Handout, ATIA, January 2026

LET'S LOOK BEFORE WE LEAP

A thoughtful approach to AI and AAC

Alyssa Hillary Zisk, PhD, AAC Research Team Lead, AssistiveWare

David Niemeijer, PhD, CEO and Founder, AssistiveWare

Summary

Introduction

“Artificial Intelligence,” or AI, is a hot topic in the world and specifically in AAC, but “AI” implementations vary widely. We took a broad view to identify different types of “AI,” their uses in AAC, and considerations around these uses. The uses of machine learning techniques for brain-computer interfaces, personal voices, and large language models (LLMs) differ – but all fall under “AI.” By addressing these cases separately, we identified key considerations.

First, we consider authenticity. Professionals express concerns about authorship: LLMs may alter or invent content. AAC users have concerns about message style and tone (Valencia et al., 2023), but also about undue judgment regarding authorship (Holyfield & Williams, 2025).

Second, we consider privacy. “AI” models often involve cloud processing in their creation, adaptation, and/or use. This has trade-offs with authenticity: models tuned with specific user data (neural signals, audio, or text output) can be more effective and authentic, but they then include that data and may reveal it in unintended ways (Valencia et al., 2023).

Third, we consider barriers to learning. A literate user can check LLM-generated text and make decisions about speed, effort, and tone. However, much like calculator use elides the math skills touse a calculator effectively, LLM use elides the language skills to check and edit LLM output.

Finally, we consider availability and accessibility – some applications of “AI” are becoming ubiquitous. Others may be touted as the future of AAC while remaining inaccessible (e.g., voice banking; Preece et al., 2024) or rarely available (e.g., brain computer interfaces; Sellwood et al., 2024).

Addressing “AI” in AAC effectively involves considering each of its use cases on its own merits. It also involves addressing these considerations – not as issues unique to the intersection of AAC and AI, but as applications of broader issues to this intersection.

Introducing "AI"

"AI" isn't just one thing—any time a computer system does tasks typically associated with human intelligence, that falls under artificial intelligence, or AI. It has been an area of academic research since the 1950s, including tasks such as image processing, decision-making, navigation, and natural language processing. As a given task becomes more common, sometimes people stop calling it AI.

As AI is about having computer systems complete tasks, it necessarily relies on mathematics and algorithms—instructions about what data to do math on, what math to do, and what to do with the results of the math. This math can be used for language detection, translation, spell checking, word or phrase prediction, summarization, self-driving cars, photo tagging, cancer detection, and many other tasks.

Many AI applications can encounter similar types of problems:

- AI applications rely on data, often pulled from the way things are currently done or normative data. This can lead to the enforcement of **uniformity** when minority groups are not represented in norms.
- If the data is **biased** or **unrepresentative**, the output is likely to have the same problems. This is a concern with algorithmic bias, where algorithms reproduce the same kinds of problems that existed in the training data—like ableist denials of care.
- Feeding AI-generated data into another AI system can lead to **system degradation**. This is a problem in creating new large language models.
- Because algorithms can be very **complex**, it may be difficult to identify any errors in the algorithms, even when examination of the output and/or input tells us there must be errors.
- Being able to rapidly process large amounts of data makes **surveillance** easier, posing **privacy** risks.
    - Mass surveillance can enable **censorship** on a broader scale.
- People who create or control AI systems may attempt to **manipulate outputs**.
- **Misunderstandings** of what a given algorithm does (or does not do) may lead to applications in situations where it is unfit for the purpose.

When can AI apply to AAC

Most applications of AI that are relevant to the AAC field are connected to natural language processing, a subfield of AI that lets programs interact with human languages. This includes speech recognition, voice banking, speech generation, translation, word prediction, and, of course, the large language models that are seeing more and more widespread use.

Some applications of AI to AAC will also use classification methods on other types of data (e.g., neural data for brain computer interfaces) or work with image generation as well (e.g., creating visuals).

What are some considerations for AI in AAC?

Major considerations include **authenticity**, **privacy**, **barriers to learning**, and **access and availability**.

Authenticity

Authenticity concerns include:

- **Actual inauthenticity**: Large language models are designed to produce something that statistically *looks like* a response to the input, which isn't the same thing as matching a user's intent or writing style. This is especially true when AAC users want to say statistically unlikely things—and we know AAC users tend to have unique vocabularies. It is important to note that "plausible but factually incorrect or fabricated content", commonly known as AI hallucinations, may be theoretically unavoidable with current language model designs (Cossio, 2025).

  In voice creation contexts, actual inauthenticity can also occur when systems change or introduce accents that were not present in the recorded speech or the prompt.

- **Presumptions of inauthenticity**: AAC users already face assumptions that we are not the source of our own communication. Very real problems with the accuracy of AI output are not unique to AAC, but their impact on AAC users is compounded by these existing assumptions.

- **Disagreements about what authenticity means anyway**: Not all AAC users *want* AAC voices that sound like our speaking voices. For some, it's not necessarily clear what our speaking voice would sound like (Preece et al., 2024). For others, it's just not the goal—an AAC user may decide not to pursue this option for a variety of reasons, including gender (Zisk, 2021), having come to identify with a specific computer generated voice as an authentic representation of themself (Bell, 2025; Preece et al., 2024) or simply not liking to listen to recordings of ones own voice.

  In image generation contexts, concerns about authenticity can overlap with discussions about iconicity, where a generalized or stylized representation covers an entire category (a picture of a general dog for the word dog) versus specificity, where one example covers a category (a picture of the family dog represents the word dog). AI images often trend towards specificity, which may be seen as more authentic, but there are reasons an AAC system might use an iconic representation for the general category and specific representations for items within the category (a picture of a general dog for the word dog and a picture of the family dog for the dog's name).

Privacy

Privacy concerns with AI in AAC include concerns in training data, input data, off-device processing during use, and the accidental reveal of information in outputs.

- Because AAC users often have unique vocabularies and writing patterns, there is a trade-off between **authenticity** and **privacy**: providing one's own prior transcripts or some facts about oneself to train a language model *can* address some of the concerns around actual inauthenticity, but at the cost of user privacy.
  - Any information that is statistically present in the training data for a large language model can be revealed in the text output of that model, including at times when that wasn't intended.
- Some large language models take user input as further training data, either for all users or for users who do not pay to opt out. This is a privacy concern, whether or not it improves authenticity for any specific user.
- Applications that involve speech recognition on an ongoing basis (e.g., with an activation phrase) are, to a certain extent, **always listening**. They have to be, in order to always detect the activation phrase—but what else might they record when not in use?
- **Context-aware** boards have to detect the context somehow—this may be from listening or from detecting and processing locations.

Barriers to learning

Barriers to learning include situations where AI use makes it **harder to use previously learned skills** (e.g., if AI edits a board, that will affect motor planning just as surely as it would if a person edited a board—but now we may not know it happened). It also includes situations where AI use leads to the **weakening of existing skills** (Budzyń et al., 2025) or **failure to retain** what was just written (Kosmyna et al., 2025). Barriers to learning also covers situations where AI use elides skills that are important to language development: vocabulary, grammar, writing, and editing skills are both needed to use AI effectively *and* skipped over when using AI without checking its output—which people without disabilities absolutely do, up to and including in academic publishing (França & Monserrat, 2024; Glynn, 2024). None of these barriers to learning are specific to AAC use or users, but many of them interact badly with existing problems facing AAC users.

When AAC users are not consistently taught literacy (Williams, 2025), the writing and editing skills needed to make useful AI prompts, check if AI output is close enough to the intent to be useful, and make any needed edits to AI prompts or outputs may not be present yet. Given the retention problems with AI outputs, using AI in AAC is unlikely to teach those needed skills. This is particularly a concern for **emerging** AAC users and in **educational contexts**: independent and expert AAC users are unlikely to be immune to AI-related barriers to learning, which affect the general population, but do and should have the right to make decisions about the relevant trade-offs for themselves.

Access and availability

Access and availability are concerns around **cost**, availability of needed **supports or resources**, and **transparency** about AI features or their absence.

- In **voice banking**, cost has been decreasing substantially for users with access to certain kinds of devices, using certain languages, and with sufficient speaking ability (or with relatives who have sufficient speaking abilities, sufficiently similar voices, and a willingness to bank their voices for use). However, outside those parameters, voice banking often remains an expensive proposition.

- Effective **brain-computer interfaces** are high-cost devices, often not commercially available yet, and only usable within research contexts. They also involve significant set-up and maintenance, and many of them involve surgery as part of the set-up. Failures in brain-computer interface use are also often described as BCI illiteracy rather than as problems in system design that still need to be addressed—as it stands, not everyone can use a brain-computer interface, and in many cases, the people who find the systems don't work for them are exactly the people who would benefit most—if it worked.

- Finally, it is not always clear when AI features are or are not being used—companies have been known to introduce new AI features without communicating this clearly to users or offering clear opt-in or opt-out options (Germain, 2025). This can make it difficult to use AI features *and* make it difficult to avoid them—regardless of preference, the lack of transparency makes it harder to act on that preference.

All technologies involve trade-offs

With AAC, trade-offs involving **speed** and **informativeness** (Hoag et al., 2004), **relevance** (Bedrosian et al., 2003), and **brevity** (McCoy et al., 2007) are not new—AI improvements simply add options to consider with different combinations around different user priorities.

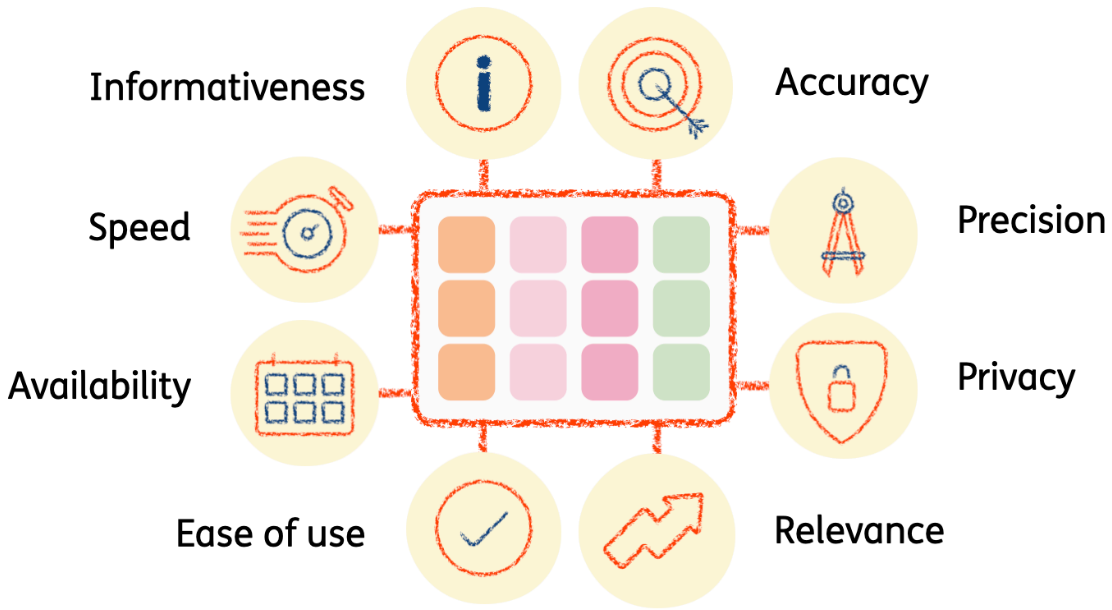

Who may benefit from large language models in message creation?

The primary beneficiaries of large language models in message creation are people for whom trading the motor effort involved in message creation for the cognitive effort of message checking is beneficial, and who are fully literate. In practice, this is likely to mean AAC users with significant motor disabilities—while nondisabled people often perceive AI tools as making them faster by reducing typing time, prompting, waiting, and checking time often take up more time than the reduction in direct typing time frees. When the typing time and effort for a given message length increase, this calculation could change, making the reduction in typing more beneficial.

Who may benefit from other applications of AI in AAC?

Many users are already benefitting from **voice banking** and other forms of personal voice creation. People who can provide an appropriate audio source or who can provide an appropriate prompt for an instruction-based voice (and can get a produced voice that follows the instructions as an output) can benefit here.

‘**Fix’ functions** meant to catch typos or missed inflections can benefit AAC users who make typos these functions will catch or are otherwise making messages that aren’t *quite* what they want to say but are close enough for the system to accurately categorize the intent—as long as the AAC user has the literacy skills to recognize whether or not the system got it right.

Brain-computer interfaces are primarily relevant for AAC users who have motor disabilities such that alternative access is relevant (a potentially wider group than typically assumed, but still not all AAC users) and access to the needed supports to set up and use them.

References

Bedrosian, J. L., Hoag, L. A., & McCoy, K. F. (2003). Relevance and speed of message delivery trade-offs in augmentative and alternative communication. Journal of Speech, Language, and Hearing Research, 46(4), 800-817.

Bell, S. A. (2025, January 9). Stephen Hawking's Eternal Voice. Nautilus. https://nautil.us/stephen-hawkings-eternal-voice-1176962/

Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H. O., ... & Mori, Y. (2025). Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. The Lancet Gastroenterology & Hepatology.

Cossio, M. (2025). A comprehensive taxonomy of hallucinations in Large Language Models. arXiv preprint arXiv:2508.01781.

França, T. F., & Monserrat, J. M. (2024). The artificial intelligence revolution... in unethical publishing: Surfing the liquid world and its intangible economy. J. Gen. Physiol., 156, e202413654.

Germain, T. (2025, August 22). YouTube secretly used AI to edit people's videos. The results could bend reality. BBC Future. https://www.bbc.com/future/article/20250822-youtube-is-using-ai-to-edit-videos-without-permission

Glynn, A. (2024). Suspected undeclared use of artificial intelligence in the academic literature: an analysis of the Academ-AI dataset. arXiv preprint arXiv:2411.15218.

Hoag, L. A., Bedrosian, J. L., McCoy, K. F., & Johnson, D. E. (2004). Trade-offs between informativeness and speed of message delivery in augmentative and alternative communication. Journal of Speech, Language, and Hearing Research, 47(6), 1270-1285.

Holyfield, C., & Williams, K. (2025, March 15). Predictive Text: Who Controls the Conversation? ASHA LeaderLive. https://leader.pubs.asha.org/do/10.1044/leader.FTR1.30032025.FAAC-predictive-text.36/full/

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X. H., Beresnitzky, A. V., ... & Maes, P. (2025). Your brain on chatgpt: Accumulation of cognitive debt when using an ai assistant for essay writing task. arXiv preprint arXiv:2506.08872.

McCoy, K. F., Bedrosian, J. L., Hoag, L. A., & Johnson, D. E. (2007). Brevity and speed of message delivery trade-offs in augmentative and alternative communication. Augmentative and Alternative Communication, 23(1), 76-88.

Preece, J., Sullivan, E., Tams-Gray, F., & Pullin, G. (2024). Making my voice and owning its future. Medical Humanities, 50(4), 624-634.

Sellwood, D., McLeod, L., Williams, K., Brown, K., & Pullin, G. (2024). Imagining alternative futures with augmentative and alternative communication: a manifesto. Medical Humanities, 50(4), 620-623.

Valencia, S., Cave, R., Kallarackal, K., Seaver, K., Terry, M., & Kane, S. K. (2023, April). "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 (pp. 1-14).

Williams, B. (2025). Reading, writing, and AAC are fundamental. Augmentative and Alternative Communication, 1-3.

Zisk, A. H. (2021, 30 June). "Why are you using a male voice?" *CommunicationFIRST*. https://communicationfirst.org/why-are-you-using-a-male-voice/