

Krisp Voice AI Lab: Revolutionizing Voice Technology

◆ This title was summarized by AI from the post below.



Arto Minasyan

2mo



Eight years ago, we started with a simple obsession: make voice AI work better than human ears expect

That obsession became Krisp's Voice AI Lab — one of the world's leading voice AI research centers, quietly built in Yerevan, Armenia.

From 2018 to 2024, we shipped seven technologies now used by tens of millions worldwide — noise cancellation, on-device turn-taking for voice agents, speech transcription, and more. Technologies most people use without ever knowing Krisp is behind them.

Then 2025 happened.

In a single year, we shipped three best-in-class breakthroughs:

- Accent Conversion — sound like yourself, only clearer
- Accent Understanding — AI that truly comprehends every accent
- Live Voice Translation — real-time, natural, human

This isn't luck. It's what happens when you concentrate rare talent around a singular mission.

Armenia has produced world-class mathematicians and engineers for generations. Our Voice AI Lab is where that intellectual tradition meets frontier AI research — one of the most focused collections of Voice AI talent on the planet.

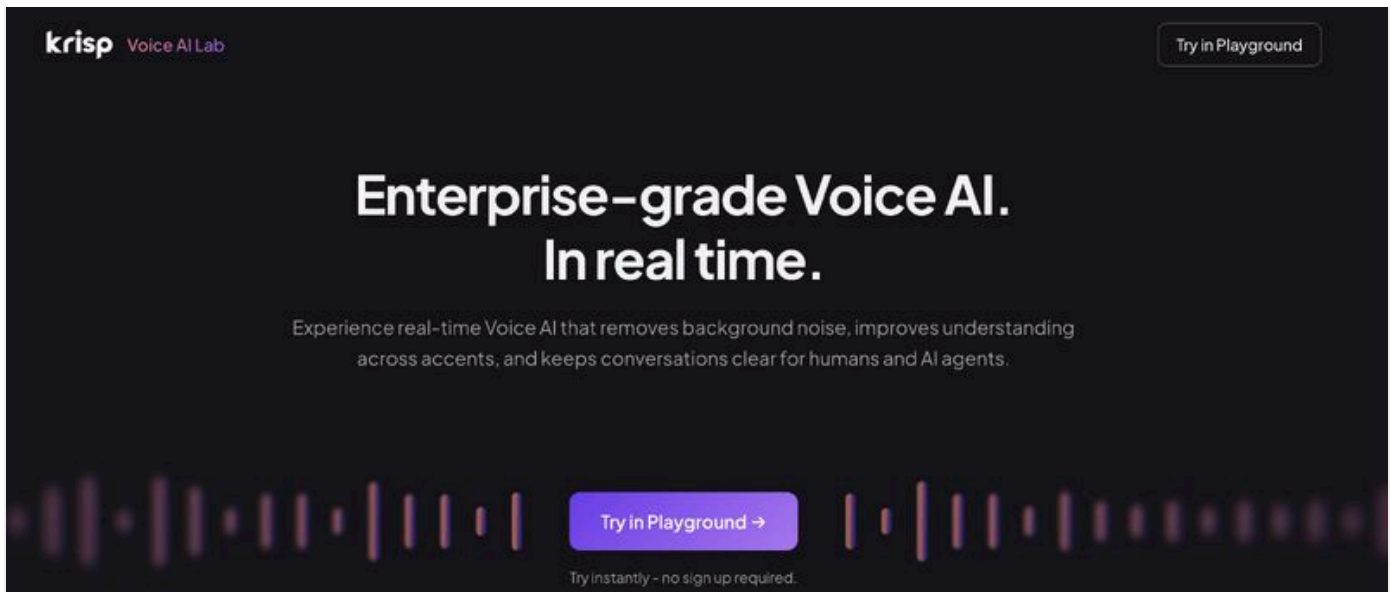
We just launched our Voice AI Lab page — experience the technologies firsthand and see why we believe Krisp is building the most compelling voice AI stack in the world.

Link is in the comment

PGR2026-00033

Krisp Technologies, Inc. v. Sanas.AI Inc.

Patent Owner Exhibit 2027



359 · 9 Comments

Like

Comment

Share



Arto Minasyan

2mo



<https://lab.krisp.ai/>

Like · Reply | 4 Reactions



Jebin Justus

2mo



moving from canceling noise to understanding accents is a profound shift toward cognitive justice...it shows that the true value of a Global Engine isn't just efficiency, but the democratization of clear communication...

Like · Reply | 1 Reaction

Mary Gabrielyan

2mo

Have been using Krisp for the past 6 years and cannot recommend it enough 🧡💙

Like · Reply

Konstantin Musatov

2mo

Incredible work! Focusing rare talent on a single mission clearly pays off. It's great to see such innovation coming out of Armenia.

Like · Reply

Artashes Hovhannisyan

2mo

Dear Artavazd, is Armenian language available? I am deaf, it would be benefit. [Arto Minasyan](#)

Like · Reply

Hemanth Vijay

2mo

[Arto Minasyan](#) Krisp is a great product. I've used it before for noise cancellation. Excited to check out the latest features!

Like · Reply

Aram Ghazaryan

2mo

Keep the rocket 🚀 !!

Like · Reply

Tigran Bayburtsyan

2mo

Nice! Having a playground is great!

Like · Reply

[See more comments](#)

To view or add a comment, [sign in](#)

More Relevant Posts

Eduard Hambardzumyan

2mo

Many people in my network often ask what we actually do at Krisp.

Instead of trying to explain it again, I'd rather let you experience it yourself.

Give it a try and see how it can transform your productivity.

Arto Minasyan

Co-founder & President @Krisp.ai | PHD in math
2mo

Eight years ago, we started with a simple obsession: make voice AI work better than human ears expect

That obsession became Krisp's Voice AI Lab — one of the world's leading voice AI research centers, quietly built in Yerevan, Armenia.

From 2018 to 2024, we shipped seven technologies now used by tens of millions worldwide — noise cancellation, on-device turn-taking for voice agents, speech transcription, and more. Technologies most people use without ever knowing Krisp is behind them.

Then 2025 happened.

In a single year, we shipped three best-in-class breakthroughs:

- Accent Conversion — sound like yourself, only clearer
- Accent Understanding — AI that truly comprehends every accent
- Live Voice Translation — real-time, natural, human

This isn't luck. It's what happens when you concentrate rare talent around a singular mission.

Armenia has produced world-class mathematicians and engineers for generations. Our Voice AI Lab is where that intellectual tradition meets frontier AI research — one of the most focused collections of Voice AI talent on the planet.

We just launched our Voice AI Lab page — experience the technologies firsthand and see why we believe Krisp is building the most compelling voice AI stack in the world.

Link is in the comment

3

Like

Comment

Share

To view or add a comment, [sign in](#)

Marianna Obrist

2mo

Fantastic news and well done [Shu Zhong](#) for shaping a new research direction on human-AI perceptual alignment beyond mainstream audio-visual modalities 🧠

[UKRI Interdisciplinary Textiles Circularity Centre UCL Computer Science](#)

Shu Zhong

Data Scientist @ Microsoft | PhD candidate at UCL

2mo · Edited

📄 Our paper "TouchAI: Exploring human-AI perceptual alignment in touch through language model representations" is finally out in the International Journal of Human-Computer Studies (with [Elia Gatti](#), [Youngjun Cho](#) and [Marianna Obrist](#)).

🧠 We asked a simple question: do language model representations align with how people perceive materials through touch?

Short answer: not really.

🔍 We found a clear gap between how humans feel materials and how AI represents them. The mismatch is not random, and shows strong perceptual biases.

👥 We intentionally worked with non-expert participants, because AI should be usable by everyone.

💡 Non-experts are often assumed to struggle to describe touch (and smell). Interestingly, we found people can describe materials in consistent ways (e.g. cattle leather or faux fur), but LLMs still struggle to recognise what material those descriptions refer to.

The issue is not the humans, it's the AI.

Current LLMs simply do not know much about how things feel.

🤖 As multimodal AI expands, underexplored senses like touch still have a long way to go.

Check the paper 📄 in the comment

[UKRI Interdisciplinary Textiles Circularity Centre](#), [UCL Computer Science](#)

[#AI](#) [#HCI](#) [#MultimodalAI](#) [#Touch](#) [#HumanAI](#) [#PerceptualAlignment](#) [#Alignment](#)

TouchAI: Exploring human-AI perceptual alignment in touch through language model representations

[sciencedirect.com](https://www.sciencedirect.com)

41

Like

Comment

Share

To view or add a comment, [sign in](#)**VoxBridge AI**

181 followers

2mo

Why Human Speech Data Still Powers AI

As artificial intelligence continues to evolve, one thing remains clear:
AI systems are only as good as the data they learn from.

High-quality human speech data plays a critical role in building accurate voice technologies, speech recognition systems, and multilingual AI models.

At VoxBridge AI, we focus on helping AI teams access structured, multilingual speech datasets through:

-  Speech data collection
-  Audio transcription
-  Multilingual linguistic teams
-  Quality-controlled datasets

Our goal is simple:

To bridge the gap between human language and machine intelligence.

As demand for diverse speech data continues to grow, collaboration between human contributors and AI teams becomes more important than ever.

If you're working on speech AI, voice technology, or multilingual datasets, we'd love to connect.

[#ArtificialIntelligence](#) [#SpeechData](#) [#AITrainingData](#) [#MachineLearning](#) [#Transcription](#) [#VoiceAI](#)

Like

Comment

Share

To view or add a comment, [sign in](#)

Yiwen Shao

2mo

Excited to give a talk at the Conversational AI Reading Group this week.

I'll discuss the evolution of audio foundation models, focusing on our work in Auden

(<https://lnkd.in/gXWKmYPt>) — exploring how large-scale audio encoders can serve as the backbone of Large Audio Language Models.

We'll cover:

- Audio encoder pretraining
- Representation learning across audio domains
- Integrating encoders with LLMs

Happy to discuss with the community!

Pooneh Mousavi

Computer Science PhD Candidate
2mo

📅 This week at our Conversational AI Reading Group:

We're happy to welcome Yiwen Shao, who will discuss the evolution of audio foundation models:

📖 Auden: From audio encoders to Large Audio Language Model.

📅 Thursday, March 5 | 11:00 AM – 12:00 PM EST

🎤 Speaker: [Yiwen Shao](#) - [Tencent](#) AI Lab

📺 YouTube channel: <https://lnkd.in/eBq6DRKp>

🔗 Details: <https://lnkd.in/eXHe6XyC>

[#ConvAIRG](#) [#ConversationalAI](#) [#ReadingGroup](#)

Conversational AI Reading Group

poonehmousavi.github.io

21

Like

Comment

Share

To view or add a comment, [sign in](#)**SUBIR MUKHERJEE**

2mo · Edited

For years, transcription was seen as a quiet back-office task — someone listening carefully and converting speech into text. Today, that same skill sits at the heart of Artificial Intelligence.

Every voice assistant, speech-to-text engine, and conversational AI system learns from human-reviewed language data. Behind every “smart” AI response is an enormous amount of carefully transcribed, annotated, and validated speech. In many ways, transcriptionists have become the silent architects of AI.

Modern transcription is no longer just about typing what we hear. It now involves understanding context, speaker intent, accents, background noise, conversational flow, and linguistic nuances — the very elements that help train and refine AI systems.

AI can process speech at scale, but it still relies on human judgment to detect subtle errors, improve contextual accuracy, and maintain data quality. That intersection between human listening and machine learning is where the real progress happens.

After more than two decades working across medical, legal, and AI-driven audio data projects, I find this evolution fascinating. The profession is quietly transforming from transcription to language intelligence and AI data quality.

Perhaps the future title is not transcriptionist at all...perhaps it is AI Language Specialist.

And those who have spent years mastering listening, context, and accuracy are already preparing the data that powers the next generation of intelligent systems.

[#ArtificialIntelligence](#)[#SpeechToText](#)[#Transcription](#)

[#DataAnnotation](#)[#AIQuality](#)

5

Like

Comment

Share

To view or add a comment, [sign in](#)**FileMarket Labs Inc.**

1,496 followers

2mo

Why Voice AI Still Struggles with Accents?

Voice AI sounds impressive. 🎤

Until it misunderstands you.

Accent bias in ASR systems is still very real.

Many models are trained heavily on:

- Standard American or British English
- Neutralized pronunciation
- Controlled recording environments

But the real world sounds different.

It sounds like:

- Regional dialects
- Code-switching mid-sentence
- Background noise
- Fast speech, slow speech, emotional speech
- Words shaped by culture, not textbooks

When someone switches between languages in one sentence,
or speaks English with a Nepali, Indian, Nigerian, or Filipino accent...

Accuracy drops.

Not because the speaker is unclear.

But because the dataset wasn't diverse enough.

Voice AI doesn't struggle with intelligence.
It struggles with representation.
And emotional tone? That's another layer entirely.
Detecting sarcasm, urgency, frustration, hesitation that requires more than phonetics.
It requires nuance. 🌐

Have you ever been misunderstood by a voice assistant?

8 · 1 Comment

Like

Comment

Share

To view or add a comment, [sign in](#)

The Brighter Side of News

541 followers

2mo

Study finds the brain detects subtle differences between AI and human speech even when people cannot consciously tell them apart.

Is AI voice technology becoming indistinguishable from human speech?

thebrighterside.news

1

Like

Comment

Share

To view or add a comment, [sign in](#)

Florea.AI

147 followers

2mo

🌸 A new PNAS paper introduces a concept worth sitting with: "The benefits and dangers of anthropomorphic conversational agents."

Researchers from the University of Sydney and University of Washington argue that the latest LLMs — GPT-4, Claude, Gemini — haven't just gotten smarter. They've become convincing mimics of human communication, often outperforming average humans in persuasive writing, empathetic responses, inferring user emotions and intent, and role-play and conversation matching.

Studies cited show GPT-4 "significantly outperforms human participants" in persuasiveness across demographics. LLM-generated text now consistently passes the Turing test.

The paper's central tension:

These systems produce human-like communication without genuine empathy, understanding, or moral inhibition, making them powerful tools for both benefit and manipulation.

The authors call this risk "anthropomorphic seduction" — users trusting AI as they would a human, without realizing the difference.

💡 Building AI that communicates like a human is no longer the hard part. Building AI that does so responsibly with transparency about what it is and isn't that's the work. This paper gives language to something we think about every day.

<https://lnkd.in/g-x4ggig>

[#ArtificialIntelligence](#) [#AI](#) [#LLM](#) [#AIEthics](#) [#ResponsibleAI](#) [#Tech](#) [#FutureOfAI](#)

The benefits and dangers of anthropomorphic conversational agents | PNAS

pnas.org

2

Like

Comment

Share

To view or add a comment, [sign in](#)

Piyush Chakerwarty

2mo

🔥 The real workflow shift is happening somewhere else. 🔥

While exploring [Sarvam](#), I realized how powerful this combination is:

Speech

Vision

Translation

Text generation

Individually, they are useful.

But when they work together in a workflow, something interesting happens.

Systems can suddenly feel smooth, it:

- Understand images
- Listen to spoken questions
- Translate languages instantly
- Respond back in natural voice

That's the real idea behind multimodal AI.

Not just smarter models.

But smarter workflows.

This carousel explains how the entire flow works together.

Over the coming days, I'll also share:

- Practical product ideas
- MVPs developers can build today
- Why this matters especially for India

If you're building in AI, this stack is worth watching.

[#AI](#) [#SarvamAI](#) [#MultimodalAI](#) [#AIIndia](#) [#AIProducts](#) [#GenAI](#) [#AIStartups](#)

8 · 4 Comments

Like

Comment

Share

To view or add a comment, [sign in](#)

38,584 followers

[364 Posts](#) · [2 Articles](#)

View Profile

Follow

More from this author

Product of the Day on Product Hunt: Epic Guide on How to Launch
Arto Minasyan · 6y

We Made \$214,500 and Set a New Record on Appsumo Deals

Arto Minasyan · 6y

Explore related topics

Voice AI Technology Trends

Voice AI Industry Expansion

Effects of AI Voice Technology

Speech Recognition Innovations

The Importance of Voice in AI Development

AI-Driven Voice Interaction Models

Explore content categories

- Career
- Productivity
- Finance
- Soft Skills & Emotional Intelligence
- Project Management
- Education
- Show more

© 2026

About

Accessibility

User Agreement

Privacy Policy

Your California Privacy Choices

Cookie Policy

Copyright Policy

Brand Policy

Guest Controls

Community Guidelines

Language