Generative AI as Communication Prosthesis: Design Tensions and Principles for Adaptive Support in Apraxia of Speech

Omotayo Emmanuel Omoyemi
University of Derby, United Kingdom
etayo25@gmail.com

Abstract: Apraxia of Speech (AoS) is a motor speech disorder characterised by inconsistent articulation, disrupted prosody and effortful initiation. Existing augmentative and alternative communication (AAC) tools and emerging AI-driven speech systems treat the speaker as a stable signal source, yet AoS produces unstable, variable input that breaks this assumption. This short paper reframes generative AI not as a productivity tool nor as an AAC accelerator, but as a communication prosthesis: an adaptive partner that infers intent from unreliable speech while preserving the user's agency, voice and trust. We articulate four HCI design tensions specific to AoS (input variability vs. response consistency; agency preservation vs. assistance density; fluent assistance vs. calibrated uncertainty; adaptive autonomy vs. explainable control) and propose five design principles for adaptive AI support. The paper aims to provoke conversation in the HCI community about how large language models should behave when the user's voice is itself the contested object.

Keywords: *Apraxia of Speech* (AoS), *Generative AI*, *Communication prosthesis*, *Adaptive interaction*, *Explainability, Augmentative and alternative communication* (AAC), *Human-centred AI.*

1. INTRODUCTION

This is Apraxia of Speech (AoS) is a motor speech disorder in which the neural planning and sequencing of speech movements is impaired while language comprehension and muscle strength remain intact (Duffy, 2020; ASHA, 2024). Its defining clinical feature is inconsistency: the same target word may be produced correctly, then distorted, then groped for on successive attempts, with token-to-token variability in articulation and prosody (Iuzzini-Seigel et al., 2017; RCSLT, 2024). For the speaker, this means that the act of being heard is unreliable; for any listening system, whether human or machine, the signal that arrives is fundamentally unstable.

Recent generative AI research has begun to explore large language models (LLMs) as communication partners for people with communication disabilities (Valencia et al., 2023; Mao et al., 2025; Cai et al., 2024). These systems can expand telegraphic input into sentences, repair disfluent text and scaffold conversation in near real time. Almost all of this work, however, is grounded in aphasia (a language disorder) or in anarthria/ALS (where speech is absent or weak but planning is intact). The HCI literature on AoS itself is sparse, even though AoS sits at the intersection where these two assistive paradigms break down. The speaker has language to say, and the motor system to attempt it, but the bridge between intent and articulation is intermittent.

This short paper argues that designing generative AI for AoS requires a different conceptual frame. We extend the notion of communication prosthesis (generative AI as assistive authoring infrastructure rather than predictive text) into the AoS context, where the prosthesis must operate adaptively on an unstable input stream while protecting the user's voice. We make three contributions:

- A reframing of generative AI for AoS as an adaptive communication prosthesis, distinct from AAC for anarthria and from LLM-based aphasia repair.
- An articulation of four HCI design tensions thatemerge when established concerns around agency, trust and explainability are confronted with the moment-to-moment motor-planning variability characteristic of AoS.
- Five proposed design principles that translatethese tensions into requirements for adaptive intervention, confidence gating, agency preservation, voice fidelity and readable adaptation, together with a research agenda for their empirical investigation.

The paper is conceptual and human-centred; it does not propose a clinical intervention, and is offered in the spirit of a “prequel” paper that motivates future empirical work.

© Omoyemi. Published by BCS Learning and Development Ltd. Proceedings of BCS HCI 2026, UK

2. BACKGROUND AND RELATED WORK

2.1 The HCI problem space of Apraxia of Speech

Clinically, AoS is treated through high-intensity motor learning guided therapy emphasising repetition, structured feedback and staged progression (Ballard et al., 2015; Maas et al., 2008). From an HCI standpoint, however, AoS has received little direct attention. Most assistive-technology work has focused either on people with stable but inarticulate speech (e.g., dysarthria, ALS) or on language-level repair (e.g., aphasia). AoS challenges both framings: the user's intended utterance is well-formed, but its realisation is unreliable across attempts, and the unreliability is the disorder itself rather than a noise channel to be filtered out.

2.2 AI-assisted AAC and conversational repair

Valencia et al. (2023) showed that LLM-generated suggestions can reduce keystroke burden for AAC users but are limited by a lack of personal voice. Mao et al. (2025) used design probes to show that people with aphasia value AI that supports conversational grounding through visual verification, grammar support and error correction, but expressed unease when suggestions diverged from their intent. Cai et al. (2024) and Weinberg et al. (2025) extended this line of work to multimodal LLM-driven AAC, while Gaines and Vertanen (2025) demonstrated that character-level LLM adaptation can substantially improve input rates for non-speaking users. Sennott et al. (2019) frame these advances within a broader principle that AAC must augment, not replace, the user's communicative agency.

2.3 Generative AI for disordered speech and self-expression

Cave and Bloch (2020) and Nathanson (2017) examine voice banking and personalised synthetic voices for people with progressive or acquired speech impairments, foregrounding the question of whose voice is being heard and the moral weight of speaking with one's own voice. Dangol et al. (2025) explored AI-supported home speech practice and reported that users and clinicians valued guided choices over opaque model outputs. Together this body of work suggests that generative AI for communication disability is moving from rate enhancement to voice and identity infrastructure, but it has not yet engaged with the specific motor-planning unreliability of AoS

2.4 Gap addressed by this paper

Existing AI-AAC systems implicitly assume one of two input regimes: stable but limited (e.g., a small switch set, a typed keyboard) or absent (eye-gaze, brain signals). AoS sits between them. The user has speech, often intelligible, but its reliability fluctuates within and across utterances, moods, fatigue states and conversational partners. A system that responds identically to a fluent and a groped attempt either over-corrects fluent speech, eroding agency, or under-supports breakdowns, leaving the user stranded. This paper addresses that gap by treating variability itself as the design surface, and asking how an LLM-based partner should adapt to it.

Our contribution does not claim agency, transparency, explainability or trust calibration as new HCI concerns in themselves. Rather, we show that AoS changes how these established concerns must be operationalised. In conventional AI-assisted communication, personalisation typically concerns stable preferences, vocabulary or interaction style. In AoS, the required level of assistance may change from one utterance to the next because the reliability of speech production itself fluctuates. Consequently, agency becomes a problem of dynamically controlling intervention intensity; transparency becomes an obligation to explain why support has changed at a particular moment; and trust calibration becomes inseparable from uncertainty about recovered communicative intent. The novelty of the framework lies in organising these concerns around moment-to-moment speech variability and translating them into adaptive interaction principles specific to AoS.

3. GENERATIVE AI AS A COMMUNICATION PROSTHESIS FOR AOS

3.1 From predictive text to prosthesis

Disability scholarship has long warned that prostheses are not neutral substitutes but devices that carry assumptions about which bodies and voices count as normal, often producing narratives that depend on the disability they claim to repair (Mitchell and Snyder, 2000). We take this caution as a starting point, not a barrier: framing generative AI as a communication prosthesis makes its normative pressures visible and contestable, rather than hiding them inside the language of prediction.

Following prior framing of generative AI as assistive authoring infrastructure, we conceptualise an AoS support system as a communication prosthesis: a layer that couples user input, conversational context and model generation to produce an utterance that the user accepts as their own. Unlike text-prediction frames, this view foregrounds three commitments. First, the system mediates an act of expression, not a typing task. Second, the prosthesis is worn by a person: it must fit the individual's linguistic style, fluctuating reliability and communicative goals. Third, it must be removable, since the user must retain the ability to communicate without it, even imperfectly.

3.2. Three modes of adaptive support

Generalising from prior assistive authoring work to AoS, an adaptive prosthesis operates in three modes whose appropriate use depends on the moment-to-moment state of the speaker.

Repair mode: When the user produces a target word with phoneme distortions or groping, the system proposes an aligned form between what was attempted and what the model believes was intended, and offers it to the user for confirmation before any output is committed. The semantic content is preserved (Mao et al., 2025); only the surface form is repaired, and the speaker remains the author of the utterance.

Scaffold mode: When the user produces partial or telegraphic input (a keyword, a syllable, a gesture), the system offers a small set of plausible continuations from which the user selects. This mirrors Dangol et al.'s (2025) finding that selectable choices, rather than autonomous outputs, sustain transparency and user accountability.

Generative mode: When speech is not currently available (for example, during severe initiation failure or fatigue), the system produces a fuller utterance from minimal cues, personalised to the user's linguistic style. This mode involves the greatest agency risk and therefore the strongest safeguards (Section 5).

Critically, mode selection may emerge from a combination of user preferences, contextual cues and model estimates of reliability, rather than being fixed by either the user or the system alone. This makes adaptive interaction the central HCI problem.

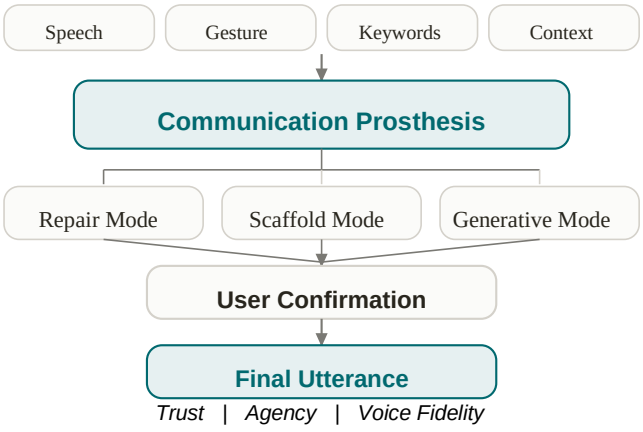

Figure 1: *The communication prosthesis. Multimodal inputs feed an adaptive layer that selects among repair, scaffold and generative modes. User confirmation gates the final utterance, whose desired qualities are trust, agency and voice fidelity.*

4. FOUR HCI DESIGN TENSIONS

4.1 T1: Input variability vs. response consistency

AoS speech is non-stationary: the same speaker may articulate fluently for one phrase and grope on the next (Iuzzini-Seigel et al., 2017). Most automatic speech recognition systems are trained on stable speech and degrade on disordered input (Jiang et al., 2024). A communication prosthesis must therefore estimate the reliability of the current utterance and tune its level of intervention accordingly. This reframes adaptation away from generic personalisation (static user model) toward moment-to-moment calibration of intervention intensity.

4.2 T2: Agency preservation vs. assistance density

Dense assistance is tempting: a fluent generated sentence from a keyword saves time. But research on human-AI co-writing shows that the level of scaffolding directly shapes users' sense of authorship and ownership (Dhillon et al., 2024). For AoS users, whose speech is already renegotiated with listeners many times a day, an over-eager prosthesis risks substituting the speaker rather than supporting them.

Weinberg et al. (2025) describe this as the risk of LLMs making jokes *for* users instead of *with* them. Designers must therefore treat agency as a first-class metric, not an afterthought (Satyanarayan and Jones, 2024).

4.3 T3: Fluent assistance vs. calibrated uncertainty

LLMs frequently appear confident when wrong (Zhang et al., 2024). For communication-disability contexts, the cost of miscalibrated trust is asymmetric: an undertrusted system is ignored, an overtrusted one speaks for the user. Recent HCI work shows that transparent uncertainty communication and explanation rationales substantially improve trust calibration in human-AI collaboration (Bansal et al., 2021; Naiseh et al., 2023). For an AoS prosthesis, this implies uncertainty must be surfaced at utterance level (for example, “I'm not sure I heard you correctly. Did you mean…?”) rather than hidden inside model logits.

4.4 T4: Adaptive autonomy vs. explainable control

Because the prosthesis silently chooses between repair, scaffold and generative modes, it makes implicit judgements about the user's reliability in the current moment. These judgements can be experienced as paternalistic if unexplained. Drawing on XAI work in clinical decision support (Panigutti et

al., 2022; Ghassemi et al., 2021), we argue that adaptive support requires readable adaptation traces: short, faithful explanations of why the system changed mode, what cue it used, and how the user can override it. Explanation here is not for diagnosis but for accountability to the speaker.

5. PROPOSED DESIGN PRINCIPLES

The principles below were developed by translating each identified tension into a corresponding interaction requirement. Rather than treating established HCI values such as agency, transparency and trust as abstract goals, we ask what system behaviour would be required to preserve those values when speech reliability changes within and across utterances. The resulting principles should therefore be understood as design propositions for subsequent participatory and empirical validation rather than as empirically established guidelines.

Synthesising across the four tensions, we propose five design principles for adaptive generative AI support in AoS. Voice fidelity (P4) is treated as a cross-cutting response to the agency and generative-mode risks that recur across T2–T4, rather than mapping one-to-one onto a single tension.

P1. Variability as a first-class signal: Treat moment-to-moment speech reliability, rather than a static diagnostic label, as the variable that drives system behaviour. Surface uncertainty information in forms that are meaningful to users and communication partners.

P2. Confidence-gated intervention: Increase the depth of generative support only as the system's confidence in having recovered the user's intended meaning increases. When confidence is low, abstain or ask, rather than guess (Wen et al., 2025; Zhang et al., 2024).

P3. Agency-preserving defaults: Default to the lightest viable mode (repair before scaffold before generative). Make the user's last verified utterance, not the model's preferred completion, the canonical output.

P4. Voice fidelity over fluency: Personalise the model to the user's idiolect, including their idiomatic and stylistic preferences, so that generative output is recognisable as theirs rather than as a generic LLM voice (Hwang et al., 2025; Nathanson, 2017).

P5. Readable adaptation: When the system shifts mode or substitutes content, provide a concise, faithful explanation that the user can interrogate or override (Mitchell et al., 2019; WHO, 2024).

These principles align with the EU AI Act's expectation of human oversight for high-risk AI used in health and assistive contexts (European Union, 2024) and with the WHO's (2024) guidance on large multimodal models in care.

6. RESEARCH AGENDA

We pose five questions to motivate future HCI work on adaptive AI for AoS.

RQ1. How can a system estimate per-utterance reliability for an AoS speaker in real time, and what representations of that estimate are interpretable to speakers and partners?

RQ2. What metrics capture communicative agency for AoS users beyond rate and intelligibility (for example, perceived authorship, conversational repair initiation, and willingness to attempt difficult utterances)?

RQ3. How should an adaptive prosthesis transition between repair, scaffold and generative modes in a way that feels coherent and non-paternalistic?

RQ4. What forms of explanation make adaptive mode changes accountable to the speaker without imposing additional cognitive load?

RQ5. What are the long-term effects of prolonged prosthetic use on the user's own speech motor learning, social identity, and listener perception?

Methodologically, these questions invite participatory design with people with AoS and speech-language pathologists, longitudinal field studies, and adversarial evaluation of LLMs against AoS-typical input variation.

7. LIMITATIONS AND POSITION

This is a position paper. We do not present an evaluated system; we sketch a conceptual frame and a set of HCI questions. AoS is heterogeneous (acquired vs. childhood, isolated vs. co-occurring with aphasia or dysarthria), and any deployed prosthesis will require disaggregated study. We have also written primarily from an HCI perspective; clinical, ethical and disability-studies perspectives, including the risk that prostheses impose normative speech (Mitchell and Snyder, 2000), must shape any system that follows this argument. Finally, the focused literature synthesis used to derive the tensions (see Appendix A) was intended to support conceptual framework development rather than provide systematic evidence of prevalence or effect; consequently, the proposed tensions and principles require validation with people with AoS, speech-language pathologists and other communication partners.

8. CONCLUSION

Generative AI is widely framed as advanced predictive text. For Apraxia of Speech, that framing is inadequate. The disorder makes the user's voice itself the unstable object that any AI partner must work with. We have argued for a communication prosthesis view of generative AI in AoS: an adaptive, agency-preserving, trust-calibrated and explainable partner that intervenes only as much as the moment requires. More broadly, we hope this paper encourages the HCI community to view

communication disabilities not simply as noisy inputs to be corrected, but as opportunities to rethink how AI should support expression, agency and identity.

9. USE OF AI IN PREPARING THIS PAPER

In line with the conference policy, the author declares that generative AI tools were used to support (a) literature discovery and source triage, (b) drafting and refining language in the introduction, related work, and design principles sections, and (c) consistency checking against earlier work by the author. All conceptual claims, the framing of AoS-specific design tensions, the design principles, and the research agenda were formulated and verified by the author. AI tools were not used for data collection or empirical analysis, as this paper is conceptual. All cited sources were verified by the author.

APPENDIX A. LITERATURE AND FRAMEWORK DEVELOPMENT

Literature identification: As this is a conceptual position paper rather than a systematic review, the literature was identified through a focused, iterative search intended to map the intersection of Apraxia of Speech, AI-mediated communication and human-centred AI. Searches combined terms relating to “apraxia of speech”, “augmentative and alternative communication”, “AAC”, “generative AI”, “large language models”, “disordered speech”, disability”, “agency”, “human-AI “trust”, “communication collaboration”, “explainability” and “adaptive interfaces”. Priority was given to peer-reviewed HCI, accessibility, speech language pathology and AI literature addressing AI-mediated communication, user agency, personal voice, adaptive support or uncertainty.

Selection criteria: Sources were retained where they contributed directly to at least one of three areas: (1) characteristics of AoS relevant to interactive system design; (2) AI or computational support for people with communication disabilities; or (3) established HCI concerns arising when AI mediates human expression. Literature concerned solely with clinical diagnosis or treatment efficacy, without implications for interaction design, was used only where necessary to establish the characteristics of AoS.

Derivation of the design tensions: The four tensions were developed through comparative synthesis rather than statistical or systematic-review coding. We examined where established assumptions in AI-assisted AAC and human-AI interaction become problematic when applied to the defining variability of AoS. Recurring conflicts were grouped into four design tensions: variability versus consistency of system response; assistance versus user agency; model confidence versus appropriately calibrated user trust; and adaptive behaviour versus the user's ability to understand and contest that adaptation. The five proposed design principles were subsequently formulated as actionable design responses to these tensions.

AoS-specific problem	Existing HCI/AI assumption	Resulting tension	Design response
Speech reliability changes between utterances	Systems generally assume relatively stable input	T1: Input variability vs. response consistency	P1 Variability as first-class signal
Users can often communicate independently but intermittently require support	More assistance is often treated as beneficial	T2: Agency preservation vs. assistance density	P3 Agency-preserving defaults
Intended meaning may be uncertain even when generation is fluent	LLM output can appear authoritative	T3: Fluent assistance vs. calibrated uncertainty	P2 Confidence-gated intervention
Support level may change dynamically	Adaptive systems can alter behaviour opaquely	T4: Adaptive autonomy vs. explainable control	P5 Readable adaptation
Generated language can become more fluent than the user's natural expression	Generation tends to optimise linguistic fluency	Cross-cutting: voice fidelity vs. normative fluency	P4 Voice fidelity over fluency

Table A1. *Mapping from AoS-specific interaction problems and established HCI/AI assumptions to the four design tensions (T1–T4) and the corresponding proposed design principles (P1–P5). The final row reflects voice fidelity as a cross-cutting concern rather than a one-to-one mapping.*

REFERENCES

ASHA (2024) *Acquired Apraxia of Speech: Practice Portal.* American Speech-Language-Hearing Association. Available at: https://www.asha.org/practice-portal/clinical-topics/acquired-apraxia-of-speech/ (Accessed: 13 June 2026).

Ballard, K.J., Wambaugh, J.L., Duffy, J.R., Layfield, C., Maas, E., Mauszycki, S. and McNeil, M.R. (2015) ‘Treatment for acquired apraxia of speech: a systematic review of intervention research between 2004 and 2012’, *American Journal of Speech-Language Pathology*, 24(2), pp. 316–337. https://doi.org/10.1044/2015_AJSLP-14-0118.

Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M.T. and Weld, D.S. (2021) ‘Does the whole exceed its parts? The effect of AI explanations on complementary team performance’, in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York: ACM, pp. 1–16. https://doi.org/10.1145/3411764.3445717.

Cai, S., Venugopalan, S., Tomanek, K., MacDonald, B., Jiang, M., Brenner, M., Nelson, J., Tu, T., Morris, M.R. and Green, J.R. (2024) ‘Using large language models to accelerate communication for users with severe motor impairments’, *Nature*

Communications, 15(1), 10168. https://doi.org/10.1038/s41467-024-53873-3.

Cave, R. and Bloch, S. (2020) 'Voice banking for people living with motor neurone disease: views and expectations', *International Journal of Language & Communication Disorders*, 56(1), pp. 116–129. https://doi.org/10.1111/1460-6984.12588.

Dangol, A., Lewis, A., Suh, H., Hong, X., Meadan, H., Fogarty, J. and Kientz, J.A. (2025) '"I want to think like an SLP": a design exploration of AI-supported home practice in speech therapy', in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. New York: ACM. https://doi.org/10.1145/3706598.3713986.

Dhillon, P.S., Molaei, S., Li, J., Golub, M., Zheng, S. and Robert, L.P. (2024) 'Shaping human–AI collaboration: varied scaffolding in LLM co-writing', in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. New York: ACM. https://doi.org/10.1145/3613904.3642134.

Duffy, J.R. (2020) *Motor Speech Disorders: Substrates, Differential Diagnosis, and Management*. 4th edn. St. Louis: Elsevier.

European Union (2024) *Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, Article 14, Human oversight. https://artificialintelligenceact.eu/article/14/.

Gaines, D. and Vertanen, K. (2025) 'Adapting large language models for character-based augmentative and alternative communication', in *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 15273–15291. https://aclanthology.org/2025.findings-emnlp.826/.

Ghassemi, M., Oakden-Rayner, L. and Beam, A.L. (2021) 'The false hope of current approaches to explainable artificial intelligence in health care', *The Lancet Digital Health*, 3(11), pp. e745–e750. https://doi.org/10.1016/S2589-7500(21)00208-9.

Hwang, A.H.C., Liao, Q.V., Blodgett, S.L., Olteanu, A. and Trischler, A. (2025) '"It was 80% me, 20% AI": seeking authenticity in co-writing with large language models', *Proceedings of the ACM on Human-Computer Interaction*, 9(CSCW1). https://doi.org/10.1145/3711020.

Iuzzini-Seigel, J., Hogan, T.P. and Green, J.R. (2017) 'Speech inconsistency in children with childhood apraxia of speech, language impairment, and speech delay', *Journal of Speech, Language, and Hearing Research*, 60(5), pp. 1194–1210. https://doi.org/10.1044/2016_JSLHR-S-15-0184.

Jiang, M., Tomanek, K., Tobin, J., MacDonald, B., Cattiau, J., Yu, P., Brenner, M., Cave, R., Seaver, K., Casado, M., Green, J.R. and Venugopalan, S. (2024) 'Learnings from curating a trustworthy, well-annotated, and useful dataset of disordered speech', *Proceedings of Interspeech 2024*, pp. 4863–4867. https://doi.org/10.21437/interspeech.2024-578.

Maas, E., Robin, D.A., Austermann Hula, S.N., Freedman, S.E., Wulf, G., Ballard, K.J. and Schmidt, R.A. (2008) 'Principles of motor learning in treatment of motor speech disorders', *American Journal of Speech-Language Pathology*, 17(3), pp. 277–298. https://doi.org/10.1044/1058-0360(2008/025).

Mao, L., Lee, J.H., Shah, Y.F. and Valencia, S. (2025) 'Design probes for AI-driven AAC: addressing complex communication needs in aphasia', *arXiv preprint* arXiv:2504.09435. https://arxiv.org/abs/2504.09435.

Mitchell, D.T. and Snyder, S.L. (2000) *Narrative Prosthesis: Disability and the Dependencies of Discourse*. Ann Arbor: University of Michigan Press.

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D. and Gebru, T. (2019) 'Model cards for model reporting', in *Proceedings of the Conference on Fairness, Accountability, and Transparency*. New York: ACM, pp. 220–229. https://doi.org/10.1145/3287560.3287596.

Naiseh, M., Al-Thani, D., Jiang, N. and Ali, R. (2023) 'How the different explanation classes impact trust calibration: the case of clinical decision support systems', *International Journal of Human-Computer Studies*, 169, 102941. https://doi.org/10.1016/j.ijhcs.2022.102941.

Nathanson, E. (2017) 'Native voice, self-concept and the moral case for personalised voice technology', *Disability and Rehabilitation*, 39(1), pp. 73–81. https://doi.org/10.3109/09638288.2016.1139193.

Panigutti, C., Beretta, A., Giannotti, F. and Pedreschi, D. (2022) 'Understanding the impact of explanations on advice-taking: a user study for AI-based clinical decision support systems', in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. New York: ACM. https://doi.org/10.1145/3491102.3502104.

RCSLT (2024) *Position paper on Childhood Apraxia of Speech (CAS)*. London: Royal College of Speech and Language Therapists. https://www.rcslt.org/.../CAS-Position-Paper-2024.pdf.

Satyanarayan, A. and Jones, G.M. (2024) 'Intelligence as agency: evaluating the capacity of generative AI to empower or constrain human action', *An MIT Exploration of Generative AI*. https://doi.org/10.21428/e4baedd9.2d7598a2.

Sennott, S.C., Akagi, L., Lee, M. and Rhodes, A. (2019) 'AAC and artificial intelligence (AI)', *Topics in Language Disorders*, 39(4), pp. 389–403. https://doi.org/10.1097/TLD.0000000000000197.

Valencia, S., Cave, R., Kallarackal, K., Seaver, K., Terry, M. and Kane, S.K. (2023) '"The less I type, the better": how AI language models can enhance or impede communication for AAC users', in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. New York: ACM. https://doi.org/10.1145/3544548.3581560.

Wen, B., Yao, J., Feng, S., Xu, C., Tsvetkov, Y., Howe, B. and Wang, L.X. (2025) 'Know your limits: a survey of abstention in large language models', *Transactions of the Association for Computational Linguistics*, 13, pp. 529–556. https://doi.org/10.1162/tacl_a_00754.

WHO (2024) *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*. Geneva: World Health Organization. https://www.who.int/publications/i/item/9789240084759.

Weinberg, T.M., Kadoma, K., Gonzalez Penuela, R.E., Valencia, S. and Roumen, T. (2025) 'Why so serious? Exploring timely humorous comments in AAC through AI-powered interfaces', in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. New York: ACM. https://doi.org/10.1145/3706598.3714102.

Zhang, M., Huang, M., Dong, S., Zhou, Y. and Qiu, X. (2024) Calibrating the confidence of large language models by eliciting fidelity', in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 2935–2948. https://doi.org/10.18653/v1/2024.emnlp-main.173