

Private, real-time AI for high-stakes conversations

An on-device copilot that listens, sees your screen, remembers your documents and gives structured answers while the conversation is still happening. Live today on macOS.

Listen

- Dual-channel speech: microphone and meeting audio
- Noise filtering
- Live question detection

See

- Screen capture and visual OCR
- Problem statements, diagrams, documents

Remember

- Uploaded documents
- Role and presentation context
- Conversation history

Assist

- Structured answers
- Code and complexity analysis
- Talking points and follow-ups

Nothing leaves the device. Floating, always-on-top overlay. Seed round: \$3.4M. pairhint.com

THE PROBLEM

Live conversations move too fast for traditional AI chat

Questions never stop. By the time you have typed a prompt, the moment has passed.

TOO SLOW

Switch, type, wait

You cannot pause a live conversation to prompt a chatbot and wait for its reply.

TOO SCATTERED

The answer is everywhere

It is in your documents, on the shared screen, or in what was said earlier. Chat sees none of it.

TOO RISKY

Private data goes out

Sending private audio and screens to third parties is a risk in sensitive conversations.

The result: doctors, salespeople, lawyers, engineers and other professionals face their hardest live moments without help, or trust a cloud tool they should not.

PairHint: real-time, context-aware help while the conversation is happening, 100% on-device.

One continuous loop: observe, understand, assist

1

Observe

- Mic and meeting audio on two channels
- Screen captured and read with OCR
- Everything is captured locally

2

Understand

- Detects the question in live speech
- Pulls the right passage from your documents
- Uses session history and role context

3

Assist

- Streams a structured answer mid-conversation
- Key points, code, complexity, follow-ups
- Talking points ready to say

Built for privacy and speed

- Speech recognition and language-model inference run on the user's own CPU and GPU
- Floating, always-on-top overlay: no tab switching
- Nothing leaves the device

PROFESSIONAL USE CASES

Built for the moments that can't wait

Questions arrive live with no time to switch tools, type a prompt, and wait.

1 Technical Interviews

Get structured answers, code solutions, and complexity analysis while you are still talking.

2 Customer Demos

Keep product details, technical specs, and your own context documents one glance away.

3 System-Design Discussions

Read diagrams and problem statements from your screen with live OCR, then get instant talking points.

4 Sales & Support Calls

Follow the conversation on two audio channels (mic & meeting) and keep follow-up answers ready to say.

5 Legal & Law Consultations

Lawyers: analyze contracts, statutes, case law, and client testimony with 100% local, privilege-safe data privacy.

6 Medical & Healthcare

Doctors: instant clinical reference, patient history insights, and medical terminology support with zero cloud data transmission risk.

7 Executive & Board Briefings

Review financial spreadsheets, slide decks, and strategic metrics on the fly with live contextual AI guidance.


8 Management Consultancy & Strategy

Formulate strategic frameworks, synthesize market intelligence, and deliver recommendations during live advisory sessions.

9 Client Support & Account Management

Track client account history, service SLA terms, and resolve complex enterprise inquiries in real time.

From setup to live answer in one desktop workflow



01 Context and control

Upload a PDF as document context, add role and session notes, pick the speech and AI engines, and hide the window from screen share with Ghost Mode.

02 Live speech + AI

Two-channel transcription, question detection, then a one-line answer, key points and an extended explanation. Switch engines or freeze an answer.

03 History and export

Q&A mode or the full timeline, labelled Speaker / Host or You. Export to PDF to debrief. History stays on your device.

Live assistance using your conversation, screen and documents

1

Context

- Uploaded documents as persistent context
- Role and presentation notes
- Session history: the conversation so far
- Pulls the right passage for each question

2

Listen

- Microphone and meeting audio on two channels
- Noise filtering and live waveforms
- Live question detection
- AUTO multi-engine speech-to-text

3

See

- Screen capture on demand
- Visual OCR
- Problem statements, diagrams, documents
- Ghost Mode hides the window from screen share

4

Assist

- One-line answer first
- Key points and extended explanation
- Code, algorithms and complexity
- Talking points and follow-ups ready to say

Then: history and export

Q&A items, the full dual-channel transcript and PDF export, all stored on your device.

Our core technology brings together live audio, screen content and documents to deliver relevant answers.

Local AI has just become good enough to feel live

Models

- Small and mid-size language models now give strong answers on consumer hardware
- Open-weight models run locally, with no per-token cloud bill

Hardware

- Apple silicon and consumer GPUs make on-device inference practical in real time
- PairHint runs on macOS 13+, Apple silicon and Intel

Privacy

- Buyers and employers increasingly reject sending raw audio and screens to the cloud
- No cloud round trip: lower latency, simpler data-protection reviews

Honest trade-off and plan

- Local inference needs a capable machine: 16 GB RAM minimum; 24 GB+ recommended.
- Privacy-first by default; optional hybrid mode planned for lower-spec and enterprise needs.

NATIVE CROSS-PLATFORM DESKTOP ARCHITECTURE

PairHint is designed for multiple desktop platforms, but is currently released only on macOS because of limited development resources.

One AI assistant. Multiple industries.

Designed for live conversations across professional and everyday situations.

TODAY

Multi-domain product

- Real-time conversation assistance
- Sales, support and meetings
- Interviews and presentations
- Audio, screen and document context

NEXT

Broader adoption

- B2C and professional subscriptions
- Team and business deployments
- Legal, healthcare and education
- Industry-adapted workflows

LATER


AI platform

- Enterprise integrations
- Embedded and white-label AI
- Licensing and partnerships
- Potential GENTA-powered products

One core technology, many use cases. PairHint is not limited to developers or any single profession.

A multi-billion-dollar opportunity across live-conversation roles

ANNUAL MARKET VALUE IN USD (JOBS x \$240 PER USER PER YEAR)



Circles not to scale

TAM | TOTAL ADDRESSABLE MARKET

\$12.5B

52M jobs in six live-conversation roles, four regions

SAM | SERVICEABLE ADDRESSABLE MARKET

\$6.0B


~25M jobs in sales, consulting, law and technical support

SOM | SERVICEABLE OBTAINABLE MARKET

\$72M

Year-6 revenue target (400K paying users at year end): 1.2% of SAM, 0.6% of TAM

TAM BY REGION (USD BILLIONS, ESTIMATED)



Total ~\$12.5B = ~52M jobs x \$240

How this market is estimated

- Value = jobs x \$240 per year (blended \$20/month, the same price assumption as the six-year projection). It is a theoretical ceiling, not a forecast.
- Jobs: USA from BLS counts for six roles (2024-25). Europe, Asia and Middle East are PairHint estimates: US counts x 1.2, 3.5 and 0.2 (assumed white-collar workforce ratios).
- SOM = Year-6 gross revenue from the projection (about 300K average paying users). PairHint is macOS-only today; to be refined with ILOSTAT data; legal uses need validation.

Sources: BLS Occupational Outlook Handbook 2024 and Employment Projections 2025 (US). Regional figures are PairHint estimates, not reported statistics; roles map to BLS sales representatives (wholesale and manufacturing), management analysts, lawyers, computer user support specialists, customer service representatives and software developers.

Private help, in the moment, with your context

	Chat AI	Meeting assistants	PairHint
Live audio awareness	User-provided	Meeting-focused	Dual-channel capture
Screen context	Manual upload	Varies	Screen analysis + OCR
Personal documents	Files / prompts	Limited	Persistent session context
Technical output	General	General	Code + explanations
Privacy model	Cloud	Cloud	On-device
Interface	Separate window	Meeting UI	Floating overlay

The edge is integration: live audio, screen and private documents in one on-device assistant that answers during the conversation.

Subscriptions today. Teams and enterprise next.

1

Consumer (today)

- Pro subscription for individual professionals
- One-time session-hour top-ups
- Self-serve: download, sign in, start a session
- Prices: pairhint.com/pricing

2

Teams (next)

- Seat-based plans
- Shared workflows and context
- Security and administration
- Pilots with companies that train, hire or support

3

Platform (later)

- Enterprise deployments
- Domain copilots: medical, legal, sales, field support
- Licensing / SDK
- Embedded and white-label real-time AI

Why the model works


- Software-first and light on services; local inference keeps cloud compute costs low, so each new user adds little marginal cost
- Recurring subscription revenue plus usage-based top-ups. Planning assumption: 85% net revenue after fees and delivery costs

Year 6 target: 400,000 paying users and \$72M revenue

Key assumptions

- Multi-domain market: an estimated 52M jobs across 6 live-conversation roles in the USA, Europe, Asia and the Middle East, including sales, consulting, law and technical support
- PairHint is macOS-only today: the reachable base is a fraction of that
- Paying users at year end: 1K, 5K, 20K, 75K, 200K, 400K
- Blended \$20/month per paying user
- Net revenue = 85% of gross, after fees, refunds, hosting, support
- Launch within 5 months of funding: year 1 is partial
- Teams, enterprise and platform revenue not included

GROSS AND NET REVENUE (\$M)



6-year cumulative: \$120.2M gross revenue, \$102.2M net revenue.

TEAM

Built end to end by an experienced founder



Nassim MOUALEK

Founder · Senior Software Developer

20 years of experience

Sole builder of PairHint and

GENTA

Also the developer of

naviptv.com · castmypitch.com

EXPERIENCE

20 years of experience at Uberall, SFR, ENGIE and more.

EXPERTISE

AI, Big Data, Cloud, and native cross-platform applications.

PRODUCT

Designed and built PairHint end to end: desktop app, floating overlay, session and document management.

AI PIPELINE

Dual-channel speech-to-text, screen OCR, context memory and on-device LLM inference, tuned for latency and privacy.

RESEARCH

Leads GENTA research: one model designed to process images, speech and text, remember information and create new content.

The longer-term bet: GENTA

GENTA is our longer-term AI research project. It brings together images, speech and text, with multiple levels of memory and tools for reasoning and content creation.

Learns how the world changes. JEPA predicts what comes next from images, speech and text.

Connects what it sees, hears and reads. Information is combined so the model can use it together.

Supports multiple memory levels. Working and short-term context, plus long-term and very long-term memory.

Checks calculations. Symbolic reasoning uses explicit mathematical operations to improve reliability.

Built for broader use. One architecture for many tasks, with parallel training across multiple processors.

Learns in parallel. Multiple synchronized instances train together and share learning updates, designed to accelerate learning.

JEPA

world model

Joint latent

one shared space

**Tiered
memory**

extensible


System-2

symbolic reasoning

PairHint is the near-term business. GENTA is the longer-term research opportunity; further investment would support training and validation.

GENTA architecture: one model, many abilities

One model processes images, speech and text, supports multiple levels of memory, and trains on 65 initial tasks, extensible to new ones. Its generation capabilities span text, audio, images, 3D objects and video.



GENTA: sees, hears, reads, remembers, thinks and creates

An early research model designed to understand images, speech and text, remember information, solve problems and create new content. Shown: image-to-3D generation.

Step	Task	Loss	IQ	Mastery / repeats	Gradient	Step time (s)	Worker RAM
52209/91000	GEN3D	0.0001	71.7	MTL 59.5% R+1146 D+7	8.98	0.54	0.25 GB · system 89%

TRAINING SKILLS

☐ 08 init_vid_aud
☐ 09 init_vid_aud_txt
☒ ► 10 gen3d
☐ 11 gen3d_abo
☐ 51 vocab
☐ 52 math_basic
☐ 53 stt_math
☐ 54 ocr_math
☐ 55 stt_ocr_math
☒ Full cycle
☐ Follow running task

Preview & training details

All metrics & task IQ



Input

Please [time_frame=0] Generate 3D PBR mesh for: A real 3D CAD model of

Document

please [time_frame=0] generate 3d pbr mesh for: a real 3d cad model of a chair x mit csail pix3d dataset. reproduce the earlier segment, if you can.

Generated output

generate 3d pbr mesh for: a real 3d cad model of a chair x mit csail pix3d dataset. lthbox: 0.41m x 0.82m x 0.44m

Expected target

Generate 3D PBR mesh for: A real 3D CAD model of a chair from MIT CSAIL Pix3D dataset. lthbox: 0.41m x 0.82m x



Volume 100%

Processes

Images, speech and text

Remembers

Working, short-term, long-term and very long-term

Creates

Text, audio, images, 3D and video

Learns

65 initial tasks, extensible

Trains in parallel

Multiple processors working together

Supports multiple levels of memory: working, short-term, long-term and very long-term. Trains on 65 initial tasks, extensible to new ones.

PairHint is the first application. GENTA can go further.

PairHint delivers real-time assistance. GENTA aims to bring the AI abilities behind it together in one compact model, with uses well beyond a single product.

PairHint today

Several AI models work together to listen, see, remember context and answer in real time.

Coordinating them adds integration work, processing overhead and delays.



GENTA: the development goal

One compact model for perception, memory, reasoning and response generation.

The aim: faster responses, lower running costs and easier deployment on local devices.

A foundation for more than PairHint

PairHint provides a concrete first use case. The same GENTA architecture could support education, professional assistants and creative tools that generate text, speech, images, 3D and video.

PairHint creates the product opportunity. GENTA could become a reusable AI foundation.

A \$3.4M round to build a repeatable business

40%

Product and AI · \$1.36M

- Latency and reliability
- Inference optimization
- Hardware coverage
- Quality measurement

35%

Go-to-market · \$1.19M

- Customer acquisition
- Target communities
- B2B pilots
- Conversion experiments

25%

Trust and scale · \$0.85M

- Security architecture
- Enterprise readiness
- Analytics
- Platform preparation

What this round unlocks

- Go-to-market in under 5 months once the round is funded
- 1,000 paying users by the end of year 1
- A path to 400,000 paying users by year 6 (see projection)