SHARE: Synthesizing Heterogeneous Autism-support Records into Evidence-based Recommendations

Saumya Chauhan*, **Mila Hong***

California Institute of Technology
1200 E. California Blvd., Pasadena, CA 91125
saumya.s.chau@gmail.com, milazoehwork@gmail.com

Abstract

Supporting children with Autism Spectrum Disorder (ASD) requires highly individualized knowledge. However, critical information is often dispersed across documents such as Individualized Education Plans (IEPs), diagnostic assessments, and caregiver notes. Thus, we propose SHARE (Synthesizing Heterogeneous Autism-support Records into Evidence-based Recommendations), a framework that combines diverse autism-related documents into a concise, actionable set of recommendations for caregivers of children with ASD. Feedback is generated using OpenAI's large language model API, grounded in user-provided evidence with optional web-based extensions for missing details, and citation-linked. After caregivers attempt and then rate recommendations, SHARE uses a Bayesian bandit algorithm with Upper Confidence Bound (UCB) re-ranking to refine future advice. While previous work mostly focuses on drafting static goals, SHARE additionally combines LLM-generated recommendations, caregiver feedback, and interpretable ranking into a pipeline that can adapt over time.

Introduction

ASD (Autism Spectrum Disorder) is characterized by challenges in social, nonverbal, and repetitive communication and behavior [1] . It is estimated to affect 1 in 127 children globally (World Health Organization 2025), and 1 in 31 children in the US[1]. Over the past 20 years, there has been an almost 300% increase in ASD diagnoses, largely credited to (1) an expansion in the definition of ASD and (2) an increase in ASD screening for children between 18–24 months as promoted by public health programs (Public Health On Call 2025). With this increase in ASD diagnoses, further intervention and early childhood support are crucial.

However, scalable and homogeneous care for children with ASD is challenging, as ASD is characterized by clinical heterogeneity across affected genes, behavioral attributes, and developmental projections (Siller 2021). Therefore, ASD support necessitates a thorough, individualized understanding of a child's needs.

In practice, vital information such as sensory triggers, accommodations, and progress milestones are often dispersed across various documents, notes, or people's memories. This fragmentation leads to time-consuming onboarding, missed information, and low continuity of care, especially as staff turn over (Sulek et al. 2017).

Related Work

Previous ASD research largely focuses on developing diagnostic tools, such as at-home multimodal evaluations (Abbas et al. 2020). Beyond diagnostics, both academic and commercial applications have concentrated on Individualized Education Plans (IEPs): structured documents outlining a child's special needs, accommodations, and services (U.S. Department of Education 2000). Many efforts explore how LLMs can generate or summarize IEPs (Monsha 2025; Waterfield et al. 2025; Gómez 2025). In parallel, bandit algorithms, particularly UCB-style methods, are used in large-scale recommender systems to balance exploration and exploitation (Yi et al. 2023). SHARE builds on the aforementioned work by applying this mathematically grounded approach to autism-related document synthesis. More specifically, SHARE derives advice across multimodal documentation and delivers them to audiences beyond educators, such as caregivers, who often do not have enough to time read every document.

Our Approach

We present SHARE, a pipeline that converts historical ASD-related records into cited recommendations for caregivers. SHARE works in three steps: (1) it accepts common document formats or zip files, including IEPs, diagnostic reports, teacher notes, and caregiver feedback; (2) cleans and splits the text into smaller chunks with stable IDs; and (3) generates recommendations using OpenAI API's large language model with enforced citations. Missing information is flagged, and caregivers can further refine feedback in a simple Gradio-based UI.

SHARE combines multi-document, citation-grounded recommendations with an adaptive ranking process. To accomplish this, we use a lightweight bandit strategy (adapted from UCB-style selection in large-scale recommendation systems (Yi et al. 2023)) that updates probabilities based on caregiver ratings and balances exploration of new suggestions with reinforcement of effective ones.

*These authors contributed equally.

[1]Autism Speaks, https://www.autismspeaks.org/what-autism

System Architecture

Data Source Ingestion

SHARE first ingests user-uploaded documents, including diagnostic assessments, teacher observations, caregiver feedback, and/or IEPs. Providing more documents creates a more holistic understanding of a child's needs, triggers, and behaviors. Files can be uploaded in DOCX, TXT, or PDF format, as well as in a .zip folder with documents of these types. Next, text is parsed with the `pypdf` library for page-level extraction and `python-docx` for educator notes. Finally, SHARE standardizes whitespace, removes personal identifiers (when chosen by the user), and normalizes formatting to prepare text for downstream processing.

Chunking and Stable IDs

Next, the preprocessed text is tokenized using OpenAI's `tiktoken` library (`cl100k_base`). The text segmented into windows of about 750 tokens and each segment is assigned a stable identifier that acts as a citation, helping trace each recommendation back to its source.

Recommendation Generation

Through a Gradio interface launched from a colab notebook[2], caregivers can trigger recommendation generation using the OpenAI API with `gpt-4.1-mini`. Outputs must contain (1) a brief summary, (2) 5–10 recommendations with at least one citation in the form `(source: file#chunkNN)`, and (3) a "missing information" section. No recommendation is provided without cited evidence. If enabled, a second pass with `web_search` fills gaps with web citations of the form `(web: Title, URL)`.

Post-Processing and UI

After recommendations are generated, SHARE saves the result as Markdown and JSON files. We use a parser to extract citated lines into a schema with IDs, text, and version history. In the Gradio UI, caregivers can edit and rate recommendations on a scale from 1–10. These ratings help refine and update the top-k recommendations selected by the ranking algorithm. Results can be exported as a `.zip` bundle or used to restart the cycle of re-rating and re-ranking on the refined set.

Feedback loop. We use ratings to inform both recommendation re-ranking and the selection of rewrite targets: low-mean items with sufficient ratings and the current top-K by UCB.

Ranking Algorithm

We use a ranking algorithm to convert the user-generated ratings into Bayesian posteriors. Each recommendation starts with a uniform prior as modeled by a Beta distribution:

$$\text{prior} \sim \text{Beta}(\alpha{=}1,\ \beta{=}1).$$

[2]The SHARE pipeline (Colab notebook, sample inputs, and exports) is available online: https://drive.google.com/drive/folders/1VRlK3cCSdp7fWttepfl1l96YZJl0DojV

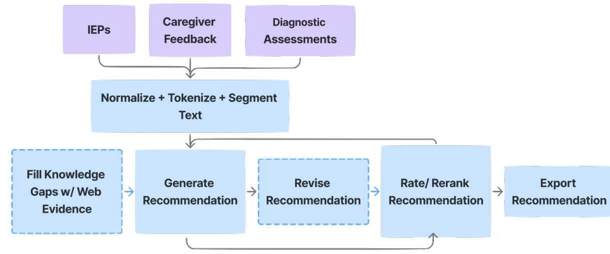

Figure 1: **System pipeline.** Document ingestion, chunking, generation with citations, re-ranking, rewriting, and exporting.

We convert each rating $r \in \{1, \ldots, 10\}$ into a fractional score $p = r/10$ and update the posterior parameters as

$$\alpha \leftarrow \alpha + p, \quad \beta \leftarrow \beta + (1 - p).$$

Each recommendation has a stable ID. Caregiver ratings update a $\text{Beta}(\alpha, \beta)$ posterior per ID. When an item is rewritten, it retains the same ID, keeping its posterior; when the system introduces a new item, it starts from $\text{Beta}(1, 1)$.

We then define μ as an estimation of the usefulness or effectiveness of a recommendation to the user.

$$\mu = \frac{\alpha}{\alpha + \beta}$$

Thus, as high ratings increase α, the recommendation's current μ or usefulness increases. Likewise, as low ratings increase β, the recommendation's current μ decreases. Since users do not have to rate every item, but we want to encourage the exploration of untried or less-rated recommendations, we also include an Upper Confidence Bound (UCB):

$$\text{score} = \mu + \sqrt{\frac{2 \ln T}{\alpha + \beta}},$$

where T is the total number of ratings across all items. As the user rates more and more recommendations, the second term decreases, giving new/rarely rated items a higher score.

Conclusion and Future Work

We presented SHARE, a framework that synthesizes diverse autism-support records into citation-grounded recommendations that adapt to caregiver feedback. In future work, we aim to incorporate video from multiple settings as an additional modality and use embedding-based clustering with high-rated recommendation examples to generate more diverse and targeted recommendations.

References

Abbas, H.; Garberson, F.; Liu-Mayo, S.; Glover, E.; and Wall, D. P. 2020. Multi-modular AI Approach to Streamline Autism Diagnosis in Young Children. *Scientific Reports*, 10.

Gómez, S. B. 2025. AI for Family Advocacy and Learning: Making Individualized Education Plans Accessible. https://innovate-us.org/ai-for-family-advocacy-and-learning-making-individualized-education-plans-accessible/. Accessed: 2026-01-07.

Monsha. 2025. How to Generate IEP Goals Using AI. https://monsha.ai/blog/how-to-generate-iep-goals-using-ai. Accessed: 2026-01-07.

Public Health On Call. 2025. Is There an Autism Epidemic? https://publichealth.jhu.edu/2025/is-there-an-autism-epidemic. Accessed: 2026-01-07.

Siller, M. 2021. Editorial: Individualizing Interventions for Young Children With Autism: Embracing the Next Generation of Intervention Research. *Journal of the American Academy of Child and Adolescent Psychiatry*, 60(6): 680–682.

Sulek, R.; Trembath, D.; Paynter, J.; Keen, D.; and Simpson, K. 2017. Inconsistent staffing and its impact on service delivery in ASD early-intervention. *Research in Developmental Disabilities*, 63: 18–27.

U.S. Department of Education. 2000. A Guide to the Individualized Education Program. https://www.ed.gov/sites/ed/files/parents/needs/speced/iepguide/iepguide.pdf. Accessed: 2026-01-07.

Waterfield, D. A.; Coleman, O. F.; Welker, N. P.; Kennedy, M. J.; McDonald, S. D.; and Cook, B. G. 2025. IEPs in the Age of AI: Examining IEP Goals Written with and Without ChatGPT. *Journal of Special Education Technology*.

World Health Organization. 2025. Autism. https://www.who.int/news-room/fact-sheets/detail/autism-spectrum-disorders. Accessed: 2026-01-07.

Yi, X.; Wang, S.-C.; He, R.; Chandrasekaran, H.; Wu, C.; Heldt, L.; Hong, L.; Chen, M.; and Chi, E. H. 2023. Online Matching: A Real-time Bandit System for Large-scale Recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems*, RecSys '23, 403–414. New York, NY, USA: Association for Computing Machinery.