EXECUTIVE SUMMARY

Hear *their* Heart

Communication Without Speech

The founder's thesis, the strongest opportunity, and the case for investigation, in brief. The complete 22-part research report is available as a companion document.

Contents

Executive Summary edition · Hear Their Heart

A note from the founder

Hear Their Heart was not started because someone spotted a market opportunity in assistive technology. It started with frustration.

My son has severe autism and struggles to communicate verbally. I have watched communication breakdowns turn into real distress and difficult behavioral episodes, simply because the people around him could not always understand what he was trying to say. Watching that happen to your own child is one of the hardest things a parent can go through.

We use an iPad running LAMP-based AAC software, and I am genuinely grateful it exists. It gives my son a way to be understood that he would not otherwise have. The clinicians, researchers, and companies documented throughout this report built something real, and their work created the foundation everything below stands on.

But I am also frustrated by how far the rest of consumer technology has moved while the communication experience available to children like my son has stayed largely the same: static grids of symbols, generic graphics, manually configured vocabulary sets, interfaces that look almost identical from one child to the next.

We are living through one of the largest technological shifts in history. Artificial intelligence can now generate images, video, speech, and adaptive digital experiences almost instantly and cheaply. Meanwhile, many families supporting non-speaking or minimally speaking children are still working with tools whose underlying interaction model has barely changed in twenty years.

The question this report is built to answer is not “is AAC helpful.” The evidence in the pages that follow says clearly that it is. The real question is why modern AI has not yet been pushed hard into this field to make communication tools more personal, more visual, more adaptive, and more responsive to the individual child in front of them, and where the actual opportunity is to change that.

That is the gap Hear Their Heart exists to close. Section “THE THESIS,” immediately following, lays out the case in full.

Founder, Hear Their Heart

A note on scope and confidence

This report reflects extensive live web research (over 200 individual search/fetch operations across five research streams) conducted in September 2026, prioritizing peer-reviewed literature, government/clinical sources, primary company sources, and patent databases. It is not a substitute for professional legal, regulatory, financial, or clinical due diligence; every section flags explicitly where a professional (patent counsel, FDA regulatory counsel, a claims-data analytics vendor) should be engaged before acting. Several explicit data gaps are called out throughout (e.g., no published national Medicaid claims data broken out for AAC-specific billing codes, no independent sizing of the SLP-therapy-services market) rather than papered over with invented figures. Two specific instances of likely AI-generated, unverifiable market-report content were caught and flagged during this research (a fabricated-looking "$50 million Series B" claim about Lingraphica, and a general caution about a ~7x spread across vendor market-size reports); a reminder that search-engine-indexed content about this sector should be treated with active skepticism, not taken at face value.

THE THESIS: AI PERSONALIZATION AS THE NEXT ERA OF ASSISTIVE COMMUNICATION

This section states the central argument of this report in plain terms before the reader goes deep into 22 parts of supporting research. It is written with more conviction than the sections around it, on purpose, and every specific claim in it is either backed by evidence found elsewhere in this report or is explicitly flagged as HYPOTHESIS. Part 9 (AI and the Next Generation of AAC), Part 11 (Opportunity Map), Part 20 (Product Concepts), and Part 21 (Why Now) all connect back to the argument made here.

The gap, stated plainly

Every part of this report that touches technology history tells some version of the same story. Blissymbolics, PECS, Minspeak, PCS symbols, GoTalk, dedicated speech-generating devices, Proloquo2Go: each was a genuine advance for its era (§2). But look at the underlying interaction model across nearly all of them, including the market leaders sold today: a fixed grid of symbols, drawn from a shared library used by every other child on the same app, arranged by a clinician or parent, navigated by touch or gaze. The core representation layer of AAC (§1.2) has been remarkably stable since the 1980s. What has changed since then is mostly packaging: the symbols moved from paper to a dedicated device, then from a dedicated device to a tablet. The symbols themselves, and the fact that they are generic rather than personal, has not fundamentally moved.

Meanwhile, outside this field, artificial intelligence has spent the last four years learning to generate photorealistic images, coherent short video, natural speech, and adaptive, personalized digital experiences on demand, at falling cost (§13.8). That capability has been aimed at advertising, entertainment, productivity software, search, and social media. It has not, by and large, been aimed at one of the oldest and most fundamental human problems there is: loving someone and not always understanding what they are trying to tell you.

This report's opportunity map (§11) already identifies device abandonment, slow communication rate, generic vocabulary, and the inability to express personality as the field's best-documented, most consistently reported problems (§4.2.1, §14). Those are not new findings. What is new, and what this section argues, is that the specific tool needed to attack several of them at once, personalization at the level of image, video, and context rather than just text prediction, has only existed in usable form for roughly the last two to three years.

The reframe

The wrong question is “does AAC work.” The evidence says it does; the field’s own peer-reviewed literature and the lived experience behind this report both say so.

The right question is: given what generative AI can now do elsewhere, why does a communication board for a non-speaking child still look like a communication board built in 2005, just running on newer glass? And what happens to outcomes, engagement, and independence once that stops being true?

The seven-part opportunity, evidence by evidence

Hear Their Heart’s founding thesis is that AI personalization should become a foundational layer of assistive communication technology, not a bolt-on feature. That thesis breaks into seven concrete product directions. Each is treated honestly below: what is realistic today, what already exists as research or a competing product, and what remains a genuine open bet.

1. AI-personalized visual communication cards, built around the child’s real life. Instead of a generic drawing standing in for “grandma,” the card reflects an image connected to the child’s actual grandmother. Instead of a generic school icon, the interface uses something recognizable from the child’s actual school. Instead of a generic snack picture, the system shows the foods the child actually eats. This is a direct, evidence-grounded extension of a finding already documented in this report’s competitive research (§8.3): today’s symbol libraries (PCS, SymbolStix, Widgit, ARASAAC) are static, professionally illustrated sets shared across every user, and the market’s own patent landscape shows this static-library model has barely moved since the 1990s. Uploading a plain family photo as a custom icon is old and already supported by most major apps; that is not the innovation. The innovation is using generative image models to produce a *consistent, stylized, repeatable* likeness that stays recognizable across dozens of future uses, and to do this fast enough that a parent or therapist can build it in minutes rather than hours. This is a genuinely open bet, not a settled one: a focused research pass for this report (§9.10) found the underlying image and video technology already mature enough to do this today (consumer tools including Google’s Gemini “Nano Banana,” OpenAI’s GPT Image 2, and Midjourney’s character-reference feature can now hold a specific person’s likeness consistent across many generated images from a handful of reference photos), yet found no AAC or autism-support product that has built this capability, and found that PRC-Saltillo, the field’s leading dedicated-device maker, shipped AI image generation into its apps in February 2026 running on Apple Intelligence and explicitly configured it to refuse to generate images of real people at all. The market leader has the platform and the distribution and has chosen not to solve this problem. That is the opening.

2. An AI visual prompt builder for routines and transitions. A parent, teacher, or therapist should be able to rapidly generate a visual sequence for getting dressed, using the bathroom, going to the

doctor, brushing teeth, handling a schedule change, or visiting somewhere unfamiliar, in a visual style that is actually familiar and engaging to that specific child. This report's technology research already found the enabling capability proven in peer-reviewed form: the ASSETS 2024 study on GPT-4V-generated visual scene displays found SLP evaluators rated AI-generated communication content as comparable to human-made output (§9.4, §11 Concept 6). What that study did not yet test at scale is exactly this use case, personalized multi-step routine sequences rather than single boards, which is the more ambitious and more clinically valuable version of the same underlying capability.

3. Personalized short social and instructional videos. "What happens when we go to the dentist." "How we get ready for school." "How we ask for a break." "What happens when plans change." This is a natural extension of Social Stories, the well-established, decades-old clinical intervention pioneered by Carol Gray, into video generated around an individual child's actual routines, environment, and interests. This is not a replacement for a clinician's judgment about what story a specific child needs; it is a way to collapse the hours currently spent searching the internet, hand-illustrating, and printing a custom story down to minutes, so that far more children get far more of an intervention already proven to work. This is not empty white space: commercial AI Social Story generators already exist (GrowTale, Ella Kids, Glint, Easy-Peasy.AI, §9.10), and they already personalize by description, a boy in a green shirt, standing in a park. None of them was found to personalize by *identity*: an image or video built from an actual photo of this specific child, this specific dentist's office, this specific sibling. That distinction, generic-but-described versus real-and-recognizable, is the entire gap this item is built to close, and it is a narrower, more defensible claim than "nobody does AI social stories," which would simply be false.

4. Dynamic, context-aware interfaces that reduce friction without taking over the choice. The goal is never for AI to decide what a person wants to say; this report treats that line as close to sacred, and the reasons why are laid out at length in §9.9. The goal is to make the *right* vocabulary easier to reach depending on whether the child is at school, at home, at a restaurant, in the car, at a doctor's appointment, playing outside, or getting ready for bed, the same way a smartphone keyboard learns what you tend to type without ever putting words in your mouth. Genuine peer-reviewed research already exists in this exact direction: SPICA's personal-memory architecture and AACessTalk's real-time contextual vocabulary recommendation (both presented at CHI/IUI in 2025, §9.3) prove the core mechanism works and, in AACessTalk's case, prove it can *increase* a child's sense of agency in conversation rather than reduce it. Neither has shipped as a commercial product yet.

5. Personalization around the child's actual interests. A child fascinated by trains should not receive the same visual communication experience as a child fascinated by dinosaurs, space, or a favorite show. This sounds simple, and it is the one item on this list with the least direct academic literature behind it in this research pass; it is closer to product design common sense than a technical breakthrough. It is included here because every parent and clinician quoted in this report's

user-voice research (§14) describes generic, one-size-fits-all vocabulary as a real barrier to engagement, and interest-based personalization is the most direct, lowest-technical-risk way to address that.

6. A caregiver copilot: describe the situation, get the supports. A parent should be able to write, in plain language, "my son has a new dentist appointment tomorrow and gets very anxious in unfamiliar places," and receive back a coordinated bundle: a visual schedule, a short personalized story, relevant communication cards, questions to prepare for, and symbols likely to matter during the visit. Today, building that bundle means searching the internet, manually creating graphics, editing templates, printing materials, and configuring several disconnected apps, exactly the uncompensated hours-long burden this report's teacher and parent journeys document in detail (§5.2, §5.4). Collapsing that into one integrated workflow is a caregiver-time problem with a caregiver-time solution, not a research problem; the individual generation capabilities it depends on (items 2 and 3 above) are the harder parts, and they are already within reach. Here too the pieces already exist separately and the bundle does not: GrowTale and Ella already generate a story from a description, Glint and HeyOtto already generate a visual schedule from a description, and this report's research (§9.10) found no product that turns one description into a coordinated kit spanning schedule, story, and communication cards together.

7. A personal communication profile and an AI-generation-compatible visual library. With real consent and privacy protection, a portable profile carrying a child's preferences, routines, triggers, interests, effective supports, and important personal vocabulary would also solve this report's single most-repeated cross-stakeholder complaint: that a child's hard-won progress resets at every transition, new school, new therapist, new device (§5.7, §11 Concept 2). Paired with a personalized visual library built for AI generation rather than a static licensed symbol set, this becomes the infrastructure layer the other six items run on, not a separate product.

Why now, specifically

The last generation of assistive communication digitized the communication board: paper became a tablet screen, a recorded button became a synthesized voice. The underlying visual vocabulary stayed generic and shared. The next generation can personalize the entire communication environment, and it can do that now, specifically, because of a stack of capabilities that did not exist together even three years ago:

- **Multimodal AI** that can look at a photo and understand what is in it well enough to generate clinically appropriate content from it, proven in peer-reviewed form only since 2023 to 2024 (§9.4, §21).
- **Generative image and video models** capable of producing coherent, personalized visual content fast and cheaply, a capability that simply did not exist at usable quality or cost five years ago.

- **On-device and low-cost inference**, with industry analysts projecting inference cost for large models to fall more than 90 percent by 2030 (§13.8), which matters directly for a product a child may use hundreds of times a day and which, for privacy reasons, should avoid sending a vulnerable child's data to the cloud by default.
- **A newly-published ethical design framework** for exactly this problem, built by HCI researchers and AAC self-advocates together between 2023 and 2026 (§9.9): suggest rather than substitute, stay visible to the communication partner, personalize from the individual's own history rather than generic web text, and keep the human in control of consent and override at every step. This framework, cited throughout this report, is what makes it possible to pursue this thesis aggressively without repeating the authorship mistakes this field has already lived through once, in the facilitated communication and Rapid Prompting Method controversies documented in §10.
- **Accessible, already-owned hardware.** Nearly every family in this report's target population already owns a tablet or smartphone capable of running these tools; the marginal hardware cost of this thesis is close to zero, in sharp contrast to the $5,000 to $8,500 cost of a dedicated device (§6.6, §7.4).
- **Identity-consistent generation became a solved technical problem only in the last cycle of model releases, and the relevant privacy rulemaking landed at almost exactly the same time.** OpenAI's Sora 2 "Cameo" (fall 2025), Google's Gemini "Nano Banana," and OpenAI's GPT Image 2 all demonstrate consumer-grade, reference-photo-conditioned identity consistency for the first time as of 2025 to 2026 (§9.10), and the FTC's amended COPPA rule, finalized April 2025 and reaching full compliance deadline in April 2026, was the first US privacy rulemaking to explicitly grapple with AI-generated likenesses and biometric identifiers for children (§9.10). A company building this thesis in 2026 is not waiting on either the technology or the regulatory ground to stabilize; both arrived within the past year.

No single one of these was the missing piece five years ago. Together, as of the last two to three years, they are.

What this thesis is not, and the tensions it must respect

Two things need to be said plainly, in the same spirit of intellectual honesty the rest of this report tries to hold to.

First, established AAC science carries a real, well-evidenced principle that a personalization-first product must respect rather than override: consistent symbol placement and appearance supports motor planning and communication speed, which is the entire reason systems like Minspeak exist (§2.1) and the entire reason core-vocabulary boards are built to stay stable over time. A product that regenerates a child's board every day in the name of personalization, rather than personalizing content while keeping layout and access methods stable, would be working against six decades of

clinical evidence, not with it. The right design target is personalized *content* inside a stable, clinically sound *structure*, not novelty for its own sake.

Second, using generative AI on real photographs of a child's actual family members, not just the child, raises a privacy question broader than standard children's-data rules address: a grandmother, sibling, or teacher whose image is used to generate a communication card never consented to a technology company processing their likeness. This is not a hypothetical concern this report is raising on its own; a dedicated research pass (§9.10) found that the FTC's own April 2025 COPPA rulemaking record explicitly discusses AI-generated avatars of children and concedes it would be difficult for an operator to even determine whether a given avatar was generated from a real child's image, and found no regulatory guidance, professional-association position, or case law anywhere that addresses the narrower and equally real question of a *non-child third party's* likeness (a grandmother, a sibling) processed for a minor's benefit. This report's existing AI-ethics research (§9.9) already establishes the core discipline this thesis must extend to images and video, not only to text: transparency, consent, and control at every step, now applied to a new and more sensitive category of data where the law itself has not yet caught up. Any company pursuing this thesis needs a privacy and consent architecture built for that reality from day one, not retrofitted after launch, and should expect to need outside counsel to help define it rather than assume an existing template covers it.

Neither tension is a reason to abandon the thesis. Both are reasons to build it carefully, with the same clinical humility this report asks of every other opportunity in it (§9.9, §22).

The investor argument

Hundreds of billions of dollars are being invested in artificial intelligence right now. Most of that capital is aimed at advertising, entertainment, productivity, search, coding, and social media, genuinely valuable problems, but not urgent ones in the way this one is. Meanwhile, millions of families are living with one of the most fundamental human problems there is: loving someone and not always understanding what they are trying to say. In an era when AI can generate a photorealistic video from a sentence, that should feel unacceptable, not inevitable.

This is not a pitch for another autism app. The opportunity this report identifies, and the one Hear Their Heart intends to pursue, is an intelligent communication and visual-support platform capable of becoming genuinely personal to the individual using it: the child's own people, own places, own interests, own routines, expressed through AI capabilities that, until the last two or three years, simply did not exist. The remainder of this report, its history, its market map, its money flows, its patent landscape, and its opportunity scoring, exists to test that thesis against the evidence rather than simply assert it. Where the evidence pushes back, this report says so directly (§4.2.1, §10, §12.15, the Final Deliverable's closing "surprising discoveries"). Where it supports the thesis, as it does across Parts 9, 11, 13, 19, 20, and 21, this report says that directly too.

Research tagging key: FACT (cited), ESTIMATE (calculation/reasoning shown), CLAIM (attributed to a named speaker/company), HYPOTHESIS (analyst inference/opportunity).

A note on terminology used throughout this section: "AAC" (Augmentative and Alternative Communication) is the umbrella clinical/field term. ASHA defines AAC as "an area of clinical practice that supplements or compensates for impairments in speech-language production and/or comprehension, including spoken and written modes of communication" and divides it into **unaided** (body-only: gesture, sign, facial expression) and **aided** (uses an external tool, ranging from no-tech picture boards to high-tech speech-generating devices) modalities. (FACT; ASHA Practice Portal: AAC)

There is active, unresolved disagreement in the field over "non-verbal" vs. "non-speaking" vs. "minimally speaking" vs. "AAC user"; addressed in a dedicated subsection of Part 1.

PART 21: WHY NOW: THE FIVE STRONGEST OPPORTUNITIES

1. AI Communication Co-Pilot. Domain-adapted transformer techniques for AAC-specific short-utterance prediction were only published in 2024–2025 (PrAACT; Gaines & Vertanen, EMNLP 2025 Findings); agentic personal-memory architectures for AAC (SPICA) and dual-sided caregiver/child modeling (AACessTalk) were presented at CHI/IUI 2025; this entire technical foundation is under two years old. Simultaneously, the field's own ethical framework for how to do this responsibly (suggestion-not-substitution, transparency, personalization, user control) has only crystallized in the 2023–2026 HCI literature (Valencia et al. CHI 2023 through the CHI 2026 "Three Praxes Framework"). Five years ago, neither the technical capability nor the ethical design consensus existed; ten years ago, the underlying LLMs themselves did not exist in usable form. On-device inference costs are simultaneously collapsing (Gartner projects >90% cost reduction by 2030), making a privacy-preserving, always-available version of this product newly feasible rather than a cloud-dependent compromise.

2. Webcam-only/software eye-gaze access. Tobii; the incumbent with the dominant patent position in dedicated eye-tracking hardware; has itself pivoted toward AI-driven, webcam-based gaze estimation (Tobii Nexus) only recently, implicitly validating that the hardware-free approach is now technically viable at consumer camera resolution and processing power. Open academic research (WebEyeTrack, 2025) demonstrates few-shot, on-device personalization is achievable without dedicated infrared hardware. This was not possible even five years ago at acceptable accuracy; it is possible now because of the same broader computer-vision advances (efficient on-device models, better camera hardware baked into every consumer tablet) that are separately being validated at massive scale by the automotive driver-monitoring industry ($2.72B market in 2024, pushing RGB-camera gaze algorithms to commodity cost).

3. Pediatric personalized voice banking. General-purpose voice cloning only reached commercially mature, low-data-requirement quality around 2023–2025; but every major commercial vendor (ElevenLabs explicitly) excludes children's voices from self-serve use, creating a clear, current, and likely temporary regulatory/trust vacuum. The specific technical approach needed for children who never developed typical speech (residual-vocalization blending) was pioneered by VocaliD but that company was absorbed into a non-AAC-focused platform (Veritone) in 2022, leaving the pediatric-specific niche freshly vacant. Academic research on cultural/linguistic identity in AAC-generated voices is dated 2026 in this research; the field is actively working out how to do this well right now, not revisiting settled territory.

4. AI board/visual-scene-display generation. The enabling capability; a multimodal model that can look at a photograph and generate SLP-quality, contextually appropriate communication content; simply did not exist before GPT-4V-class models arrived in 2023. The first rigorous test of this specific application (the ASSETS 2024 study) is barely two years old. Before this, "generate a custom board from a photo" was not a deferred feature; it was not a solvable problem at all, since it required genuine scene understanding paired with clinical vocabulary judgment that only recent multimodal LLMs can approximate.

5. Nonprofit-hybrid/PBC AI-native platform. Public Benefit Corporation statutes have matured (Delaware, California, New York, and others have had PBC law in force for roughly a decade or more) into a well-understood, investor-legible legal structure, rather than a novel experiment. CoughDrop has now operated its cross-subsidy pricing model long enough (since 2014, through a 2023 acquisition) to be treated as validated precedent rather than a hypothesis; the report can point to a real, still-operating example rather than a theoretical model. Post-pandemic normalization of subscription SaaS procurement in K-12 education (accelerated 2020–2023) has made school districts meaningfully more comfortable with recurring software-subscription purchasing than they were a decade ago, lowering the go-to-market barrier for exactly this business model in exactly this buyer segment.

PART 20: PRODUCT CONCEPTS, RANKED

The concepts below are described in full, with hardware/software requirements, IP strategy, pricing, and five-year potential, in the extensive companion report.

RANK	CONCEPT	RATIONALE
1	**AccessAll** (Concept 7)	The flagship expression of this report's core thesis; the only concept that both resolves the revenue-vs-access tension by design and fully embodies the seven-part personalization vision (personalized imagery, routine and story generation, context awareness, caregiver coaching) under one mission-aligned platform.
2	**CoPilot** (Concept 1)	Best "why now," most novel and defensible technically, viable as a standalone feature/licensing play or the technical core of AccessAll.
3	**StoryBuilder** (Concept 11)	Directly answers the founder's sharpest frustration: generic content versus a child's real life. No direct existing competitor was identified doing exactly this, subject to the confirmation noted in the research notes; high urgency, genuine novelty.
4	**SituationKit** (Concept 12)	Cheapest and fastest concept to prototype once its underlying capabilities exist; the most immediately understandable feature for a caregiver, and likely the platform's primary front door.
5	**ClearPath** (Concept 5)	Addresses the single highest-leverage, best-documented bottleneck (funding/documentation) with the clearest near-term revenue model and shortest sales cycle.
6	**HomeVoice** (Concept 10)	Directly serves the best-evidenced, most universal unmet need (caregiver training gap) with a genuine academic-research head start.
7	**GazeFree** (Concept 4)	Highest social-impact-per-dollar for the most access-constrained population (CP/Rett), riding a genuine "why now" cost curve.
8	**SceneBoard** (Concept 6)	Strong, low-risk, high-frequency feature; likely most valuable bundled into #1 or #2 rather than standalone.
9	**SpeakClear** (Concept 8)	Strong defensible niche, but narrower addressable population and harder ML problem (data-scarce ASR).
10	**OwnVoice** (Concept 3)	Meaningful and dignity-affirming, but the VocaliD precedent suggests this is more durable as a feature/partnership than a standalone company.
11	**TrendGuard** (Concept 9)	Genuinely novel category, but execution is gated by vendor-data-partnership risk outside the company's own control.
12	**PathPass** (Concept 2)	Highest pure social impact if it succeeds as a standard, but weakest as a venture-scale business on its own; best pursued as a nonprofit/consortium initiative or a trust-building feature of #1.

FINAL DELIVERABLE

The 10 biggest insights

1. **AAC is a stack, not a product.** Representation, vocabulary organization, output, access method, hardware, and software are independently swappable layers (§1.2); most real innovation; historically and today; happens at exactly one layer while inheriting the rest from incumbents, which explains both the durability of incumbents and the opportunity in narrow, layer-specific plays.
2. **Four or five different parties use, choose, prescribe, pay for, and approve payment for the same device** (§6.1); this fragmentation, not underlying technology cost, is the primary driver of price, delay, and even product design (a "locked" $8,000 dedicated tablet and a $250 consumer-app-plus-iPad combination can be functionally identical; the price gap is a reimbursement-code artifact, not an engineering one, §6.2).
3. **Simplicity and durability have repeatedly beaten sophistication commercially:** GoTalk, PECS, and arguably Proloquo2Go's original 2009 breakout (more a platform-timing story than an AAC-science one) all won by being fast to learn and rugged, not by being the most advanced system available (§2.2–2.3).
4. **Device/therapy abandonment (30–50%, stable across decades) is the industry's actual central failure mode**, and it is driven overwhelmingly by training, onboarding, and caregiver-support gaps; not by device capability (§4.2.1, §18); meaning most "build better technology" theses target the wrong lever.
5. **The single most-repeated friction point across every stakeholder journey** (child, parent, SLP, teacher, district, hospital, payer) **is the absence of a portable, interoperable communication profile** that travels with the child; no company has solved this, and it is largely a coordination/standard-setting problem, not a hard technical one (§5.7).
6. **Patent risk in this space is real, asymmetric, and concentrated in two places:** Tobii's ~850+-filing eye-tracking portfolio and PRC/Semantic Compaction Systems' demonstrated willingness to litigate Minspeak-style icon-encoding IP against a two-person competitor (§8.1–8.2); while the AI-prediction-for-AAC layer itself looks comparatively open as of 2026, a window that will likely close as big-tech generative-AI patenting broadens (§8.6).
7. **Reimbursement rules, not silicon, set hardware prices:** the CMS requirement that a tablet-based SGD be "locked" to function solely as a speech-generating device at time of issue is why a functionally ordinary tablet costs 8x more once it is DME-billable (§6.2).
8. **AI's biggest documented risk here is not hypothetical:** "putting words in a nonspeaking person's mouth" is the direct continuation of the Facilitated Communication/RPM/Spelling-to-

Communicate controversy, which has already caused real, documented harm (false abuse allegations in the 1990s) when unvalidated communication-authorship claims were trusted for a population that often cannot self-report being misrepresented (§10.1–10.4, §9.9).

9. **Independent, high-quality research converges on the same root problems from different angles**; Yau et al. (2024, autism-focused, Australia) and Berenguer et al. (2022, general AAC, 19 studies/297 parents) reach the same five-theme structure (training, stigma, resources, engagement, device fit) via entirely different populations and methods, which is unusually strong grounds for prioritization confidence in a field otherwise full of noisy, conflicting market claims (§4.2.1).
10. **The market-sizing "problem" is itself a signal, not just a limitation:** vendor reports disagree by ~7x, no Medicaid claims data is published for AAC-specific HCPCS codes, and the SLP-therapy-services market (likely larger than the device market) has never been sized; meaning a company that builds genuinely better first-party outcomes data would hold an information advantage that currently belongs to no one, including the incumbents (§7.0, §7.7–7.8).

The 10 biggest unsolved problems

1. No portable, standardized communication-profile/interoperability layer across devices, providers, and life transitions.
2. Device/therapy abandonment stuck at 30–50% for decades despite AAC being formally "evidence-based practice."
3. The communication-rate/speed gap between AAC output and real-time spoken conversation, which drives social exclusion and being underestimated.
4. No good way to give a child who never had typical speech (or is losing it) an authentically personal synthetic voice; commercial voice-cloning vendors explicitly exclude children.
5. A structural SLP/teacher training gap (75% of surveyed SLPs lack AAC continuing-education access; teachers report being excluded from device selection yet expected to be daily experts).
6. Multilingual/culturally appropriate vocabulary remains a bolt-on afterthought, not a first-class design constraint, across nearly all commercial systems.
7. Medical-necessity documentation and prior authorization; the single highest-leverage bottleneck in the whole system; remains manual, unautomated, and fragmented across 50+ distinct state Medicaid processes.
8. No rigorous, broadly agreed-upon framework yet exists for how much AI should shape or generate a nonspeaking child's message without compromising authorship; the field is actively inventing this norm in real time (2023–2026 HCI literature).
9. Intent determination for early/emergent and multiply-disabled communicators remains clinically hard, and is structurally the same unresolved problem underlying the FC/RPM/S2C controversy; no scalable, validated, non-facilitator-dependent tool exists to help here.

10. No one has sized the actual Medicaid/insurance dollar flow into AAC with real claims data, meaning every company and researcher in this space is operating with an incomplete picture of the public-payer economics that dominate it.

The 10 strongest business opportunities

1. A personalized AI visual and video communication platform (personalized imagery, routine and Social-Story-style video generation, context-aware vocabulary, caregiver coaching, all under one product).
2. AI Communication Co-Pilot (personalized prediction and phrase generation, authorship-preserving), the platform's predictive-text component.
3. Personalized AI Social-Story and routine-video generator (Concept 11), the platform's most novel and least-contested component.
4. Portable Communication Profile / interoperability standard.
5. Pediatric personalized voice-banking platform.
6. Webcam-only/software-only eye-gaze access layer.
7. AI-powered medical-necessity documentation and appeals automation.
8. AI board/visual-scene-display auto-generation for teachers and SLPs.
9. Unified voice-in/voice-out app for CAS/dysarthria.
10. District-level AAC outcomes/abandonment-risk analytics.

The 5 opportunities to investigate first

1. **Personalized AI Social-Story and routine-video generator (StoryBuilder, Concept 11):** the single most direct expression of this report's founding thesis, the least contested by existing competitors of anything on this list, and the fastest way to test whether personalization measurably changes outcomes and engagement before committing to the full platform.
2. **AI Communication Co-Pilot:** strongest technical "why now" for the predictive-text layer specifically, most direct alignment with responsible-AI-for-children design principles, viable via multiple business models.
3. **AI-powered medical-necessity/appeals automation:** fastest path to real revenue and the clearest, best-documented pain point, with the least clinical-validation burden; a useful funding bridge while the platform's harder, more novel components mature.
4. **Webcam-only eye-gaze access layer:** highest social-impact-per-dollar for the most access-constrained population, riding a real, current cost-curve shift.
5. **The full personalized AI platform (AccessAll):** the vehicle that eventually combines 1 through 4 into the single highest-ceiling outcome; the highest-capital, longest-timeline option of the five,

and one that should be approached only after de-risking the StoryBuilder and CoPilot components individually.

The single opportunity that deserves the most investigation: A Personalized AI Communication and Visual-Support Platform

- **Core insight:** For four decades, AAC's underlying visual vocabulary has been generic and shared across every user of a given app; the child adapts to the technology's library of symbols rather than the technology adapting to the child's actual life. Generative AI has, in the last two to three years, become capable of producing personalized images, video, and adaptive digital experiences fast enough and cheaply enough to invert that relationship. This report's own research finds the specific enabling pieces already proven in peer-reviewed form (GPT-4V-generated visual scene displays rated as SLP-comparable, §9.4; contextual AI vocabulary recommendation shown to increase, not decrease, a child's sense of agency, §9.3), but finds no identified company has yet combined personalized imagery, personalized routine and Social-Story video, context-aware vocabulary, and caregiver coaching into one coherent product built specifically for children with complex communication needs.
- **Proposed product:** A platform combining five capabilities under one authorship-preserving design discipline: personalized visual communication cards built around the child's real people, places, and interests (Concept 1's imagery layer); an AI visual prompt builder for routines and transitions (Concept 6); a personalized Social-Story and routine-video generator (Concept 11); a caregiver copilot that turns a plain-language description of a situation into a coordinated support kit (Concept 12); and a predictive, personalized phrase-suggestion layer (Concept 1's CoPilot component). Every AI-generated item is a suggestion an adult reviews before it reaches the child, consistent with the authorship framework in §9.9, extended here from text to images and video.
- **Exact user:** Children roughly ages 4 to 14 with complex communication needs, autism, cerebral palsy, apraxia, Down syndrome, and related conditions primarily, starting with families and clinics already using an Open-Board-Format-compatible app (CoughDrop) or willing to pilot a new one, and starting product development with the single highest-conviction, lowest-competition component (StoryBuilder, Concept 11) before building out the full platform.
- **Payer:** Initially families, on a cross-subsidized subscription consistent with the CoughDrop precedent (§12.13); schools and clinics on a site-license basis; Medicaid and insurance via the E2511 software HCPCS code as a validated, later-stage expansion once clinical evidence exists.
- **Existing competitors:** No identified company combines these five capabilities today. The closest individual pieces, confirmed by dedicated research for this report (§9.10), are Cboard's Cbuilder (board generation only, open-source), a small but real category of generic AI social-story and visual-schedule generators (GrowTale, Ella Kids, Glint, HeyOtto, Easy-Peasy.AI, all personalizing by description rather than by real identity), and PRC-Saltillo's own AI image feature

(shipped February 2026 on Apple Intelligence), which is the single most direct evidence of incumbent intent in this exact direction, and which explicitly refuses to generate images of real people.

- **Why competitors haven't solved it:** The incumbent hardware and software makers are fundamentally reimbursement-infrastructure companies whose moat and revenue model is the DME and insurance-billing relationship built up over decades (§6.2 through §6.4), not generative-AI product velocity, and PRC-Saltillo's own February 2026 feature shows the clearest evidence available: when the field's leading AAC vendor finally shipped generative image AI, it chose a safety-first guardrail, refusing to generate images of real people at all, over solving the harder consent and personalization problem, likely a reasonable, defensible choice for a large incumbent with brand risk to protect, but one that leaves the capability unclaimed (§9.10). The dedicated AI social-story vendors found in this research (GrowTale, Ella Kids, Glint, HeyOtto) are small, general-audience products built for the broader autism-parent market, not AAC-integrated, and personalize by description rather than identity, likely because identity-consistent generation of real minors and their relatives is a harder consent problem than a small team wants to take on first. Open-source and mission-driven AAC players (CoughDrop, Cboard) have shipped real, fast-moving individual AI features but lack the capital for a coordinated, clinically-validated multi-capability platform. Big-tech accessibility research (Google's Project Relate and Project Euphonia) has repeatedly produced strong research without durably commercializing it. No party in this market currently has both the AI capability and the sustained, mission-driven motivation to solve the consent problem and build this specific combination.
- **Technological advantage:** A proprietary, consented, longitudinal dataset of what actually works for a given child, imagery, stories, and phrasing that measurably improved communication and reduced distress, which is a genuinely defensible moat (§19.8) that generic foundation-model access cannot replicate on its own.
- **Estimated development difficulty:** Medium to high for the full platform; low to medium for the recommended first component (StoryBuilder). The individual generation techniques are published and reproducible; the harder work across the whole platform is clinical-template design, consent architecture for imagery of real people, and the caregiver and therapist-facing review workflow, not novel model research.
- **Regulatory risk:** Low to moderate if positioned as non-diagnostic communication-assistance and caregiver-support software, consistent with the posture of existing AAC apps and existing Social-Stories practice; this requires confirmation from actual FDA regulatory counsel, and separately requires real legal review of consent and likeness rights for AI-generated imagery of identifiable children and family members before any public launch, a distinct and less-precedented question than standard children's software privacy compliance.
- **Clinical-validation requirements:** SLP- and behavior-analyst-rated content-quality studies following the ASSETS 2024 evaluator-in-the-loop methodology; anxiety-reduction and routine-comprehension outcome measures for the Social-Story component specifically, given Social

Stories' established evidence base sets a real bar to meet, not just a format to imitate; words-per-minute and authenticity measures for the predictive-text component, per the original Co-Pilot proposal.

- **Potential IP:** The clinically-templated content-generation pipelines that keep AI output inside evidence-based structures (Social Stories conventions, core-vocabulary motor-planning consistency) rather than free-form generation; the consent and likeness-rights architecture for generating imagery of real identifiable children and family members, an area with little existing precedent to copy; the proprietary longitudinal usage-and-outcome dataset described above.
- **Estimated pricing:** $10 to $25 per month for families, $500 to $2,000 per school building per year for site licenses, consistent with CoughDrop and general education-SaaS benchmarks found in this research.
- **Rough market opportunity:** Layers onto the roughly $260 million per year US pediatric AAC device-and-software bottom-up TAM (§7.4) as a premium tier or successor product, with meaningful upside into the unserved roughly half of the pediatric complex-communication-needs population (§7.7) if personalization measurably improves adoption and reduces the 30 to 50 percent abandonment rate documented throughout this report (§4.2.1, §18.3).
- **Customer-acquisition strategy:** Partner-first distribution through the CoughDrop and Open Board Format ecosystem and SLP and AAC-specialist professional networks (ISAAC, USSAAC, ASHA's AAC special-interest group), and specifically through autism-focused early-intervention and behavior-analysis practices for the StoryBuilder component, where Social Stories are already a trusted, familiar clinical tool; explicitly avoiding an app-store-only consumer launch given this is a clinician-gatekept market (per the Verbally failure lesson, §18.3).
- **First prototype:** A StoryBuilder pilot: a caregiver describes an upcoming situation for their child, the system generates a short personalized routine story, and a small group of SLPs and behavior analysts rate the generated content's clinical appropriateness and the family's engagement with it against their own hand-built Social Stories.
- **First 20 people/organizations to interview:** AAC self-advocates endever* (via CommunicationFIRST), Beth Moulam, and Emma Zurcher-Long; CoughDrop's Brian Whitmer; AssistiveWare's David Niemeijer; the Gaines and Vertanen (Michigan Tech) and SPICA and AACessTalk research teams; behavior analysts and SLPs experienced in building Social Stories by hand, to learn exactly what clinical structure a generation pipeline must respect; five or more parents of current AAC users across different diagnoses; two to three special-education directors; staff at CommunicationFIRST and the Autistic Self Advocacy Network; a pediatric-hospital AAC clinic team such as Gillette Children's; and a disability-tech investor such as Adaptation Ventures or Remarkable.
- **Assumptions that must be validated:** that personalized, AI-generated visual and video content measurably improves engagement, comprehension, or anxiety outcomes compared to generic materials, not just novelty appeal that fades; that clinicians and behavior analysts will trust and

adopt AI-generated content inside an intervention as established and evidence-sensitive as Social Stories; that a workable consent and likeness-rights model exists for generating imagery of real family members, not only the primary child user; that families and schools will pay for this on top of an AAC app they already own.

- **Reasons this idea might fail:** the clinical-template constraint needed to keep generated content safe and evidence-aligned could limit personalization enough to blunt the product's central advantage; the third-party consent question could prove a harder legal and trust barrier than anticipated, slowing launch; an incumbent or a well-resourced new entrant could ship a comparable feature first once the underlying multimodal-generation capability becomes common knowledge; the clinician-gatekept sales cycle could be too slow for a typical venture-fund timeline.
- **What evidence would convince us to abandon it:** a rigorous pilot showing generated content does not measurably improve engagement or outcomes over well-made generic or manually-personalized materials once the novelty wears off; or clinicians and behavior analysts consistently judging AI-generated Social-Story-style content as clinically inappropriate or unsafe despite genuine, good-faith template constraints.
- **What evidence would convince us to aggressively pursue it:** a pilot showing meaningfully better engagement, comprehension, or anxiety-reduction outcomes than generic materials; genuine enthusiasm and willingness to use it expressed by behavior analysts and SLPs who currently build Social Stories by hand, the group with the most at stake in judging whether this is a real clinical tool or a novelty; and a workable, legally sound consent model for generating imagery of real family members.

Surprising discoveries that challenge the assumptions behind this entire project

1. **The dominant, best-documented problem in the field is not a technology problem at all.** Device abandonment (30–50%, stable for decades) is driven overwhelmingly by training, onboarding, and support gaps; meaning a large share of "build better AAC technology" venture theses, including several attractive-looking ones in this very report, are aimed at a secondary lever, not the primary one. Any new entrant should treat onboarding, caregiver coaching, and longitudinal engagement as core product surface area, not a nice-to-have.
2. **The single most price-defining fact in the hardware market is a Medicare reimbursement rule, not manufacturing cost or R&D sophistication.** "Build a cheaper/better device" is a far smaller opportunity than it appears, because the price gap between a locked $8,000 DME-billable tablet and an equivalent $250 consumer setup is a coding artifact, not an engineering one.

3. **A sophisticated, decades-old professional consensus (ASHA/APA/AAIDD) already firmly rejects several communication-authorship approaches (FC, RPM) that remain beloved by an emotionally invested, organized community of families and self-advocates.** Any AI-assisted communication company enters an unusually polarized trust environment where being perceived as "just another unvalidated authorship technology" is a real, live reputational risk; not a fringe concern to footnote.
4. **The most technically sophisticated and best-funded new entrants in adjacent technology (BCI companies; Synchron's $345M raised, Neuralink, Precision Neuroscience) are essentially irrelevant to the near-term pediatric AAC opportunity**, despite dominating tech-press coverage of "the future of AAC." They are adult-paralysis-focused, invasive, pre-commercial, and plausibly a decade or more from any responsible pediatric application; meaning the most-hyped technology in this space is among the least actionable for a company focused on children today.
5. **Open-source and volunteer-run projects have shipped genuinely novel AI features faster than well-funded incumbents**; Cboard's Cbuilder (AI board generation) went live in January 2026, ahead of any equivalent feature identified from Tobii Dynavox, PRC-Saltillo, or AssistiveWare in this research. This suggests capital intensity is not the binding constraint in this market; distribution, clinical trust, and reimbursement infrastructure are; a genuinely counterintuitive finding for anyone assuming a well-funded startup automatically out-innovates the field.
6. **The market-sizing literature for this sector is unusually unreliable** (vendor reports disagree by roughly 7x) **with major, acknowledged gaps** (no published Medicaid claims data for AAC-specific billing codes, no sizing of the SLP-therapy-services market that likely dwarfs the device market). Any pitch or plan built on a specific "$X billion market" claim lifted from a vendor report should be treated with real skepticism; and the more legitimate opportunity may lie in the unmeasured, currently-unserved population (roughly doubling the addressable market per §7.7) rather than in taking share of the well-measured, already-served one.

This synthesis (Parts 11, 20, 21, 22, and the Final Deliverable) was produced by the lead research/writing agent, drawing on and cross-referencing the five parallel research streams (A: History/Taxonomy/People/Organizations; B: Market/Competitors/Patents/Startups/Failures; C: Pediatric Populations/User Journeys/Controversies/User Voice; D: Money Flows/Market Size/Business Models; E: AI/Adjacent Industries/Build-vs-Buy/BCI Patents). All opportunity scores and rankings in Parts 11 and 20 are the analyst's own HYPOTHESIS-level judgment calibrated against the cited evidence, not independently benchmarked figures; a disciplined caveat this entire report has tried to apply consistently, per the research standards set out in the original brief.