corti.ai

The frontier AI lab for healthcare & life sciences

JUNE 2026

Meet Corti & egs

Corti was founded in 2017 to prove AI can meaningfully augment healthcare

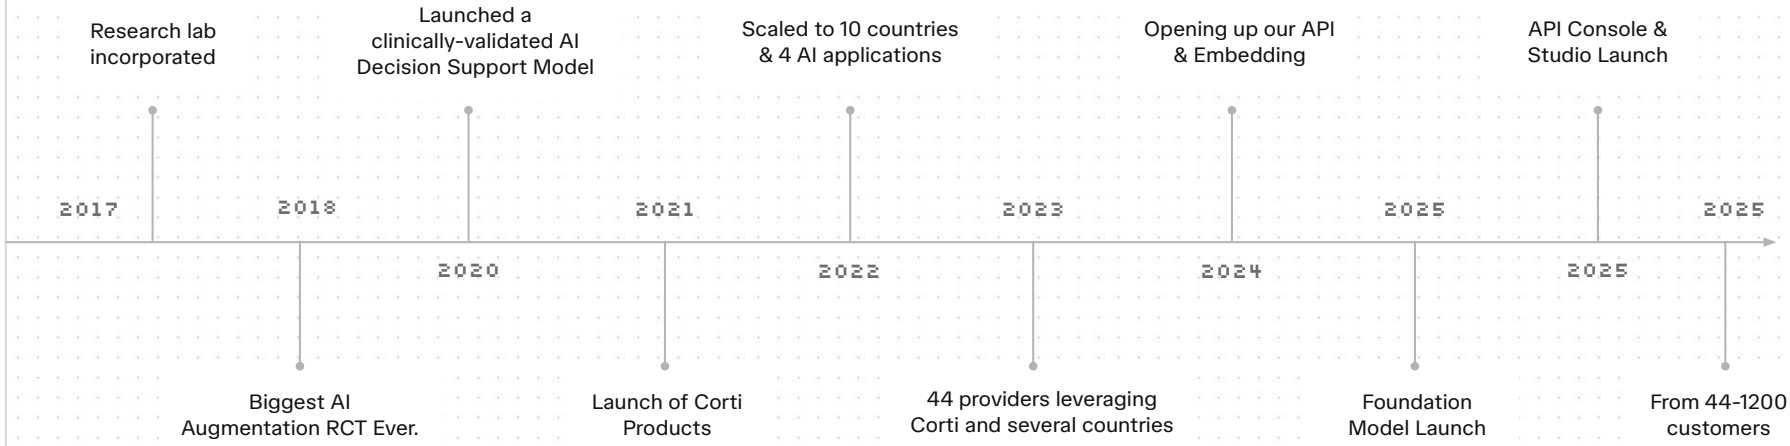

PRODUCT

The only full-stack AI platform built for the constraints of healthcare

Ecosystem	Developer-Built Applications
Agents	Experts · Pre-Configured Agents
Builder Enablement	AI Studio · SDKs · APIs · UI Components · Embedded Apps
Core Products	Speech-to-Text API · Symphony for Medical Coding · Agentic Framework · Text Generation API
Foundation Platform	Symphony · FactsR Reasoning · Code Like Humans
Infrastructure	Sovereign Cloud · Clinical Guardrail Architecture
Compliance	ISO:27001 · HIPAA COMPLIANT · GDPR · AICPA SOC 2 TYPE 2 · NHS DTAC · NHS DSPT · BSI C5 · DORA · ISAE 3000 · CJIS · NIS2

egs & corti

Bundles & Third Party Workflows

- Sonic Labs (Corti Assistant Embedded)
  - dns.comfort
  - dns.comfort.ki.connect
- Speech Processing Solutions
  - SpeechLive

Powered by corti.....

Welche Systeme werden unterstützt?

PVS:

Abasoft EVA
Aesculap (CH)
Agora (CH)
Albis
AOris
Arkandus
CARW/Penta
Compumed M1
CSPmed (A)
Data-AL
Data-Vital/David X
Duria, Duria²
EasyMed
EL-Win
ImageSystems
Indication
InnomedChoice (A)
Interarzt
KIWI
Medavis
MedicalOffice
Medistar
Medys
MedStreaming (USA)
Nephro
Pegamed
Poesy
Polypoint (CH)
Profimed
Q-Med
Quincy
S3
Smarty
T2Med
TabletOffice
Tomedo
TriaMed (CH)
Turbomed
Vitomed (CH)
Webmed (A)
WinRadiolog
x.comfort
x.concept
x.Isynet

weitere Systeme auf Anfrage ..

KIS:

Briselat
CGM Clinica
ClinicControl
(Reha-Net)
IMedOne
ITR
Medico//s
Nexus
Phoenix (CH)
u.v.m.

Market Feedback and Requirements

Why integration matters?!

- Seamless Integration into **familiar** workflows (Dictation or EHR)
- **No** “Copy & Paste” to get structured data into the EHR
- “Real Integration”& “Less Wrapping and UI Emulations”

  → CA Embedded provides a **robust, consistent, and secure interface** for parent applications to control and interact with embedded Corti Assistant.

Corti Assistant - Platform Strategy

Offering	Definition	Status
Embedded iFrame or WebView2	Direct iframe integration embedded inside EHR host platforms to sit right in the active frame. (Browser Application)	Focus
Android Application	Secure MDM-deployable build optimized for dedicated enterprise hospital tablets and clinical mobile devices.	Focus
iOS Application	Native iPad & iPhone client built for bedside rounds. Fully sandbox-compliant background recording algorithms.	Focus
• Standalone . NET & Browser Application Standalone - Corti Assistant • Apple Watch App		EOS(L)

Corti Assistant Embedded

Dictation: Preset voice commands & Punctuation Control

Template assembler - writing style customization at the section level

Template assembler - format customization at the section level

Dictation: Enabling interim text for supported languages

Template Assembler - Structural Customization

New UI language: Norwegian, Danish, Swedish, Dutch, Italian

Warning in case of muted microphone

Template Assembler

CA Embedded - Latest Improvements and Features

Preset voice commands & Punctuation Control

Description: Users currently interrupt their dictation flow to manually switch sections. This roadmap item introduces preset voice commands and the option to toggle spoken punctuation on/off for a smoother, more flexible dictation experience, which also sets the stage for a future standalone dictation workflow.

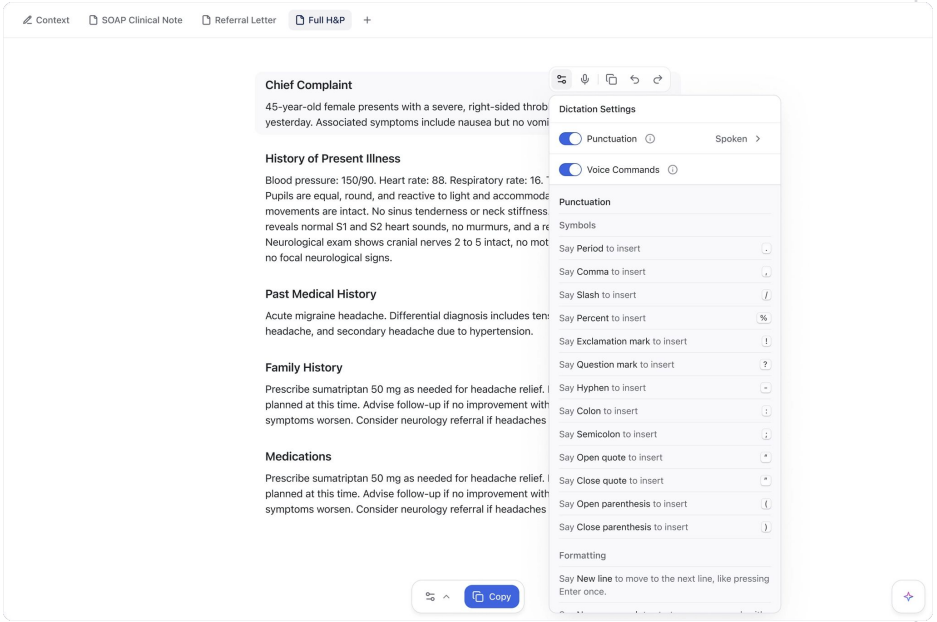

Template assembler - **writing style** customization @ section level

Description: This initiative focuses on giving clinicians greater control over their documentation by enabling customization of template section writing style. By introducing flexible style options—based on the most common user requests—clinicians can tailor templates to better match their communication preferences, improve clarity, and ensure outputs align with their specific clinical and organizational standards.

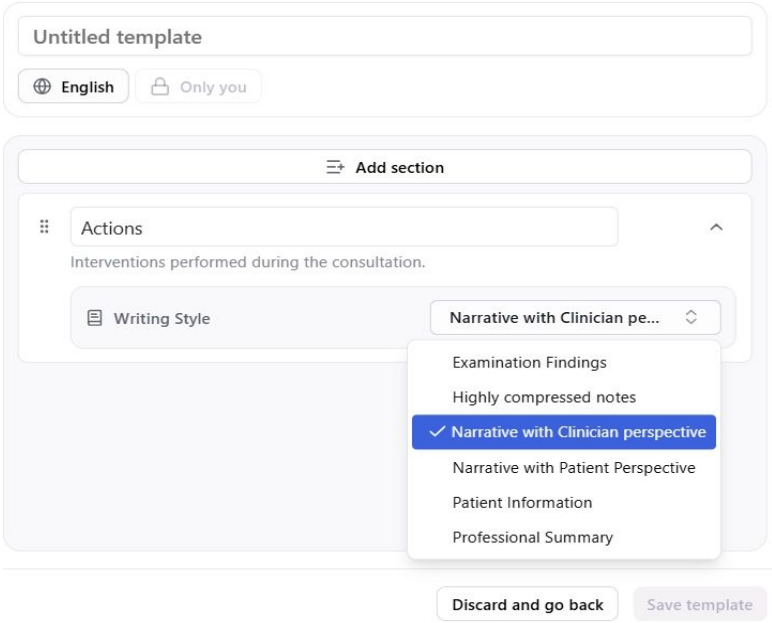

Template assembler - format customization @ section level

Description: This initiative focuses on giving clinicians greater control over their documentation by enabling customization of template section formatting. By introducing flexible formatting options - based on the most common user requests- clinicians can tailor templates to better match their workflows, improve readability, and ensure outputs align with their specific clinical and organizational needs.

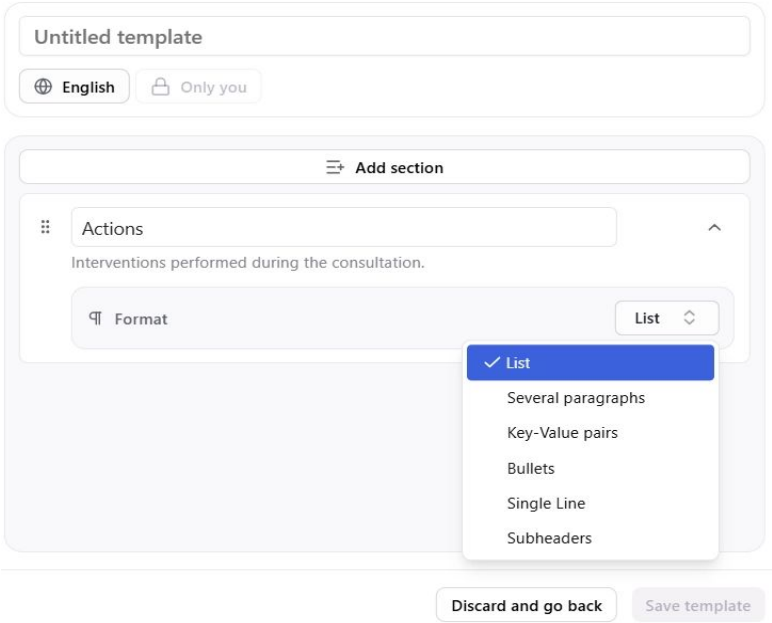

Template Assembler - Structural Customization

Description: Medical professionals share many core documentation needs, but differ in how they organize and structure their notes. This project focuses on enabling clinicians to personalize their documentation experience through predefined customization options that affect the structure of their template. These settings are applied at the user level, allowing individual preferences to be reflected while maintaining a consistent underlying framework.

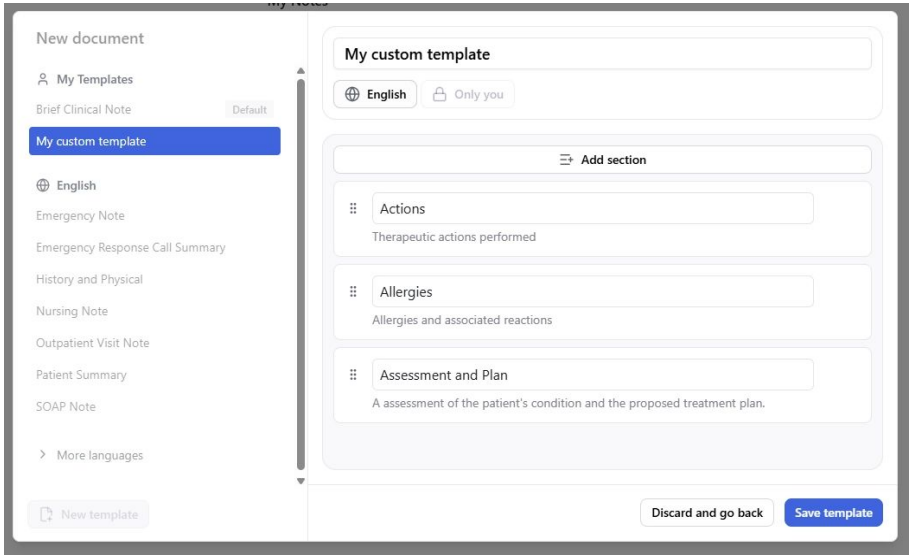

Enabling interim text for supported languages

Description: Building on capabilities available in the Corti ASR, this feature enables interim text during dictation. Assistant users will see partial transcriptions appear in real time while they are speaking, rather than waiting for the final text to appear. This provides quicker visual confirmation that their dictation is being captured and creates a more responsive and reassuring dictation experience.

UI Theming & Component Visibility in Assistant Embedded

Description: Enables customers to align the embedded Corti experience with their own brand by customizing the application's color theme. This first step toward a fully configurable UI allows host platforms to adjust accent and primary colors so the embedded experience visually blends with their environment - delivering a seamless, native look and feel for end users while minimizing implementation effort for partners.

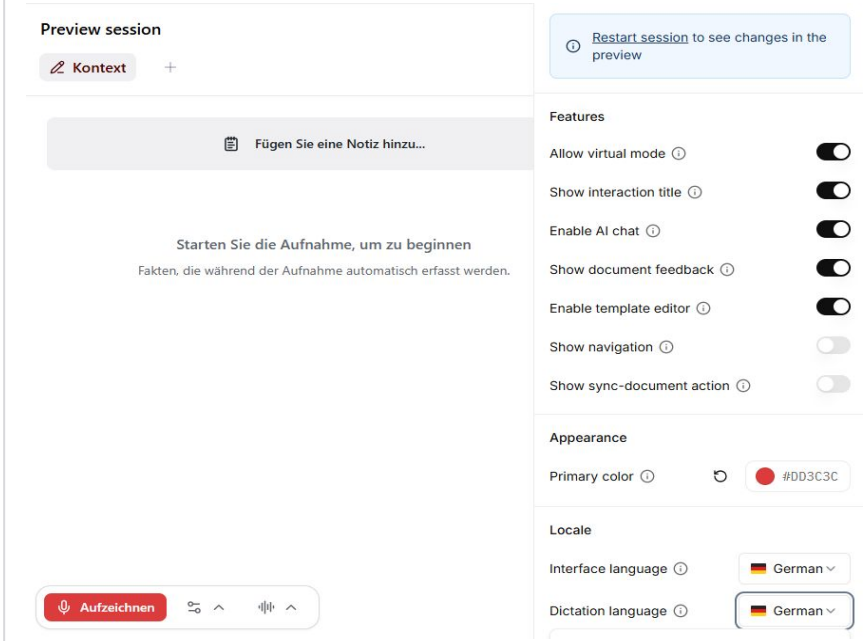

Warning in case of muted microphone

Description: Introduces an in-app notification that alerts users when a session recording is started with a muted microphone, helping prevent silent or incomplete recordings. This enhancement improves user experience, reduces entries in the product hazard log and enables faster root-cause identification when sessions report no audio data.

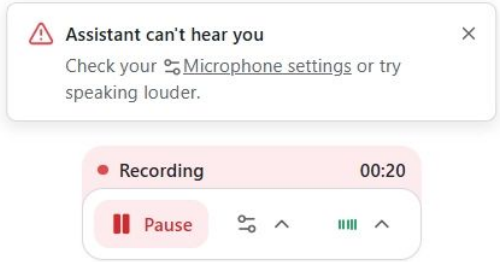

Language Support

Description: Additional Languages for Ambient Documentation and ASR

CA Embedded - WIP

Ability to Set Different Microphones for Dictation vs. Ambient Recording

Configure separate microphones for dictation and ambient recording - or use the same microphone for both. This makes it easy to use a directional mic to minimize background noise when dictating, and a mic optimized for capturing room audio during ambient sessions, ensuring the best possible recording quality in every situation.

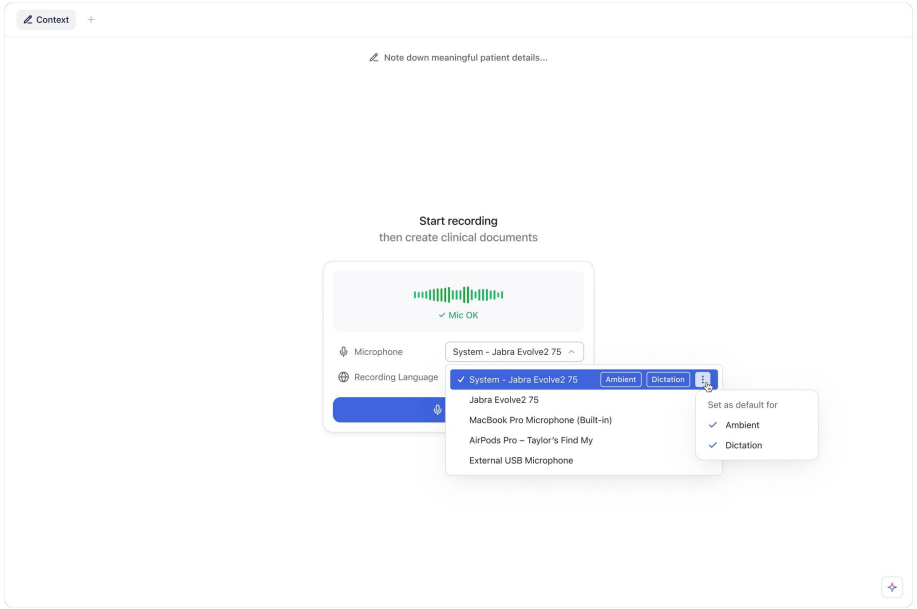

In-product Guidance for Microphone Set-Up

Description: This initiative aims at adding in-app guidance to help users set up their microphone for the best possible ambient recording experience. Users will receive recommendations on choosing the most suitable microphone from available devices, along with tips on placement and room setup to help capture both sides of the conversation more clearly and improve transcript quality downstream.

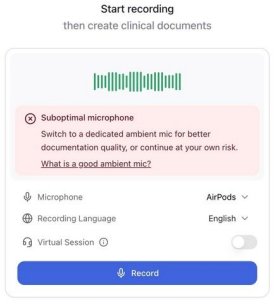

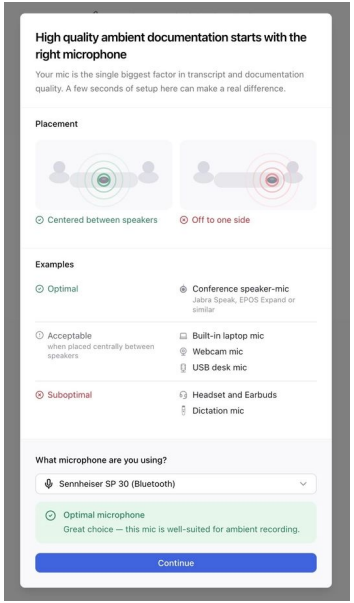

CA Embedded - Future Roadmap

CA Embedded - Roadmap

- Improved Ambient Documentation Worklow
- Agents in Assistant
- Dynamic Clinical Documentation Gernation (Base Section)
  Blank and structured empty templates
- Customer Voice Trigger Phrases

.... and much more :-)

egs & corti

Bundles & Third Party Workflows

- Sonic Labs (Corti Assistant Embedded)
  - dns.comfort
  - dns.comfort.ki.connect

- Speech Processing Solutions
  - SpeechLive

Powered by corti.....

corti.ai

Preset voice commands & Punctuation Control

Description: Users currently need to manually interact with the interface to move between sections while dictating in documents or the Context tab, which can interrupt the flow of input. This roadmap item introduces a set of preset voice commands designed to create a smoother, more seamless dictation experience by reducing the need for those manual actions. Additionally, to accommodate varying documentation styles, we will enable users to choose whether they want spoken punctuation on or off during dictation, providing greater flexibility and control over how their dictated text is formatted. These enhancements also lay the groundwork for a future standalone dictation workflow, enabling even more efficient and uninterrupted voice-driven documentation.

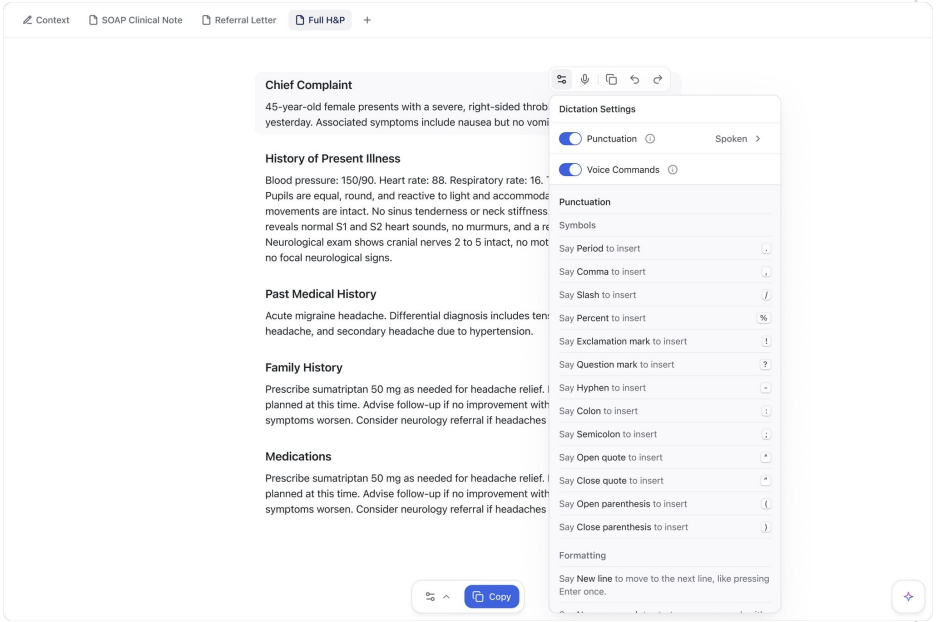

In certain industries hallucination risk have asymmetric negative consequences

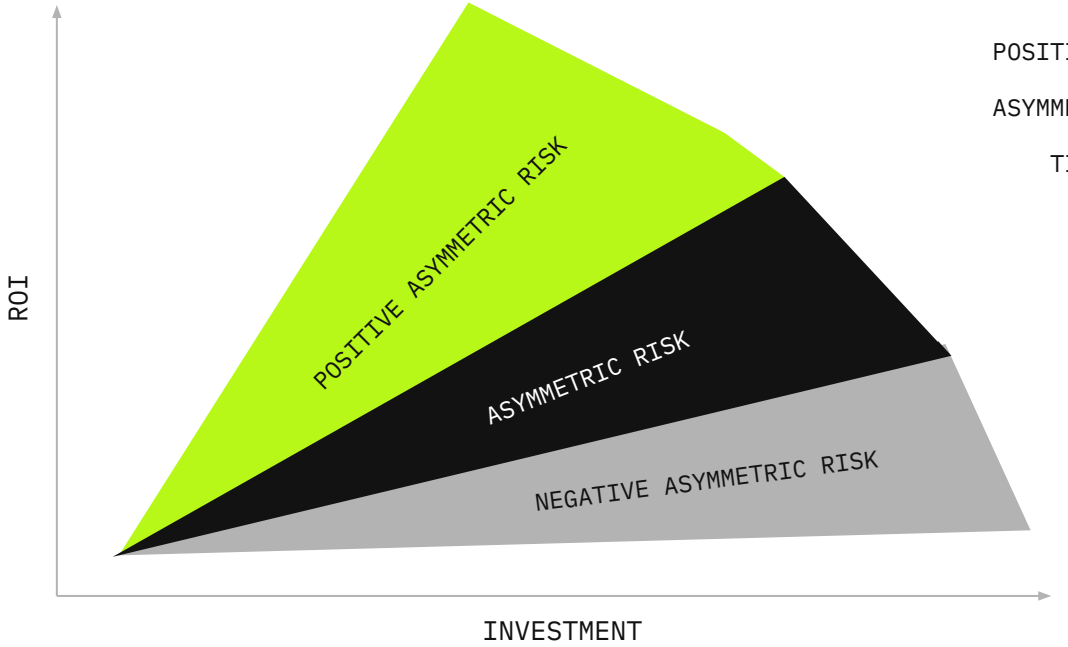

POSITIVE ASYMMETRIC RISK: E.G., AD GENERATION

ASYMMETRIC RISK: E.G., PREDICTIVE MAINTENANCE

TIVE ASYMMETRIC RISK: E.G., STEERING AN AIRPLANE

In vertical domains you compete on constraint mastery not compute & scaling laws

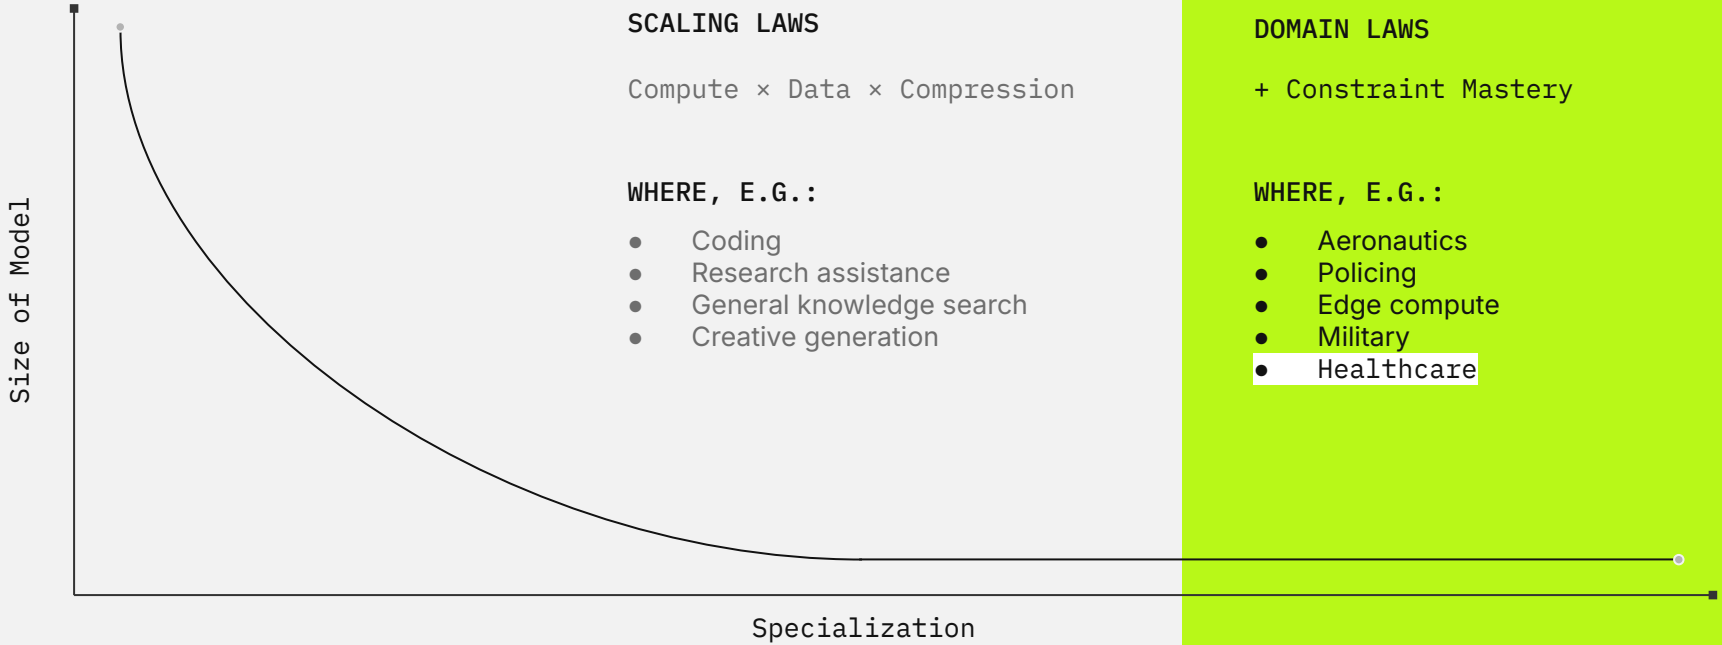

SCALING LAWS

Compute × Data × Compression

WHERE, E.G.:

- Coding
- Research assistance
- General knowledge search
- Creative generation

DOMAIN LAWS

+ Constraint Mastery

WHERE, E.G.:

- Aeronautics
- Policing
- Edge compute
- Military
- Healthcare

MEET CORTI

The only Frontier AI Lab built for clinical reality, European-founded with R&D at our core

250,000 Citations

IN THE R&D TEAM

$100 Million

VENTURE-BACKED FUNDING

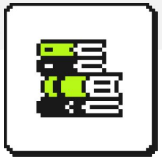

25+ Publications

PROVING THE TECHNOLOGY

5.5m Monthly Patients

INTERACTIONS SUPPORTED

We are the only AI lab dedicated to healthcare.

Corti is a vertical AI lab building the leading developer platform for healthcare and life sciences, powered by a compliant, accurate, and governable AI Stack.

We offer healthcare-native models, orchestration, auditability, and compliance primitives, delivered via APIs, SDKs, and web components.

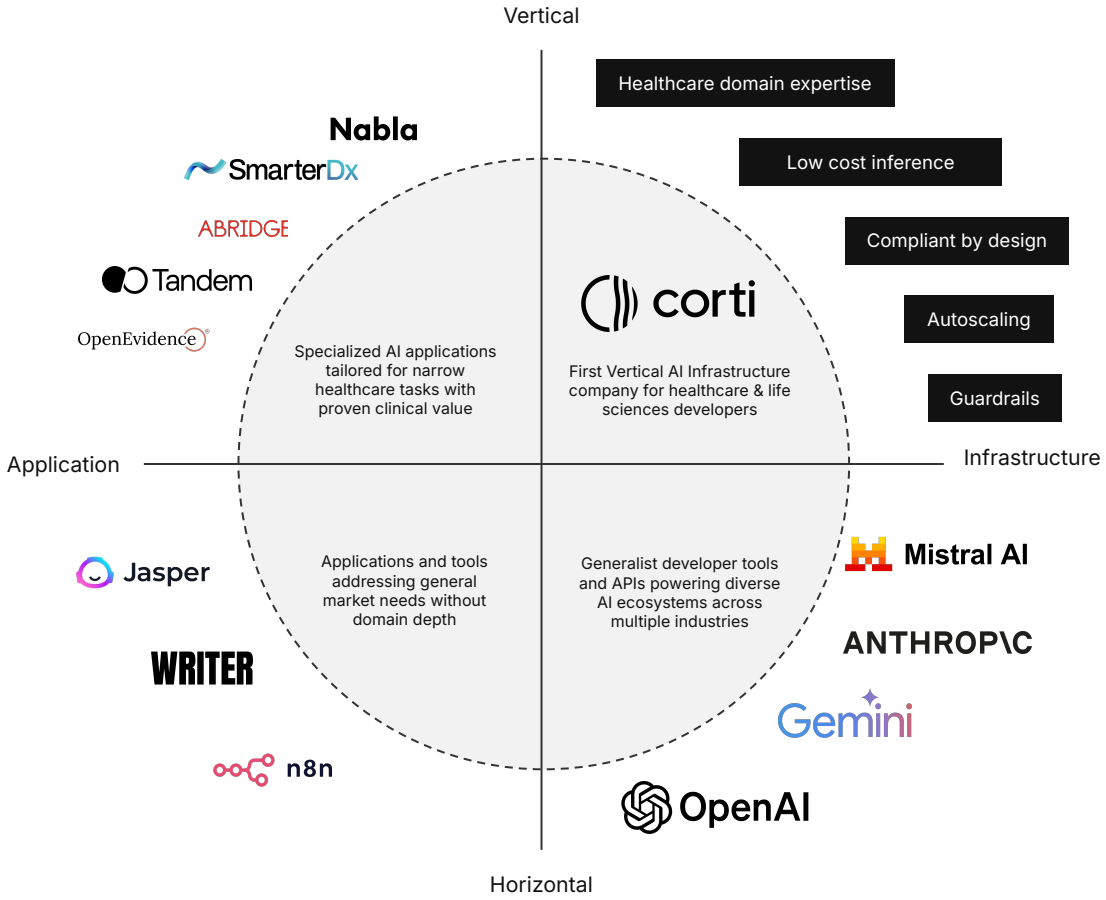

Healthcare needs AI that is built to thrive in clinical reality.

COMPARE

Building vertically means optimizing for specific domain law challenges

	Corti	Anthropic	OpenAI
Excellent World Model	✓	✓	✓
Healthcare Benchmarks	✓	✓	✓
HIPAA Compliance	✓	✓	–
Healthcare Specific Models	✓	–	–
Sovereign Hosting in Europe	✓	–	–
Healthcare Specific Guardrails	✓	–	–
SLAs & Performance Guarantees	✓	–	–
AI Act Compliance	✓	–	–
Boots on the ground to help deploy	✓	–	–

We are committed to peer-reviewed research

+369 h-index

+248k citations

+938 i-index

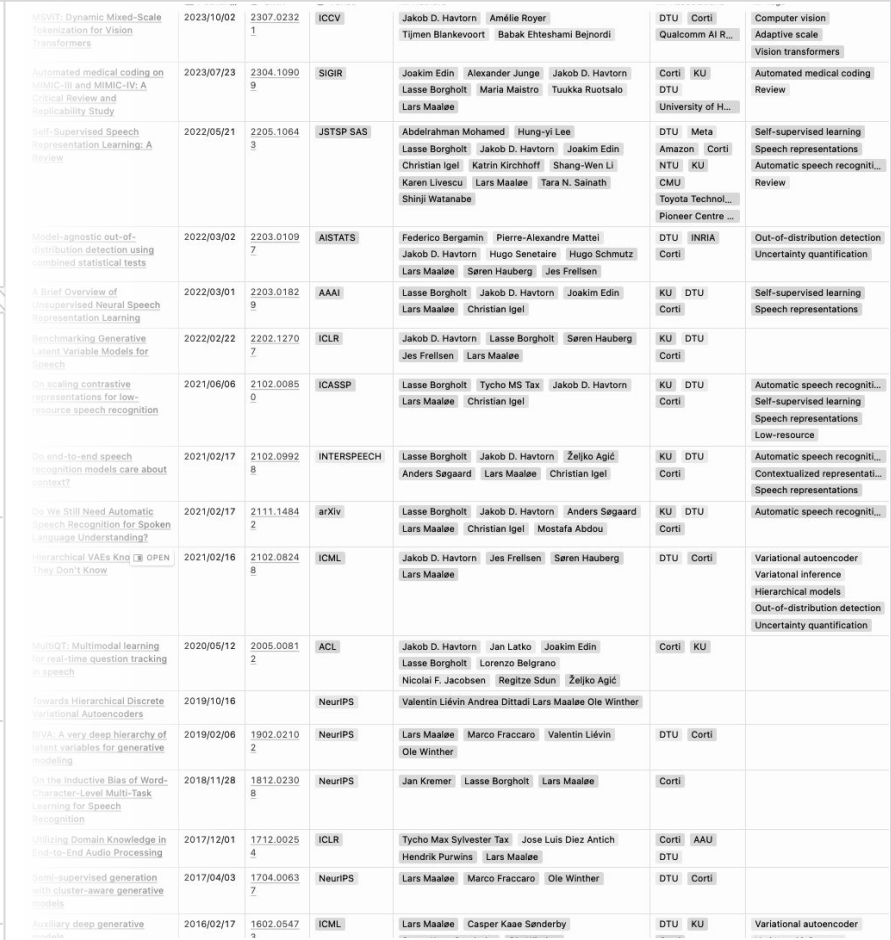

MSViT: Dynamic Mixed-Scale Tokenization for Vision Transformers	2023/10/02	2307.02321	ICCV	Jakob D. Havtorn, Amélie Royer, Tijmen Blankevoort, Babak Ehteshami Bejnordi	DTU, Corti, Qualcomm AI R...	Computer vision, Adaptive scale, Vision transformers
Automated medical coding on MIMIC-III and MIMIC-IV: A Critical Review and Replicability Study	2023/07/23	2304.10909	SIGIR	Joakim Edin, Alexander Junge, Jakob D. Havtorn, Lasse Borgholt, Maria Maistro, Tuukka Ruotsalo, Lars Maaløe	Corti, KU, DTU, University of H...	Automated medical coding, Review
Self-Supervised Speech Representation Learning: A Review	2022/05/21	2205.10643	JSTSP SAS	Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, Tara N. Sainath, Shinji Watanabe	DTU, Meta, Amazon, Corti, NTU, KU, CMU, Toyota Technol..., Pioneer Centre ...	Self-supervised learning, Speech representations, Automatic speech recogniti..., Review
Model-agnostic out-of-distribution detection using combined statistical tests	2022/03/02	2203.01097	AISTATS	Federico Bergamin, Pierre-Alexandre Mattei, Jakob D. Havtorn, Hugo Senetaire, Hugo Schmutz, Lars Maaløe, Søren Hauberg, Jes Frellsen	DTU, INRIA, Corti	Out-of-distribution detection, Uncertainty quantification
A Brief Overview of Unsupervised Neural Speech Representation Learning	2022/03/01	2203.01829	AAAI	Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Lars Maaløe, Christian Igel	KU, DTU, Corti	Self-supervised learning, Speech representations
Benchmarking Generative Latent Variable Models for Speech	2022/02/22	2202.12707	ICLR	Jakob D. Havtorn, Lasse Borgholt, Søren Hauberg, Jes Frellsen, Lars Maaløe	KU, DTU, Corti	
On scaling contrastive representations for low-resource speech recognition	2021/06/06	2102.00850	ICASSP	Lasse Borgholt, Tycho MS Tax, Jakob D. Havtorn, Lars Maaløe, Christian Igel	KU, DTU, Corti	Automatic speech recogniti..., Self-supervised learning, Speech representations, Low-resource
Do end-to-end speech recognition models care about context?	2021/02/17	2102.09928	INTERSPEECH	Lasse Borgholt, Jakob D. Havtorn, Željko Agić, Anders Søgaard, Lars Maaløe, Christian Igel	KU, DTU, Corti	Automatic speech recogniti..., Contextualized representati..., Speech representations
Do We Still Need Automatic Speech Recognition for Spoken Language Understanding?	2021/02/17	2111.14842	arXiv	Lasse Borgholt, Jakob D. Havtorn, Anders Søgaard, Lars Maaløe, Christian Igel, Mostafa Abdou	KU, DTU, Corti	Automatic speech recogniti...
Hierarchical VAEs Knq OPEN They Don't Know	2021/02/16	2102.08248	ICML	Jakob D. Havtorn, Jes Frellsen, Søren Hauberg, Lars Maaløe	DTU, Corti	Variational autoencoder, Variational inference, Hierarchical models, Out-of-distribution detection, Uncertainty quantification
MultiQT: Multimodal learning for real-time question tracking in speech	2020/05/12	2005.00812	ACL	Jakob D. Havtorn, Jan Latko, Joakim Edin, Lasse Borgholt, Lorenzo Belgrano, Nicolai F. Jacobsen, Regitze Sdun, Željko Agić	Corti, KU	
Towards Hierarchical Discrete Variational Autoencoders	2019/10/16		NeurIPS	Valentin Liévin Andrea Dittadi Lars Maaløe Ole Winther		
BIVA: A very deep hierarchy of latent variables for generative modeling	2019/02/06	1902.02102	NeurIPS	Lars Maaløe, Marco Fraccaro, Valentin Liévin, Ole Winther	DTU, Corti	
On the Inductive Bias of Word-Character-Level Multi-Task Learning for Speech Recognition	2018/11/28	1812.02308	NeurIPS	Jan Kremer, Lasse Borgholt, Lars Maaløe	Corti	
Utilizing Domain Knowledge in End-to-End Audio Processing	2017/12/01	1712.00254	ICLR	Tycho Max Sylvester Tax, Jose Luis Diez Antich, Hendrik Purwins, Lars Maaløe	Corti, AAU, DTU	
Semi-supervised generation with cluster-aware generative models	2017/04/03	1704.00637	NeurIPS	Lars Maaløe, Marco Fraccaro, Ole Winther	DTU, Corti	
Auxiliary deep generative	2016/02/17	1602.05473	ICML	Lars Maaløe, Casper Kaae Sønderby	DTU, KU	Variational autoencoder

Production-grade by default

Healthcare-native infrastructure means safety, compliance, and reliability are built in, not bolted on.

Foundation models

Trained on clinical data. Purpose-built for healthcare accuracy.

Sovereign Cloud

Data residency for the strictest regulated markets (EU, Nordics, NHS).

Research Pedigree

Published at NeurIPS, ICML, ICLR, ACL, npj Digital Medicine.

Scale

Trusted for 1M+ patient interactions every week.

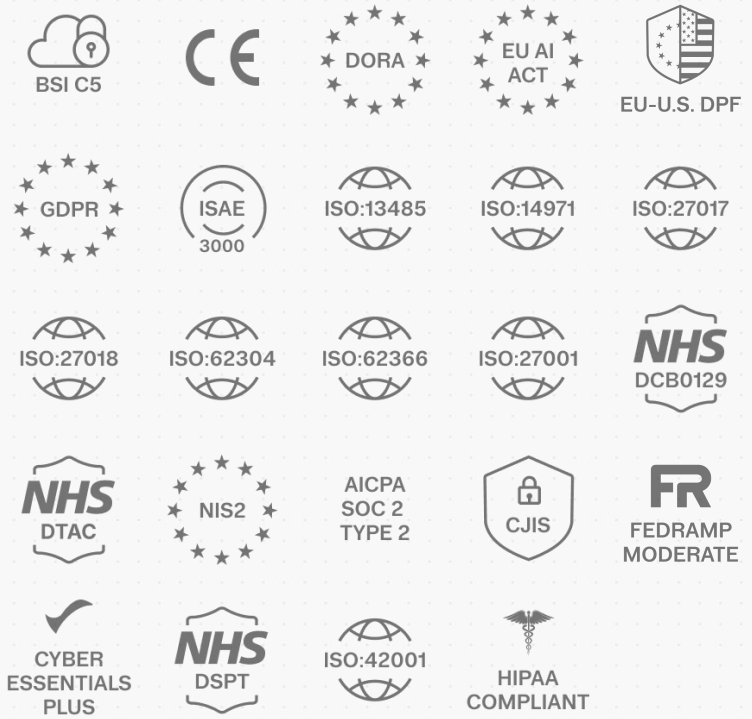

Corti Assistant - 6 ways of providing it...

EOL

Standalone .NET & Browser App

High-performance native Windows app. Global hotkeys, offline session caching, and low-latency audio device mapping.

Apple Watch App

Ultra-discreet background scribing. One-tap screenless capture preserves patient eye contact and care focus.

?

Browser App

Zero-footprint SaaS web interface. Instantly secure clinical scribing on standard workstation browsers with WebRTC.

Available

Embedded iFrame or WebView2

Direct iframe integration embedded inside EHR host platforms to sit right in the active frame.

Android App

Secure MDM-deployable build optimized for dedicated enterprise hospital tablets and clinical mobile devices.

iOS Application

Native iPad & iPhone client built for bedside rounds. Fully sandbox-compliant background recording algorithms.

Corti Assistant - 6 ways of providing it...

EOL

Standalone .NET

High-performance native Windows app. Global hotkeys, offline session caching, and low-latency audio device mapping.

Apple Watch App

Ultra-discreet background scribing. One-tap screenless capture preserves patient eye contact and care focus.

?

Browser App

Zero-footprint SaaS web interface. Instantly secure clinical scribing on standard workstation browsers with WebRTC.

Available

Embedded iFrame or WebView2

Direct iframe integration embedded inside EHR host platforms to sit right in the active frame.

Android App

Secure MDM-deployable build optimized for dedicated enterprise hospital tablets and clinical mobile devices.

iOS Application

Native iPad & iPhone client built for bedside rounds. Fully sandbox-compliant background recording algorithms.

PRODUCT

8 major launches in 12 months

Every capability a healthcare vertical AI lab needs, shipped and validated in production.

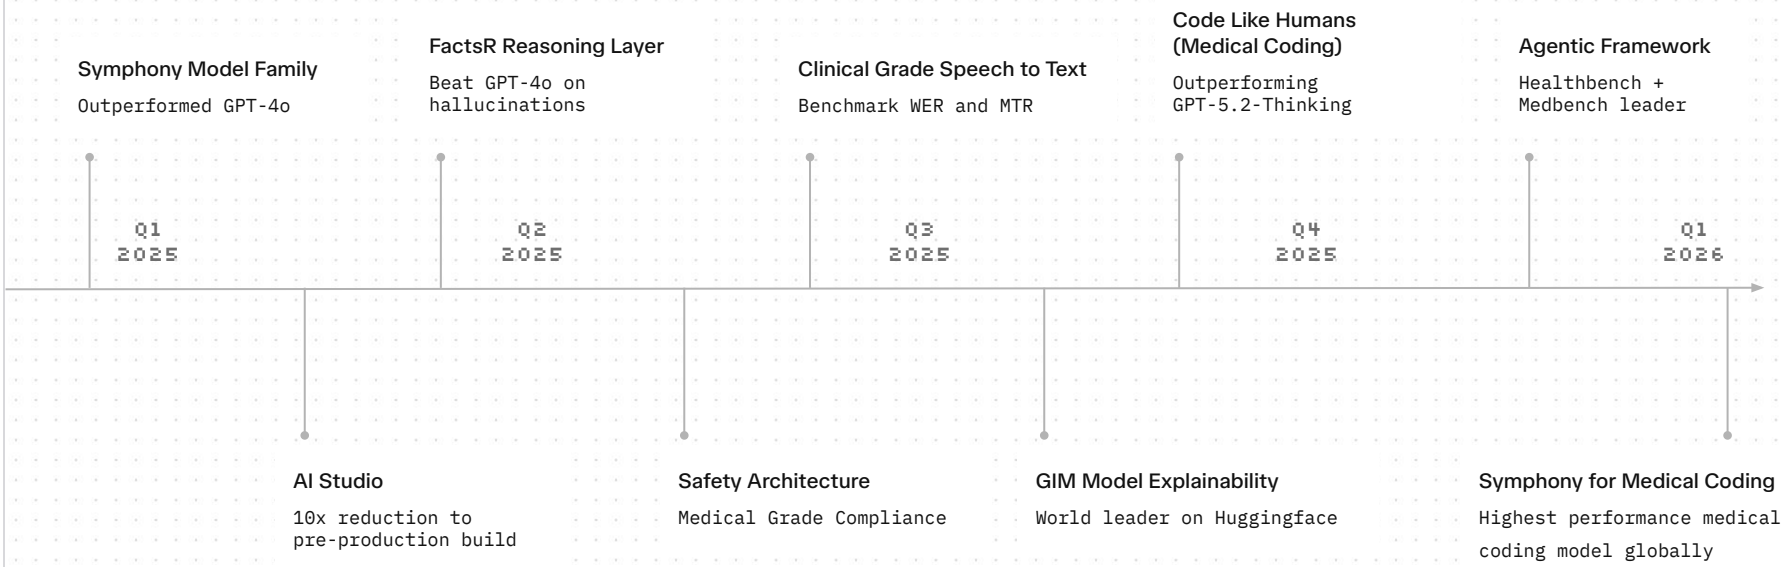

PRODUCT

Agentic Framework: Build AI agents that take on work in healthcare & life-sciences

Make multi-agent systems production ready.

- → ORCHESTRATION
- → AUDITABILITY
- → CONTEXT
- → COMPLIANCE
- → GUARDRAILS
- → SPEED

EXPLORE MORE

Production grade

Governance, audit trails, and parameter-level traceability built into the architecture.

Governed orchestration

Build complex, composable flows while Corti handles context, memory, and A2A coordination.

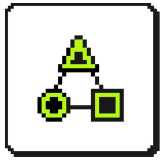

Open & extensible

MCP-compatible. Connect your own tools and external systems while inheriting governance and guardrails.

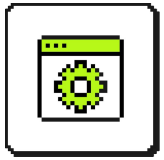

Domain expertise

Governed, compliance-ready architecture with healthcare expertise and domain reasoning built in.

PRODUCT

Symphony for Medical Coding: The world's most accurate medical coding model

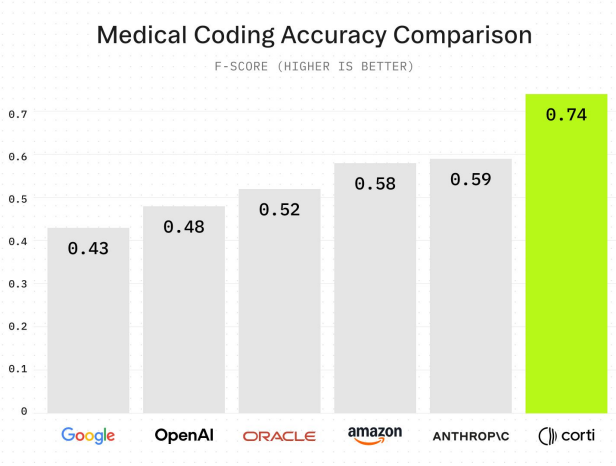

EXPLORE MORE

Best-in-class accuracy

Applies coding rules natively and decides what should and shouldn't be coded. Fewer false positives, lower denial rates, less rework.

Auditable

Every prediction includes source text, primary code, ranked alternatives, and evidence. Chain of custody in every API response.

Global

One integration point to U.S. and international coding standards (ICD-10-CM, ICD-10-PCS, CPT, and ICD-10-International).

Evergreen & Fast Updates

No accuracy window after system updates. Not reliant on retraining cycles. Coding guidelines in → code predictions our ***very fast*** given our agentic model setup.

API-native

Single API. Structured output. No implementation project, no professional services.

PRODUCT

Speech-to-Text: The most accurate clinical ASR for healthcare builders

THE SPEECH RECOGNITION LAYER FOR THE AGENTIC AGE IN HEALTHCARE.

EXPLORE MORE

Highly accurate

Unmatched medical term recall. Captures critical terminology where general ASR fails.

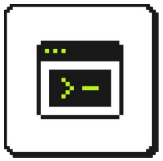

Developer control

API-first and fully modular. Customize to fit your workflow and embed directly into your stack.

Agentic-ready

Built to meet healthcare data requirements out of the box. No compliance retrofitting required.

Compliant on day 1

Structured outputs designed for downstream AI agents.

PRODUCT

Text Generation: Clinical reasoning infrastructure that de-risks LLMs at operational scale

1. Automate documentation
2. Reason across data
3. Produce structured data

EXPLORE MORE

Reasoning built in

Higher-order clinical reasoning layer that general-purpose LLMs don't give you out of the box.

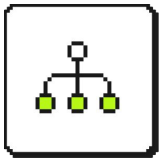

Auditable orchestration

Corti handles context pre-loading, output constraints, and clinical alignment so your team doesn't have to build it from scratch.

Reduced risk

Built for environments where safety is critical and margins are tight.

Lower cost

Reusable reasoning steps and structured outputs cut the marginal cost of every task you run through the engine.

PRODUCT

Full stack control down to the GPU

A unified stack that connects your multi-agent system directly to optimized GPU execution, giving developers full control over performance and enabling teams to focus entirely on workflow design rather than complex infrastructure.

Scale models and workloads without re-architecting your system

Control performance bottlenecks with visibility from API to GPU

Run securely in hosted or sovereign cloud environments

Benchmark-proven across every layer of the stack

Major benchmark wins in the last 6 months, outperforming OpenAI, Anthropic, Amazon Health, Oracle Health, Microsoft Health and Google Health.

SYMPHONY FOR MEDICAL CODING

F-SCORE (HIGHER IS BETTER)

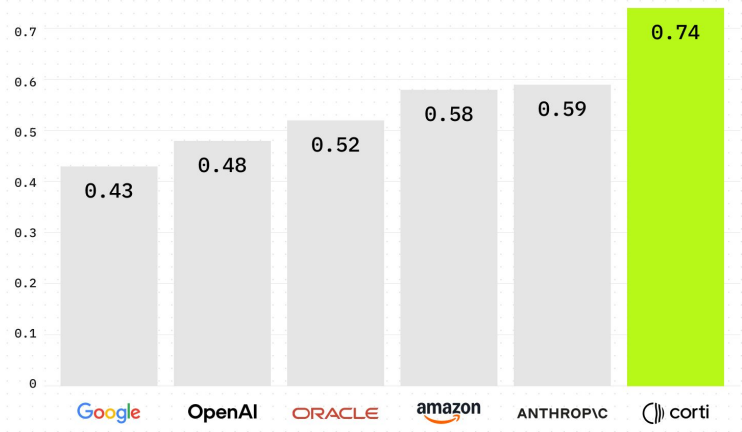

CLINICAL-GRADE SPEECH TO TEXT

WORD ERROR RATE (LOWER IS BETTER)

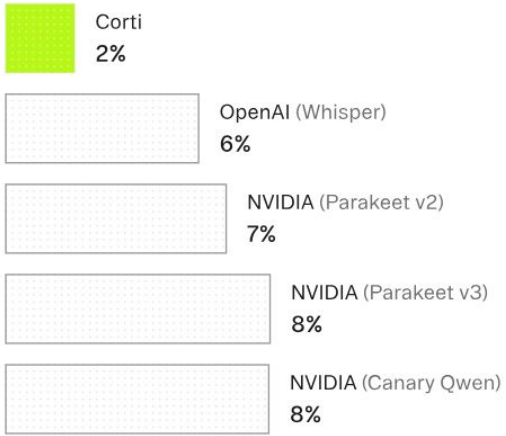

PRODUCT

Benchmark-proven across every layer of the stack

Major benchmark wins in the last 6 months, outperforming OpenAI, Anthropic, Amazon Health, Oracle Health, Microsoft Health and Google Health.

CODING

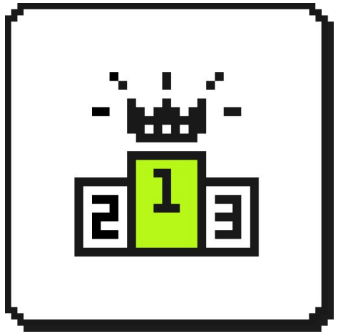

#1 ACI and MDACE medical coding benchmarks.

Proven in real-life healthcare across specialties.

HEALTHBENCH

#1 OpenAI HealthBench dataset.

Most powerful clinical reasoning, full stop.

SD BENCH

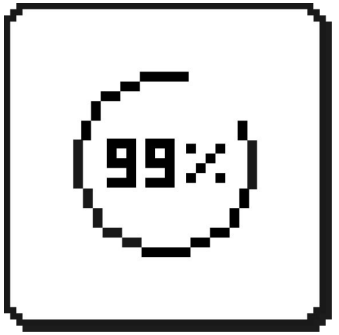

(WIP) promising results toward becoming the #1

on the medical superintelligence benchmark.

Developer use cases

USE CASE

3 generations of products built on Corti

We see three distinct generations of products being built on Corti that map to the core workflows of healthcare.

GEN 1

Scribes

Point-and-click solutions to augment human workflows.

GEN 2

Agents

Point-and-click agents added to the workflow to fulfill jobs.

GEN 3

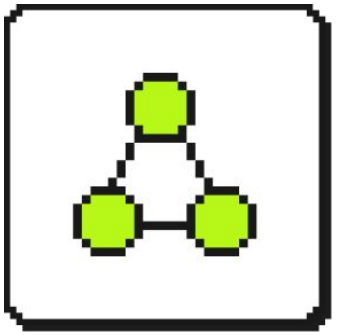

Multi-Agents

Autonomous headless multi-agent systems do end-to-end work.

USE CASES

Use cases spanning the entire patient journey:

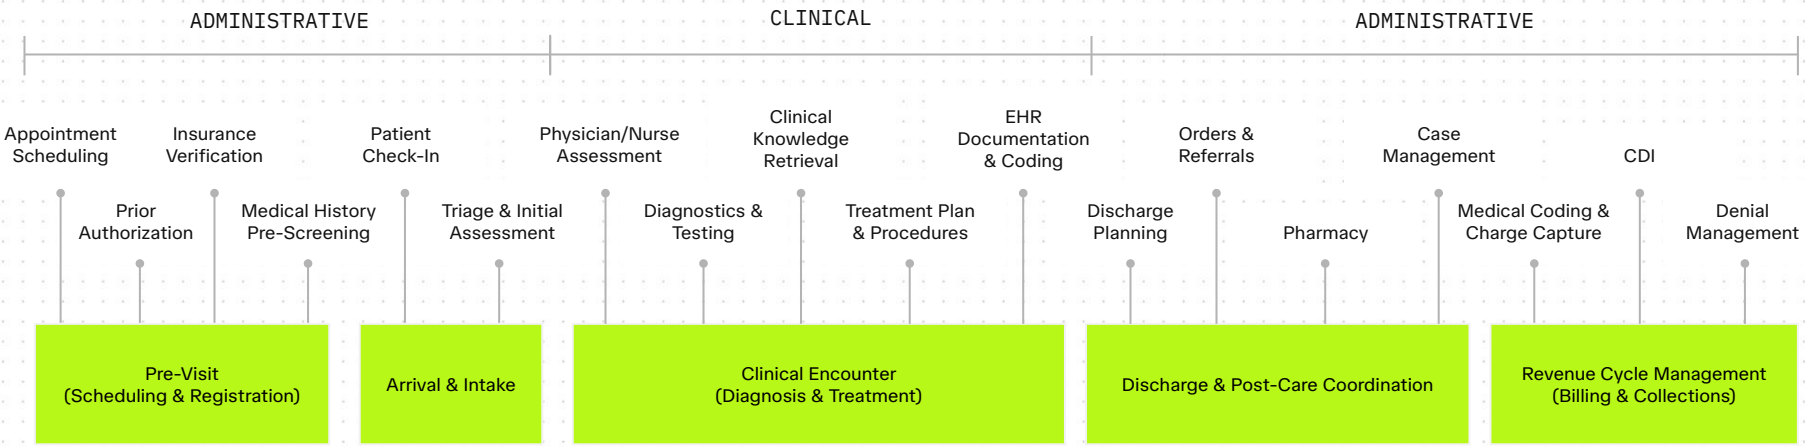

USE CASES

Developers across +50 specialties are building next generation healthcare apps on Corti.

Ambient Scribe + Dictation

$3Bn TAM by 2033 & $600M revenue in 2025

AGENTS

Rev Cycle management automation

$70Bn TAM by 2030, 24% CAGR

Clinical decision support

$11Bn TAM by 2030

AGENTS

Prior authorization

$5Bn subset of AI in RCM

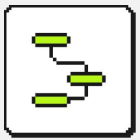

Patient navigation + triage

$5Bn TAM by 2030

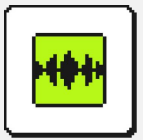

(COMING SOON)

Voice agents

$3Bn TAM by 2030

USE CASES

EXPLORE MORE

Corti API endpoints: powering seamless AI integration.

Interactions

An interaction represents a complete patient session, enabling developers to create, update, and manage interactions with a unique ID for real-time processing.

Recordings

Supports audio file uploads for post-session processing. Once linked to an interaction, Corti AI generates transcripts, facts, and structured documents.

Streams

Enables live transcription and real-time fact extraction. A WebSocket connection streams audio input to Corti AI, which responds instantly with transcripts and clinical findings.

Transcripts

Converts recorded interactions into accurate, structured text for review and integration into medical documentation.

Facts

Extracts key clinical details such as symptoms, vitals, and diagnoses in real time, ensuring relevance and precision in documentation.

Templates

Allows organizations to customize document structure, writing style, and formatting for consistency and clarity in medical reporting.

Documents

Generates structured clinical notes, referral letters, and patient summaries using transcripts and predefined templates. Multiple output formats are supported per interaction.

Codes

Uses AI-powered coding to generate ICD-10, CPT, and SNOMED codes for accurate billing and compliance. Custom models can be trained for client-specific documentation needs.

Agents

Modular artificial intelligence system for software developers to build advanced AI agents that perform high-quality clinical and operational tasks.

corti.ai