

International Multidisciplinary Journal of Emerging Technologies and Applications (IMJETA)

Vol. 1, No. 3, pp. 68-78, June 2026

Received June 9, 2026; Revised June 12, 2026; Accepted June 25, 2026

Published June 30, 2026

Architectural Proposal for a Hybrid Recommender Engine for Accessible Educational Content

Luis David Jiménez Martínez^[0009-0002-4950-5176] and Francisco Javier Álvarez Rodríguez^[0000-0001-6608-046X]

Centro de Ciencias Básicas, Universidad Autónoma de Aguascalientes, 20131, Aguascalientes, Mexico

al221273@edu.uaa.mx, francisco.alvarez@edu.uaa.mx

Abstract. This paper presents an architectural proposal for a modular and adaptive hybrid recommender system designed to improve equitable access to PDF-based educational content for students with visual impairments. The system overcomes the limitations of traditional recommenders, which optimize rankings solely by thematic relevance while ignoring format-related barriers, by implementing a closed-loop approach that prioritizes technical accessibility prior to ranking the educational material. The architecture is structured into four independent layers: automated ingestion via multimodal artificial intelligence, dynamic student profiling, a sequential hybrid recommendation engine, and a semantic presentation interface. The technical viability of the proposal was successfully validated through a functional Python prototype integrated with the Gemini 2.5 Flash model from Google AI Studio, evaluating a repository of real-world documents split between the fields of algebra and programming. Experimental results demonstrated that the system accurately resolves the cold-start problem through an initial technical questionnaire and effectively mitigates the occurrence of false positives in practice, thanks to a reverse feedback loop that automatically updates functional security profiles and restricts resources reported with access flaws.

Keywords: Inclusive Recommender System; Digital Accessibility; Visual Impairment; Artificial Intelligence; Accessible Educational Content.

1. Introduction

Equitable access to digital educational materials constitutes one of the greatest challenges in contemporary higher education, especially for the student population with visual impairments (Seale, 2014). Over the last decade, Learning Object (LO) repositories have grown exponentially; however, a critical amount of these PDF-based resources is published without complying with elementary digital inclusion guidelines, such as the PDF/UA standard (ISO, 2014) or the Web Content

Accessibility Guidelines (WCAG) (W3C, 2024). The assessment of the degree of accessibility of these documents has traditionally been performed manually by experts, a process that is costly, inefficient, and prone to human error, making the mass auditing of digital academic catalogs impossible (Jiménez Martínez, 2025). As a consequence, students with visual impairments recurrently face severe access barriers, where screen readers and other assistive technologies are unable to interpret content due to latent structural flaws within the files (Vigo et al., 2013). In a broader global context, resolving these technical barriers directly contributes to the United Nations 2030 Agenda for Sustainable Development, specifically aligns with Sustainable Development Goal 4 (SDG 4: Quality Education), which advocates for equal access to all levels of education for persons with disabilities, and SDG 10 (SDG 10: Reduced Inequalities), which aims to empower and promote the social and economic inclusion of all individuals irrespective of their functional conditions.

The evaluation of digital document accessibility is an escalating challenge, as recent literature identifies a modern accessibility crisis in scholarly repositories. Over the last decade, shifts in publishing models and platform infrastructures have paradoxically contributed to a decline in document standard compliance (Kumar & Wang, 2024). Empirical studies across Latin American educational repositories confirm that a massive percentage of PDF assets fail basic accessibility testing, establishing severe barriers for regional inclusion (Acosta-Vargas et al., 2017). To mitigate this, initial efforts in the state of the art focused on building dedicated plugins to assist in the automated generation of structured elements (Zulfiqar et al., 2020), as well as implementing specialized frameworks like A11yPDF to bridge the inclusion gap (Aljedaani et al., 2024). More recently, benchmarking efforts have shifted toward evaluating automated testing and Large Language Model (LLM)-based validation tools to establish reliable datasets for automatic verification (Kumar et al., 2025). However, these evaluators function as standalone diagnostic tools, creating a clear gap for architectures that seamlessly couple automated structural auditing with dynamic user delivery.

To mitigate the cognitive overload involved in manually searching for legible materials, the use of educational recommender systems has been proposed (Tarus et al., 2018). In parallel, research in special education has highlighted the importance of adapting digital resources according to the functional requirements of learners with disabilities (Drigas & Ioannidou, 2013). Nonetheless, traditional recommendation approaches based on collaborative or content-based filtering suffer from severe structural deficiencies when applied to inclusion contexts (Burke, 2017). To address special needs, contemporary research has explored user interface adaptation (Loitsch et al., 2017), context-aware service selection frameworks (Namoune et al., 2022), context-aware recommender systems for learning (Verbert et al., 2012), and personalization incorporating Point of Interest (POI) accessibility levels or monitoring models for neurodivergent populations (Deng et al., 2021; Mauro et al., 2020; Santos et al., 2017). Despite these advancements, existing inclusive recommenders rarely evaluate the structural and legibility constraints of

deep document layouts in real-time. Furthermore, the critical phenomenon of "Cold Start" occurs where the system cannot generate accurate suggestions without prior user history (Pazzani & Billsus, 2007). Even more alarmingly, standard recommenders optimize algorithms exclusively based on thematic relevance or implicit tastes (Mehrabi et al., 2021). This algorithmic bias causes a highly relevant pedagogical resource to be placed at the top of the ranking, completely ignoring whether the file is technically legible or accessible for the specific functional condition of a particular student (Naghiaei et al., 2022).

To simultaneously resolve assessment automation and suggestion personalization, a paradigm shift toward algorithmic fairness is required. This implies designing architectures that prioritize technical access viability over generic market preferences (Ekstrand et al., 2019). Under this premise, the present research proposes the architecture of an adaptive and inclusive recommender system, structured in a modular and sequential manner across four autonomous operational layers. The proposal stands out by integrating an intelligent evaluation engine that audits documents without manual intervention, and a hybrid recommendation engine that processes cross-functional constraints.

The scientific and innovative value of this proposal lies in overcoming the limitations of existing systems through a feedback-driven approach (Salehie & Tahvildari, 2009). The system not only generates a dynamic model of the student by isolating their non-negotiable visual requirements, but also enables telemetry channels and inverse semantic interactions. In this way, direct user feedback regarding the practical usability of the resource allows the system to learn from false positives and refine algorithmic security thresholds in real time. To validate the technical and operational viability of the described architecture, a functional prototype was designed and implemented in Python, which delegates complex structural evaluation to foundational multimodal vision-language models (VLM) and executes hierarchical classification through a combination of symbolic inference engines and multicriteria probabilistic models.

2. Architectural Methodology

The methodology of this research is grounded in the design and interconnection of a software architecture oriented toward inclusion and algorithmic fairness. The proposed system transcends traditional thematic accuracy metrics to prioritize the technical viability of educational resources, resolving access barriers in an automated and adaptive manner.

The architecture of the inclusive recommender system has been designed under a modular and sequential approach, enabling the integration of artificial intelligence for technical assessment and a hybrid recommendation engine for experience personalization. The core objective is to transform a repository of conventional PDF documents into a personalized selection of educational resources that guarantee the inclusion of the user with visual impairments, as shown below in Figure 1.

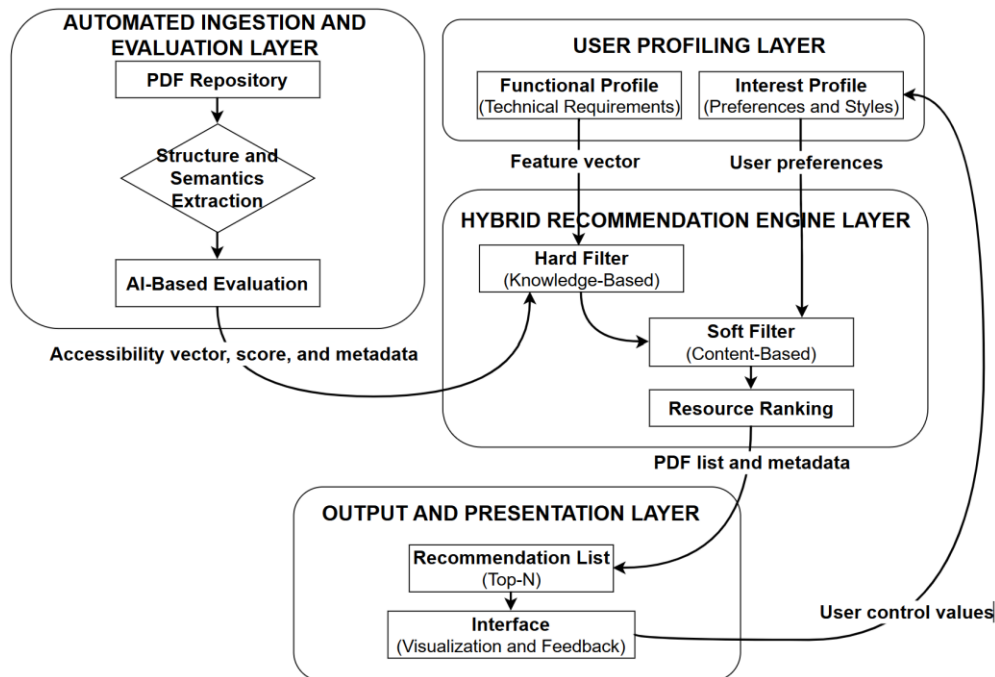


Figure 1. Recommender system architecture diagram

2.1. Automated Ingestion and Evaluation Layer

This layer functions as an automated data pipeline that processes PDF files to generate inclusion metadata without manual intervention (see Figure 2), consolidating the results into the following structure:

- Reception and Preprocessing Module: Validates the PDF files from the dataset and extracts their basic metadata.
- Structural Feature Extraction Module: Analyzes tags and the semantic architecture of the document.
- Artificial Intelligence Assessment Engine: Validates thirteen accessibility criteria (such as PDF/UA (ISO, 2014), hierarchies, reading order, and contrast) via natural language processing and extracts titles and short descriptions.
- Accessibility Score and Metadata Generator: Calculates the final accessibility percentage based on compliance with the previous criteria.

Record = {PDF_ID, [D_0, D_1, ..., D_12], Accessibility_Score, Topic_ID, Title, Description}

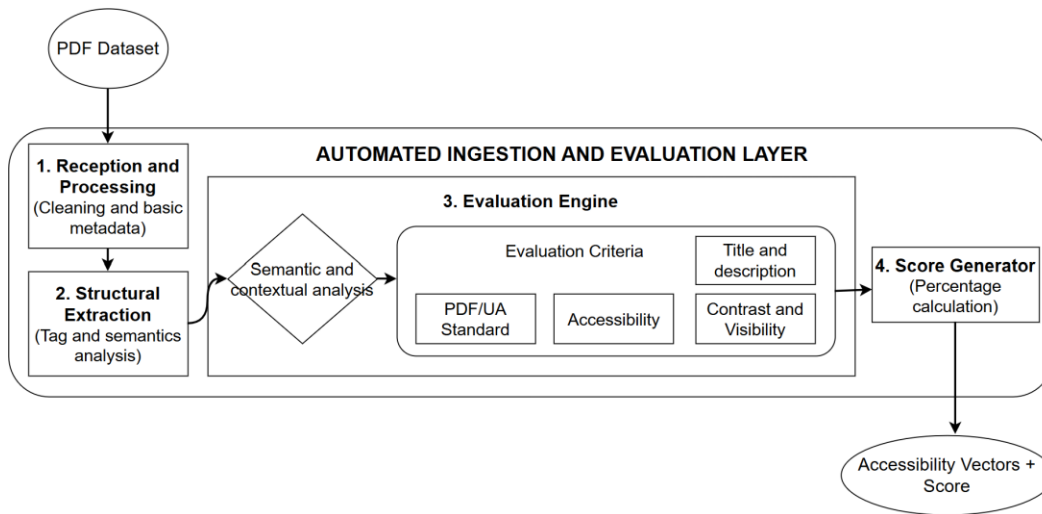


Figure 2. Internal process diagram of the automated ingestion and evaluation layer

2.2. User Profiling Layer

This layer creates, manages, and updates the user model through three submodules that dynamically feed the profile database (see Figure 3), emitting an updated vector in a consolidated format:

- Data Acquisition Module: Collects visual acuity and assistive software usage through an initial questionnaire to resolve the cold-start problem (Pazzani & Billsus, 2007).
- Functional Profile Module: Acts as a hard filter by centralizing user limitations and calculating the minimum required accessibility threshold.
- Interest Profile Module: Registers the topic of interest, captures student satisfaction (1-5), and adds failed resources to an exclusion list (Knijnenburg et al., 2012).

U-Model Package = {USER_ID, [(Topic_ID, Visual_Acuity, Min_Required_Score, Assistive_Req, Contrast_Pref, PDF_ID [Satisfaction_Scale, Accessibility_Validation]), Blocklist[]]}

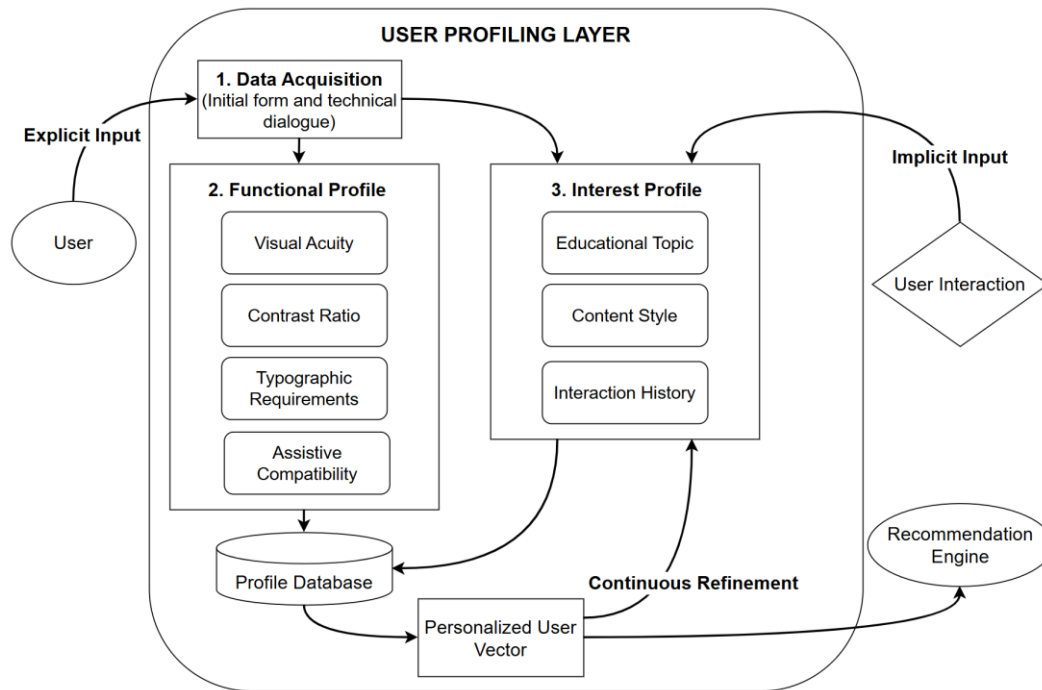


Figure 3. Block diagram of the user profiling layer

2.3. Hybrid Recommendation Engine Layer

As the system's decision center, this layer cross-references Layer 1 tagged documents with Layer 2 student profiles through a sequential double-filtering pipeline (see Figure 4):

- Knowledge-Based Filtering Submodule (Hard Filter): Contrasts the PDF accessibility vector against the user's functional profile via a rule engine, discarding non-viable materials to isolate "Accessible Candidates".
- Content-Based Filtering Submodule (Soft Filter): Hierarchizes the approved candidates through a probabilistic model that calculates resource utility by combining the AI score with the student's satisfaction history.
- Classification and Ranking Module: Orders results by relevance and selects the best N documents (Top-N) for deployment.

Top-N Ranking = {USER_ID, Topic_ID, [(Rank_Position, PDF_ID, Accessibility_Score, Calculated_Similarity, Title, Description)]}

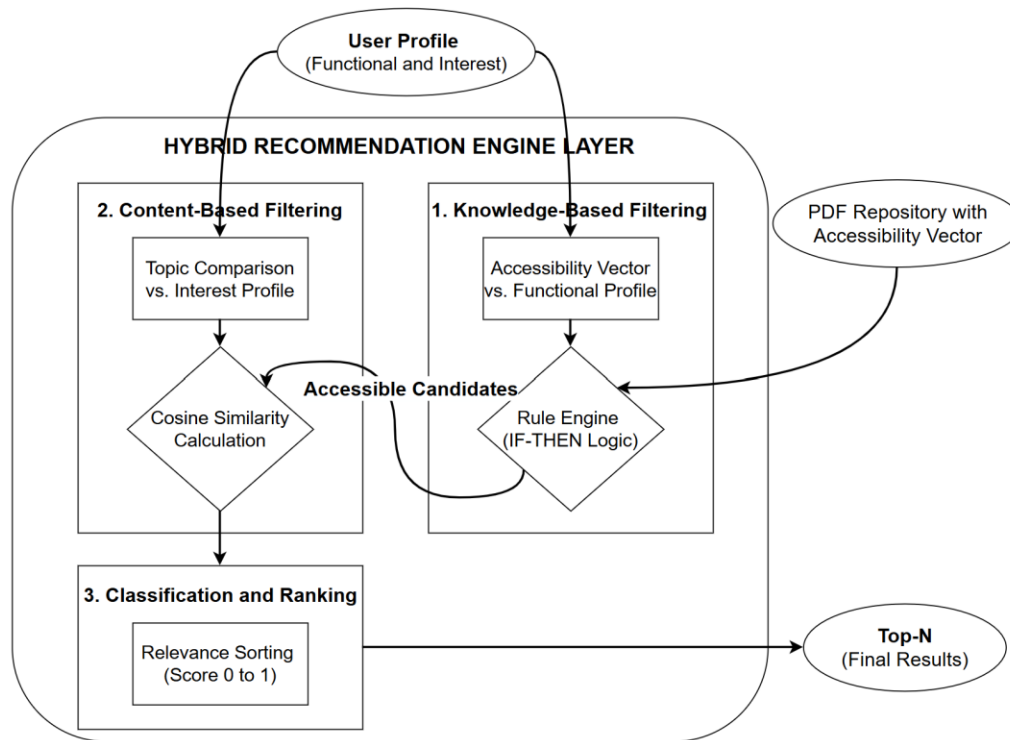


Figure 4. Double-filtering diagram of the hybrid recommendation engine

2.4. Output and Presentation Layer

This final phase of the architecture receives the hierarchized ranking of PDF documents from the preceding layer and manages the semantic interaction environment with the user. Its purpose is to capture explicit student feedback to effectively close the adaptive loop of the platform. It is divided into two operational components that result in two data structures:

- Results Rendering Module: Displays the final recommendations list (Top-N), semantically deploying the title, short description, and accessibility score obtained in Layer 1.
- User Feedback Module: Captures the student's rating regarding the utility and accessibility of the resource. It sends these control values to Layer 2 to refine preferences and future ordering criteria.

Interface Input = {USER_ID, Topic_ID, [(Rank_Position, PDF_ID, Accessibility_Score, Calculated_Similarity, Title, Description)]}

Control Telemetry = {USER_ID, PDF_ID [Satisfaction_Scale, Accessibility_Validation, Blocklist[]]}

3. Results

This section presents the results obtained from the implementation and experimental validation of the proposed architecture. To demonstrate the system's

technical viability, a functional prototype was developed in the Python programming language and subjected to a real-time case study simulation using a closed interaction loop.

3.1. Experimental Environment and Case Study

The Layer 1 pipeline delegates computer vision processing to the Google AI Studio API using the Gemini 2.5 Flash model, reducing response times to a few seconds. Layer 2 adaptive logic and Layer 3 recommendation engines were coded natively in Python for a lightweight execution environment. The test scenario utilized eight real-world PDF documents, evenly split between algebra and programming topics.

3.2. Cold-Start Mitigation

The first phase of experimentation evaluated the system's response upon adding a new user to the platform (USER_01), simulating a mild visual impairment condition. By executing explicit interaction through the Layer 2 interface, the following control variables were entered into the initial questionnaire: Topic_ID = 1 (Algebra), Visual_Acuity = 2 (Mild), and Assistive_Req = 0 (Not dependent on a screen reader). As a result of this profiling, the Layer 2 logical inference engine automatically calculated a strict security threshold of Min_Required_Score = 0.85.

Upon processing this context, the engine immediately filtered out the programming files due to a thematic discrepancy with the requested interest. Subsequently, it contrasted the logical rules against the four remaining algebra files, discarding any resource whose accessibility score fell below the minimum threshold or that presented a null value in the D12 dimension (Interoperability with Assistive Technologies), thereby isolating the "Accessible Candidates" subset. For the approved materials, the AI then estimated their final proximity index by mathematically weighting the technical quality score against the user's default initial satisfaction factor. The final semantic rendering from Layer 4 deployed the ordered listing (Top-N), validating that the system is capable of providing useful and inclusive recommendations from the very first interaction, effectively resolving the cold-start problem.

3.3. Feedback Validation

To verify the adaptive nature of the architecture and the mitigation of false positives in practice, a complete reverse feedback loop cycle was simulated. While reading the options on the Layer 4 interface, the user selected a specific resource and explicitly reported experiencing severe reading difficulties on their device, marking the document with a status of "Not Accessible". Upon dispatching this control event, the system demonstrated the convergence of its closed-loop design through immediate, sequential actions executed natively in Layer 2. Layer 2 intercepted the faulty document identifier (PDF_ID) and permanently indexed it within the student's in-memory exclusion list (Blocklist[]) to veto its recurrence in the future. Furthermore, the functional profiling submodule automatically recalculated the filter's strictness level—likewise based on the satisfaction scale assigned by the

user—dynamically raising the visual security threshold (Min_Required_Score) as an algorithmic protection measure against prior evaluation failures.

When an instantaneous catalog update was requested, the Layer 3 recommendation engine reprocessed the pipeline. The results in the validation console confirmed the success of the architecture, as the reported resource was completely omitted from the pipeline, and the Layer 4 interface reconfigured itself to display a new selection of academic alternatives shielded by a higher accessibility standard for the student.

4. Conclusion

The architecture proposed in this research demonstrates that it is possible to efficiently automate the structural assessment of accessibility in educational documents and couple it with a recommendation engine that adapts to the needs of students with visual impairments. Through experimentation with the functional prototype, it was verified that approaches based on logical rules and probabilistic decision models resolve the deficiencies of common recommenders, achieving a secure filtering process that prevents suggesting unreadable materials. The most valuable aspect of the system lies in its closed-loop operation, which allows the platform to learn directly from the practical experience expressed by the student.

Although a formal field-test protocol was not deployed, preliminary usability feedback from a user with visual impairments confirmed the functional viability of the interactive prototype. However, certain technical limitations must be acknowledged for future work. First, the Automated Ingestion and Evaluation Layer exhibits an operational dependence on the Google AI Studio infrastructure, which could introduce latency variations or cost scaling challenges in massive educational deployments. Second, since this architecture is strictly designed to evaluate and filter existing digital assets without altering their original file structures, documents with severe degradation or low-quality scans will be automatically discarded by Layer 1 due to non-compliance with minimal accessibility benchmarks.

5. References

- Acosta-Vargas, P., Luján-Mora, S., Acosta, T., & Salvador-Ullauri, L. (2017). Accesibilidad de documentos PDF en repositorios educativos de Latinoamérica. *Enfoque UTE*, 8(1), 162-177. <https://dialnet.unirioja.es/servlet/articulo?codigo=10207290>
- Aljedaani, W., Rudhravaram, S. K., Chintham, A., Habib, A., & Eler, M. M. (2024). A11yPDF: Bridging the gap to inclusive PDFs. *Proceedings of the 21st International Web for All Conference*, 37-38. <https://doi.org/10.1145/3677846.3677859>
- Burke, R. (2017). Multisided fairness for recommendation. In *Proceedings of the Workshop on Fairness, Accountability and Transparency in Machine Learning*. <https://doi.org/10.48550/arXiv.1707.00093>

- Deng, L., Rattadilok, P., & Xiong, R. (2021). A machine learning-based monitoring system for attention and stress detection for children with autism spectrum disorders. *Proceedings of the 3rd International Conference on Intelligent Medicine and Health*, 1–5. <https://doi.org/10.1145/3484377.3484381>
- Drigas, A., & Ioannidou, R. E. (2013). Special Education and ICTs. *International Journal of Emerging Technologies in Learning (ijET)*, 8(2), 41–47. <https://doi.org/10.3991/ijet.v8i2.2514>
- Ekstrand, M. D., Burke, R., & Diaz, F. (2019). Fairness and discrimination in retrieval and recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '19)* (pp. 1403–1404). Association for Computing Machinery. <https://doi.org/10.1145/3331184.3331380>
- ISO. (2014). *Document management applications – Electronic document file format enhancement for accessibility – Part 1: Use of ISO 32000-1 (PDF/UA-1)* (ISO Standard No. 14289-1:2014). International Organization for Standardization. <https://www.iso.org/standard/64599.html>
- Jiménez Martínez, L. D. (2025). *Diseño y desarrollo de una plataforma web de aprendizaje en línea integrando aula invertida y adaptada para usuarios débiles visuales* [Tesis de maestría, Universidad Autónoma de Aguascalientes]. Repositorio Institucional UAA. <http://bdigital.dgse.uaa.mx:8080/xmlui/handle/11317/3289>
- Knijnenburg, B. P., Willemsen, M. C., Gantner, Z., Soncu, H., & Newell, C. (2012). Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4–5), 441–504. <https://doi.org/10.1007/s11257-011-9118-4>
- Kumar, A., Padath, T., & Wang, L. L. (2025). Benchmarking PDF accessibility evaluation: A dataset and framework for assessing automated and LLM-based approaches for accessibility testing. *ASSETS '25: Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility*, 1–24. <https://doi.org/10.1145/3663547.3746380>
- Kumar, A., & Wang, L. L. (2024). Uncovering the new accessibility crisis in scholarly PDFs: Publishing model and platform changes contribute to declining scholarly document accessibility in the last decade. *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, 1–16. <https://doi.org/10.1145/3663548.3675634>
- Loitsch, C., Weber, G., Kaklanis, N., Votis, K., & Tzovaras, D. (2017). A knowledge-based approach to user interface adaptation from preferences and for special needs. *User Modeling and User-Adapted Interaction*, 27(3–5), 1–20. <https://doi.org/10.1007/s11257-017-9196-z>
- Mauro, N., Ardissono, L., & Cena, F. (2020). Personalized recommendation of PoIs to people with autism. *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, 1–8. <https://doi.org/10.1145/3340631.3394845>

- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Naghiaei, M., et al. (2022). *CPFair: Personalized consumer and producer fairness re-ranking for recommender systems*. arXiv. <https://doi.org/10.48550/arXiv.2204.08085>
- Namoune, A., Sen, A. A., Tufail, A., Alshantiti, A. M., Nawaz, W., & Benrhouma, O. (2022). A two-phase machine learning framework for context-aware service selection to empower people with disabilities. *Sensors*, 22(14), Article 5142. <https://doi.org/10.3390/s22145142>
- Pazzani, M. J., & Billsus, D. (2007). Content-based recommendation systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The adaptive web* (pp. 325–341). Springer. https://doi.org/10.1007/978-3-540-72079-9_10
- Santos, F., de Almeida, A., Martins, C., Gonçalves, R., & Martins, J. (2017). Using POI functionality and accessibility levels for delivering personalized tourism recommendations. *Computers, Environment and Urban Systems*, 77, 1–15. <https://doi.org/10.1016/j.compenvurbsys.2017.08.007>
- Salehie, M., & Tahvildari, L. (2009). Self-adaptive software: Landscape and research challenges. *ACM Transactions on Autonomous and Adaptive Systems*, 4(2), Article 14, 1–42. <https://doi.org/10.1145/1516533.1516538>
- Seale, J. (2014). *E-learning and disability in higher education: Accessibility research and practice* (2nd ed.). Routledge. <https://www.routledge.com/E-learning-and-Disability-in-Higher-Education-Accessibility-Research-and-Practice/Seale/p/book/9780415629416>
- Tarus, J. K., Niu, Z., & Mustafa, G. (2018). Knowledge-based recommendation: A review of ontology-based recommender systems for e-learning. *Artificial Intelligence Review*, 50(1), 21–48. <https://doi.org/10.1007/s10462-017-9539-5>
- Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnić, I., & Duval, E. (2012). Context-aware recommender systems for learning: A survey and future challenges. *IEEE Transactions on Learning Technologies*, 5(4), 318–335. <https://doi.org/10.1109/TLT.2012.11>
- Vigo, M., Brown, J., & Conway, V. (2013). Benchmarking web accessibility evaluation tools: Measuring the harm of sole reliance on automated tests. In *Proceedings of the International Cross-Disciplinary Conference on Web Accessibility (W4A '13)*. Association for Computing Machinery. <https://doi.org/10.1145/2461121.2461124>
- W3C. (2024). *WCAG 2 overview*. World Wide Web Consortium. <https://www.w3.org/WAI/standards-guidelines/wcag/>
- Zulfiqar, S., Arooj, S., Umar Hayat, S., Shahid, S., & Karim, A. (2020). Automated generation of accessible PDF. *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, 1–3. <https://doi.org/10.1145/3373625.3418045>