

INTERVIEWWHISPERER

User Guide

Every hotkey. Every workflow. One reference.

A real-time AI assistant for the moments that matter most —
built to make your performance sharper, not stealthier.

What is InterviewWhisperer?

InterviewWhisperer is a real-time AI assistant that runs as an on-screen overlay during your interviews and high-pressure conversations. It listens to the question you're being asked, looks at what's on your screen when you need it, and streams a personalized answer back in under a second — drawn from the documents you uploaded, the conversation so far, and the question itself.

It's designed for one thing: **helping you perform at your best in moments where your mind might blank, your stack of facts is too deep to recall instantly, or the question is more technical than your prep covered.** You stay the candidate. The tool just makes sure nothing you already know slips away under pressure.

REAL-TIME

Sub-second answers based on the question being asked, the documents you've uploaded, and the full session history.

CONTEXT-AWARE

Reads your resume, the job description, technical papers — anything you upload at the start. Every answer is personalized.



JUST FOR YOUR EYES

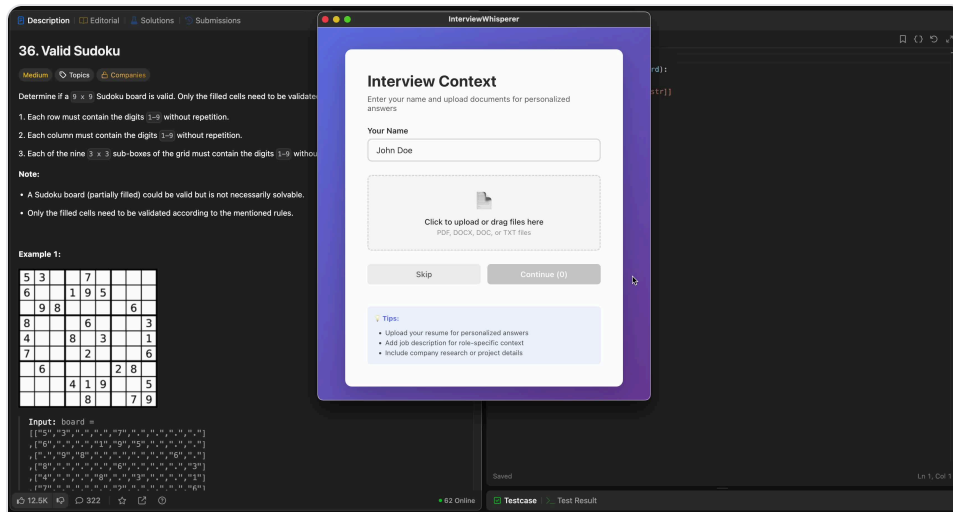
The panel can hide from screen-share tools so the answer stays visible to you while staying out of the recording.

What this guide covers: every keyboard shortcut, when to use each one, and the most common workflows. The cheat sheet on page 4 is designed to be printed and kept beside your monitor.



Getting Started

Three steps from launch to your first answer. Most users are set up in under two minutes.

- 1 Launch InterviewWhisperer**
Open the app. You'll land on the **Interview Context** window. Enter your name — the AI uses it when introducing you in answers.
- 2 Upload context documents**
Drag in your resume, the job description, any prep notes, even research papers. PDF, DOCX, DOC, or TXT. The more context you provide, the more personalized every answer becomes.
- 3 Hit Continue and start your session**
The overlay appears and starts listening. When the interviewer asks something, press   and the answer streams in within a second.
















The Interview Context window. Enter your name, drop in documents, hit Continue.

If you only remember one hotkey:   — this is your fast-answer key, the one you'll use dozens of times per session.







Every Hotkey

 — print and keep beside your monitor















CORE ACTIONS

  ;	Submit the question — answer streams in ~1s	  W	Screenshot answer — coding, diagrams, anything visual
   A	Append clipboard contents to context buffer	  	Submit the full clipboard buffer as one question
   X	Clear the clipboard buffer to start fresh		

VISIBILITY & PRIVACY

  H	Toggle the response panel visibility	 	Hide or show the entire window
  I	Private overlay mode — stays out of screen share		










DISPLAY

   M	Cycle answer length: brief, normal, detailed	   T	Toggle light and dark theme
 ]	Increase response font size	  [Decrease response font size
 	Scroll response up	 	Scroll response down

WINDOW POSITION

   Arrows	Nudge the window 10 pixels at a time	  R	Reset window to default position
---	--------------------------------------	---	----------------------------------

SESSION

   E	Export full transcript of the session	   C	Clear all responses on screen
   Q	Quit the app		

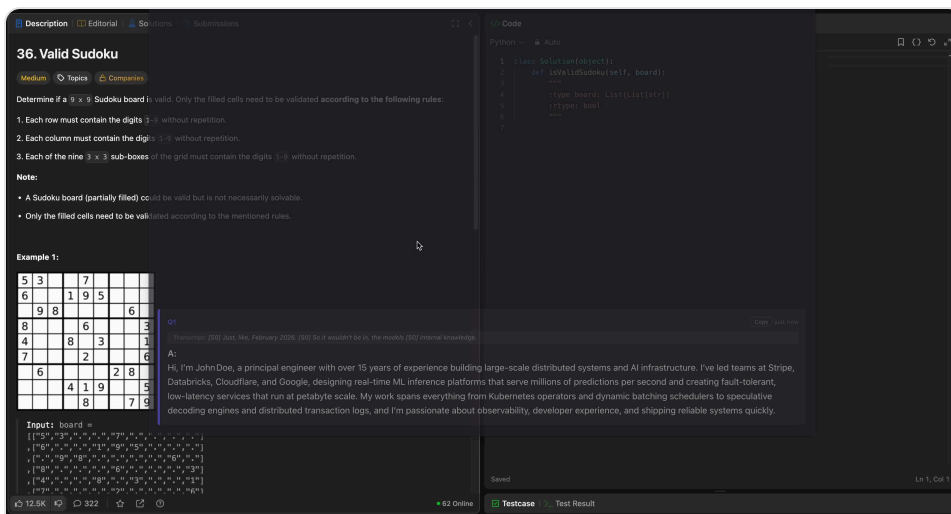
Audio questions: the ⌘ ; flow

For everything the interviewer says out loud — behavioral, system design verbal, motivation questions, follow-ups — the workflow is the same.

1 · Listen, then submit



The app is always listening once you start the session. When the interviewer finishes their question, press ⌘ ; and a personalized answer streams in within about a second. The answer draws on your uploaded documents, the questions asked earlier in this session, and everything that's been transcribed so far.



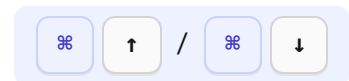
After ⌘ ; — the answer streams in personalized from your resume.

2 · Adjust on the fly



If the interviewer wants more depth — "walk me through the impact" — press ⌘ ↑ M to cycle to detailed mode. Press again for brief mode when you just need a one-line pivot. The mode persists for the next answer.

3 · Review past answers



Newer answers appear at the top of the panel — that's intentional, because your camera is up there too, and reading the latest answer keeps your eyes naturally near the lens. To scroll back through older answers in the session, use ⌘ ↑ and ⌘ ↓.

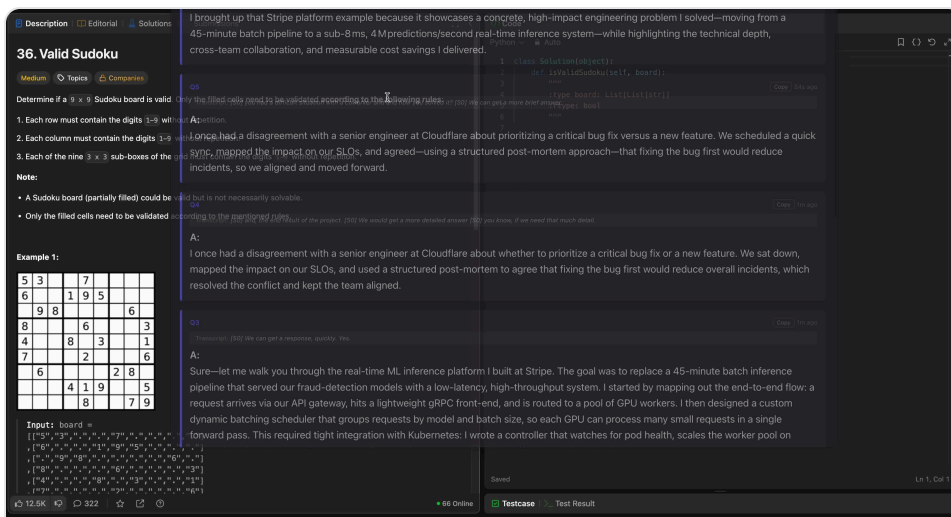
Visual problems: the ⌘ W flow

For anything on your screen the AI needs to see — coding problems, system design diagrams, math, architecture prompts — use the screenshot workflow.

1 · First press: queue a screenshot

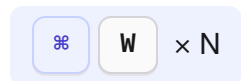


The first time you press ⌘ W, the app captures what's on screen and adds it to the screenshot queue. You'll see "1 screenshot queued" in the panel — but nothing is sent yet. This lets you stage multiple screenshots before sending.



A queued screenshot shows in the panel, ready to send.

2 · For long problems: queue more



If the problem spans multiple screens — like a long self-attention prompt with separate windows for the question, the constraints, and the starter code — press ⌘ W each time to keep queuing screenshots. The queue counter ticks up. The vision model will see all of them together when you send.

3 · Final press: send everything

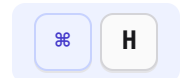


One more press of ⌘ W sends all queued screenshots to the vision model in a single call. The full solution streams back — code, explanation, time complexity, edge cases. Follow-up questions from the interviewer? Use ⌘ ; again — the model has the full visual context now.

Three ways to hide

InterviewWhisperer has three different visibility controls, and each one is for a different moment.

Toggle the response panel



The standard "hide the answers" shortcut. The panel disappears; the rest of the app keeps running. Press again to bring it back. Use this when you want to clear your view briefly to focus on the interviewer.

Hide the entire window



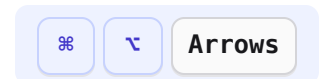
A more complete hide — the whole InterviewWhisperer window vanishes. Useful when you need a totally clean screen briefly. Press again to restore.

Private overlay mode



Toggles the overlay's private mode so it stays out of screen-share and screen-recording tools like Zoom, Meet, and Teams. The panel remains visible to you on your own screen. This is the mode you'll want active during a live interview.

Move the panel anywhere



Hold ⌘ \ and tap any arrow key to nudge the panel 10 pixels at a time. Drag it to wherever it fits your setup. If it ever ends up off-screen or somewhere awkward, ⌘ R snaps it back to default.

Getting the most out of it

Practice with mock interviews first. The muscle memory matters. Run two or three full mock interviews before a real one so the hotkeys feel automatic, not deliberate.

Upload everything you'd review the night before. Resume, JD, prep notes, system design study sheets, technical papers — the model uses them all. The more thorough your context upload, the more personalized your answers.

Treat ⌘ Shift M as a tone dial. Brief for quick pivots ("can you summarize your experience?"). Normal for typical questions. Detailed when they want the full breakdown — projects, end-to-end design, deep behavioral.

Use ⌘ I before joining the call. Enable private overlay mode *before* you connect to Zoom/Meet/Teams. That way the panel is configured correctly from the start.

The panel is positioned for eye contact. New answers appear at the top, near your camera. Reading the latest answer keeps your gaze where you want it. Don't move the panel down to a corner where you'd have to look away to read it.

Keep ⌘ R in mind. If you ever lose the panel — off-screen, behind another window — ⌘ R resets its position to default.

Export the session after every real interview. ⌘ Shift E saves the full transcript — every question and answer. Use it as context for the follow-up interview, or to review what you said.

Use ⌘ Shift A for multi-part prompts. If the interviewer pastes a long problem in chat, copy each part, then ⌘ Shift A appends it to the buffer. When you've got all the pieces, ⌘ Shift ↵ submits the combined buffer as one question.

Troubleshooting

Quick fixes for the most common issues.

Audio isn't being transcribed.

Check microphone permissions in *System Settings* → *Privacy & Security* → *Microphone* and confirm InterviewWhisperer is allowed. Restart the app after granting access.

The response panel isn't appearing.

Press ⌘ H in case the panel is toggled hidden. If it's still missing, press ⌘ R to reset position — it may have moved off-screen.

Hotkey isn't responding.

Make sure InterviewWhisperer is running (check the menu bar). Some hotkeys can conflict with macOS or other apps — try the alternate where one exists (e.g., ⌘ ` instead of ⌘ H).

Answer feels off-topic.

The transcription may have missed part of the question. Listen for the full question, then press ⌘ ; right after they finish speaking — that's the cleanest moment.

Screenshot answer is missing context.

If the problem spans multiple screens or panels, queue more screenshots with additional ⌘ W presses before sending. One final press sends everything together.

App won't quit normally.

Use ⌘ Shift Q to fully quit InterviewWhisperer. Closing the window only hides it.

Still stuck? Email support@interview-whisperer.app. Include what you were trying to do, what happened instead, and your macOS version. We respond within one business day.

One last thing

InterviewWhisperer was built to be there when your mind blanks, when the question is more specific than your prep, when the moment matters and you need to be sharp. You're still the one in the chair. The tool just makes sure nothing you already know slips away under pressure.

Practice with it. Run mocks. Get the hotkeys into your fingers. Then when the real interview comes, the muscle memory takes over and the tool fades into the background — exactly where it belongs.

Good luck out there.