

Factually: Exploring Wearable Fact-Checking for Augmented Truth Discernment

CHITRALEKHA GUPTA, HANJUN WU, PRAVEEN SASIKUMAR, SHREYAS SRIDHAR, PRIAMBUDI BAGASKARA, and SURANGA NANAYAKKARA, Augmented Human Lab, National University of Singapore, Singapore




Fig. 1. (a) A person using Factually in an everyday setting (b) Factually integrated with a smartwatch, (c) A companion mobile app.

Wearable devices are transforming human capabilities by seamlessly augmenting cognitive functions. In this position paper, we propose a voice-based, interactive learning companion designed to amplify and extend cognitive abilities through informal learning. Our vision is threefold: (1) to enable users to discover new knowledge on-the-go through contextual interactive quizzes, fostering critical thinking and mindfulness; (2) to proactively detect misinformation, empowering users to critically assess information in real time; and (3) to provide spoken language correction and prompting hints for second language learning and effective communication. As an initial step toward this vision, we present Factually — a proactive, wearable fact-checking system integrated into devices like smartwatches or rings. Factually discreetly alerts users to potential falsehoods via vibrotactile feedback, helping them assess information critically. We demonstrate its utility through three illustrative scenarios¹, highlighting its potential to extend cognitive abilities for real-time misinformation detection. Early qualitative feedback suggests that Factually can enhance users' fact-checking capabilities, offering both practical and experiential benefits.

CCS Concepts: • **Human-centered computing** → **Haptic devices**; **Mobile devices**.

Additional Key Words and Phrases: Assistive Augmentation, Fact-Checking, Wearable Assistant

1 INTRODUCTION

Misinformation has become an unavoidable challenge in the digital age, with bite-sized falsehoods spreading rapidly across social media platforms and casual conversations. Fact-checking in real-time is inherently inconvenient, often socially awkward, and prone to being forgotten. Despite the rise of automated fact-checking systems, existing solutions typically require active engagement, such as searching for questionable claims, which disrupts natural conversations.

This paper was presented at the 2025 ACM Workshop on Human-AI Interaction for Augmented Reasoning (AIREASONING-2025-01). This is the authors' version for arXiv.

¹Concept Video Link

Authors' address: Chitralekha Gupta, chitralekha@nus.edu.sg; Hanjun Wu, michelle@ahlab.org; Praveen Sasikumar, praveen@ahlab.org; Shreyas Sridhar, shreyas@ahlab.org; Priambudi Bagaskara, bagas@ahlab.org; Suranga Nanayakkara, suranga@ahlab.org, Augmented Human Lab, National University of Singapore, Singapore.

To address this gap, we present *Factually*, a wearable live fact-checking system designed to augment an individual's ability to discern truth seamlessly.

Factually integrates with everyday wearable devices such as smartwatches or rings to provide real-time, discreet feedback through vibrations when potentially false information is detected. The aim of *Factually* aligns with the goals of assistive augmentation [12] where it enhances cognitive capabilities in naturalistic settings, blending immediacy and social integration into an intuitive system. Our work builds on the broader vision of using wearable technologies to amplify human abilities, with a specific focus on tackling misinformation.

In this paper, we describe the design of *Factually*, its implementation using existing fact-checking mechanisms, and its potential to transform how individuals engage with information. Through initial qualitative demonstrations, we highlight *Factually*'s promise as an effective tool for fostering truth-centered behaviors and enhancing human cognition in socially integrated ways.

2 RELATED WORK

Factually builds upon a growing body of research in misinformation detection, assistive augmentation [12], and wearable technologies, extending these fields by addressing the challenges of real-time fact-checking in social contexts.

Several key studies have advanced our understanding of misinformation, particularly the cognitive and social processes that underpin its spread and persistence. Lewandowsky et al. have made significant contributions to this field, exploring why individuals believe misinformation and the challenges of correcting it. For instance, [5] demonstrated that while explicit warnings can reduce the impact of misinformation, they often fail to eliminate its lasting effects. Similarly, Blank's work on "double misinformation" [1] shows how misinformation can confuse people and distort their understanding of events. Blank's earlier work [2] on the social construction of memory further explains how misinformation impacts the formation of false memories. These impactful contributions, alongside studies like [6] work on the "continued influence effect" and [4] analysis of misinformation dynamics online, emphasize the complex interaction of psychological and social factors in the spread of misinformation and provide valuable insights into strategies for mitigation.

2.1 Misinformation Detection and Cognitive Augmentation

Kozyreva et al. [7, 8] discuss interventions aimed at combating misinformation and emphasize the importance of timely cognitive support. Pennycook et al. [9] examine the role of cognitive reflection in discerning truth and demonstrate that nudges to encourage reflective thinking can mitigate the spread of misinformation. Through *factually*, we align with these principles by providing real-time fact-checking assistance, and employ subtle, tactile cues that prompt introspection and self-correction without disrupting conversations.

2.2 Wearables for Misinformation Detection

Wearable technologies have increasingly been employed as tools for augmenting human cognitive and sensory abilities. The Wearable Reasoner explored the use of explainable AI to enhance rational decision-making by providing verbal feedback through a wearable device [3]. While effective in delivering justifications for decisions, its focus on audio-based interactions raises practical challenges in socially sensitive scenarios. *Factually* addresses these limitations by relying on non-verbal tactile feedback, ensuring seamless integration into conversations. *Factually* builds on the foundation of *assistive augmentation* [12], which envisions wearable technologies as transformative tools for expanding human capabilities.




Fig. 2. Overview of the Technical Implementation.

2.3 Real-time Fact Checking

Fact-checking systems have traditionally focused on web-based or multimedia platforms, with limited attention to conversational contexts. For example, Rashkin et al. [10] propose a language model-based system to detect false claims in online text. While effective in structured environments, such solutions are less applicable to dynamic, interpersonal interactions. Factually bridges this gap by operating as a wearable assistant capable of on-the-fly fact verification during conversations. Similarly, Setty et al. [11] introduce Factiveverse, a live fact-checking tool for online applications. While Factiveverse targets media platforms, Factually focuses on augmenting individual users' cognitive abilities in real-world, conversational scenarios.

3 FACTUALLY

3.1 Design considerations in the Assistive Augmentation space

The design of the *Factually* is inspired by the principles of assistive augmentation, as outlined by Tan et al. [12]. They introduce assistive augmentation as a paradigm that enhances human capabilities along two primary dimensions: *ability* (expanding cognitive and sensory capacities) and *integration* (seamlessly embedding augmentation into everyday life). We design *Factually* so that it embodies these dimensions in the context of misinformation detection and truth discernment, in the following way:

Ability Dimension: *Factually* augments human cognitive capabilities by providing real-time feedback about potentially false information. It empowers users to discern truth more effectively during conversations, addressing the cognitive overload often associated with manual fact-checking. The system's vibrational feedback mechanism enables users to extend their perceptual limits, allowing for quick and intuitive recognition of misinformation. This aligns with the ability dimension of assistive augmentation, where tools amplify human cognition without replacing human judgment.

Integration Dimension: The integration dimension emphasizes the importance of embedding augmentation seamlessly into users' daily lives. *Factually* achieves this through its wearable design, which uses discreet tactile feedback instead of disruptive auditory or visual cues. Whether integrated into a smartwatch or a smart ring, *Factually* ensures that users remain engaged in social interactions without drawing attention to the system's operation. Its subtle design makes it particularly suitable for naturalistic settings, such as casual conversations or collaborative environments, where overt fact-checking would be impractical or socially awkward.

3.2 System Design

Our proof-of-concept system consists of the following components, as shown in Figure 2:

- **Wearable Interface:** Factually integrated into common wearable device (e.g. smartwatch). Vibrotactile feedback is used to indicate when a statement is flagged as potentially false, ensuring subtle and socially acceptable interactions.
- **Fact-Checking Backend:** The system leverages a combination of large language models (LLMs) and web-based resources to evaluate the truthfulness of statements in real time. While the current implementation uses state-of-the-art fact-checking tools as a proof of concept, this mechanism can be replaced with more sophisticated models as they become available.
- **Real-Time Processing:** Audio input is transcribed and analyzed for potentially false claims. If flagged, Factually sends a vibration alert to the wearable device. Users can then access additional context about the flagged statement through a connected mobile application, if desired.

3.3 Proof-of-concept Scenarios

To demonstrate the potential of *Factually*, we developed three proof-of-concept scenarios (Figure 3)².

3.3.1 Scenario 1: Health Misinformation. In this scenario, we assess the role of Factually in detecting health-related misinformation in a live conversation. This scenario has two granddaughters, Emma and Grace, who are taking out the rubbish while discussing their grandmother’s medication.

Emma asks Grace if she has given their grandmother Neurontin. Grace responds that she hasn’t and questions the purpose of the medication. Emma confidently states that it is used to lower blood pressure.

At this point, *Factually* vibrates, alerting the users to a potential error. The device provides a tactile cue, prompting the wearer to check the accompanying feedback: Neurontin is a medication used to help manage epileptic seizures. With this information, Grace points out the error, avoiding a potential health risk.

This scenario demonstrates *Factually*’s ability to extend perceptual capabilities by identifying incorrect statements and providing accurate information in real time. By doing so, it enhances user awareness and helps prevent potentially harmful misunderstandings.

3.3.2 Scenario 2: Social Conversations. In this scenario, we assess the role of factually in a casual social setting. Two friends are having lunch and discussing Taylor Swift. One claims that Taylor Swift can speak Chinese, while the other disagrees, leading to a playful bet: the loser will pay for lunch.

The first individual finds a YouTube video appearing to show Taylor Swift speaking Mandarin. However, *Factually* vibrates subtly, allowing the user to access the feedback: while Taylor Swift does not speak Mandarin, there is a highly realistic deepfake depicting her doing so. Armed with this information, the user refutes the claim, shifting the outcome of the bet.

This scenario illustrates how *Factually*’s wearable and tactile design supports its seamless integration into social interactions. By aligning with familiar gestures, such as glancing at a wrist or adjusting a wearable device, *Factually* remains inconspicuous while providing timely and actionable information.

²Concept Video Link




Fig. 3. Use-Case Scenarios of Factually. (A) Scenario 1: health-related misinformation detection during a conversation between two grand daughters. (B) Scenario 2: misinformation detection during a casual lunch conversation. (C) Scenario 3: misinformation detection during learning.

3.3.3 Scenario 3: Everyday Learning. In this scenario, we evaluate the role of *Factually* in learning. A parent is helping their child with homework. The child asks if dinosaurs lived alongside humans. The parent confidently responds affirmatively.

Factually vibrates, prompting the parent to consult the device’s feedback: dinosaurs went extinct 65 million years before humans appeared. Armed with this new information, the parent corrects their statement, encouraging a more accurate understanding of history.

This scenario highlights how *Factually* encourages users to question the validity of their claims and practice mindfulness in speech. By promoting introspection and self-correction, *Factually* fosters a truth-centered identity and critical thinking.

3.4 Initial User Reactions

We presented *Factually* to 10 potential users to get their initial qualitative feedback. Overall, participants found *Factually* intuitive and useful, particularly appreciating its non-intrusive nature. They expressed interest in using such a system in their daily lives, citing its potential to enhance critical thinking and mindfulness. However, they also suggested improvements in the vibrotactile feedback’s specificity and the ability to customize the system’s sensitivity to different topics.

4 LIMITATIONS AND FUTURE WORK

The current prototype of *Factually* employs general-purpose large language models and web-based resources, which may not always produce domain-specific or highly accurate results. Future implementations would incorporate specialized fact-checking models trained on misinformation datasets to enhance reliability.

Factually relies on cloud-based fact-checking using large language models and web resources, but network latency can delay real-time feedback. Future versions could address this by incorporating edge computing or on-device inference for faster response times.

While *Factually*’s discreet vibrotactile feedback minimizes social disruptions, broader and long-term usability studies are needed to assess its practicality in diverse social contexts to gauge its social acceptability as well as its impact on

Chitralkha Gupta, Hanjun Wu, Praveen Sasikumar, Shreyas Sridhar, Priambudi Bagaskara, and Suranga Nanayakkara

user’s cognitive load. Investigating the ethical implications of real-time fact-checking systems, such as privacy concerns and the potential for misuse, will be vital to ensure responsible deployment.

REFERENCES

- [1] Hartmut Blank, Anu Panday, Ross Edwards, Ewa Skopicz-Radkiewicz, Violet Gibson, and Vasudevi Reddy. 2022. Double misinformation: Effects on eyewitness remembering. *Journal of Applied Research in Memory and Cognition* 11, 1 (2022), 97.
- [2] Hartmut Blank, Eva Walther, and Simon D Isemann. 2016. The past is a social construction: susceptibility to social influence in (mis) remembering. In *False and distorted memories*. Psychology Press, 65–81.
- [3] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2020. Wearable Reasoner: towards enhanced human rationality through a wearable device with an explainable AI assistant. In *Proceedings of the Augmented Humans International Conference*. 1–12.
- [4] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the national academy of Sciences* 113, 3 (2016), 554–559.
- [5] Ullrich KH Ecker, Stephan Lewandowsky, and David TW Tang. 2010. Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & cognition* 38 (2010), 1087–1100.
- [6] Hollyn M Johnson and Colleen M Seifert. 1994. Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of experimental psychology: Learning, memory, and cognition* 20, 6 (1994), 1420.
- [7] Anastasia Kozyreva, Stephan Lewandowsky, and Ralph Hertwig. 2020. Citizens versus the internet: Confronting digital challenges with cognitive tools. *Psychological Science in the Public Interest* 21, 3 (2020), 103–156.
- [8] Anastasia Kozyreva, Philipp Lorenz-Spreen, Stefan M Herzog, Ullrich KH Ecker, Stephan Lewandowsky, Ralph Hertwig, Ayesha Ali, Joe Bak-Coleman, Sarit Barzilai, Melisa Basol, et al. 2024. Toolbox of individual-level interventions against online misinformation. *Nature Human Behaviour* (2024), 1–9.
- [9] Gordon Pennycook and David G Rand. 2021. The psychology of fake news. *Trends in cognitive sciences* 25, 5 (2021), 388–402.
- [10] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 conference on empirical methods in natural language processing*. 2931–2937.
- [11] Vinay Setty et al. 2024. LiveFC: A System for Live Fact-Checking of Audio Streams. *arXiv preprint arXiv:2408.07448* (2024).
- [12] Felicia Fang-Yi Tan, Chitralkha Gupta, Dixon Prem Daniel Rajendran, Pattie Maes, and Suranga Nanayakkara. 2025. Assistive Augmentation: Fundamentally Transforming Human Ability. *Interactions* 32, 1 (2025), 22–27.