

Title	An AI tutorial for speech and language therapists: Translating concepts from the AI literature into accessible knowledge and clinically relevant applications
Authors	Oliveira-Buckley, Ana;O'Sullivan, Barry;McKean, Cristina;Frizelle, Pauline
Publication date	2026-02-06
Original Citation	Oliveira-Buckley, A, O'Sullivan, B, McKean, C & Frizelle, P 2026, 'An AI tutorial for speech and language therapists: Translating concepts from the AI literature into accessible knowledge and clinically relevant applications', International Journal of Language and Communication Disorders, vol. 61, no. 2, e70201, pp. 1-17. https://doi.org/10.1111/1460-6984.70201
Type of publication	Article (peer-reviewed)
Link to publisher's version	https://www.scopus.com/pages/publications/105029500571 - 10.1111/1460-6984.70201
Rights	© 2026 The Author(s). International Journal of Language & Communication Disorders published by John Wiley & Sons Ltd on behalf of Royal College of Speech and Language Therapists. This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made. - https://creativecommons.org/licenses/by-nc-nd/4.0/
Download date	2026-09-15 05:22:00
Item downloaded from	https://hdl.handle.net/10468/18966



University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

DISCUSSION OPEN ACCESS

An AI Tutorial for Speech and Language Therapists: Translating Concepts From the AI Literature Into Accessible Knowledge and Clinically Relevant Applications

Ana Oliveira-Buckley¹  | Barry O'Sullivan² | Cristina McKean³  | Pauline Frizelle¹ 

¹Department of Speech and Hearing Sciences, University College Cork, Cork, Ireland | ²School of Computer Science and IT, University College Cork, Cork, Ireland | ³Department of Education, University of Oxford, Oxford, UK

Correspondence: Ana Oliveira-Buckley (anabuckley@ucc.ie)

Received: 19 August 2025 | **Revised:** 19 December 2025 | **Accepted:** 19 January 2026

Keywords: artificial intelligence | child | language therapy | speech therapy

ABSTRACT

Background: Artificial Intelligence (AI) is increasingly discussed as a tool that can support speech and language therapy (SLT). However, clinical adoption of AI requires improved AI literacy among clinicians. AI is a rapidly evolving and often inconsistently defined field that can be difficult to navigate. Despite the definition provided by the EU AI Act, AI terminology can feel abstract for non-technical readers.

Aims: To provide a foundational understanding of AI tailored for SLTs, by translating complex concepts into accessible language and organising them across three levels: (i) *AI techniques* (how AI works); (ii) *AI capabilities* (what AI can do) and (iii) *clinical applications* (how AI can support SLT).

Methods: This tutorial is informed by foundational AI literature, established AI taxonomies, relevant SLT literature and regulatory and ethical guidelines. Clinical analogies are used to explain technical concepts, with additional technical detail signposted where relevant. Existing and conceptual examples illustrate the relevance of AI across paediatric SLT practice.

Main contribution: This tutorial provides: (i) a clinician-focussed interpretation of the EU AI Act definition; (ii) an organisation of key AI concepts into techniques, capabilities and clinical applications; (iii) a production-line model for mapping clinical needs to AI design choices and (iv) a practice-focussed discussion of ethical and regulatory considerations.

Conclusion: AI is best understood as a set of techniques that enable specific capabilities, which in turn support clinical applications. This tutorial promotes the safe, ethical and accountable use of AI as a tool that can support rather than replace clinicians.

WHAT THIS PAPER ADDS

What is already known on this subject

- Current Artificial Intelligence (AI) literature is typically designed for technical audiences, making it difficult for clinicians to interpret. This can hinder the effective and responsible integration of AI into clinical practice.

What this paper adds to the existing knowledge

- This tutorial provides a clinician-focussed explanation of AI, structured across three levels: (i) *AI techniques* (how AI works); (ii) *AI capabilities* (what AI can do) and (iii) *clinical applications* (how AI supports practice) in paediatric speech and language therapy. It also addresses key challenges, ethical considerations and regulatory requirements relevant to clinical contexts.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2026 The Author(s). *International Journal of Language & Communication Disorders* published by John Wiley & Sons Ltd on behalf of Royal College of Speech and Language Therapists.

What are the potential or actual clinical implications of this work?

- This tutorial lays the groundwork for informed engagement with emerging AI tools. It prepares clinicians to evaluate how different *AI techniques* and *capabilities* may support core clinical tasks (e.g., assessment, therapy planning and delivery).

1 | Introduction: The Challenge of Deciphering AI

Artificial Intelligence (AI) is expected to impact every corner of modern life, including healthcare. Liss and Berisha (2020) propose that AI can enhance the profession of speech and language therapy by facilitating more efficient clinical documentation, providing assistive technologies for clients, permitting objective assessment and personalising practice. Yet for many clinicians, AI remains an ambiguous concept and understanding what it is and how it applies to SLT, can feel like trying to solve a jigsaw puzzle without a reference image.

Currently, there is no universally agreed-upon definition of AI, with the term often described as technology that mimics human intelligence (e.g., Cordella et al. 2024). However, as human intelligence itself is not fully understood, this explanation remains inherently vague. Misconceptions are also common. For instance, many associate AI solely with modern tools such as ChatGPT or personalised social media feeds, overlooking older technologies and broader applications (Mitchell 2020).

Amid this lack of clarity, Regulation (EU) 2024/1689 on AI, referred to as the EU AI Act (European Parliament and Council of the European Union 2024) introduced a formal definition that is increasingly recognised as a gold standard within the AI research community. According to the Act:

“An AI system means a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.”

This definition provides a valuable foundation, emphasising a system’s ability to infer from input data and generate meaningful outputs. However, for non-technical audiences, it can feel abstract and difficult to interpret without additional context.

The challenge in understanding AI is further compounded by the existing literature on AI in SLT, which often does not define AI or describes it in a fragmented or overly restrictive manner. For instance, Fang et al. (2023) conducted a systematic review of AI in story writing without clarifying what they defined as AI. Other papers allude to “technological developments” without sufficient explanation of the nature of these developments, or focus on one or two clinical applications only (e.g., Cordella et al. 2024).

Beyond definitions, numerous classification systems have emerged over time, reflecting different perspectives and priorities across disciplines. In 2020, the EU reviewed 35 existing AI frameworks and found that most were either technically inaccurate (e.g., misclassifying AI capabilities) or incomplete, often focusing solely on underlying AI techniques and overlooking the functions AI is designed to perform (European Institute of Innovation and Technology 2020). In response, the EU developed its own AI taxonomy, which is widely accepted within the AI research community, alongside frameworks such as the Association for the Advancement of Artificial Intelligence (AAAI)-20 keyword classification (Association for the Advancement of Artificial Intelligence n.d.). These taxonomies (discussed in the following section) play important roles in structuring research and guiding policy. However, they cater to policymakers and technical experts, posing interpretation challenges for clinicians.

A recent survey by Aggarwal et al. (2025) found that only a small proportion of SLTs and audiologists were aware of the clinical applications of AI, highlighting the need for greater awareness and training. Cordella et al. (2024) made an important contribution with their tutorial on Machine Learning (ML) and its applications in the *diagnosis* and *recovery prediction* of adults with aphasia. However, the tutorial does not address the broader scope of AI or its potential application with paediatric populations or intervention.

To the best of our knowledge, no publication has yet aimed to provide a broad, accessible overview of AI tailored to SLTs and to contextualise its application across paediatric speech and language therapy. This tutorial aims to fill this gap by: (i) highlighting emerging gold standards within the AI community as reference points; (ii) translating key concepts for SLTs and, where necessary, reframing them in a more intuitive manner; (iii) exploring existing and potential clinical applications of AI in paediatric populations and (iv) addressing key limitations, ethical concerns and regulatory issues.

Box 1. A Note on Terminology and Technical Detail

Throughout this tutorial, technical precision is balanced with accessibility. The goal is to provide an entry point to AI for SLTs, rather than an exhaustive technical overview. To maintain readability, additional technical details and/or nuanced discussion have been moved to supplementary materials. Readers wishing to explore the more technical or theoretical aspects of AI are encouraged to consult the list of suggested reading and resources in Supplementary Material 1, as well as the broader AI literature.

2 | Translating Formal Definitions to Practical Understanding

AI is not a new concept. The term was introduced by John McCarthy at the seminal AI conference held at Dartmouth College in 1956 (Mitchell 2020). It was at this event that McCarthy and other key pioneering researchers formally established AI as a distinct field of research, laying the foundation for decades of technological advancement (Thompson 2025), but without consensus on a definition.

As previously discussed, the EU AI Act introduced a formal definition that is increasingly recognised as a gold-standard within the AI research community. However, while technically robust, this definition may feel abstract for clinicians unfamiliar with AI terminology. We propose a clinician-focussed explanation of AI that unpacks the key components of the EU's definition, rephrases them in more accessible terms and adds a clear clinical perspective.

2.1 | A Clinician-Focussed Explanation of AI

AI encompasses a set of evolving technologies (*AI techniques*) that are used to build systems designed to simulate human-like abilities (*AI capabilities*). AI systems take in information (e.g., test scores, speech and language samples), process it and produce an output. This output might be a prediction (e.g., how likely a child is to make progress with a therapy approach), a recommendation (e.g., which activities might be most helpful), new content (e.g., personalised therapy materials), or a decision (e.g., whether to adjust the difficulty of a task). In speech and language therapy, such outputs can be translated into *clinical applications* spanning the full scope of practice, from administrative tasks, to assessment, diagnosis, therapy planning and therapy implementation.

There are three broad types of AI systems: (i) fixed rule-based systems that follow human-set instructions and behave in the same way in every scenario; (ii) more flexible data-driven systems (i.e., ML systems) that learn patterns from data and, over time, may adapt to new situations and (iii) hybrid systems that use a combination of both. AI systems also differ in (i) the level of independence once deployed, with some always requiring human input and others making adjustments independently and (ii) the capacity to generate new content. In SLT, AI can be integrated into digital tools (e.g., therapy apps that adjust task difficulty in real time) or physical tools (e.g., robots that interact with children during therapy sessions).

In essence, AI is not a single, uniform technology. It exists along a spectrum of AI techniques, from traditional, rule-based methods to advanced, data-driven approaches. Many AI systems combine elements from across this spectrum to deliver the AI capabilities required for a specific clinical application. The concepts introduced in this tutorial are explored in greater detail throughout the paper. Of note, these concepts are not fixed or mutually exclusive. For instance, clinical applications often draw on multiple overlapping capabilities, each of which may be supported by different AI techniques; a dynamic explored in more detail later.

3 | Interpreting AI Taxonomies From a Clinician Lens

The EU AI taxonomy (European Institute of Innovation and Technology 2020) supports researchers and policy makers by categorising AI across five key dimensions: (i) *industries*—where AI is applied; (ii) *enterprise functions*—the role of AI in different departments within a business, such as marketing and sales; (iii) *locations*; (iv) *AI capabilities*—what AI can do; and (v) *enabling technology types*—technologies that AI depends on, such as data storage. While this structure offers insight into AI capabilities, it does not clarify the clinical relevance of these capabilities or explain the techniques that underpin AI.

The AAAI-20 Keyword classification system (AAAI, n.d.) also provides a valuable structure for AI researchers by organising AI topics under 20 keywords (e.g., ML, Natural Language Processing), each with its own list of subtopics. However, it assumes familiarity with AI terminology, presenting subtopics in alphabetical order without a clear organisational structure (e.g., subheadings that distinguish between the methods used to train ML systems from the tasks they perform). This lack of structure limits the classification's value for non-technical audiences.

3.1 | A Clinician-Oriented Classification Structure

To build a bridge between the expert AI and SLT communities, we present a practical classification structure informed by established AI taxonomies. This structure reorganises AI concepts into two core levels and introduces a third level specific to clinical practice:

- **Level 1: AI techniques** (*how* AI works)
- **Level 2: AI capabilities** (*what* AI can do)
- **Level 3: Clinical applications** (*how* AI can be used in clinical practice).

We use human and clinical analogies to explain AI concepts in Levels 1 and 2. In Level 3, we discuss existing and potential clinical applications of AI to illustrate how AI has been, or could be, integrated into SLT practice. Of note, these levels reflect a logical dependency, where AI techniques power AI capabilities, which in turn enable clinical applications. This relationship is illustrated in Section 7 using a production line model to demonstrate how these components can interact within an AI-driven clinical tool.

4 | Level 1: AI Techniques (how AI works)

AI techniques refer to the core technologies that underpin AI systems. A widely accepted distinction in the AI literature (e.g., Russell and Norvig 2021) divides AI techniques into two core categories: (i) *symbolic AI* (rule-based AI) and (ii) *subsymbolic AI* (data-driven AI). For clarity, we refer to symbolic AI as *traditional AI* and subsymbolic AI as *ML* throughout this paper.

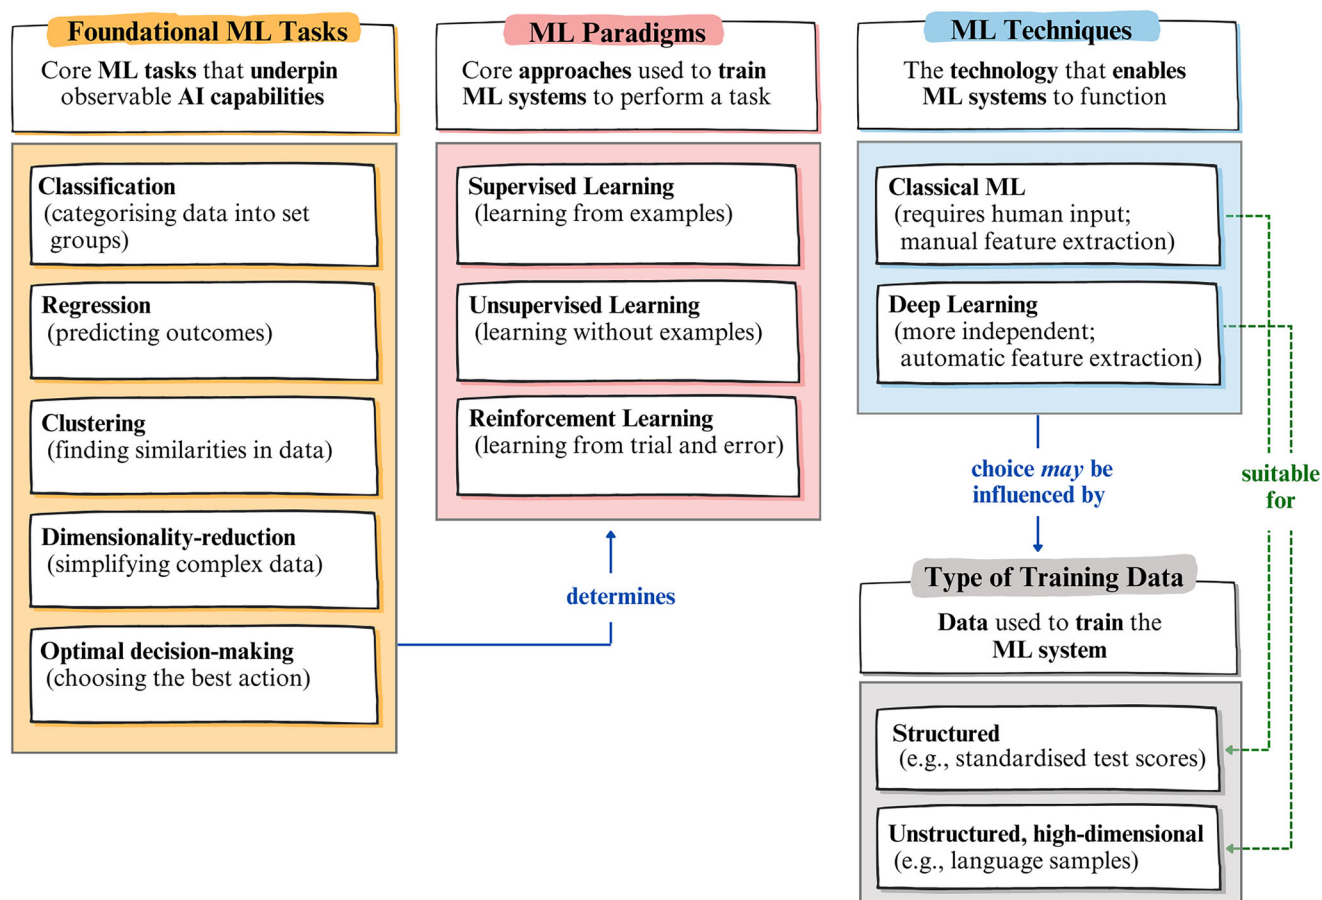


FIGURE 1 | Simplified overview of ML, highlighting the relationship between ML paradigms, foundational ML tasks, ML techniques and types of training data.

4.1 | Traditional AI (Symbolic AI): The Rule Follower

Traditional AI is “The rule-follower”: it does exactly what it is programmed to do (Bernard, 2021). These systems rely on clear, pre-defined instructions such as “If X happens, do Y” (Wooldridge 2020) (e.g., a therapy app that follows the rule “If the child selects the wrong answer, prompt them to try again”). This approach mirrors explicit grammar teaching, where an SLT explains a rule (e.g., “If something already happened, add -ed”) and the child applies it (Finestack 2018).

Traditional AI is effective for predictable, easy to define tasks. However, it cannot adapt or generalise beyond its programmed rules, making it less useful in ambiguous or unpredictable contexts such as natural communication (Russell and Norvig 2021; Mitchell 2020). This is where ML offers a powerful alternative.

4.2 | ML (Subsymbolic AI): The Learner

The key premise of ML is simple: when a concept is too complex to define, teach it by showing (Kubat 2021). For instance, while trying to explain what a person’s voice sounds like would be challenging, playing a few audio recordings would be much more effective. This makes ML “The Learner”: it analyses large

volumes of data, identifies patterns and applies that knowledge to new, unseen data (Russell and Norvig 2021). This mirrors implicit grammar intervention, where children are exposed to grammatical structures in conversation and language-based tasks and gradually internalise and generalise them without direct instruction (Finestack 2018).

The AI community recognises numerous subfields and technical distinctions within ML. However, this level of detail is beyond the scope of this paper. Instead, we build on “The Learner” concept and explain ML through three practical elements: (i) the core tasks it performs (i.e., *foundational ML tasks* that form the basis of broader AI capabilities); (ii) the ways it is trained to perform those tasks (i.e., *ML paradigms*) and (iii) the AI techniques that enable the learning (i.e., the *ML techniques*) (see Figure 1). For a more detailed technical overview of ML terminology tailored to SLTs, see Cordella et al. (2024).

4.2.1 | Foundational ML Tasks

Russell and Norvig (2021) outline common types of tasks (also referred to as *learning problems*) that ML systems are designed to perform. These tasks (illustrated with SLT-related examples in Figure 2) include: (i) classifying data (i.e., *classification*); (ii) predicting outcomes (i.e., *regression*); (iii) grouping data (i.e., *clustering*); (iv) simplifying data (i.e., *dimensionality reduction*) and

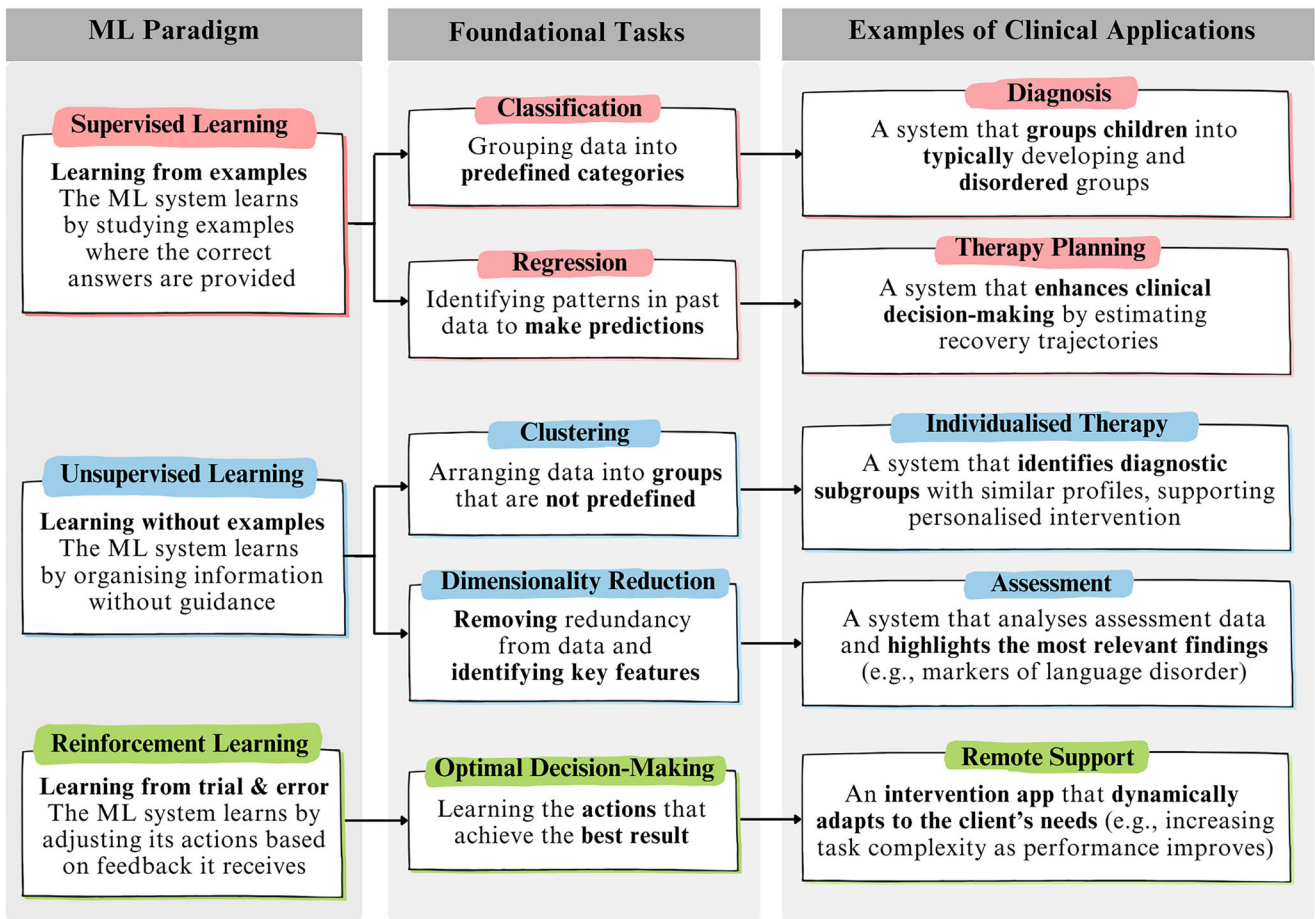


FIGURE 2 | Applied mapping of *ML paradigms* to the foundational *ML tasks* they can perform, with examples of *clinical applications*. It builds on Figure 1 by offering more detail to reinforce understanding and clinical relevance.

(v) choosing the best action (i.e., *optimal decision-making*).¹ These foundational tasks underpin all observable AI behaviours in real-world applications (i.e., AI capabilities such as computer vision), either individually or in combination, depending on what the AI system is designed to do.

4.2.2 | ML Paradigms: How ML Systems Are Trained

Before an ML system can be used in practice, it must undergo a training phase in which it is provided with a large dataset (referred to as *training data*) from which it learns to complete a specific ML task. The quality, quantity and diversity of this data significantly impact how well the system performs and how effectively it generalises to new, unseen data (Kubat 2021); a point we return to in Section 8.

There are three main ways to train machines: *supervised learning*, *unsupervised learning* and *reinforcement learning*. These are referred to in the literature as *forms of learning* (e.g., Russell and Norvig 2021), or *learning paradigms* (e.g., Bernard 2021). Albeit not a universally agreed term, we adopt the latter for its intuitive framing of how ML systems learn from data. The choice of *learning paradigm* depends on the task to be performed (Bernard 2021). Figure 2 illustrates common mappings between these ML paradigms and ML tasks.

Supervised learning (learning from examples): Here, the machine learns to classify data from examples, where both the input data (i.e., the data fed into the machine), and the correct output for each input (i.e., the expected answer) are provided. This is referred to as *labelled data* (Russell and Norvig 2021). This would be like a child learning vocabulary by looking at multiple images of cats and dogs, each clearly labelled (e.g., “These are cats. These are dogs”). Over time, the child would learn to recognise the distinguishing features (e.g., cats have whiskers) and begin to generalise this knowledge to identify new, unlabelled cats and dogs.

Unsupervised learning (learning without examples): In unsupervised learning, the machine is provided with input data, but no examples of expected answers to guide its learning (referred to as *unlabelled data*) (Russell and Norvig 2021). The machine independently identifies patterns in the data, potentially uncovering unexpected relationships or groupings (Kubat 2021). This would be like giving a child a mixed set of pictures of cats and dogs, without labels, and asking them to sort the images in a way that makes sense to them. A child might group animals by size, by colour or by other unanticipated features.

Reinforcement learning (learning from trial and error): Reinforcement learning is used in dynamic or interactive scenarios where the goal is to discover the most effective strategy for

TABLE 1 | Key differences between classical ML and deep learning, with relevance to SLT practice.

Distinguishing category	Classical ML	Deep (machine) learning
Complexity of algorithms	Simple algorithms: Statistical models or shallow artificial neural networks (see Supplementary Material 2 for an explanation of artificial neural networks).	Complex algorithms: Large and multilayered (deep) artificial neural networks.
Data feature extraction	Manual: System relies on humans to define relevant data features.	Automatic: System independently identifies relevant data features.
Suitable data type	Structured: e.g., numerical test scores, demographic data (e.g., age, sex).	Unstructured: e.g., voice recordings and visual data such as observations of non-verbal communication.
Generative capabilities	No: Analyses data only.	Yes: Can analyse data and create new content (e.g., text).
Example of clinical application	Analysis of numerical test scores such as standardised assessment scores.	Transcription of ecologically valid data such as speech and language samples.

achieving a desired outcome. The machine performs an action, receives feedback in the form of *rewards* (positive reinforcement) or *penalties* (negative reinforcement), and gradually identifies which actions lead to better results (Russell and Norvig 2021). This mirrors how a child learns to use fluent speech in the Lidcombe Program; receiving *positive verbal contingencies* for stutter-free speech and *corrective contingencies* for moments of stuttering (Onslow et al. 2017). With repeated feedback, the child learns which speaking behaviours increase fluency.

4.2.3 | ML Techniques

Once a ML paradigm is selected, the final step is to choose the specific ML technique to power the AI system. Broadly speaking, ML techniques fall into two main categories based on complexity: *classical ML* and *deep learning*. Established taxonomies (e.g., AAAI-20) refer to classical ML simply as ML, only explicitly labelling deep learning. However, to people new to AI, this may suggest that deep learning exists outside of ML. Therefore, to avoid potential confusion for the lay reader, we adopt the terms *classical ML* and *deep learning*, also found in some literature (e.g., Bernard 2021; IBM 2021). The type of training data available may influence the choice of ML technique (and vice versa).

Classical ML consists of relatively simple algorithms and relies on human experts to specify which elements of the data (i.e., features) the system should use to perform a task (i.e., *manual feature extraction*) (Dargan et al. 2020). These techniques are well suited to structured data that is easy to define and organise (e.g., numerical or categorical data such as standardised scores and demographic information). Hence, classical ML is a suitable choice for a tool designed to predict language disorders based on age and test scores.

Deep learning, in contrast, uses more complex algorithms that learn from large sets of data with limited human input, automatically identifying relevant features (i.e., *automatic feature extraction*) (Mitchell 2020). This makes deep learning ideal for handling unstructured, high-dimensional data that is “messy”

and difficult to quantify (e.g., speech recordings, language samples, or informal clinical observations). For instance, a system built to transcribe or analyse speech must interpret prosody, intelligibility, articulation and other nuanced features that are difficult for a human programmer to define, but can be learned directly from raw audio using deep learning (see Supplementary Material 2 for a more detailed overview of classical ML and deep learning).

Accordingly, deep learning is particularly valuable in SLT, where much of the clinically meaningful data, such as natural speech and language samples, is complex, multi-dimensional and unstructured. Another distinctive feature of deep learning is its ability to create entirely new content (i.e., *Generative AI*), such as personalised therapy materials. See table 1 for a summary of the key distinctions between classical ML and deep learning.

4.2.4 | Interpretability: How ML Systems Explain Their Decisions

A key limitation of ML is the difficulty in understanding exactly how the ML system generates its outputs; referred to as the *black box problem*. This lack of *interpretability* (i.e., the ability to inspect how the system works) and/or *explainability* (i.e., the ability to understand why a specific decision was made) (Russell and Norvig 2021) poses a major barrier in clinical contexts, where in line with the Health and Care Professions Council (2024), decisions must be transparent, defensible and accountable.

This challenge is pronounced in deep learning models, whose multiple layers of mathematical computations and automatic feature extraction make their decision process difficult to interpret (Champendal et al., 2023). Interpretability and explainability tends to be less of an issue in classical ML, particularly when the system is trained using supervised learning and structured data. In these cases, the model’s decision-making process is often based on relatively simple and traceable logic (Murdoch et al. 2019). To address the black box problem, researchers are developing:

TABLE 2 | Alignment between EU AI taxonomy of capabilities and reframed categories for speech and language therapy.

AI Capability (EU Taxonomy)	AI Capability (SLT Perspective)
Interaction intelligence ▪ Computer vision ▪ Computer audition ▪ Computer linguistics	Perception & communication
Analytic intelligence ▪ Forecasting ▪ Discovery ▪ Planning & search	Clinical reasoning
Motion & creative Intelligence ▪ Robotics & control ▪ Creativity	Movement Creativity

- i. **Hybrid systems** that combine the strengths of deep learning with the transparency of more interpretable systems (for a review of their use in clinical decision systems, see Kierner et al. 2023).
- ii. **Explainable AI (XAI)** techniques that aim to make the decision-making process of *ML* systems more interpretable (see Dwivedi et al. 2023). At present, this remains an evolving field, particularly in relation to deep learning (Hosain et al. 2024).

5 | Level 2: AI Capabilities (What AI Can Do)

Next, we turn to AI capabilities, that is, what AI can do. As previously noted, these observable AI behaviours are grounded in the foundational *ML* tasks discussed in section 4.2.1. Consider ChatGPT, for instance, that appears to engage in a genuine conversation with its users. Its “conversational ability” results from training the AI system to predict the next word in a sequence of words, based on patterns learned from vast datasets of labelled examples. In other words, ChatGPT’s apparent human-like communication is simply a machine executing a prediction task (i.e., regression) using supervised learning (Briganti 2024).² This point warrants an important clarification: despite the term “intelligence,” AI systems do not possess intelligence in the human sense. They cannot think, understand, or reason. Instead, they mimic aspects of intelligence through mathematical computations (Mueller et al. 2025).

The EU AI taxonomy classifies AI capabilities into three overarching categories: (i) *interaction intelligence* (i.e., computer vision, computer audition and computer linguistics); (ii) *analytic intelligence* (i.e., forecasting, discovery and planning and search) and (iii) *motion and creative intelligence* (i.e., robotics and control and creativity). In this paper, we retain the underlying EU concepts but reframe them using terminology and examples that are more intuitive for SLTs, enabling clearer links between AI capabilities and clinical applications (see Table 2):

- i. **Clinical reasoning:** Reasoning, problem-solving, planning and making predictions and decisions (e.g., supporting therapy planning, or suggesting diagnoses).

- ii. **Perception and communication:** Interpreting visual, auditory, or other sensory input (e.g., perceiving images, gestures, tone or facial expressions), and engaging in natural language tasks (e.g., producing, interpreting or transcribing natural speech and language).
- iii. **Movement:** Physically interacting with the environment, including navigating space, manipulating objects, or using gestures to communicate. While this may overlap with non-verbal aspects of perception and communication, we present it as a separate domain to reflect the broader and growing role of robotics in SLT and educational contexts (e.g., social robots or robotic tutors).
- iv. **Creativity:** Creating content (e.g., making personalised therapeutic resources).

Each domain is further specified in Table 3 into defined clinical abilities that we have directly mapped to corresponding AI terminology, associated functions, and SLT applications.

Most AI capabilities can, in principle, be supported by any of the three core AI techniques (traditional AI, classical *ML*, or deep learning), individually or in combination. In practice, however, newer AI techniques generally achieve greater effectiveness, particularly for capabilities involving perception and communication. For example, while computer vision was once driven by traditional AI and classical *ML*, its performance has advanced dramatically with the introduction of deep learning (Russell and Norvig 2021).

6 | Level 3: Clinical Applications

Finally, we move to the third level, clinical applications, which we categorise into key areas of SLT practice, namely (i) *assessment*; (ii) *diagnosis*; (iii) *therapy planning*; (iv) *therapy implementation* and (v) *administrative duties*. For each area, we present a combination of tested, proposed and/or conceptual applications. The limitations of these clinical applications are addressed in Section 8. Of note, where we use the term “could”, we are referring to conceptual applications suggested by the authors, that are considered technically feasible. Our aim is not to systematically or critically appraise existing research, but rather to highlight key insights from the literature and propose possible directions for SLT innovation.

6.1 | Assessment

AI-powered assessment tools offer scalable, efficient and accessible solutions that facilitate large-scale screening, remote evaluations and automated data analysis. While clinician expertise remains invaluable, these technologies can reduce clinician workload and make ecologically valid assessments more feasible.

One particularly promising application is automated assessment, where AI-powered tools *could* support large-scale screening of communication difficulties. In Ireland and the UK, national health services offer free hearing and vision screenings for children in their first year of formal education, conducted by school nurses, to detect impairments that can have lasting consequences if untreated. Despite DLD affecting approximately 8% of children (i.e., two children per classroom of 30) and having significant long-term impacts (Norbury et al., 2016), there is no equivalent national screening programme for language difficulties. Scaling language assessment is arguably more challenging than vision and hearing screenings due to the specialist skills required. However, AI-powered, gamified apps may help bridge this gap by enabling teachers or nurses to conduct large-scale screenings with minimal training. These tools *could* assess communication skills, flag risk factors and promote child engagement through interactive tasks.

Technology-assisted large-scale screening has already shown promise, particularly for receptive language, where children can show understanding through simple, structured responses (e.g., pointing to an image or animation). For instance, Frizelle et al. (2023) demonstrated that children as young as five could independently complete the Test of Complex Syntax–Electronic (TECS-E) (powered by traditional AI), with results comparable to clinician-supported administration. The tool was later norm-referenced for classroom use, with up to 25 children completing it simultaneously using tablets and headphones (Frizelle et al. 2025).

Assessing expressive language, on the other hand, requires more advanced methods. Speech recognition can serve as a foundational step by converting children's spoken language into text. In fact, emerging evidence suggests that speech recognition (refer to Table 3 for a reminder of AI capabilities) can produce more accurate transcriptions than clinician-produced transcriptions and is less affected by speech rate (Fox et al. 2021). Transcriptions can be subsequently analysed using Natural Language Processing (NLP), which has shown potential in evaluating syntax, vocabulary, discourse structure and linguistic complexity (Al-Ali et al. 2024). Of note, recent studies indicate that NLP systems powered by deep learning can automate the scoring of narrative macrostructure (Jones et al. 2019) and sentence formulation tasks (Wang et al. 2020). Ongoing research is also exploring the use of ML-driven NLP for early childhood language screening (e.g., Zhang et al. 2020).

AI *could* also support more ecologically valid assessment practices such as Language Sample Analysis (LSA), which can be underused due to time demands, scoring variability and clinician training requirements (Lammert et al. 2024). ML-powered transcription and analysis *could* make LSA more viable in routine practice, particularly for less experienced clinicians.

Future AI implementation efforts *could* involve web-based platforms, provided by national health services, that leverage ML-driven speech recognition, NLP and data mining to support early identification of communication difficulties. In this case, caregivers *could* upload audio or video samples collected via smartphones or AI-powered monitoring devices such as the Language Environment Analysis (LENA) system (see Gilkerson and Richards 2020)³.

Beyond linguistic analysis, integrating computer vision and affective computing capabilities (underpinned by deep learning), into gamified screening apps, for example, *could* offer a more holistic view of a child's communication profile, capturing non-verbal aspects such as facial expressions, gestures and tone of voice. At present, computer vision is being explored in instrumental assessments, including automated imaging for videofluoroscopy (Girardi et al. 2023) and velopharyngeal competence analysis (Corneford et al. 2024).

6.2 | Diagnosis

Once assessment data is collected, whether remotely or in-person, it can be processed by ML systems that leverage data mining capabilities to support diagnostic decisions. The choice of ML paradigm (i.e., supervised vs. unsupervised learning) is particularly relevant here, as it shapes the level of diagnostic detail achieved. Supervised learning can be used to train ML systems to classify children as having, or not having, a speech and language disorder (e.g., Justice et al. 2019; Lammert et al. 2024). Unsupervised learning, in contrast, can be employed to uncover hidden patterns in the assessment data, clustering children into diagnostic subgroups with shared characteristics (e.g., Stevens et al. 2019). The latter approach enables a more nuanced understanding of disorder profiles, which *could* guide more client-specific interventions.

6.3 | Therapy Planning

Assessment and diagnostic data *could* inform AI-driven therapy planning, with forecasting and planning capabilities playing a key role. Traditional AI systems (e.g., the SPELTA system that recommends therapy exercises based on predefined rules; Robles-Bykbaev et al. 2016) have demonstrated how rule-based models can generate consistent and transparent therapy recommendations. However, ML systems offer greater potential for personalisation and adaptability.

ML-driven forecasting *could* support clinical decision-making by estimating recovery trajectories (e.g., Cordella et al. 2024), supporting decisions such as waitlist prioritisation. ML-powered planning and routing *could* propose therapy goals tailored to a child's individual needs (ranging from broad, long-term objectives to specific session targets).

Additional ML-driven capabilities, such as NLP and data mining, *could* further refine dynamic and personalised therapy planning by analysing therapy session data, tracking progress and identifying subtle patterns in a child's response to intervention. Beyond analysis, Generative AI *could* contribute directly to content

TABLE 3 | Human capabilities mapped onto AI terminology.

Domain of human ability	Clinical ability	Corresponding AI terminology	Examples of specific functions	Examples of clinical applications
Perception & communication	Interpret visual information	Computer vision: AI that perceives and interprets visual information	- Image classification - Object detection - Visual tracking - Semantic segmentation (i.e., understanding the relationship between two objects (Chai et al. 2021)) - Image restoration	- Facial recognition - Assistive technology for the visually impaired
	Perceive and respond to emotional cues	Affective computing: AI that detects and reacts to human emotions	- Emotion recognition - Sentiment analysis	- Therapy bots - Emotion-aware bots
	Transcribe speech	Speech recognition: AI that converts speech to text	- Transcription - Voice command recognition	- Transcription services - Voice assistants (e.g., Siri)
	Use speech	Speech synthesis: AI that artificially produces human-like speech	- Conversion of text to speech	- AI-generated voices for AAC devices or virtual therapy assistants
	Interpret and use language	Natural language processing (NLP): AI that interprets, understands and communicates using human language	- Language modelling - Question answering - Text summarisation - Text generation - Information retrieval - Language analysis (e.g., lexical, semantic, syntactic and/or pragmatic analysis)	- Chatbots/Agents/Virtual therapy assistants - Lexical, semantic, syntactic and/or pragmatic analysis of language samples
Motion	Move and use gestures	Robotics: AI that physically moves and interacts with the surrounding physical environment	- Non-verbal communication (i.e., gestures, sign language) - Object manipulation - Physical interaction with surrounding environment	- Use of robots for sign language or social interaction with Autistic children
Clinical reasoning	Foresee	Forecasting: AI that predicts past, present or future outcomes based on past data	- Predicting future outcomes	- Prediction of therapy progress
	Investigate	Data mining/discovery: AI that analyses large datasets and uncovers patterns in data	- Clustering - Anomaly detection - Uncovering patterns in assessment data	- Diagnosis - Identification of diagnostic subgroups - Automated analysis of assessment data
	Make informed decisions	Planning, routing & scheduling: AI that supports decision making by suggesting a series of actions (<i>planning</i>), finding the most efficient path (<i>routing</i>), and/or coordinating resources (<i>scheduling</i>)	- Decision optimisation - Planning/recommendations - Game playing	- Personalised therapy plans based on assessment data - Interactive therapy apps or virtual assistants that adapt to the child's level

(Continues)

TABLE 3 | (Continued)

Domain of human ability	Clinical ability	Corresponding AI terminology	Examples of specific functions	Examples of clinical applications
Creativity	Innovate	Generative AI: AI that generates entirely new content such as text, images, video, and audio	- Generate new content	- Therapeutic resources generated by AI-powered tools such as ChatGPT (for text generation), DALL-E (for image generation), and/or Jukebox (for music generation) - Clinical documentation drafted using AI-powered tools such as Heidi.

creation, such as formulating personalised therapy goals in real-time. Rakap and Balikci (2024) explored this in the context of IEP goal development for pre-schoolers with autism and found that ChatGPT significantly improved the quality of goals written by special education teachers.

6.4 | Therapy Implementation

ML-supported speech and language therapy *could* help therapists to generate tailored communications, therapy ideas and materials, adapt interventions in real time and extend support beyond the clinic through remote and adaptable AI-powered tools.

6.4.1 | Therapy Individualisation and Improved Efficiency

Generative AI *could* assist SLTs in efficiently personalising therapy sessions by brainstorming innovative therapeutic tasks that are tailored to the client's interests and needs (see example from an educational context in Fan et al. (2024)). It *could* also generate customised materials on demand, reducing preparation time and helping to maximise within-session dosage. The latter is an important factor in intervention efficiency, particularly in the face of growing SLT waiting lists. In May 2024, over 75,000 children in England were waiting for speech and language therapy, representing an unprecedented high (Kirby 2024).

Dosage control *could* be achieved through careful prompt engineering,⁴ whereby therapists *could* create AI-generated resources by specifying a child's interests, age, ability level and therapy targets, as well as the required number of repetitions of each target. For instance, a dinosaur enthusiast might receive a personalised storybook featuring themselves and their friends, embedding phonological, syntactic, or narrative targets at a controlled frequency that supports optimal exposure. This aligns with Frizelle et al. (2021), who highlight the importance of sessions rich in “active ingredients” (e.g., varied visual representations and linguistic input, repeated exposure to therapy targets, etc.) to improve therapy efficiency.

Generative AI tools *could* also support real-time therapy, enabling interactive and immersive experiences. For example, a child

practising specific speech or language targets at sentence level *could* use image generation tools (e.g., DALL-E) to illustrate their sentence, reinforcing the connection between language and meaning. For narrative macrostructure targets, AI video generators (e.g., Synthesia or HeyGen) *could* bring a child's spoken or written story to life. ML-driven speech synthesis tools *could* offer another layer of engagement, converting the child's written text into audio with varied voices, making practice more playful and motivating.

6.4.2 | AI as a Therapeutic Tool

SLTs *could* help young people with DLD build the skills to use Generative AI systems, like LLMs, as a practical, long-term support tool for navigating both academic and everyday language demands.

For example, ChatGPT *could* serve as a personalised learning assistant, simplifying complex texts, rephrasing difficult concepts, or generating accessible summaries. With SLT guidance, a child or teenager *could* learn to ask targeted questions to clarify vocabulary, extract key ideas from reading assignments, or organise their written work more effectively.

Generative AI *could* also promote independence in communication. Young people *could* use ChatGPT to draft text messages, organise thoughts, or prepare for social interactions through simulated conversations. Speech recognition capabilities may further support those with written language difficulties, allowing them to dictate ideas that ChatGPT can convert into well-structured text.

6.4.3 | Accessible, Adaptable and Remote Support

AI systems, such as therapeutic apps powered by deep learning, offer promising ways to extend therapy outside the clinic. These tools should never replace therapists, but they *could* offer a practical bridge in care, ensuring that children receive meaningful therapeutic support when direct SLT intervention is unavailable. Traditionally, during therapy gaps (whether between sessions, across therapy blocks, or while awaiting services) parents/caregivers and educators rely on home programs with

activity suggestions. However, adapting these tasks to a child's level can be difficult. Even with in-clinic demonstrations of how to adjust tasks, generalising this knowledge without specialist training is often challenging.

ML-driven therapy apps with chatbot functionality, especially those integrating speech recognition, NLP, speech synthesis, planning and Generative AI, *could* help address this gap. These tools *could* adjust task difficulty in real time and provide immediate feedback based on the child's performance. With additional capabilities such as computer vision and affective computing, such systems *could* interpret emotional cues and offer encouragement when frustration is detected, simplify tasks when needed, or increase complexity when a child is ready for more.

To date, research has largely focussed on robots designed to support social communication in autistic children (e.g., Marino et al. 2020; Doğan and Çolak 2024; Fachantidis et al. 2020). However, most rely on traditional AI, limiting adaptability, and their high cost restricts use in home settings. Therapy apps may offer a more scalable and cost-effective alternative. While robots benefit from multimodal interaction (e.g., gestures, eye-gaze tracking, etc.) to create natural and immersive experiences, therapy apps *could* simulate similar effects through animated avatars, emoji-based feedback, highlighted text, speech synthesis and interactive animations. These features may enhance engagement, reinforce learning and offer responsive support.

Belpaeme et al. (2018) outline key principles for designing social robots as second-language tutors, many of which can inform the design of SLT therapy apps/chatbots. These include: (i) presenting the AI tool as a peer-like companion rather than an instructor, to boost child engagement and (ii) incorporating personalisation and emotionally responsive interaction.

6.4.4 | Personalised and Multilingual Communication

Generative AI *could* enhance communication with parents and caregivers by adapting information to different literacy levels and language backgrounds. Many parents struggle with dense written reports, making it difficult to fully understand their child's needs and progress. Generative AI *could* improve information accessibility by: (i) simplifying complex information into parent-friendly summaries using plain-language, images or videos; (ii) integrating visuals into appointment notifications to reduce reliance on text and (iii) generating short, narrated videos to ensure that key information remains accessible to parents with low literacy or visual impairments. For instance, tools such as Nuance's DAX Copilot, used in healthcare to generate multilingual patient instructions, *could* be adapted for SLT.

6.5 | Administrative Duties

AI-powered tools have the potential to streamline administrative tasks allowing clinicians to dedicate more time to direct client care (Farmer et al. 2025). Some of these tools have already demonstrated success in healthcare, suggesting their potential for SLT.

6.5.1 | Automated Documentation and Report Generation

Deep learning-driven speech recognition (to transcribe voice notes) and NLP capabilities *could* assist SLTs in generating written progress reports, summarising therapy sessions, and drafting therapy notes. In healthcare, AI tools such as Microsoft's Dragon Copilot and Heidi are already automating clinical documentation.

6.5.2 | Caseload Management

Tools like Ambient AI Scribes, have been implemented in healthcare, to automate appointment management and reduce administrative burdens. Similar AI systems *could* be introduced to SLT settings to manage scheduling, automate reminders and streamline follow-up communications, helping to reduce missed appointments and improve service coordination. ML-powered virtual assistants/chatbots *could* handle routine inquiries from parents, such as explaining appointment details or providing general therapy-related information. This technology is already well established. For example, Salesforce's "Agentforce" platform is used by organisations like Heathrow Airport to automate and personalise customer support.

7 | Connecting the Three Levels: A Practical Illustration

To consolidate the concepts discussed throughout this paper, we present a conceptual production line model that illustrates the relationships among the three levels of our AI classification. The model presents the three levels in reverse order to reflect the real-life process of designing AI tools for clinical use, starting with identifying a clinical need and working backwards to determine appropriate AI techniques. It also highlights the flexible and dynamic interplay between levels, where a single clinical application may require multiple AI capabilities, and each capability may be supported by different AI techniques (see core structure of the model in figure 3).

In Figure 4, we introduce two additional layers to support practical application of the model: (i) *tool objective*—to clarify the clinical function of the AI tool (e.g., generate therapeutic materials) and (ii) *specific requirements*—to detail the step-by-step tasks the system must perform so that each one can be clearly mapped to the relevant AI Capabilities.

Figure 5 applies the extended model to a hypothetical AI tool designed to identify children with DLD from spoken language samples. The example shows how hybrid systems can combine different AI techniques depending on the clinical priority (e.g., speed versus transparent decision-making). The tool performs three main tasks: (i) transcription; (ii) linguistic analysis and (iii) classification of the analysed sample (as indicative of DLD or not). Deep learning is suitable for the first two tasks, where the priority is efficient and accurate processing of complex, unstructured and high-dimensional data. For the final classification step, however, classical ML may be preferable. It offers greater transparency and interpretability, allowing specific linguistic markers (e.g.,




FIGURE 3 | Core structure of the clinician-oriented classification model, outlining the three interrelated levels.

number and type of syntactic errors) to be directly linked to diagnostic outcomes, enabling SLTs to understand and justify the system's diagnostic decision.

8 | Practical, Ethical and Regulatory Considerations

SLTs should view AI-driven applications as supports for clinical judgment; not substitutes. Their integration into clinical practice, particularly those based on ML, raises practical, regulatory and ethical questions that require careful consideration.

Can I trust the system's response? As noted earlier, it can be difficult to interpret how a ML-driven system produces its output. The level of interpretability should align with the tool's purpose. For high-impact applications (e.g., diagnosis or therapy planning), SLTs should prioritise interpretable approaches, such as traditional AI, classical ML algorithms trained with supervised learning, or hybrid systems that balance performance and transparency. For lower-impact applications (e.g., generating therapy materials), clinicians may accept the use of less interpretable systems, provided clinical judgment is retained.

Special caution is required when using Generative AI tools like ChatGPT to source information or support clinical decisions. These tools do not “understand”; they predict the next likely word and may produce convincing but inaccurate responses (i.e., hallucinations). SLTs must cross-check outputs against trusted sources and apply professional judgment.

Was the right data used to train the system? SLTs must critically assess whether the data used to train the system reflects both the *client population* and the *clinical context*. The former is a particular challenge for a relatively small field like SLT. Most available speech datasets reflect adult populations, overlooking the variability of paediatric speech (Gohider and Basir 2024), which limits the development of child-specific tools (Usha and Alex 2023). With regard to clinical context, existing evidence shows that AI systems can reproduce and amplify biases present in their training data, including negative stereotypes related to race and ethnicity (e.g., Ayoub et al. 2024). Specific to SLT, Lewis et al. (2025) identified cultural, disability and linguistic biases in AI-generated materials (e.g., defaulting to White-presenting characters, portraying disability negatively and privileging English). The authors caution that such biases risk perpetuating inequities in SLT, particularly for culturally and linguistically diverse children with disabilities, and highlight the importance of improving AI literacy among clinicians and




FIGURE 4 | A step-by-step guide to the extended production line model, detailing how to move from *tool objective* to the selection of *AI techniques*.

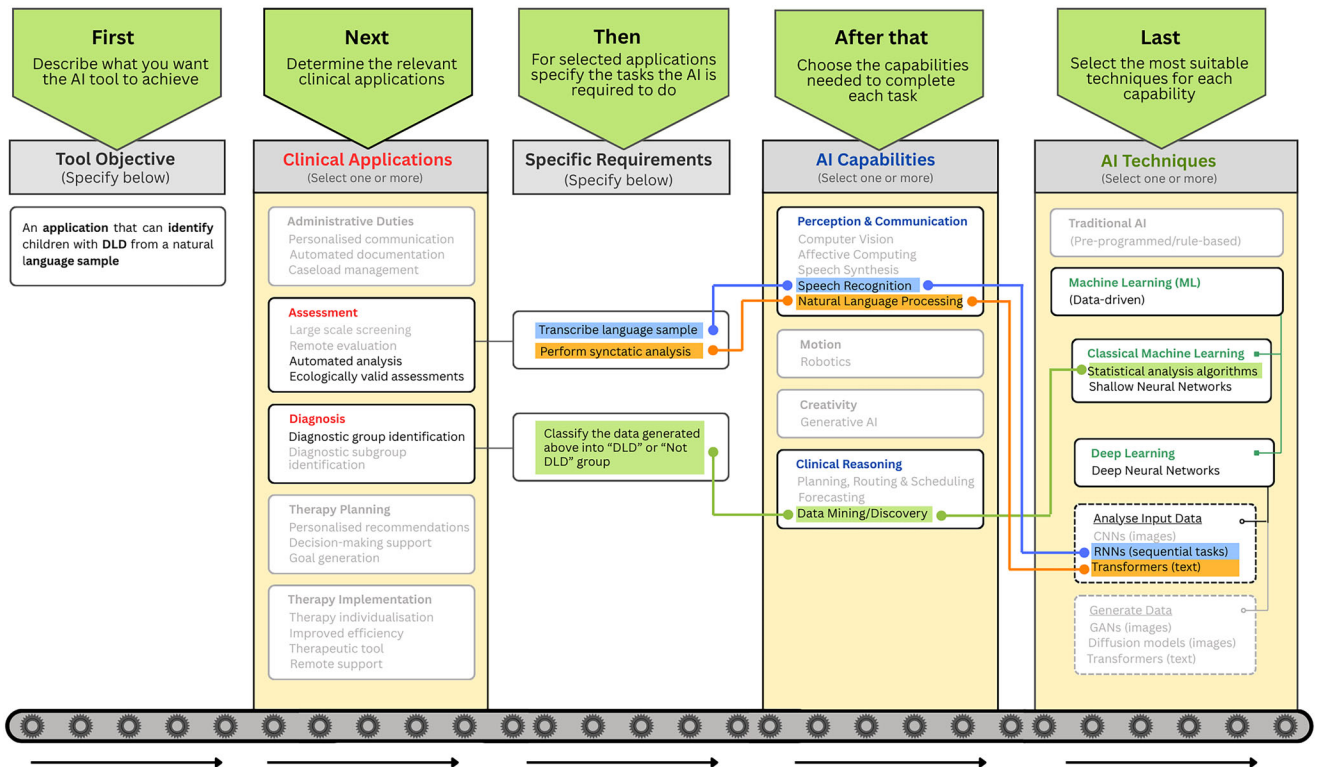


FIGURE 5 | Application of the production line model to a hypothetical AI tool. The tool's objective is to identify DLD from spoken language samples. Within *clinical applications*, the focus is on assessment and diagnosis. *Specific requirements* include transcribing and analysing the sample, followed by classifying the analysed sample. These tasks map to *AI capabilities* such as speech recognition, NLP and clinical reasoning. Appropriate *AI techniques* are then selected: *deep learning* would support efficient transcription and analysis, while *classical ML* could support more transparent, interpretable classification.

using structured protocols to evaluate AI-generated materials for bias.

With regard to clinical context, ML systems often underperform when used outside the context they were trained on. For instance, Tobin et al. (2024) found that speech recognition models trained on read speech were less accurate when applied to conversational speech in children with speech disorders. However, accuracy improved when the models were trained directly on conversational data.

Is there evidence to support the use of this tool? Few AI tools for SLT have undergone rigorous testing. A recent review by Deka et al. (2025) on the use of AI in speech interventions found limited evidence comparing AI-driven intervention to traditional interventions. Clinical uptake of AI should remain cautious until stronger evidence emerges.

How can I ensure ethical use of AI? Jobin et al. (2019) identify core principles for ethical AI, including: (i) *transparency and informed consent*—clients must understand and consent to the use of AI; (ii) *privacy and data protection*—systems must safeguard sensitive data and comply with GDPR; (iii) *beneficence*—AI must enhance care quality; (iv) *responsibility*—clinicians remain accountable for all decisions and (v) *justice*—systems must avoid perpetuating bias and support equitable access, especially in light of risks of digital exclusion. The environmental cost of training and deploying large-scale AI models must also be considered. LLMs, in particular, require substantial computational resources and energy consumption, raising concerns about sustainability and responsible use (Ji and Jiang 2025). Recent work on *green AI* explores potential mitigation strategies, including (i) the design of more energy-efficient AI systems (i.e., *green-in AI*) and (ii) the use of AI to support environmentally sustainable practices (i.e., *green-by AI*) (Bolón-Canedo et al. 2024).

What are the regulatory requirements? Healthcare and education applications of AI are classified as *high-risk* by the EU AI Act and must comply with strict standards of data quality, fairness, transparency and accountability (see Supplementary Material 3 for further information on the EU's classification of AI systems). This aligns with the broader aim of promoting *trust-worthy AI* (i.e., AI that is lawful, ethical and robust) (High-Level Expert Group on Artificial Intelligence 2019).

9 | A Final Thought: AI or IA?

Fears that “AI will replace humans” are widespread; however, many experts promote a vision of human-machine collaboration rather than replacement (e.g., Dégallier-Rochat et al. 2022). Interestingly, Doug Engelbart proposed renaming AI as Intelligence Augmentation (IA) to better reflect this goal (Russell and Norvig 2021). The former CEO of IBM, Ginni Rometty, echoes this viewpoint, noting that: “*These are technologies to augment human intelligence. In fact, I don't like the word AI...because AI says replacement of people...and that is not what we're talking about. By and large we see a world where this is a partnership between man and machine... It's a tool that helps, and it makes you as a professional do a better job.*” (Rometty 2017, 09:05). This perspective is crucial for clinicians: AI can streamline practice,

but decisions remain in the hands of SLTs to ensure service delivery is ethical, evidence-based, clinically appropriate and emotionally responsive.

10 | Conclusion

As AI becomes more embedded in paediatric speech and language therapy, SLTs need a clear understanding of what it is, what it can do and how it can support practice. This tutorial addresses this need by translating key AI concepts into accessible, clinically relevant knowledge for SLTs and organising them across three core levels: AI techniques (how AI works), AI capabilities (what AI does) and clinical applications.

Our focus has been on paediatrics; however, many of the AI techniques and principles discussed are also applicable to adult contexts. For instance, (i) ML paradigms such as unsupervised learning could be used to identify subgroups associated with different types of dementia; (ii) Generative AI could be used to create communication supports for adults with aphasia and (iii) ML-driven therapy apps could support intervention for adults with acquired language disorders.

AI's potential for enhancing clinical efficiency, personalisation and access to care must be balanced with careful consideration of challenges around interpretability, ethics, regulation and clinical rigour. Ultimately, AI is a tool to enhance clinical practice; it's safe and effective use will always depend on the expertise, judgement and oversight of clinicians.

Advancing the role of AI in speech and language therapy, particularly ML, depends on the availability of large, diverse and high-quality datasets that reflect real clinical contexts. SLTs, researchers and AI developers must collaborate to build these datasets and co-design tools that are ethical and clinically meaningful.

Funding

This research is part of a PhD funded by Irish Research Council (IRC). Project ID: GOIPG/2024/4514

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

Endnotes

¹ *Optimal decision-making* is not a term explicitly adopted by Russell and Norvig (2021). However, it aligns with their description of decision-making under uncertainty and provides an intuitive label for SLTs.

² This example is simplified for clarity. Large Language Models can be trained using multiple *AI paradigms and tasks*.

- ³ LENA has known limitations, such as difficulty distinguishing speech-related from non-speech related vocalisations (Cristia et al. 2021), but it demonstrates the potential of AI-enabled home monitoring systems.
- ⁴ Prompt engineering is the process of carefully composing the instructions (referred to as prompts) provided to a Generative AI system to elicit the best possible response. Clear, well-crafted prompts help the system better interpret the intended query, improving the accuracy and usefulness of its response (Meskó 2023).

References

- Aggarwal, K., R. Ravi, and K. Yerraguntla. 2025. "Artificial Intelligence and Patient Care: Perspectives of Audiologists and Speech-Language Pathologists." *Intelligence-Based Medicine* 11: 100214. <https://doi.org/10.1016/j.ibmed.2025.100214>.
- Al-Ali, H., E. Soares, G. Leal, et al. 2024. "Towards Automated Game-Based Early Screening for Language Disorder." In *Proceedings of the 2024 CLASP Conference on Multi-Modality and Interaction in Language Learning* 26–31. Association for Computational Linguistics.
- Association for the Advancement of Artificial Intelligence. n.d. AAAI-20 Keywords. AAAI. <https://aaai.org/conference/aaai/aaai-20/aaai20keywords/>.
- Ayoub, N. F., K. Balakrishnan, M. S. Ayoub, T. F. Barrett, A. P. David, and S. T. Gray. 2024. "Inherent Bias in Large Language Models: A Random Sampling Analysis." *MCP: Digital Health* 2, no. 2: 186–191. <https://doi.org/10.1016/j.mcpdig.2024.03.003>.
- Belpaeme, T., P. Vogt, R. van den Berghe, et al. 2018. "Guidelines for Designing Social Robots as Second Language Tutors." *International Journal of Social Robotics* 10, no. 3: 325341. <https://doi.org/10.1007/s12369-018-0467-6>.
- Bernard, E. 2021. *Introduction to Machine Learning*. Wolfram Media Inc.
- Bolón-Canedo, V., L. Morán-Fernández, B. Cancela, and A. Alonso-Betanzos. 2024. "A Review of Green Artificial Intelligence: Towards a More Sustainable Future." *Neurocomputing* 599: 128096. <https://doi.org/10.1016/j.neucom.2024.128096>.
- Briganti, G. 2024. "How ChatGPT Works: A Mini Review." *European Archives of Oto-Rhino-Laryngology* 281, no. 4: 1565–1569. <https://doi.org/10.1007/s00405-023-08337-7>.
- Chai, J., H. Zeng, A. Li, and E. W. Ngai. 2021. "Deep Learning in Computer Vision: A Critical Review of Emerging Techniques and Application Scenarios." *Machine Learning With Applications* 6: 100134. <https://doi.org/10.1016/j.mlwa.2021.100134>.
- Champendal, M., H. Müller, J. O. Prior, and C. S. Dos Reis. 2023. "A Scoping Review of Interpretability and Explainability Concerning Artificial Intelligence Methods in Medical Imaging." *European Journal of Radiology* 169: 111159. <https://doi.org/10.1016/j.ejrad.2023.111159>.
- Cordella, C., M. J. Marte, H. Liu, and S. Kiran. 2024. "An Introduction to Machine Learning for Speech-Language Pathologists: Concepts, Terminology, and Emerging Applications." *Perspectives of the ASHA Special Interest Groups* 10, no. 2: 432–450. https://doi.org/10.1044/2024_PERSP-24-00037.
- Cornefjord, M., J. Bluhme, A. Jakobsson, et al. 2024. "Using Artificial Intelligence for Assessment of Velopharyngeal Competence in Children Born With Cleft Palate With or Without Cleft Lip." *The Cleft Palate-Craniofacial Journal* 62, no. 10: 1684–1694. <https://doi.org/10.1177/10556656241271646>.
- Cristia, A., M. Lavechin, C. Scaff, et al. 2021. "A Thorough Evaluation of the Language Environment Analysis (LENA) System." *Behavior Research Methods* 53, no. 2: 467–486. <https://doi.org/10.3758/s13428-020-01393-5>.
- Dargan, S., M. Kumar, M. R. Ayyagari, and G. Kumar. 2020. "A Survey of Deep Learning and Its Applications: A New Paradigm to Machine Learning." *Archives of Computational Methods in Engineering* 27, no. 4: 1071–1092. <https://doi.org/10.1007/s11831-019-09344-w>.
- Dégallier-Rochat, S., M. Kurpicz-Briki, N. Endrissat, and O. Yatsenko. 2022. "Human Augmentation, Not Replacement: A Research Agenda for AI and Robotics in the Industry." *Frontiers in Robotics and AI* 9: 997386. <https://doi.org/10.3389/frobt.2022.997386>.
- Deka, C., A. Shrivastava, A. K. Abraham, S. Nautiyal, and P. Chauhan. 2025. "AI-Based Automated Speech Therapy Tools for Persons With Speech Sound Disorder: A Systematic Literature Review." *Speech, Language and Hearing* 28, no. 1: 2359274. <https://doi.org/10.1080/2050571X.2024.2359274>.
- Doğan, S., and A. Çolak. 2024. "Social Robots in the Instruction of Social Skills in Autism: A Comprehensive Descriptive Analysis of Single-Case Experimental Designs." *Disability and Rehabilitation: Assistive Technology* 19, no. 2: 325–344. <https://doi.org/10.1080/17483107.2022.2087772>.
- Dwivedi, R., D. Dave, H. Naik, et al. 2023. "Explainable AI (XAI): Core Ideas, Techniques and Solutions." *ACM Computing Surveys* 55, no. 9: 1–33. Article 194. <https://doi.org/10.1145/3561048>.
- European Institute of Innovation and Technology. 2020. *Creation of a Taxonomy for the European AI Ecosystem*. https://eit.europa.eu/sites/default/files/creation_of_a_taxonomy_for_the_european_ai_ecosystem_final.pdf.
- European Parliament, & Council of the European Union. 2024. *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144, and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)*. Official Journal of the European Union, L, 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>.
- Fachantidis, N., C. K. Syriopoulou-Delli, I. Vezyrtzis, and M. Zygopoulou. 2020. "The Effectiveness of Socially Assistive Robotics in Children With Autism Spectrum Disorder." *International Journal of Developmental Disabilities* 66, no. 2: 113–121. <https://doi.org/10.1080/20473869.2018.1495391>.
- Fan, H., G. Chen, X. Wang, and Z. Peng. 2024. "LessonPlanner: Assisting Novice Teachers to Prepare Pedagogy-Driven Lesson Plans With Large Language Models." In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* 1–20. Association for Computing Machinery. <https://doi.org/10.1145/3654777.3676390>.
- Fang, X., D. T. K. Ng, J. K. L. Leung, and S. K. W. Chu. 2023. "A Systematic Review of Artificial Intelligence Technologies Used for Story Writing." *Education and Information Technologies* 28: 14361–14397. <https://doi.org/10.1007/s10639-023-11741-5>.
- Farmer, R. L., A. B. Lockwood, A. Goforth, and C. Thomas. 2025. "Artificial Intelligence in Practice: Opportunities, Challenges, and Ethical Considerations." *Professional Psychology: Research and Practice* 56, no. 1: 19–32. <https://doi.org/10.1037/pro0000595>.
- Finestack, L. H. 2018. "Evaluation of an Explicit Intervention to Teach Novel Grammatical Forms to Children With Developmental Language Disorder." *Journal of Speech, Language, and Hearing Research* 61, no. 8: 2062–2075. https://doi.org/10.1044/2018_JSLHR-L-17-0339.
- Fox, C. B., M. Israelsen-Augenstein, S. Jones, and S. Laing Gillam. 2021. "An Evaluation of Expedited Transcription Methods for School-Age Children's Narrative Language: Automatic Speech Recognition and Real-Time Transcription." *Journal of Speech, Language, and Hearing Research* 64, no. 9: 3533–3548. https://doi.org/10.1044/2021_JSLHR-21-00096.
- Frizelle, P., A. Buckley, T. Biancone, et al. 2023. "How Reliable Is Assessment of Children's Sentence Comprehension Using a Self-Directed App? A Comparison of Supported versus Independent Use." *Journal of Child Language* 51, no. 6: 1395–1423. <https://doi.org/10.1017/S0305000923000545>.
- Frizelle, P., A. Oliveira-Buckley, T. Biancone, et al. 2025. "The Understanding of Complex Syntax in Children From 5 to 9 Years, Using a Novel Assessment Approach—The Test of Complex Syntax-Electronic (TECS-E)." *International Journal of Language & Communication Disorders* 60, no. 4: e70079. <https://doi.org/10.1111/1460-6984.70079>.

- Frizelle, P., A.-K. Tolonen, J. Tulip, C.-A. Murphy, D. Saldana, and C. McKean. 2021. "The Influence of Quantitative Intervention Dosage on Oral Language Outcomes for Children With Developmental Language Disorder: A Systematic Review and Narrative Synthesis." *Language, Speech, and Hearing Services in Schools* 52, no. 2: 738–754. https://doi.org/10.1044/2020_LSHSS-20-00058.
- Gilkerson, J., and J. A. Richards. 2020. *A Guide to Understanding the Design and Purpose of the LENA System (Technical Report LTR-12)*. LENA Foundation. https://www.lena.org/wp-content/uploads/2020/07/LTR-12_How_LENA_Works.pdf.
- Girardi, A. M., E. A. Cardell, and S. P. Bird. 2023. "Artificial Intelligence in the Interpretation of Videofluoroscopic Swallow Studies: Implications and Advances for Speech–Language Pathologists." *Big Data and Cognitive Computing* 7, no. 4: 178. <https://doi.org/10.3390/bdcc7040178>.
- Gohider, N., and O. A. Basir. 2024. "Recent Advancements in Automatic Disordered Speech Recognition: A Survey Paper." *Natural Language Processing Journal* 9: 100110. <https://doi.org/10.1016/j.nlp.2024.100110>.
- Health and Care Professions Council. 2024. *Standards of Conduct, Performance and Ethics*. <https://www.hcpc-uk.org/standards/standards-of-conduct-performance-and-ethics/>.
- High-Level Expert Group on Artificial Intelligence. 2019. *Ethics Guidelines for Trustworthy AI*. European Commission. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.
- Hosain, M. T., J. R. Jim, M. F. Mridha, and M. M. Kabir. 2024. "Explainable AI Approaches in Deep Learning: Advancements, Applications and Challenges." *Computers and Electrical Engineering* 117: 109246. <https://doi.org/10.1016/j.compeleceng.2024.109246>.
- IBM. 2021. *What is Machine Learning (ML)?* IBM Think. <https://www.ibm.com/think/topics/machine-learning>.
- Ji, Z., and M. Jiang. 2025. "A Systematic Review of Electricity Demand for Large Language Models: Evaluations, Challenges and Solutions." *Renewable and Sustainable Energy Reviews* 225: 116159. <https://doi.org/10.1016/j.rser.2025.116159>.
- Jobin, A., M. Lenca, and E. Vayena. 2019. "The Global Landscape of AI Ethics Guidelines." *Nature Machine Intelligence* 1: 389–399. <https://doi.org/10.1038/s42256-019-0088-2>.
- Jones, S., C. Fox, S. Gillam, and R. B. Gillam. 2019. "An Exploration of Automated Narrative Analysis via Machine Learning." *PLoS ONE* 14, no. 10: e0224634. <https://doi.org/10.1371/journal.pone.0224634>.
- Justice, L. M., W. Y. Ahn, and J. A. R. Logan. 2019. "Identifying Children With Clinical Language Disorder: An Application of Machine-Learning Classification." *Journal of Learning Disabilities* 52, no. 5: 351–365. <https://doi.org/10.1177/0022219419845070>.
- Kierner, S., J. Kucharski, and Z. Kierner. 2023. "Taxonomy of Hybrid Architectures Involving Rule-Based Reasoning and Machine Learning in Clinical Decision Systems: A Scoping Review." *Journal of Biomedical Informatics* 144: 104428. <https://doi.org/10.1016/j.jbi.2023.104428>.
- Kirby, J. 2024. "76,000 children in England Waiting for Speech and Language Therapy." *The Independent*. <https://www.independent.co.uk/news/uk/england-nhs-england-community-government-b2584094.html>.
- Kubat, M. 2021. *An Introduction to Machine Learning*. 3rd edn. Springer Nature Switzerland AG.
- Lammert, J. M., A. C. Roberts, K. McRae, L. J. Batterink, and B. E. Butler. 2024. "Early Identification of Language Disorders Using Natural Language Processing and Machine Learning: Challenges and Emerging Approaches." *Journal of Speech, Language, and Hearing Research* 68, no. 2: 705–718. https://doi.org/10.1044/2024_JSLHR-24-00515.
- Lewis, A., A. Dangol, H. Suh, A. Olszewski, J. Fogarty, and J. A. Kientz. 2025. "Exploring AI-Based Support in Speech-Language Pathology for Culturally and Linguistically Diverse Children." In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (1–19). Association for Computing Machinery. <https://doi.org/10.1145/3706598.3714131>.
- Liss, J., and V. Berisha. 2020. August. *How Will Artificial Intelligence Reshape Speech-Language Pathology Services and Practice in the Future?* ASHA Journals Academy. <https://academy.pubs.asha.org/2020/08/how-will-artificial-intelligence-reshape-speech-language-pathology-services-and-practice-in-the-future/>.
- Marino, F., P. Chilà, S. Trusso Sfrazzetto, et al. 2020. "Outcomes of a Robot-Assisted Social-Emotional Understanding Intervention for Young Children With Autism Spectrum Disorders." *Journal of Autism and Developmental Disorders* 50, no. 6: 1973–1987. <https://doi.org/10.1007/s10803-019-03953-x>.
- Meskó, B. 2023. "Prompt Engineering as an Important Emerging Skill for Medical Professionals: Tutorial." *Journal of Medical Internet Research* 25: e50638. <https://doi.org/10.2196/50638>.
- Mitchell, M. 2020. *Artificial Intelligence: A Guide for Thinking Humans*. Pelican.
- Mueller, J., L. Massaron, and S. Diamond. 2025. *Artificial Intelligence*, (3rd ed.). John Wiley & Sons.
- Murdoch, W. J., C. Singh, K. Kumbier, R. Abbasi-Asl, and B. Yu. 2019. "Definitions, Methods, and Applications in Interpretable Machine Learning." *Proceedings of the National Academy of Sciences* 116, no. 44: 22071–22080. <https://doi.org/10.1073/pnas.1900654116>.
- Norbury, C. F., D. Gooch, C. Wray, et al. 2016. "The Impact of Nonverbal Ability on Prevalence and Clinical Presentation of Language Disorder: Evidence From a Population Study." *Journal of Child Psychology and Psychiatry, and Allied Disciplines* 57, no. 11: 1247–1257. <https://doi.org/10.1111/jcpp.12573>.
- Onslow, M., M. Webber, E. Harrison, et al. 2017. *The Lidcombe Program Treatment Guide*. Lidcombe Program Trainers Consortium.
- Rakap, S., and S. Balıkcı. 2024. "Enhancing IEP Goal Development for Preschoolers With Autism: A Preliminary Study on ChatGPT Integration." *Journal of Autism and Developmental Disorders*. <https://doi.org/10.1007/s10803-024-06343-0>.
- Robles-Bykbaev, V., M. López-Nores, J. García-Duque, J. J. Pazos-Arias, and D. Arévalo-Lucero. 2016. "Evaluation of an Expert System for the Generation of Speech and Language Therapy Plans." *JMIR Medical Informatics* 4, no. 3: e23. <https://doi.org/10.2196/medinform.5660>.
- Rometty, G. 2017, January 19. *Ginni Rometty: It should be Augmented Intelligence, not Artificial* [Video]. YouTube. World Economic Forum. <https://www.youtube.com/watch?v=j98rY3vhPhE>.
- Russell, S. J., and P. Norvig. 2021. *Artificial Intelligence: A Modern Approach*, (4th ed.). Pearson.
- Stevens, E., D. R. Dixon, M. N. Novack, D. Granpeesheh, T. Smith, and E. Linstead. 2019. "Identification and Analysis of Behavioral Phenotypes in Autism Spectrum Disorder via Unsupervised Machine Learning." *International Journal of Medical Informatics* 129: 29–36. <https://doi.org/10.1016/j.ijmedinf.2019.05.006>.
- Thompson, J. K. 2025. *The Path to AGI: Artificial General Intelligence: Past, Present, and Future*. Technics Publications, LLC.
- Tobin, J., P. Nelson, B. MacDonald, et al. 2024. "Automatic Speech Recognition of Conversational Speech in Individuals With Disordered Speech." *Journal of Speech, Language, and Hearing Research* 67, no. 11: 4176–4185. https://doi.org/10.1044/2024_JSLHR-24-00045.
- Usha, G. P., and J. S. R. Alex. 2023. "Speech Assessment Tool Methods for Speech-Impaired Children: A Systematic Literature Review on the state-of-the-Art in Speech Impairment Analysis." *Multimedia Tools and Applications* 82: 35021–35058. <https://doi.org/10.1007/s11042-023-14913-0>.
- Wang, Y., E. Prud'hommeaux, M. Asgari, and J. Dolata. 2020. "Automated Scoring of Clinical Expressive Language Evaluation Tasks." In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications* (177–185). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.bea-1.18>.

Wooldridge, M. 2020. *The Road to Conscious Machines: The Story of AI*. UK, Penguin Random House.

Zhang, X., F. Qin, Z. Chen, L. Gao, G. Qiu, and S. Lu. 2020. "Fast screening for children's developmental language disorders via comprehensive speech ability evaluation-using a novel deep learning framework." *Annals of translational medicine* 8, no. 11: 707. <https://doi.org/10.21037/atm-19-3097>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Supplementary Material 1: Suggested reading and resources for readers seeking a deeper understanding AI. **Supplementary Material 2:** Classical machine learning versus deep learning: artificial neural networks explained. **Supplementary Material 3:** The EU AI Act and its risk-based regulatory framework for AI.