Context-Aware Recommendation System for Augmentative and Alternative Communication

Leonardo S. V. Filho[1] **, Beatriz P. Silva**[1]**, Taciana P. Falcão** [1]**,**
Rafael F. L. Mello [1,2] **, André C. A. Nascimento** [1]

[1] Departamento de Computação, Universidade Federal Rural de Pernambuco (UFRPE)
Recife, Brazil

[2]Cesar School, Recife, Brazil
{leonardo.vianafilho, beatriz.pereiras, taciana.pontual}@ufrpe.br
{rafael.mello, andre.camara}@ufrpe.br

Abstract. *Augmentative and Alternative Communication (AAC) systems rely on large vocabularies of pictograms, the navigation of which can be demanding, especially in spontaneous communication in different environments. We propose a context-sensitive pictogram recommendation system that integrates spatial clustering of GPS coordinates and fuzzy temporal modeling by Gaussian functions, ordering pictograms by combining temporal likelihood and frequency of use with temporal decay. We compare this probabilistic approach with a Random Forest model under the same temporal evaluation protocol, using more than 671,000 real interactions from 27 users of a real AAC application. The probabilistic model achieves a Recall@5 of 0.28 and an MRR@10 of 0.17, consistently outperforming the Random Forest and establishing an interpretable benchmark for future comparisons.*

Resumo. *Sistemas de Comunicação Aumentativa e Alternativa (CAA) baseiam-se em grandes vocabulários de pictogramas cuja navegação pode ser exigente, sobretudo na comunicação espontânea em diferentes ambientes. Propomos um sistema de recomendação de pictogramas sensível ao contexto que integra agrupamento espacial de coordenadas GPS e modelagem temporal difusa por funções gaussianas, ordenando os pictogramas pela combinação de verossimilhança temporal e frequência de uso com decaimento temporal. Comparamos essa abordagem probabilística com um modelo Random Forest sob o mesmo protocolo de avaliação temporal, usando mais de 671.000 interações reais de 27 usuários de uma aplicação real de CAA. O modelo probabilístico atinge Recall@5 de 0,28 e MRR@10 de 0,17, superando consistentemente o Random Forest e estabelecendo uma interpretável para comparações futuras.*

1. Introduction

Augmentative and Alternative Communication (AAC) systems support individuals with speech, language, or motor impairments by enabling communication through visual symbols called pictograms. The Picture Exchange Communication System (PECS) has proven effective in developing requesting skills and in promoting retention and generalization of communication abilities across different contexts in children with multiple disabilities (Alfuraih et al. 2024). Recent icon-based AAC systems have also adopted

PECS-inspired training protocols, including a structured progression from low-tech resources such as flashcards to high-tech communication interfaces (Srinivasan et al. 2022). Mobile AAC applications provide these resources on consumer devices, e.g., tablets, allowing nonverbal users to express themselves by selecting visual cards that include both standard symbols and personalized content tailored to individual needs. As these applications are deployed across diverse daily environments, their vocabularies grow considerably, creating navigation challenges that directly affect communication fluency and spontaneity (Neamtu et al. 2019).

Despite their communicative value, AAC systems commonly rely on grid-based interfaces that require users to navigate through large numbers of pictograms to locate the appropriate symbol. This interaction overhead imposes a significant cognitive load, particularly for beginning communicators and users with neurodivergent profiles (Zastudil et al. 2024). The challenge intensifies across different usage environments — such as home, school, or outdoor settings — as the set of relevant pictograms shifts according to situational context. Context-aware recommendation has emerged as a promising strategy to address this challenge, yet prior work has focused primarily on isolated modeling choices (e.g., location-only or frequency-only approaches (Garcia et al. 2016)), leaving open the question of how different recommendation paradigms compare under a unified, realistic experimental setup.

In this work, we propose a context-aware pictogram recommendation system that jointly exploits spatial and temporal situational cues together with individual usage history. In this setting, the recommendation task can be interpreted as a multiclass classification problem, where each pictogram represents a candidate class, among the available classes in the considered time window (Fig. 1).

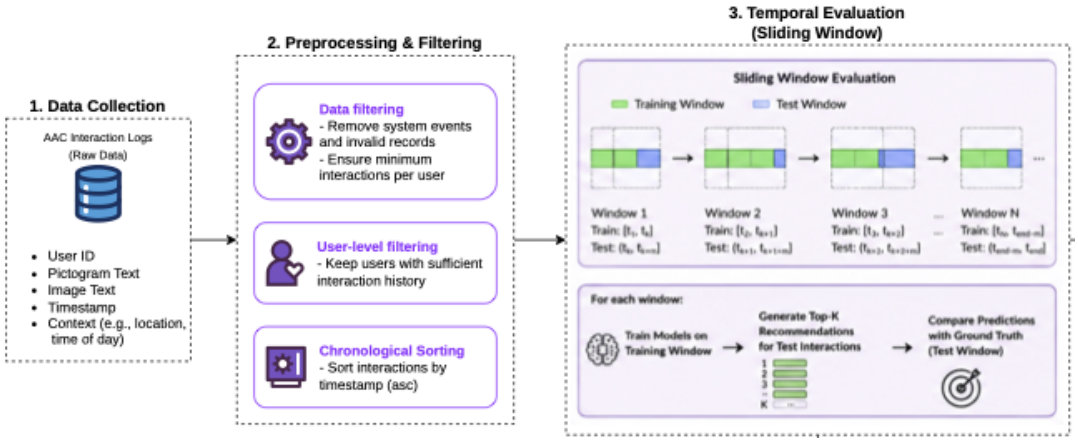


Figure 1. Overview of the proposed context-aware pictogram recommendation data preprocessing and temporal sliding-window evaluation.

The overall system architecture is illustrated in steps 1 and 2 of Figure 1. Spatial context is modeled through density-based clustering of GPS coordinates, and temporal context through fuzzy Gaussian membership functions that enable smooth transitions between periods of the day; these signals are combined with temporally decayed usage frequency to produce a final ranking score for each candidate pictogram.

To situate this approach within the broader landscape of recommendation strate-

gies, we compare it against a Random Forest model trained on the same contextual features under the same rolling temporal evaluation protocol, enabling an empirical analysis of the trade-offs between interpretable, parameter-free probabilistic reasoning and supervised learning on real interaction logs from an AAC application.

The contributions of this paper are threefold:

(i) the design of a context-aware probabilistic recommendation model that integrates spatial clustering and fuzzy temporal modeling of user behavior in AAC systems;
(ii) an empirical comparison between a probabilistic approach and a supervised learning model (Random Forest) for pictogram prediction, conducted under a unified experimental protocol with shared contextual representations; and
(iii) an analysis based on real-world interaction data from a deployed AAC application, providing insights into the design of effective and scalable recommendation systems for assistive communication.

2. Related work

AAC systems commonly rely on grid-based interfaces composed of pictograms, requiring users to navigate through multiple options to construct messages or to communicate basic needs (e.g., 'I need to go to the bathroom' or 'I am thirsty'). While effective, this interaction paradigm can impose a significant cognitive load, particularly for beginning communicators and users with neurodivergent profiles (Zastudil et al. 2024). The effort is even greater for users with motor disabilities such as cerebral palsy, for whom a single touch on the screen may demand substantial control of the muscles involved in the movement. In such contexts, reducing the number of interactions required to communicate a message is essential to support fluency and spontaneity, motivating adaptive mechanisms that prioritize relevant symbols according to the user's current situation.

Pictogram-based communication, as adopted in applications such as Livox (Neamtu et al. 2019), Proloquo2Go (AssistiveWare 2025) and Tobii Dynavox (Tobii Dynavox 2019), offers visual support that enables nonverbal users to express needs, actions, and emotions. However, as the size of pictogram vocabularies increases, locating the appropriate symbol becomes progressively more difficult, especially across different environments such as home, school, or outdoor settings. This challenge motivates the use of adaptive mechanisms capable of prioritizing relevant symbols according to the user's current situation.

Context-aware recommendation systems (CARS) have emerged as a promising strategy to address cognitive load in AAC by incorporating situational cues into the recommendation process. In AAC scenarios, contextual information such as location and time provides strong signals about users' communicative intent, enabling systems to proactively highlight relevant pictograms. Evangeline et al. (Evangeline and Moorthy 2024) emphasize the role of contextual factors such as location, time, and social context in adapting vocabulary suggestions to users' daily routines and supporting the independence of individuals with speech difficulties. A practical implementation of this context sensitivity is the AAC system of Garcia et al. (Garcia et al. 2016), which automatically adjusts the user's communicative profile according to the detected location, using GPS for outdoor areas and Bluetooth for indoor environments. More recent work explores modular and interpretable architectures that

integrate linguistic normalization, semantic retrieval, and context-based re-ranking to refine pictogram prediction under noisy user input (Evangeline et al. 2026). Together, these observations motivate recommendation strategies that leverage spatial and temporal information to reduce cognitive load while preserving the familiar structure of pictogram grids.

Prior research on pictogram prediction has evolved from statistical frequency methods toward machine learning and deep learning approaches. Pereira et al. (Pereira et al. 2022a) conducted a systematic mapping study summarizing existing techniques for next-pictogram prediction on AAC boards, concluding that methods are progressively adopting recurrent neural networks and Transformer architectures due to their ability to capture sequential temporal patterns. The study further shows that hybrid or mixed strategies combining local user data with general contextual signals tend to outperform purely local or purely general models. A related but distinct line of research has focused on sequential item prediction in AAC, through transformer-based models that predict the next pictogram or communication card from the interaction history. PictoBERT adapts BERT to next-pictogram prediction and reports gains over n-gram and knowledge-based baselines (Pereira et al. 2022b), while PrAACT extends this line by emphasizing adaptability to user-specific vocabularies (Pereira et al. 2024). This is a different problem from the one addressed here: whereas those studies focus on sequence modeling for next-item prediction, our work investigates context-aware recommendation under a controlled comparison between an interpretable probabilistic approach and a supervised model sharing the same preprocessing, contextual representation, and temporal evaluation protocol.

Despite these advances, a direct empirical comparison between interpretable probabilistic approaches and supervised learning methods for pictogram recommendation, under identical data preprocessing, contextual representation, and temporal evaluation protocol, remains largely unexplored. Garcia et al. (Garcia et al. 2016) demonstrated gains from combining location context and sentence reuse, but did not incorporate temporal modeling or compare against learning-based alternatives. This work addresses that gap by evaluating a probabilistic context-aware model alongside a Random Forest model on real interaction data from a deployed AAC application, enabling a controlled analysis of the trade-offs between these two recommendation paradigms.

3. Methods

This section describes the two on-device pictogram recommendation models proposed in this work and the temporal evaluation protocol applied to both: a parameter-free probabilistic model that exploits spatial and temporal signals from each user's interaction history, and a Random Forest supervised alternative trained on the same contextual features. Both share the data representation produced by the preprocessing and clustering steps of Section 4.

3.1. Pictogram usage feature extraction

Relative Frequency with Temporal Decay: Pictogram usage is quantified through a relative frequency augmented with temporal decay, where each interaction contributes a weight that decays exponentially with age, so recent interactions exert greater influence than older ones. Formally, let $\mathcal{H}$ denote the set of interactions in the current training

window, p_i the pictogram selected at interaction i, t_i its timestamp, and t_{now} the current inference time, i.e., the timestamp of the test interaction being predicted. The exponential weight of interaction i is

$$w_i \;=\; e^{-\lambda\,(t_{\text{now}}-t_i)}, \qquad \lambda = \frac{\ln 2}{t_{1/2}}, \tag{1}$$

where λ is the decay rate derived from a predefined half-life $t_{1/2}$. The relative frequency of a candidate pictogram p is then

$$\hat{f}(p) \;=\; \frac{\displaystyle\sum_{i\in\mathcal{H}} \mathbf{1}[p_i = p]\cdot w_i}{\displaystyle\sum_{j\in\mathcal{H}} w_j}, \tag{2}$$

where $\mathbf{1}[\cdot]$ is the indicator function. Equation 2 normalizes the decayed counts of p by the total decayed mass of the window, balancing recency against gradual forgetting of outdated patterns.

Fuzzy Temporal Modeling: Temporal usage patterns are represented using fuzzy logic applied to the hour of day. The day is partitioned into K overlapping temporal periods, each characterized by a Gaussian membership function centered at a reference hour c_k with spread σ_k. For an interaction occurring at hour $h \in [0, 24)$, the degree of membership in period k is

$$\mu_k(h) \;=\; \exp\!\left(-\frac{(h-c_k)^2}{2\,\sigma_k^2}\right), \qquad k = 1,\ldots,K. \tag{3}$$

This continuous formulation avoids the artificial boundaries of hard time discretization: an interaction at 11:50 contributes to both the *Morning* and *Noon* periods, reflecting gradual behavioral shifts. Figure 2 illustrates the six membership functions used in this work ($K = 6$, centers at $c_k \in \{0, 4, 8, 12, 16, 20\}$ h, $\sigma_k = 3$ h for all k), showing the smooth overlap between adjacent periods.

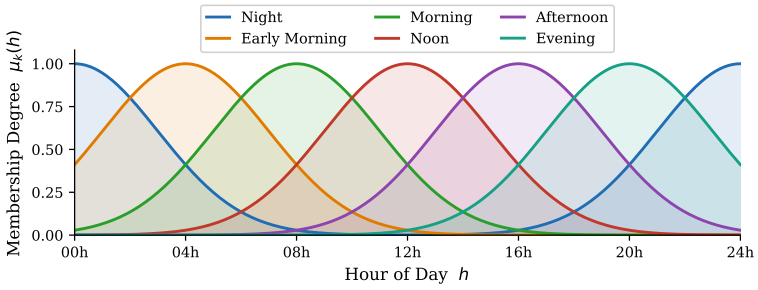


Figure 2. Gaussian membership functions defining the six temporal periods used for fuzzy temporal modeling (K=6, σ=3 h). The overlap between adjacent functions enables smooth transitions and avoids hard temporal boundaries.

Based on the training history, the membership vector $\boldsymbol{\mu}(h_i) = [\mu_1(h_i), \ldots, \mu_K(h_i)]^\top$ of each interaction i is accumulated in a weighted manner

also subject to temporal decay w_i (Eq. (1)) to construct for each pictogram p a fuzzy temporal profile.

$$\boldsymbol{\phi}(p) \;=\; \frac{\sum_{i \in \mathcal{H}} \mathbf{1}[p_i = p] \cdot w_i \cdot \boldsymbol{\mu}(h_i)}{\sum_{i \in \mathcal{H}} \mathbf{1}[p_i = p] \cdot w_i}, \tag{4}$$

which summarizes the pictogram's typical usage distribution across the periods of the day.

3.2. On-device Context-Aware Recommender Models

In this section we present the two on-device recommender models compared in this work: the proposed probabilistic ranking model and a supervised multiclass classifier used as a learning-based reference. To avoid ambiguity, throughout the paper the probabilistic model is the primary proposed contribution, while the Random Forest serves as the supervised point of comparison.

Probabilistic Ranking Model: The probabilistic recommendation model exploits historical usage patterns and temporal information without any supervised learning. User activity is modeled independently for each individual using only the interaction history available in the training data, with strict respect for temporal causality. The training window expands cumulatively until reaching a predefined maximum duration; once this limit is reached, a fixed-size sliding window is applied, continuously incorporating new interactions while discarding older ones. At inference time, the fuzzy membership vector $\boldsymbol{\mu}(h_{\text{now}})$ of the current timestamp is computed, and the temporal likelihood of each candidate pictogram p is obtained as the inner product between this vector and the pictogram's temporal profile (Eq. (4)),

$$\ell(p) \;=\; \boldsymbol{\mu}(h_{\text{now}})^{\top}\, \boldsymbol{\phi}(p), \tag{5}$$

which reduces the K-dimensional profile to a scalar. Both the relative frequency $\hat{f}(p)$ and the temporal profile $\boldsymbol{\phi}(p)$ are estimated from the interactions belonging to the user's currently active spatial cluster (Section 4.3), so that spatial context conditions the candidate statistics. The final ranking score is the product $s(p) = \hat{f}(p) \cdot \ell(p)$; candidates are sorted in descending order of $s(p)$ and the top-ranked pictograms are returned as recommendations.

Multiclass Supervised Learning Recommender: As a supervised learning alternative for comparison, we trained a Random Forest classifier (Breiman 2001) on the same contextual features used by the probabilistic model. Since the per-user vocabulary remains bounded (typically a few hundred distinct pictograms, as detailed in Section 5), it can be treated directly as a multiclass classification problem, with each unique pictogram as a distinct class. This formulation avoids a prior candidate-filtering stage: instead of a two-step retrieve-then-rank pipeline, which risks discarding relevant pictograms before scoring, the classifier produces a probability distribution over all classes in a single pass, and the resulting ranked list is used directly as the recommendation output.

The feature vector for each interaction comprises: (i) the spatial cluster identifier obtained from Section 4.3, encoded as a categorical variable via label encoding; (ii) the fuzzy Gaussian membership values for each temporal period (Section 4.2), providing a continuous representation of the time of interaction; and (iii) the relative usage frequency of each candidate pictogram within the current training window. The target label is the identifier of the selected pictogram.

4. Materials

This section describes the dataset, experimental environment, and data preparation pipeline adopted across all experiments. It also details the behavioral clustering procedure and the activity-richness criteria used to select users suitable for context-aware evaluation.

4.1. Dataset and Experimental Environment

The dataset used in this study was extracted from the Livox AAC application and consists of records of real user interactions with an AAC system. Each record contains information about the selected pictogram, the click location, the interaction timestamp, and an anonymous user identifier. All personal information was kept confidential, ensuring compliance with the Brazilian General Data Protection Law (Lei Geral de Proteção de Dados — LGPD, Law No. 13.709/2018).

The initial dataset comprised 1,672,888 records from 1,002 unique users distributed across 18 columns. Experiments were conducted in a Python 3.12.3 environment, using Pandas and NumPy for data manipulation and Scikit-learn for auxiliary algorithms including DBSCAN, BallTree, LabelEncoder, and OneHotEncoder.

4.2. Data Preparation and Preprocessing

The following columns were selected as primary features for training: `user_uuid`, `click_location`, `card_written_text`, and `event_timestamp`. Rows containing null values in any of these columns were removed.

A common preprocessing pipeline was adopted across all experiments, with the final representation adapted to each model through a configuration parameter (`--model`); the feature extraction of Section 3.1 was applied equally in all folds.

4.3. Geospatial Location Clustering

To capture recurring spatial usage patterns, we applied density-based spatial clustering via DBSCAN over each user's geographic coordinates (latitude and longitude). DBSCAN was configured with the Haversine distance metric, $\varepsilon = 400\,\text{m}$, `min_samples` = 5, and `max_cluster_points` = 43,000, enabling the identification of regions with high click concentration while labeling isolated points as noise (cluster -1).

Clustering was performed per user on the training split. For users with many interactions, spatially stratified sampling reduced the number of points per cell before running DBSCAN, lowering computational cost without sacrificing geographic diversity. Cluster labels were then propagated to the full training and test sets via centroid-based quantization: each interaction was assigned to the nearest centroid (BallTree search, Haversine distance) within the 400 m threshold, or labeled -1 otherwise.

4.4. User Filtering

To ensure that models were trained and evaluated on data from users with sufficient temporal and spatial diversity, i.e., a prerequisite for meaningfully constructing spatial clusters and fuzzy temporal profiles, we applied a sequential set of activity-richness filters. A user was retained only if they simultaneously met all of the following criteria: (i) at least 50 total interactions; (ii) interactions spanning at least 3 distinct values of the

`click_location` field; (iii) activity distributed over at least 18 distinct hours of the day; and (iv) app usage across at least 3 distinct weeks as measured by `week_order`.

The `click_location` field was further decomposed into latitude and longitude, and records with invalid geographic values were removed. After this stage, the dataset was reduced to 671,022 records corresponding to 30 users. During the experimental phase, one additional user whose interaction patterns were inconsistent with pictogram-based communication and two users whose test splits were unusable due to extreme sparsity were excluded, resulting in a final dataset of **27 users**.

It is important to note that this filtering introduces a *selection bias*: results are representative of high-engagement users whose AAC usage spans multiple spatial and temporal contexts. Generalization to occasional users or users with single-location profiles would require separate evaluation.

5. Experimental Setup

Model evaluation follows a rolling temporal validation protocol with sliding windows, preserving the chronological order of interactions and avoiding future information leakage (Ji et al. 2023). For each user, the interactions are ordered chronologically and divided into successive training and testing segments. In each iteration, the model components are reconstructed solely from the training portion, and the recommendations are generated sequentially for each event in the test set. This procedure ensures strict temporal integrity and closely mirrors real-world deployment conditions in AAC systems, where a model must make recommendations based exclusively on past behavior.

Grouping users by test-proportion band, the low band (3–15% held-out, 85–97% training) encompasses 24 of the 27 users and the mid band (23–24% held-out) the remaining 2; this distribution arises naturally from the rolling protocol, which requires a substantial historical window to initialize the contextual profiles reliably.

The window configuration was empirically selected among four candidate settings, evaluated according to a trade-off between recommendation quality and the number of users that retained sufficient historical context for reliable profile initialization.

The hyperparameters considered for the Random Forest were: `n_estimators` = 150, `max_depth` = 22, `min_samples_leaf` = 5, and `class_weight` = `balanced` to mitigate the strong class imbalance inherent in large pictogram vocabularies. All other hyperparameters were kept at Scikit-learn defaults.

For the proposed model, three components were configured independently. Spatial clustering used DBSCAN with haversine distance, `eps` = 400 m and `min_samples` = 5, with clustering input capped at $43{,}000$ points per user to bound memory usage. Fuzzy temporal modeling represented time-of-day context through six Gaussian membership functions (dawn, morning, noon, afternoon, evening, night) centered every $\sim$4h across the 24h cycle, all with $\sigma = 2.0$. Usage frequency was weighted by recency via exponential half-life decay, $w(\Delta t) = 0.5^{\Delta t/h}$, with $h = 30$ days.

To characterize the complexity of the recommendation space, Fig. 3 presents the distribution of unique pictograms used per user. Usage is highly variable, with a median of 216 unique items and upper percentiles reaching 333 (P75) and 463 (P90), highlighting the large and sparse class space of the task. These characteristics reinforce the challenges

faced by supervised models, particularly class imbalance and limited per-class observations, helping to explain the performance gap observed for the Random Forest.

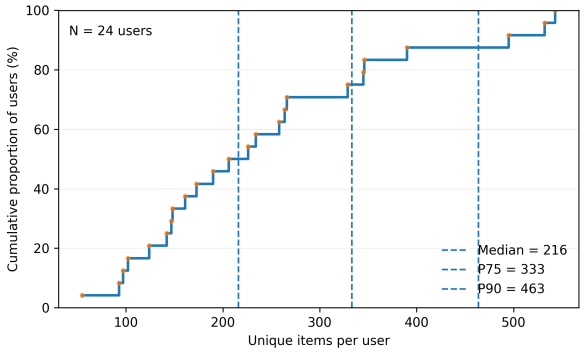


Figure 3. Empirical cumulative distribution of the number of unique pictograms used per user, illustrating the large and sparse class space of the recommendation task (median of 216 unique items, 333 at P75 and 463 at P90).

6. Results and discussion

This section reports the performance of the proposed probabilistic model and the Random Forest model under the rolling temporal evaluation protocol. Results are presented for the low test proportion band (3–15% held-out), which encompasses 24 of the 27 users and provides sufficient sample size for reliable aggregation (Table 1). Results for the mid-band group ($N = 2$) are omitted from aggregate tables due to insufficient statistical power and are discussed qualitatively.

Table 1. Accuracy, Recall and MRR (Mean $\pm$ Std) — low test proportion band ($N = 24$); higher is better (↑). *Base* is the proposed probabilistic model; *RF* the Random Forest.

Model	Stat.	Acc@1	R@3	R@5	R@10	MRR@3	MRR@5	MRR@10
Base	Mean	0.091	0.208	0.280	0.400	0.140	0.157	0.173
	Std	0.074	0.122	0.147	0.184	0.093	0.098	0.100
RF	Mean	0.044	0.111	0.161	0.254	0.073	0.084	0.097
	Std	0.055	0.110	0.143	0.180	0.077	0.084	0.088

The probabilistic model achieves Accuracy@1 of 0.091 and Recall@5 of 0.280, values that reflect the inherent difficulty of exact top-1 prediction in a recommendation space with potentially hundreds of pictogram candidates. Performance improves consistently as the recommendation cutoff increases: Recall grows from 0.208 at rank 3 to 0.400 at rank 10, and MRR@10 reaches 0.173, indicating that relevant pictograms are frequently retrieved within the upper positions of the ranked list even when they do not occupy the first slot. The standard deviations are moderate (e.g., $\pm$0.074 for Acc@1), reflecting reasonable consistency across users and evaluation windows despite differences in individual usage patterns.

The Random Forest model underperforms the probabilistic model across all metrics by a substantial margin: Acc@1 is halved (0.044 vs. 0.091) and MRR@10 drops by

44%. This is counterintuitive given its access to the same features and a richer hypothesis space, and we attribute it to three compounding factors: (i) the extreme class imbalance and long-tail usage distribution, with hundreds of candidate labels, hinder discriminative training even with class weighting; (ii) after temporal filtering, each user provides few samples per class, limiting the learning of reliable class boundaries; and (iii) the probabilistic model carries an inductive bias well-aligned with the domain, directly modeling recency and context without generalizing across classes, making it robust under data sparsity. These findings suggest that, for sparse, highly personalized AAC recommendation, interpretable probabilistic reasoning may be preferable to off-the-shelf supervised learning lacking dedicated strategies for label sparsity (e.g., embedding-based representations or hierarchical pictogram grouping).

Figure 4 illustrates the evolution of recommendation metrics over time for both the probabilistic and the Random Forest models, across all users along the time windows. The probabilistic model shows a gradual upward trend in Recall@5 and MRR as the training window expands and the contextual profiles stabilize, consistent with the cumulative learning design. The Random Forest exhibits higher variance across evaluation windows, which may reflect sensitivity to the class composition of each training fold.

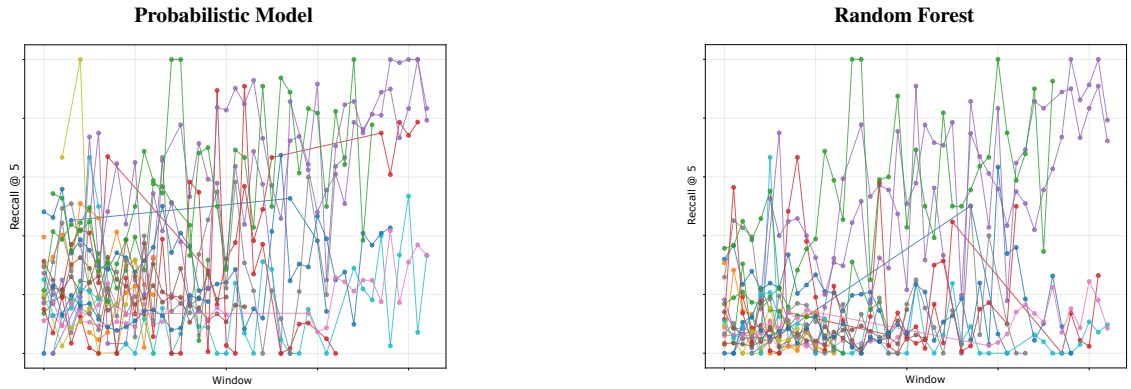


Figure 4. Evolution of recommendation performance (Recall@5) over time for the probabilistic model (left) and the Random Forest model (right), aggregated over all users along the evaluation windows.

To complement the aggregate metrics, Table 2 presents qualitative prediction examples from the final evaluation window for two representative users (16 and 28). The top-1 prediction does not always match the target, but the target frequently appears within the top-5 positions, consistent with the aggregate Recall@5. For User 28, both models repeatedly predict food-related pictograms (e.g., *peanut butter jelly sandwich*) when the target is also a food item (e.g., *crumb donut*), suggesting that this user's context is strongly associated with mealtime and that the main challenge is discriminating among semantically similar items within the same cluster — a limitation that semantic or embedding-based models could address.

7. Conclusions

This work investigated context-aware pictogram recommendation for AAC systems through a comparative evaluation of two paradigms on real interaction data from a deployed application. The proposed probabilistic model integrates spatial clustering and

Table 2. Probabilistic (Base) and Random Forest (RF) predictions in the final evaluation window. *Base Rank* and *RF Rank* give the position of the ground-truth (*Expected*) pictogram in each model's ranked list (lower is better).

User	Expected	Base Pred.	Base Rank	RF Pred.	RF Rank
16	supermercado do vovô	biscoito cookies toddy	20	bolinho bauducco	7
16	cheetos	biscoito cookies toddy	4	bolinho bauducco	4
28	crumb donut	peanut butter jelly sandwich	5	peanut butter jelly sandwich	5
28	crumb donut	peanut butter jelly sandwich	3	air mattress	8
28	chocolate donut	peanut butter jelly sandwich	3	peanut butter jelly sandwich	2

fuzzy temporal modeling to produce personalized rankings without supervised training, achieving Recall@5 of 0.280 and MRR@10 of 0.173 under a rolling temporal protocol. The comparison reveals an important empirical finding: a more expressive, supervised approach does not necessarily yield better performance in this domain, as the Random Forest underperformed across all metrics (Acc@1 of 0.044 versus 0.091). Under the data sparsity and extreme class imbalance characteristic of personalized AAC recommendation, a simple method with strong inductive bias, e.g., recency weighting and context-aware ranking, can be more effective than a general-purpose classifier trained on the same features.

Several limitations point to future directions. First, the activity-richness filtering reduced the dataset to 27 high-engagement users whose AAC usage spans multiple spatial and temporal contexts; extending the evaluation to occasional or single-location users would broaden the applicability of the findings. Second, the absence of ablation studies prevents attributing performance to individual components (spatial clustering, temporal decay, or fuzzy modeling), making their isolation an important next step. Third, the evaluation is restricted to exact pictogram matching, which may underestimate the utility of contextually plausible near-misses; since both models struggle to discriminate among semantically similar pictograms within the same contextual cluster, semantic representations (e.g., embedding-based or graph-based models) may provide complementary gains.

Acknowledgements. The authors thank the Livox Application Team for their support and collaboration. This work has also been supported by the following Brazilian research and innovation agencies: FINEP, FACEPE, CAPES and CNPq.

References

[Alfuraih et al. 2024] Alfuraih, R. K., Almalki, N. S., and AlNemary, F. M. (2024). Effectiveness of picture exchange communication system in developing requesting skills for children with multiple disabilities. *Frontiers in Psychology*, Volume 15 - 2024.

[AssistiveWare 2025] AssistiveWare (2025). Proloquo2Go: Symbol-based AAC application. `https://www.assistiveware.com/products/proloquo2go`. Accessed: 2026-04-08.

[Breiman 2001] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

[Evangeline et al. 2026] Evangeline, B., Moorthy, A. D., and Kuppa, K. S. (2026). Cprsat: A context-aware pictogram recommendation system – a speech assistive technology. *SoftwareX*, 34:102635.

[Evangeline and Moorthy 2024] Evangeline, S. B. and Moorthy, A. D. (2024). Improving communication for people with speech disabilities: Exploring context-aware recommendation systems in assistive technology. In *2024 International Conference on Intelligent Computing and Emerging Communication Technologies (ICEC)*, pages 1–6.

[Garcia et al. 2016] Garcia, L. F., de Oliveira, L. C., and de Matos, D. M. (2016). Evaluating pictogram prediction in a location-aware augmentative and alternative communication system. *Assistive Technology*, 28(2):83–92. PMID: 26479456.

[Ji et al. 2023] Ji, Y., Sun, A., Zhang, J., and Li, C. (2023). A critical study on data leakage in recommender system offline evaluation. *ACM Trans. Inf. Syst.*, 41(3).

[Neamtu et al. 2019] Neamtu, R., Camara, A., Pereira, C., and Ferreira, R. (2019). Using Artificial Intelligence for Augmentative Alternative Communication for Children with Disabilities. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11746 LNCS:234–243. ISBN: 9783030293802.

[Pereira et al. 2022a] Pereira, J., Medeiros, S., Zanchettin, C., and Fidalgo, R. (2022a). Pictogram prediction in alternative communication boards: a mapping study. In *Anais do XXXIII Simpósio Brasileiro de Informática na Educação*, pages 705–717, Porto Alegre, RS, Brasil. SBC.

[Pereira et al. 2022b] Pereira, J. A., Macêdo, D., Zanchettin, C., de Oliveira, A. L. I., and do Nascimento Fidalgo, R. (2022b). Pictobert: Transformers for next pictogram prediction. *Expert Systems with Applications*, 202:117231.

[Pereira et al. 2024] Pereira, J. A., Pereira, J. A., Zanchettin, C., and do Nascimento Fidalgo, R. (2024). Praact: Predictive augmentative and alternative communication with transformers. *Expert Systems with Applications*, 240:122417.

[Srinivasan et al. 2022] Srinivasan, S., Patel, S., Khade, A., Bedi, G., Mohite, J., Sen, A., and Poovaiah, R. (2022). Efficacy of a novel augmentative and alternative communication system in promoting requesting skills in young children with autism spectrum disorder in india: A pilot study. *Autism & Developmental Language Impairments*, 7:23969415221120749.

[Tobii Dynavox 2019] Tobii Dynavox (2019). My Tobii Dynavox. `https://www.tobiidynavox.com`. Accessed: 2026-04-08.

[Zastudil et al. 2024] Zastudil, C., Holyfield, C., Smith, J. A., Nguyen, H., and MacNeil, S. (2024). Predictive anchoring: A novel interaction to support contextualized suggestions for grid displays. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24, New York, NY, USA. Association for Computing Machinery.