



POLICY ON THE USE OF ARTIFICIAL INTELLIGENCE

A licensed doctor approves every prescription. AI supports; it never decides.

Document No.: CYD-AI-01 | Version 1.0 | Effective date: November 5, 2026 | Owner: Zebroid Private Limited | Status: Approved

1. Purpose and scope

This policy governs every use of artificial intelligence, including Zebroid AI, in the services of Consult Your Doctor and in the work of the staff of Zebroid Private Limited. Its purpose is to make sure that AI improves safety and convenience without replacing the judgement of a doctor.

2. Core principle

Every prescription is reviewed and approved by a licensed doctor, who makes the final decision. AI does not diagnose, does not prescribe and does not change treatment. An AI tool may flag a possible problem or send a reminder, but a person who is qualified to do so acts on it.

3. Permitted uses

- Checking prescriptions for possible errors, such as dose or frequency problems, duplicate medicines, known interactions and listed allergies.
- Flagging diagnostic results that fall outside the reference range, and unusual combinations of results, for review by a professional.
- Follow-up of patients after a consultation, including medicine reminders and symptom check-ins, under the Zebroid AI 12-Day Follow-up Protocol (CYD-AI-02).
- Administrative assistance, such as scheduling and drafting non-clinical messages, reviewed before sending.

4. Prohibited uses

- Making or communicating a diagnosis, choosing a medicine or dose, or telling a patient to stop or change a medicine.
- Issuing a prescription or report without the approval of the responsible professional.
- Using patient data to train or improve a model without a lawful basis, the patient's consent where required, and the approval of the data protection lead.
- Entering identifiable patient data into any public or unapproved AI tool.

5. Transparency to patients

Patients are told that AI is used to check prescriptions and to support follow-up, that it can make mistakes, and that a doctor makes every final decision. The prescription and the consent form say so in plain language.

6. Validation and monitoring

Before an AI tool is put into use, the medical lead and the technical lead test it against real and synthetic cases and record the results. After launch, the company measures how often the tool flags correctly, how often it misses a problem and how often doctors override it, and reviews these figures every month. A fall in performance, or any harm connected with the tool, is reported under CYD-CLN-03 and may lead to the tool being suspended.



7. Governance and accountability

The directors appoint an AI governance owner who keeps the register of AI tools, their purposes, vendors, versions and risks, approves changes, and reports to the directors each quarter. Vendors of AI tools must sign a data processing agreement and may not use the company's data for their own purposes.

8. Legal classification

Whether any AI function is regulated as medical device software, or as a high-risk system under the European Union's rules on artificial intelligence, is assessed in CYD-REG-01. No function that falls under such rules is put into use until the assessment is complete and any required steps have been taken.

9. Data used by AI tools

AI tools may use only the data that they need for their stated purpose. Identifiable patient data is used only in secure, approved systems, under written agreements that forbid the vendor from using the data for its own purposes or for training of its own models. Where data is used to test or improve a tool, it is anonymised or pseudonymised wherever possible, and the data protection lead must approve the use in writing. Patients are told about the use in the privacy notice, and where the law requires consent, it is obtained.

10. Levels of human oversight

Function	Who decides	Oversight
Prescription safety check	Doctor	Every flag is reviewed by the doctor before approval; the doctor may override with a recorded reason.
Report flagging	Pathologist or doctor	Flags support review only; the professional approves the report.
Follow-up check-ins and reminders	Doctor	Red-flag replies go to a human at once; the doctor reviews every alert and the final summary.
Administrative drafts	Staff member	Reviewed by the sender before sending.

11. Fairness and safety

The company tests its AI tools for different groups of patients, for example by age, sex and language, to find any group for which the tool performs worse, and corrects it or limits the use of the tool for that group. Tools are designed to err on the side of caution: where a tool is not sure, it must pass the matter to a person rather than reassure the patient.

12. Documentation and explainability

For each tool the register records its purpose, its inputs and outputs, the vendor, the version, the testing done, the known limits, the human oversight and the owner. Where a tool raises a flag, the doctor or professional can see the reason for it, in plain language, so that the flag can be judged. Changes to a tool, including model updates, are tested and approved before use and recorded in the register.

13. Vendor due diligence

Check	Done
Purpose, performance and limits documented and tested by the company	[]
Data processing agreement signed; no use of data for the vendor's own purposes	[]



Check	Done
Security review, including hosting location and sub-processors	☐
Right to audit and to receive notice of changes and incidents	☐
Regulatory status of the tool confirmed (see CYD-REG-01)	☐
Exit plan and return or deletion of data	☐

14. Incidents involving AI

An error, a missed alert, a wrong message or any unexpected behaviour of an AI tool is reported as an event under CYD-CLN-03. The AI governance owner can suspend a tool at once. The company records what happened, why, what was done and whether patients were affected, and tells affected patients and, where required, regulators.

15. Roles

Role	Responsibility under this policy
Directors	Approve the policy; decide on high-risk uses; receive quarterly reports
AI governance owner	Keeps the register; approves new tools and changes; leads reviews and incident handling
Medical lead	Validates the clinical safety of each tool; confirms the oversight arrangements
Technical lead	Builds, tests, monitors and secures the tools
Data protection lead	Reviews data use, notices and consent
Doctors and staff	Use the tools as intended; report errors and concerns

16. Review of this policy

The AI governance owner reviews this policy every year, and whenever a law, a regulator's guidance, a serious incident or a new kind of tool makes a change necessary. Changes are approved by the directors and communicated to all doctors and staff, and to patients where they affect what patients are told.

17. Training and review

Doctors and staff receive training on this policy when they join and once a year. The policy is reviewed every year and whenever a new AI function is introduced.

Approval

This document is approved by the directors of Zebrold Private Limited, the operator of Consult Your Doctor.

Hemendrah Kumar Sadamsetty

Electronically signed | November 5, 2026

Hemendrah Kumar Sadamsetty

Director, Zebrold Private Limited

Date: November 5, 2026

Indu Reddy Morthala

Electronically signed | November 5, 2026

Indu Reddy Morthala

Director, Zebrold Private Limited

Date: November 5, 2026