

# MDwAIstScheduler: A Low-Cost, Voice-Activated Device for Hands-Free Clinical Scheduling

Diego Mardian  
dmardia2@asu.edu  
Arizona State University  
Tempe, Arizona, USA

Frank Liu  
fwliu1@asu.edu  
Arizona State University  
Tempe, Arizona, USA

## Abstract

Physicians spend nearly half their workday on EHR tasks and administrative work, contributing to burnout and reducing time for direct patient care. We present MDwAIstScheduler, a low-cost, belt-worn voice assistant that allows hands-free calendar management during patient encounters. Hidden beneath a lab coat, the device avoids the eye-contact disruptions caused by visible screens or wrist-worn devices. Running on a Raspberry Pi with cloud-based speech recognition and LLM intent extraction, the system lets clinicians simply say “Schedule a follow-up with Mr. Smith next Tuesday at 2” and automatically creates the calendar event. Our demo showcases this end-to-end pipeline.

## 1 Introduction

Healthcare providers face an administrative burden that directly impacts their well-being and quality of patient care. Physicians spend only 27% of their workday in direct clinical time with patients, while 49.2% is consumed by EHR and desk work [9]. This overload is a significant driver of clinician burnout, which affects over 43% of physicians in the United States [8].

Visible technology also makes the problem worse. Studies show that EHR use during patient encounters reduces physicians’ ability to maintain eye contact and actively listen [1], and eye-gaze research confirms that physician attention to screens directly influences patient engagement [2]. Wrist-worn devices and desk-mounted systems suffer from the same issue of pulling the physician’s gaze away from the patient.

We present **MDwAIstScheduler**, a belt-worn voice assistant designed to be *invisible* during patient encounters. Hidden beneath a lab coat, the device enables hands-free scheduling without any visual distraction. A physician simply speaks a command “Schedule a meeting with Dr. Patel next Friday” and the system transcribes, interprets, and creates the calendar event automatically. Unlike commercial ambient AI scribes that focus on documentation and require expensive infrastructure [4, 10], MDwAIstScheduler targets a focused workflow (scheduling) using low-cost hardware, making it practical for smaller practices.

## 2 Belt-Worn Wearable Design

Research on wearable health devices identifies torso-worn devices, including belts, as an emerging category that enables unobtrusive operation without disrupting daily activities [6, 7]. For clinician-facing applications, this unobtrusiveness is critical.

A belt-mounted device hidden beneath a lab coat offers key advantages over alternatives:

- **No visual distraction:** The physician never looks at a screen or wrist; the patient sees no device.

- **Preserved eye contact:** Unlike EHR screens that impair communication [1], the device is invisible.
- **Hands-free operation:** Voice input requires no manual interaction, keeping hands available for patient care.
- **Natural conversation flow:** Scheduling commands are in normal speech without interrupting the encounter.



**Figure 1: The MDwAIstScheduler system, designed to be worn on the belt and hidden beneath a lab coat during patient encounters.**

## 3 System Architecture

MDwAIstScheduler runs on a Raspberry Pi 4 (\$35–75) with a USB microphone (\$2–3), keeping total hardware cost under \$100. The system shown in Figure 1 consists of three components:

(1) **Continuous Voice Capture.** The device monitors ambient audio and triggers processing when a scheduling command is detected.

(2) **Speech-to-Text.** Audio is sent to Google Speech Recognition API for transcription. Offloading to the cloud maintains high accuracy on limited hardware, consistent with hybrid edge-cloud architectures [3, 11].

(3) **LLM Intent Extraction.** Transcribed text goes to Mixtral-8x7B (via Groq API) with a prompt that extracts scheduling parameters (date, time, duration, attendee, description) as JSON. This replaces brittle rule-based parsing with flexible natural language understanding that handles varied phrasing and implicit references like “next Monday” [5]. The JSON is then used to create a Google Calendar event via its API.

## 4 User Interaction

Users can (1) speak natural-language scheduling commands to the device, (2) observe real-time transcription and intent extraction displayed on a monitor, and (3) see the resulting calendar event created in Google Calendar. The intermediate JSON output from the LLM to illustrate how natural language is mapped to these structured scheduling parameters is observable as an option. The system can also handle phrasing variations for scheduling.

## 5 Discussion and Future Work

MDwAIsScheduler demonstrates that targeted, low-cost, unobtrusive voice wearables can reduce administrative burden while preserving patient interaction quality. The belt-worn form factor represents a new approach. By remaining hidden, it eliminates the visual distractions of screen-based or wrist-worn alternatives [1].

Future work includes: (1) evaluating the system with physicians in family medicine practice to measure impact on time allocation and perceived burden, (2) integrating with EHR scheduling systems beyond Google Calendar, (3) exploring on-device LLM inference for improved privacy [11], and (4) adding multi-turn dialogue for disambiguation of ambiguous commands.

## References

- [1] Maria A Alkureishi, Wei Wei Lee, Michael Lyber, et al. 2016. Impact of Electronic Medical Record Use on the Patient-Doctor Relationship and Communication: A Systematic Review. *Journal of General Internal Medicine* 31, 5 (2016), 548–560. doi:10.1007/s11606-015-3582-1
- [2] Onur Asan and Enid Montague. 2014. Dynamic Modeling of Patient and Physician Eye Gaze to Understand the Effects of Electronic Health Records on Doctor-Patient Communication and Attention. *International Journal of Medical Informatics* 83, 3 (2014), 225–234. doi:10.1016/j.ijmedinf.2013.11.003
- [3] Alberto Cuevas et al. 2025. Real-Time Speech-to-Text on Edge: A Prototype System for Ultra-Low Latency Communication with AI-Powered NLP. *Information* 16, 8 (2025), 685. doi:10.3390/info16080685
- [4] Sameer Ghatnekar et al. 2025. Real-World Evidence Synthesis of Digital Scribes Using Ambient Listening and Generative Artificial Intelligence for Clinician Documentation Workflows: Rapid Review. *JMIR AI* 4 (2025), e76743. doi:10.2196/76743
- [5] Trent Hodgson and Enrico Coiera. 2018. Electronic Health Record Interactions through Voice: A Review. *Applied Clinical Informatics* 9, 3 (2018), 541–559. doi:10.1055/s-0038-1666844
- [6] Lu Lu, Jue Zhang, Yingchao Xie, et al. 2020. Wearable Health Devices in Health Care: Narrative Systematic Review. *JMIR mHealth and uHealth* 8, 11 (2020), e18907. doi:10.2196/18907
- [7] Yvonne O'Connor et al. 2021. Review of Wearable Devices and Data Collection Considerations for Connected Health. *Sensors* 21, 16 (2021), 5589. doi:10.3390/s21165589
- [8] Sachin J Shah, Christine A Sinsky, Karleen F Giannitrapani, Tait D Shanafelt, et al. 2025. Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout. *JAMA Network Open* 8, 2 (2025), e2460637. doi:10.1001/jamanetworkopen.2024.60637
- [9] Christine Sinsky, Lacey Colligan, Ling Li, Mirela Prgomet, Sam Reynolds, Lindsey Goeders, Johanna Westbrook, Michael Tutty, and George Blike. 2016. Allocation of Physician Time in Ambulatory Practice: A Time and Motion Study in 4 Specialties. *Annals of Internal Medicine* 165, 11 (2016), 753–760. doi:10.7326/M16-0961
- [10] Aaron A Tierney, Gonzalo Gavidia, and Christine Sinsky. 2024. Ambient Artificial Intelligence Scribes to Alleviate the Burden of Clinical Documentation. *NEJM Catalyst* 5, 3 (2024), CAT.23.0404. doi:10.1056/CAT.23.0404
- [11] Yifan Zhou et al. 2025. Intelligent Data Analysis in Edge Computing with Large Language Models: Applications, Challenges, and Future Directions. *Frontiers in Computer Science* 7 (2025), 1538277. doi:10.3389/fcomp.2025.1538277