

AI Act - How Our AI Works - English

Making medical information accessible and understandable for everyone

You're Interacting with AI

When you use the Ditto app, you're working with artificial intelligence (AI) technology that helps translate complex medical information into language you can understand. This page explains exactly how our AI works, what it can and cannot do, and how we keep your information safe.

Important: Our AI is a helpful assistant, not a doctor. It never provides medical advice, diagnosis, or treatment recommendations. We take the quality of our AI output extremely serious, but just as humans, it can make mistakes. Always remain critical of the output you receive and validate with

At Ditto, we advocate responsible and compliant use of AI. This document is set up in concordance with Article 50 of the EU AI Act in order to help our users understand how they are interacting with AI.

What Our AI Does

AI is advanced that technology that can do (certain) things that normally requires human intelligence. In the case of Ditto, we have three main steps in our AI:

1. Medical Consultation Transcription: going from an audio recording to a transcript
2. Patient-friendly Summaries: summarizing a transcript or document such that you can understand
3. Translation: translating summaries to the language you understand best

Medical Consultation Transcription

What happens: When you record a conversation with your doctor, our AI converts the audio into a written transcript.

How it works:

- Your recording is stored on your device
- The audio is temporarily uploaded to Azure (Sweden Central EU region) where most of our processing happens. Your data is encrypted during transit.

- From here, we send the encrypted audio to our partner Juvoly. They are a NEN-certified Dutch company with data centers in the Netherlands. Their Grutto AI models transform the audio into an accurate transcript of what was said. They send this back to Ditto, always encrypted.
- The recording is immediately deleted from the system of Ditto and Juvoly after processing
- A secondary cleanup process ensures no recordings remain in Azure
- The transcript is returned to your device where it is stored
- Your original audio recording remains on your device under your control

What you get: A complete written summary of your medical conversation that you can easily read and share.

Patient-friendly Summaries

What happens: Our AI reads medical conversations and documents, then creates easy-to-understand summaries.

How it works:

- Your document content is securely transmitted to Azure for AI processing (Sweden Central EU region)
- The AI analyzes your medical transcript or document
- It rewrites this information using simple, clear language at B1 level
- We run multiple quality assessments to ensure that the summary is reliable and accurate. If any of our key checks fail, we don't give you output if it does not reach our quality bar
- For the development of our AI, we only use our own and synthetic data, never yours. We do offline validation together with patients, doctors and AI experts.

What you get: Medical information explained in everyday language that makes sense to you and your family.

Translation

What happens: Our AI takes summaries that are in one language and translates it into another one

How it works:

- The AI reads your patient-level summary
- It extracts the context and current language
- It creates a translated version in the language you requested

What you get: Your summary in the language that works best for you or your loved ones

What Our AI Cannot Do

No Medical Advice

- Our AI **never** tells you what medical decisions to make
- It **never** diagnoses conditions or suggests treatments
- It **never** replaces conversations with your doctor
- It **never** provides emergency medical guidance

No Guarantee of Perfection

- AI can sometimes make mistakes in transcription or translation
- Complex medical situations might not be perfectly captured
- Technical terms might occasionally be simplified too much or not enough

No Decision-Making

- Our AI processes and translates information only
- It doesn't make clinical judgments
- It doesn't prioritize which information is most important for your health
- It doesn't provide personalized medical recommendations

How We Keep Your Information Safe

Data Protection and Processing

- Recordings are stored on your device, not permanently in the cloud
- AI processing temporarily uploads data to Azure with appropriate safeguards
- Recordings are automatically deleted from Azure after processing
- A secondary cleanup process ensures no recordings remain in Azure systems
- We comply with strict European privacy laws (GDPR/AVG)

Your Data, Your Control

- Your recordings are stored on your device - you decide when to keep or delete them
- AI-generated content is stored on your device - you decide when to keep or delete it
- We don't automatically delete anything from your device
- No permanent storage occurs in Azure or other cloud systems
- When you delete the app, your data goes with it

Advanced Security

- Your information is encrypted when stored and when transmitted
- Only you and people you choose can see your AI-generated summaries
- We use enterprise-grade security systems

Quality Control

- LLM-as-a-judge systems assess content quality and appropriateness
 - Offline evaluation setup tests safety methods using non-user data
 - Continuous monitoring validates that quality systems work as intended
 - You can always report if something doesn't look right
-

Your Control Over AI

You Decide

- **Analytics control:** You can opt out of analytics data collection in app privacy settings
- **Delete your data:** Remove recordings and all AI content through the app interface (delete all event information)
- **Initial consent:** You control your consent during app setup with clear information about data handling
- **Data stays yours:** Everything is stored on your device under your control

Your Rights

- **Know when AI is used:** We clearly mark all AI-generated content
- **Access your data:** See everything our AI has created for you
- **Correct mistakes:** Ask us to fix errors in AI-generated content
- **Get explanations:** Ask how our AI processed your specific information
- **Give feedback:** You can always give feedback on received AI output

The Technology Behind Our AI

Azure + GPT-4o (Language & Reasoning)

We use **Microsoft Azure OpenAI Service** with **GPT-4o** for core language tasks. Benefits include:

- Advanced generative & comprehension capabilities
- Enterprise-grade security, compliance, and protection

- Data isolation per customer; inputs/outputs aren't used to train Microsoft/OpenAI base models
- Encryption in transit and at rest, plus identity & network controls
- Built-in content safety filtering and abuse monitoring
- Support for healthcare compliance (e.g. via proper contracts / configurations)

Juvoly / Grutto (Speech-to-Text)

We rely on **Juvoly's Grutto (V3)** for voice transcription, especially in medical contexts. Key features:

- Trained on clinical speech and medical vocabulary for higher accuracy
- Fully hosted in Dutch data centers (sovereign infrastructure)
- Very fast and energy-efficient (e.g. ~40× faster, lower energy than comparable models)
- Zero retention of raw audio or transcripts; does not learn from users' conversations
- Integrated and certified for Dutch GP settings

Langfuse (Observability & Optimization)

We use **Langfuse** to monitor, debug, and refine our AI workflows—without storing any user-identifiable content. Highlights:

- Captures internal technical traces: API calls, latencies, token usage, tool invocations, and more
 - Enables insights into performance bottlenecks, error patterns, cost/latency trade-offs, and version regressions
 - **No user data retention** — only anonymized metadata is stored; personally identifiable content is excluded or masked
 - Supports strong governance on data retention and privacy
-

How We Ensure Quality

Before AI Processing:

- Your data is processed securely using enterprise-grade Azure infrastructure
- We use secure, encrypted connections
- We validate that your content is appropriate for AI processing

During AI Processing:

- Our AI follows strict guidelines about medical content
- It's instructed to write at B1 Dutch language level for accessibility

- It's programmed to avoid providing medical advice
- It focuses only on making existing information more understandable

After AI Processing:

- LLM-as-a-judge systems evaluate content quality and appropriateness
 - Offline evaluation validates safety methods using ground truth data (non-user data)
 - Language level validation confirms B1 Dutch accessibility targets are met
 - We use your feedback to improve future processing
-

Transparency and Trust

Open About AI Use

- We never hide that you're interacting with AI
- All AI-generated content is clearly labeled
- We explain our AI's capabilities and limitations

Continuous Improvement

- We track how well our AI performs
- We fix problems when they occur
- We update our systems based on user feedback
- We stay current with the latest AI safety practices

Legal Compliance

Our AI system complies with:

- **EU AI Act** (classified as Limited Risk AI System)
 - **GDPR/AVG** privacy regulations
 - **Dutch healthcare** privacy laws
 - **Medical device** regulations (our AI is not a medical device)
-

Working Together: You, Your Doctor, and Our AI

Your Role

- Use AI summaries as a starting point for understanding
- Always discuss important medical decisions with your healthcare provider
- Report any errors or concerns you notice
- Keep your original medical documents and recordings as primary sources

Your Doctor's Role

- Remains your primary source of medical advice and care
- Can review AI-generated summaries if you choose to share them
- Makes all medical decisions based on their professional judgment
- Can help you understand any medical information, with or without AI assistance

Our AI's Role

- Helps you understand medical information more easily
 - Saves you time by creating readable summaries
 - Makes medical documents more accessible to you and your family
 - Supports better communication about your health
-

Questions and Support

Common Questions

Q: Is my information used to train AI models?

A: No. Your personal health information is never used to train or improve AI models. This is clearly stated during app setup and covered by your consent. We only use your data to provide you with transcriptions, summaries, and translations.

Q: Can I trust AI-generated summaries?

A: Our AI summaries are designed to be helpful and accurate, but they should supplement, not replace, your original medical documents and conversations with healthcare providers.

Q: What if the AI makes a mistake?

A: You can report errors through our app, and we'll investigate. Always refer to your original medical documents and consult your healthcare provider if you notice inconsistencies.

Q: Can family members see AI-generated content?

A: Only if you specifically choose to share it with them. You have complete control over who sees your AI-processed information.

Get Help

- **Technical issues:** support@ditto.care
 - **Privacy questions:** privacy@ditto.care
 - **AI-specific concerns:** privacy@ditto.care
 - **General inquiries:** info@ditto.care
-

Our Commitment to You

We believe that understanding your health information is a fundamental right. Our AI technology is designed to:

Empower you with understandable health information

Protect your privacy with the highest security standards

Support your relationship with healthcare providers

Give you control over your own health data

Be transparent about how our technology works

Remember: You are always in control. Our AI is a tool to help you better understand your health information, but you make all the decisions about your healthcare in partnership with your medical providers.

Last updated: September 25, 2025

About Ditto Care: We're a Dutch healthcare technology company committed to making medical information accessible and understandable for everyone while maintaining the highest standards of privacy and security.