the decision provided by an AI tool, if the human review essentially amounts to adoption of the AI decision.

This could similarly remove a local government's defense in a claim against it in negligence. Typically, where a claim in negligence has been brought against a local government, a local government can argue it is exempt from liability for negligence where that liability arose from a policy decision rather than an operational decision. However, if the underlying policy decision in question can be shown to have been made in whole or in part with AI, it is unclear whether this defense would still be arguable.

Furthermore, AI usage in a decision-making process is likely to make that decision vulnerable to being overturned on judicial review. The ability to argue the reasonableness of a decision would be significantly complicated and undercut by human reasoning being absent or minimal in the decision-making in question.

V. CONCLUSION

Local governments are faced with managing the rising use of AI both externally and internally. When dealing with AI generated materials from members of the public, it is important to be wary and review for potential AI hallucinations and other inaccuracies. The volume and velocity which you may receive AI generated correspondence, as well as the assertive or potentially combative tone of this correspondence, present challenges as to the frequency, depth and tone that should be used in response. When considering incorporating AI use in your organization, you should first ensure you have policies which adequately cover all staff, elected officials, and relevant third-parties; set out training and usage guidelines; ensure human review of all AI use and output; prohibit AI use on confidential or personal information; and avoid incorporating AI into core decision-making processes.

Appendix A

Glossary of Artificial Intelligence Terms

The following terms are taken from Stanford University's Human Centered Artificial Intelligence Glossary, located at https://hai.stanford.edu/ai-definitions :

- **Artificial Intelligence (AI)** is a term coined in 1955 by John McCarthy, Stanford's first faculty member in **AI**, who described it as "the science and engineering of making intelligent machines." Today it is a broad term for computer systems that can perform tasks with human-like intelligence, such as understanding language, recognizing images, learning from data, reasoning, and making decisions. Modern **AI** often works by finding patterns in large amounts of data and using those patterns to generate predictions or responses. It can be narrow (good at a specific task) or more general-purpose, like today's large language models that can handle many tasks.

- **Agentic AI** refers to AI systems designed to act as autonomous or semi-autonomous agents: they can set or interpret goals, plan and sequence actions, use tools (like web browsers, code, or APIs), make decisions based on feedback, and adapt over time to complete tasks. Unlike a purely reactive chatbot that only responds turn-by-turn, **agentic AI** is oriented around ongoing task execution—breaking down objectives, coordinating steps, and sometimes operating with minimal human oversight within defined constraints.

- **Generative AI (or GenAI)** refers to AI systems that can create new content like text, images, music, code, or video. These systems learn patterns from training data and generate novel outputs that resemble the original data, often powered by architectures like GANs, transformers, diffusion models, and variational autoencoders. These models power applications including chatbots, code generation, and creative tools. They also raise questions about the potential for misuse including creating misinformation and deepfakes.

- **Generative Pre-Trained Transformer (GPT)** is a family of large language models developed that uses the transformer architecture to generate human-like text based on input prompts. The "pre-trained" aspect means the model is first trained on massive amounts of text from the internet to learn language patterns, grammar, facts, and reasoning, then can be fine-tuned or used directly for specific tasks like conversation, writing, translation, or code generation without task-specific training.

- **Hallucinations in AI** refers to instances where an artificial intelligence system generates information or responses that are incorrect, misleading, or entirely fabricated but presented as factual. This often happens in language models or image generation when the AI produces outputs not supported by the training data or real-world facts.

- A **Large Language Model (LLM)** is an AI system trained on massive amounts of text data to understand and generate human-like language. It uses deep learning techniques, specifically neural networks with billions of parameters, to predict and produce coherent text, answer questions, translate languages, write code, and perform various other language-based tasks.

- **Prompt Engineering** is the practice of carefully crafting instructions, or "prompts," to guide AI language models toward producing desired outputs. By adjusting the wording, structure, and context provided in a prompt, users can significantly influence the quality, style, and accuracy of the model's responses. This skill has become essential for effectively using large language models across tasks ranging from creative writing to technical problem-solving to data analysis.

NOTES

NOTES