AI Use Policy

Template for Mental Health Clinics

United States edition

A free, ready-to-adopt template for psychology and mental health practices

> **Provided by NovoPsych.** This template is a starting point, not legal advice. Have it reviewed by your own legal and professional advisers before adopting it.

How to use this template

This policy is written to be adopted largely as-is. The body (Part 3) is deliberately free of fill-in blanks so you don't have to edit it clause by clause. Everything specific to your practice lives in one place — the Schedule (Appendix A) — which you complete once.

To adopt this policy:

1. **Complete the Schedule (Appendix A)** — your practice name, the people who hold each governance role, your audit frequency, your default consent period, and your register of approved AI tools.
2. **Read Appendix B (Jurisdiction Requirements)** — this edition is tailored to your jurisdiction. Confirm the listed laws and professional-body guidance still apply (regulators update these regularly).
3. **Adapt the optional items** — if you want to name specific prohibited tools, or add practice-specific rules, the Schedule has space for it. No need to touch the body.
4. **Approve, date, and version it** using Appendix D, and brief your team.

> **A note on roles.** The policy refers to three roles — the AI Governance Lead (owns the policy and approves new tools), the Documentation Auditor, and the Incident Contact. In a small practice these may all be the same person. Map each role to a real name in the Schedule.

Part 1 — What the rules require

Two bodies of rules govern AI use in clinical practice in every jurisdiction: privacy and data-protection law, and the standards of your professional registration body. The specific laws, regulators, and guidance documents that apply to your practice are set out in Appendix B. Whatever your jurisdiction, the same themes recur:

- **Health data is specially protected.** It generally cannot be collected, used, or disclosed without a lawful basis — most often the patient's consent.
- **Sending data to a third-party AI vendor is a disclosure.** It needs a basis, appropriate security, and — where the vendor processes data overseas — specific cross-border safeguards. You remain accountable for what your vendor does with the data.
- **Patient data must not train the vendor's AI models.** Your agreement with the vendor must prohibit secondary use, including model training.

- **The practitioner stays accountable.** Regulators are unanimous: AI is a decision-support tool. Clinical judgement is never delegated to an algorithm, and the practitioner is responsible for the accuracy of the final record.
- **Consent is informed and usually tool-specific.** Patients must understand which tool is used, what it captures, where data goes, and that they can withdraw at any time without affecting their care. For tools that record a consultation, consent is obtained before recording begins.
- **Use is transparent.** Patients are told when AI is involved, and AI-assisted documents are labelled as such in the record.
- **Breaches must be reported.** Most jurisdictions have a mandatory data-breach notification regime (see Appendix B).
- **Competence is required.** Practitioners must understand the tool's capabilities, limitations, and biases before using it.

Part 2 — What clinicians worry about

Beyond formal regulation, a good policy should address these practical and clinical concerns head-on:

- **Accuracy and hallucinations.** Generative AI is probabilistic and can produce confident but wrong output — a mistranscribed medication or fabricated symptom could drive an unsafe decision if not caught.
- **Over-reliance and deskilling.** Busy clinicians may “rubber-stamp” AI notes without real review, eroding independent clinical judgement over time.
- **Loss of clinical nuance.** Text transcripts miss tone, body language, micro-expressions, and cultural context — all central to psychological assessment.
- **Medico-legal discoverability.** AI-generated notes and transcripts in a patient's file can be subpoenaed. A hallucinated detail or inaccurate risk assessment can create real liability.
- **Effect on the therapeutic relationship.** Recording or AI monitoring can change the dynamic and make some clients less willing to disclose.

A sound policy doesn't pretend these risks away — it manages them with consent, human review, vendor due diligence, and audit.

Part 3 — The policy

This is the operative policy. Adopt it as written; the details specific to your practice are recorded in the Schedule (Appendix A), and the legal and professional requirements for your jurisdiction are in Appendix B.

1. Purpose

Our practice is committed to the lawful, ethical, and responsible use of artificial intelligence (AI) to **support — not replace** — human-delivered psychological care. This policy governs how AI tools may be used in our clinical and administrative work, to protect patient privacy, uphold our professional obligations, and preserve the quality and safety of care. AI tools are adopted only where they support clinical workflows (such as documentation) and improve patient outcomes, and always under the accountability of a qualified practitioner.

2. Core principles

a) **Patient consent is required** for the use of AI tools with their personal or health information, and may be withdrawn at any time without affecting their care.

b) **Clinical decisions always remain with the practitioner.** AI is an assistant, never a decision-maker. Practitioners understand how their AI tools work, and their limitations and biases.

c) **Patient data is never used to train AI models.** We use only tools whose providers contractually guarantee this.

3. Who this applies to

This policy applies to all practitioners and clinical staff at our practice. The user of any AI tool covered by this policy is always a **qualified practitioner** (or clinical staff acting under a practitioner's direct accountability) — never the patient. Any limits on use by non-clinical administrative staff are recorded in the Schedule.

4. Regulatory and ethical framework

We use AI consistently with the privacy and data-protection law and the professional standards that apply to our practice. These are set out for our jurisdiction in Appendix B and are reviewed whenever the law or professional guidance changes.

5. Approved, restricted, and prohibited uses

Approved uses (with patient consent and clinician review) — the specific approved tools are listed in the Schedule:

a) AI scribes / ambient documentation for session transcription and note summarisation.

b) Drafting, summarising, and grammar-checking of clinical notes and correspondence.

Restricted uses (permitted only with enhanced safeguards and mandatory independent verification by the clinician):

c) Tools offering differential-diagnosis support or pattern analysis, which must always be independently verified and never relied on alone.

Prohibited uses:

d) Using AI as the sole basis for any diagnosis or any suicide or violence risk assessment.

e) Entering identifiable patient information into free, public, or non-approved AI systems (for example, consumer-tier ChatGPT, Gemini, or Copilot).

f) Using any AI tool that has not been approved through the process in clause 10.

g) Any additional tools or uses prohibited by our practice, as listed in the Schedule.

6. Privacy, security, and vendor due diligence

Before any AI tool is approved, we confirm that:

a) **A binding data agreement is in place** — a Business Associate Agreement, data processing agreement, or equivalent (see Appendix B for what your jurisdiction requires) — before any patient data is shared.

b) **The vendor contractually guarantees that patient data is never retained or used to train its own or any third-party AI models.**

c) **Data is encrypted** in transit and at rest, and stored in a known, compliant location (recorded in the Schedule).

d) **Cross-border transfers are accounted for.** Where a vendor processes data outside our jurisdiction, we apply the cross-border safeguards required by law (see Appendix B) and understand we remain accountable for the data.

e) **De-identification or redaction** is used wherever practical to minimise identifiable data sent for processing.

f) **Retention is deliberate.** Where the tool allows transcripts or other artifacts to be retained, we set a retention period consistent with our professional record-keeping and legal obligations (recorded in the Schedule), rather than defaulting to indefinite storage without a reason.

7. Informed consent and patient transparency

a) **Consent is obtained before AI use** — and, for any tool that records or processes a consultation, before recording begins. Consent is tool-specific: we explain which tool, what it captures, where data is stored and who can access it, the privacy risks, and the right to withdraw. A blanket “we use AI” consent is not sufficient.

b) **Consent is recorded** — we note the patient's consent (and any withdrawal) in their clinical record, for the consent period recorded in the Schedule.

c) **Withdrawal is honoured immediately** and without any reduction in the quality of care provided.

d) **Authorship is transparent** — AI-assisted documents are clearly labelled in the record (see the sample label in Appendix C).

8. Human review and accountability

a) **AI output that becomes clinical documentation — session summaries, notes, letters, and reports — must be reviewed, edited, and approved by the responsible clinician before it is saved to the record.**

b) **Session transcripts are a verbatim source artifact, not clinician-authored documentation.** Where our tool and retention settings keep a transcript (see the Schedule), it is stored as a source record — useful for later review and insight — and does not require line-by-line clinician review; the clinician remains responsible for any summary or note drawn from it.

c) **The clinician retains full professional, ethical, and legal accountability** for every clinical decision and for the accuracy of the final documentation. AI never carries that accountability.

9. Risk management and quality assurance

a) We acknowledge AI's limitations — hallucinations, bias, and the absence of contextual clinical judgement — and train staff to watch for them.

b) **Audit.** The Documentation Auditor reviews a sample of AI-assisted documentation, at the frequency set in the Schedule, for accuracy and completeness.

c) **Incident reporting.** Any suspected data breach, AI-generated inaccuracy, or use of an unapproved tool is reported to the Incident Contact without delay and recorded in our incident register. Data-breach notification obligations are set out in Appendix B.

10. Training and governance

a) **Training.** All clinical staff complete training on the capabilities, limitations, and privacy boundaries of any approved AI tool before using it.

b) **Approval process.** No new AI tool may be used until approved by the AI Governance Lead, who confirms it meets clauses 6 to 8 and records it in the Schedule.

c) **Ownership and review.** This policy is owned by the AI Governance Lead and reviewed at least annually, or sooner if the law or professional guidance changes.

Appendix A — Schedule (complete this)

This is the only part you must fill in. Everything the policy body refers to as practice-specific is recorded here.

Practice details and roles

Item	Detail
Practice name	
AI Governance Lead (owns policy, approves tools)	
Documentation Auditor	
Incident Contact	
Audit frequency (e.g. quarterly)	
Default patient consent period	
Transcript retention setting (e.g. delete after use / 7 years / indefinite)	
Limits on use by non-clinical staff (if any)	
Additional prohibited tools / uses (if any)	

Register of approved AI tools

Tool	Approved use	Data agreement signed?	Data location	No training on data?	Approved by / date
e.g. NovoNote	Session notes / summaries	Yes — date	Australia	Yes	

Appendix B — Jurisdiction requirements: United States

> **Currency:** Checked June 2026. The HIPAA Security Rule update (Jan 2025 NPRM) remains proposed, not final. State psychology-licensing rules may add requirements.

Privacy and data-protection law

- **HIPAA** — a vendor that creates, receives, maintains, or transmits PHI on your behalf is a business associate. A signed Business Associate Agreement (BAA) must be in place before sharing PHI.
- **No training on PHI** — the BAA must restrict the vendor to permitted purposes; it may not use PHI to train or improve its own AI models without explicit authorisation.
- **ACA Section 1557 / 45 CFR §92.210** — covered entities must make reasonable efforts to identify and mitigate discrimination from patient-care decision-support tools, including AI (in effect since 1 May 2025; enforcement posture uncertain).
- **HIPAA Security Rule** — a proposed update (Jan 2025) would strengthen cybersecurity requirements; track its progress.

Professional standards

- **APA** — “Ethical Guidance for AI in the Professional Practice of Health Service Psychology” (June 2025): AI must augment, not replace, human decision-making; the psychologist is responsible for final decisions, understanding the tool's training data, and mitigating bias. See also the APA Companion Checklist for AI-enabled tools.
- **State licensing boards** — check your state board for any additional AI, telehealth, or record-keeping requirements.

Cross-border data transfer

Where data is processed outside the US or by sub-processors, ensure the BAA binds all sub-processors to HIPAA-equivalent terms and that the minimum-necessary standard is applied.

Data-breach notification

HIPAA Breach Notification Rule: notify affected individuals and HHS (and, for large breaches, the media) without unreasonable delay and no later than 60 days from discovery.

Appendix C — Sample consent script and authorship label

The example below is written for **NovoNote**, an AI scribe whose audio is never saved, whose transcript retention the practice controls (it can be deleted after use, or kept for a set period — for example seven years or indefinitely — to support ongoing care and later insight), and whose data is not used to train AI models. Adapt the wording to your approved tool and your retention setting (recorded in the Schedule). You can give patients NovoPsych's ready-to-use resources: the NovoNote patient consent form and the NovoNote patient information sheet (both on the NovoNote security page).

Sample verbal consent script (adjust the retention sentence to match your setting):

> *"In our sessions I use an AI note-taking tool called NovoNote to help me document our work accurately, so I can focus on you rather than on writing. It listens to the session and produces a written summary, which I review and edit before it becomes part of your record. The audio is never saved. The session transcript [is deleted once the summary is written / is kept securely as part of your record for [period] to support your care]. Your information is stored securely in Australia and is never used to train AI models. You can decline, or change your mind at any time, and it won't affect your care in any way. Are you comfortable with me using it today?"*

Sample authorship label for the clinical record:

> *"Case note created with the assistance of AI (NovoNote), reviewed and adopted by [clinician name and title] on [date]."*

Appendix D — Version control

Version	Date	Approved by	Summary of changes
1.0			Initial adoption

> *Prepared by NovoPsych as a free resource for the mental health community. It reflects regulatory and professional guidance current as of June 2026 and is provided for general information only — it is not legal advice. Adapt it to your circumstances and seek your own legal and professional advice before adoption. Learn more at novopsych.com.*