

Using generative AI to respond to patient complaints

Reading time:

Dr Owen Bradfield Chief Medical Officer of MIPS

Date created: 06/05/2026



As artificial intelligence (AI) becomes more accessible and sophisticated, healthcare practitioners are relying on it to streamline administrative tasks. MIPS is increasingly seeing members use generative AI to assist them respond to patient complaints. However, this comes with medicolegal risks that must be carefully considered. This article sets out some guidance for safe use of generative AI.

What is generative AI?

Generative AI is a type of artificial intelligence that learns from patterns in text or documents to create new content that looks or sounds like it could have been created by a human. It works by:

1. Being trained on large volumes of existing information.
2. Identifying patterns in data.
3. Using probabilities to predict the next most likely word in a sentence.
4. Refining through feedback – During training, the AI compares its outputs to the real data and adjusts itself to get closer and closer to the right answer.
5. Generating new content – Once trained, it can take a prompt (like “Write me a response to a complaint from a patient”) and create a brand-new result that follows the learned patterns.

Some generative AI tools are “open source”. This means that they are made available for public use and are trained on data entered by users into the tool. Therefore, if personally identifiable information is entered into an open-source AI tool by a practitioner, it may be sent overseas, may be used to train the tool, and may unintentionally appear in outputs generated within and outside Australia. This can result in serious breaches of privacy laws.

What is personally identifiable information?

Personally identifiable information includes any health information or personal information that identifies an individual. This includes information that, when matched with existing data, may re-identify an individual.

Australian privacy laws only allow practitioners to disclose personally identifiable information about a patient without their consent in limited circumstances. The disclosure must be reasonable and must be connected to the primary purpose for which information was collected (ie the provision of healthcare). Disclosing personal information into a generative AI tool to respond to a patient's complaint may not be connected to the primary purpose for collection and may breach Australian privacy laws if disclosed without consent.

Australian privacy laws do not apply to information that is de-identified. However, merely removing a patient's name, address and/or date of birth may be insufficient to sufficiently de-identify information if it is reasonably re-identifiable. For example, if a practitioner provided health information about "a politician who died at a beach", this may not be sufficiently de-identified because a Google search may reveal the individual's identity.



Do NOT enter personally identifiable information into generative AI tools

MIPS strongly recommends that practitioners NOT enter any personally identifiable information into any commercially available AI. This means that practitioners should NOT copy and paste a complaint letter that may contain personally identifiable patient information directly into an AI tool.

Entering personally identifiable information, including complaint details, into commercial AI platforms without explicit patient consent could constitute unauthorised and unlawful disclosure of personal information. This could amount to a serious data breach.

In addition, Australian privacy laws only allow personally identifiable information to be sent overseas either with explicit patient consent, or where the information is sent to a country with privacy protections that are substantially similar to those in Australia. Again, failure to comply with these requirements could amount to a serious data breach.

Instead of using personally identifiable information when interacting with generative AI tools, practitioners should use de-identified or synthetic data. Information must be generic, gender non-binary and free from names, ages, dates of birth, place names or any other information that may re-identify you, the patient or any third party (including other health service providers).

Examples of safe and unsafe prompts are given in the following table. Fictitious examples are used (including the use of

pseudonyms) to illustrate the importance of de-identification. Any similarity to an actual person, living or deceased, is purely coincidental.

| Not safe - could breach privacy | Safe - identifying information removed |
|--|--|
| <p>Chloe* is a 9-year-old girl in Grade 4 at Richmond Primary School. She needed a filling. I explained the procedure to her mother, Helen, who is school teacher. She was very anxious about the procedure and had lots of questions. I told Helen to stop asking so many questions and suggested that Chloe's dental caries may be due to her consumption of Coca Cola. Helen became upset, withdrew consent to the procedure, left the clinic, and has now written a complaint letter. Please suggest a response to this complaint.</p> | <p>I am a dentist. The mother of a young child has complained about my communication and consent process prior to performing a dental procedure. Please suggest factual and empathic ways to respond to the relative's complaint.</p> |
| <p>Mrs Allen*, aged 91 years of age, from Sydney Hills Nursing home presented with worsening confusion. She did not have a fever or dysuria when she presented to Royal North Shore Hospital in June 2025. She subsequently died and I have been asked by the Coroner to provide a statement outlining my treatment of the deceased. Please write my statement for me.</p> | <p>How should I go about drafting a factual statement setting out the circumstances of my treatment of a patient with confusion who subsequently died of suspected sepsis.</p> |
| <p>Mr Smith* is a 46-year-old heroin addict who was involved in a motor vehicle accident in September 2025 while driving under the influence of drugs. He was taken by ambulance to Royal Melbourne Hospital with a fractured clavicle. He was discharged on opioids and benzodiazepines. I refused to continue prescribing them after 2 weeks because I was concerned about the risk of addiction. The patient has now complained to Ahpra. I don't think I did anything wrong. Draft me a response.</p> | <p>My patient with a history of an opioid use disorder is unhappy that I refused to prescribe opioids and benzodiazepines for them. I was concerned about the ongoing risk of misuse, diversion and overdose. The patient has now complained to a regulator. Help me draft a polite but firm response.</p> |

*Pseudonyms

In September 2024, an **investigation** by the Office of the Victorian Information Commissioner (OVIC) found that a child protection worker breached Victorian privacy laws by entering sensitive personal details into ChatGPT to draft a court report about a child at risk. OVIC ruled that the Department of Families, Fairness and Housing failed to take reasonable steps to protect personal information, issuing a compliance notice and requiring the department to block staff access to generative AI tools like ChatGPT. The case highlights the risks of entering personal information into public AI systems.

In another **case** in August 2024, an individual used Microsoft's Copilot AI to generate a fictional sexual harassment scenario when giving a work health and safety presentation about sexual harassment. Unbeknownst to the presenter, the AI produced a case study that included confidential details of a real sexual harassment case. The individual in the AI-generated case study was known to the presenter's audience. The incident highlights the potential for AI to inadvertently disclose personal information when generating purportedly fictitious case studies.

Always check generative AI outputs for accuracy and tone

There is no doubt that generative AI tools can quickly and easily generate outputs that are detailed and seem to say all the right things. However, you must carefully review and personalise each AI output before finalising your response. If you have withheld identifying information when entering a prompt into generative AI (as you should), the output may be generic and clinically irrelevant. It may also contain errors and omissions. You should consider what is generated to be a rough first draft only to help you structure your thoughts and to assist in finding the appropriate tone for what you want to say. Submitting an inaccurate, incomplete or insensitive response may do you a disservice. The complainant may not feel heard or may believe that you have given their complaint inadequate care and attention. This may cause further dissatisfaction, and they may escalate their complaint further (for example, by making a notification to a regulatory body).

Generative AI is not a substitute for seeking

advice and assistance from MIPS

Members should always seek advice from MIPS before finalising a response to a patient complaint, irrespective of whether they have used generative AI to draft the response. Don't be afraid to tell us if you have used generative AI to draft your early response. This helps us understand how you have engaged with the complaint and what further input may be required before submitting your response. You remain personally and professionally accountable for all communication with patients, and an inappropriate response could have legal, regulatory, or reputational consequences. MIPS can ensure that the response is accurate, defensible, and consistent with medico-legal obligations, providing an essential safeguard before the response is sent.



Key takeaways

Generative AI use will continue to increase as the technology evolves and improves. Therefore, rather than shying away from the technology, MIPS encourages practitioners to engage with AI safely, professionally, responsibly and transparently.

In summary, when using generative AI to respond to patient complaints:

- do not enter personally identifiable patient information into generative AI
- remember that merely removing names, dates of birth or addresses may be insufficient to de-identify personal information
- always check the accuracy, relevance and tone of the output
- always seek advice and input from MIPS before submitting a response to a patient complaint, irrespective of whether generative AI was used to generate your response
- tell MIPS if you have used AI to generate your preliminary draft – this will help us to help you!

At MIPS, our medico-legal expertise is built into everything we do. If you're considering your indemnity options, [get a quote today](#).

Disclaimer:

Medical Indemnity Protection Society ABN 64 007 067 281 | AFSL 301912

All information on this page is of a general nature only and is not intended to be relied upon as, nor to be a substitute for, specific legal or other professional advice.

No responsibility for the loss occasioned to any person acting on or refraining from action as a result of any material published can or will be accepted by MIPS.

You should seek legal or other professional advice before relying on any content, and practise proper clinical decision making with

regard to the individual circumstances.

Information is only current at the date initially published.

If in doubt, contact our claims and 24-hour medico-legal advice and support team on 1300 698 573.

You should consider the appropriateness of the information and read the [Member Handbook Combined PDS and FSG](#) before making a decision on whether to join MIPS.
