



Freeversity: Syllabus-Aware Personalization of Free Educational Resources

Aniruddh Dhurve, Mubashshir Khan, Aaryan Sahu, and Alok Meshram

Department of Computer Science and Data Science, Guru Nanak Institute of Engineering and Technology,
RTM Nagpur University, Nagpur, India

E-mail: (author e-mail IDs and corresponding-author designation to be added before submission)

How to Cite this Article:

Dhurve, A., Khan, M., Sahu, A. & Meshram, A. (2026). Freeversity: Syllabus-Aware Personalization of Free Educational Resources. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(10), 1-9.
<https://doi.org/10.55041/ijcope.v2i10.165>

License:

This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

© The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.



<https://doi.org/10.55041/ijcope.v2i10.165>

Abstract—Free educational resources are abundant but fragmented across providers, formats, difficulty levels, languages, and curriculum structures. This paper presents Freeversity, a lightweight prototype that converts a learner profile and an uploaded syllabus into a module-organized learning roadmap. The methodology combines an 11-question placement quiz, cumulative level filtering, content-modality weighting, language and resource filters, PDF/DOCX/TXT extraction, and a hybrid score of TF-IDF cosine similarity and tag overlap. In a descriptive case study on a 39-resource catalogue, seven roadmap modules, and one bundled sample syllabus, the heuristic assigned non-zero syllabus-match scores to 38 of 39 resources; the highest score was 46.5/100 (Khan Academy: Statistics and Probability). These are heuristic ranking scores, not probabilities or accuracy estimates. Eighteen core-logic tests passed; nine API tests could not be run because Flask was unavailable. A validation protocol based on expert relevance judgments, ranking metrics, ablations, and inter-rater agreement is specified. No learning gains or superiority over neural recommenders are claimed.

Keywords—educational recommender systems, syllabus alignment, personalized learning paths, free educational resources, TF-IDF, explainable ranking.

1. INTRODUCTION

Free online learning materials have expanded through university courseware, video platforms, open textbooks, and provider-hosted training, and the role of massive open online courses continues to evolve [34], [50]. Access alone, however, does not yield a coherent study plan: learners must decide which materials address a syllabus, at what difficulty, in which format, and in what order, and completion rates in open online courses are typically low [49]. Recommender systems for learning have been surveyed extensively as a way to

reduce this burden [15], [23], [46], [47].

Open educational resources (OER) have a specific meaning: materials in the public domain or released under an open license that permits reuse, adaptation, and redistribution [32], [37]. Free-to-access material is not necessarily openly licensed, and research on OER emphasizes efficacy, cost, and perception issues separately from licensing [40]. Because the license status of every entry in the current catalogue is unverified, this paper uses the term “free educational resources”.

Learning-path personalization has been studied in terms of learner models, resource selection, sequencing, and evaluation [16], [18], [27]. Matching a curricular need to resource content is a retrieval problem [28], [30], and recent work applies knowledge graphs and large language models (LLMs) to educational recommendation [1], [2], [3], [4], [7]. Such models add capability but also cost, data requirements, and reduced inspectability, and evidence from the wider recommendation and retrieval literature shows that simple, well-tuned baselines are often competitive and should be reported [14], [33], [35]. Freeversity therefore adopts a transparent lexical baseline.

This paper treats Freeversity as a system-design and preliminary-prototype study, not as a new ranking algorithm. Its contributions are: (1) an end-to-end architecture that integrates syllabus matching, learner level, modality preference, language, and progress persistence; (2) an exact statement of the implemented scoring and filtering logic; (3) a descriptive case study on the bundled catalogue and sample syllabus; and (4) a reproducible validation protocol. The research questions are: RQ1, how can multiple learner and resource constraints be integrated in a lightweight pipeline? RQ2, what ranking behavior does the heuristic produce on a sample syllabus? RQ3, what evidence is still needed to establish recommendation quality?



2. LITERATURE REVIEW

A. Educational Recommenders and Learning Paths

Surveys of e-learning recommenders and learning-path personalization show that quality depends on learner context, resource metadata, sequencing, and evaluation design, not on a single similarity score [15], [18], [23], [46], [47]. Integrative OER recommendation stresses learner agency [11], and Bayesian learner models have been used to match learners to open resources [21]. Learner-model research includes knowledge tracing [6], [8], [48] and open learner models that let students inspect and adjust their profiles [22]. Recent studies combine LLMs with knowledge graphs or contrastive sequence models [2], [3], and propose deterministic ranking with LLM-generated explanations [1]. Freeversity implements none of these learned components; it addresses a practical subset of the design dimensions with explicit rules.

B. Syllabus Alignment and Lexical versus Semantic Retrieval

Aligning learning outcomes to resources can be formulated as a lexico-semantic problem [30], and course descriptions have been embedded to map curricular knowledge [28]; concept extraction supports prerequisite and domain modelling [17]. Sentence-BERT and related encoders [24], [31], [36] provide semantic similarity, and multilingual encoders [13], [25] are relevant to Hindi–English catalogues. However, benchmark studies show that lexical retrieval remains a strong and robust baseline [10], [14], and that comparisons against weak baselines can overstate neural gains [35]. FAIR-oriented metadata and semantic similarity have also been applied to OER discovery [5], [39]. Freeversity uses TF-IDF with tag overlap, a classical vector-space approach described in recommender and retrieval texts [41], [43]; it is not a neural embedding method.

C. Content-Modality Preference

Freeversity lets a learner prioritize watching, reading, or building-oriented resources. Reviews and experiments report little support for the claim that matching instruction to a diagnosed learning style improves outcomes [38], [44], [45]. The system therefore treats modality only as a soft ranking preference and makes no claim that a match improves learning.

D. Explainability, Evaluation, and Ethics

Explainable recommendation research motivates exposing why an item was ranked [19]. Offline evaluation of recommenders depends on design choices such as relevance labels, sampling, and metrics [20], [29], and recommender studies have been criticised for weak baselines and poor reproducibility [33], [42]. Responsible use of AI in education also requires attention to consent, privacy, and ethics review [9], [12]. These points shape the validation protocol in Section 3.G.

TABLE I. RESEARCH AREAS AND THE BOUNDARY OF THE PRESENT WORK

Area	Representative studies	Relation to Freeversity
Learning-path and e-learning recommenders	[15], [18], [46], [47]	Informs problem dimensions; no trained learner model implemented
LLM / knowledge-graph recommendation	[1], [2], [3], [7]	Future extension; not part of the prototype
Syllabus–resource alignment	[17], [28], [30]	Motivates section-level matching (future work)
Semantic and lexical retrieval	[14], [25], [31], [35]	Sentence-embedding baseline planned, not implemented
OER and metadata	[5], [32], [37], [39]	Catalogue of free-access resources; licenses unverified
Learning styles and evaluation	[20], [33], [38], [45]	Modality is a soft preference; evaluation is descriptive

3. METHODOLOGY

A. Research Design

The study is a descriptive prototype study. It (i) specifies the pipeline and scoring logic of the Freeversity source archive, (ii) runs that implementation on its bundled catalogue and a single sample syllabus, and (iii) defines a protocol for later validation. No human participants are involved, and no ground-truth relevance labels exist; outputs are therefore reported as heuristic behavior, consistent with guidance that offline recommender results need explicit relevance judgments to be interpreted as quality [20], [42].

B. Data Sources

The catalogue holds 39 curated Data Science resources, each with title, description, URL, provider, source type, modality, curriculum module, tags, language, duration, certificate flag, downloadability flag, and level; such descriptive metadata is the kind of structured information that FAIR-oriented OER work emphasizes [5], [39]. Table II and Fig. 1 profile the catalogue.



The diagnostic quiz has 11 multiple-choice questions, and the case study uses one bundled syllabus covering Python fundamentals, probability and statistics, data wrangling, regression, and classification. Metadata completeness was checked only as a data-structure test; it does not verify that URLs resolve, resources remain free, or licenses permit redistribution.

TABLE II. CATALOGUE PROFILE (n = 39)

Attribute	Observed composition
Curriculum modules	Python Basics 6; Statistics 5; Data Cleaning 7; Data Visualization 4; Linear Regression 7; Classification 5; Applied ML Practice 5
Source type	Company 14 (35.9%); Open 12 (30.8%); University 7 (17.9%); YouTube 6 (15.4%)
Declared level	Beginner 19 (48.7%); Intermediate 16 (41.0%); Advanced 4 (10.3%)
Content modality	Watching 18 (46.2%); Reading 11 (28.2%); Building 10 (25.6%)
Language	English 36 (92.3%); Hindi 3 (7.7%)




Fig. 1. Catalogue composition by source type, declared level, content modality, and language (n = 39).

C. Pipeline

Figure 2 shows the data flow. The reference implementation uses Python with Flask, a browser-side front end, SQLite storage, scikit-learn TF-IDF [26], and pdfplumber for PDF text. A PostgreSQL path exists in code but was not independently tested, so only the SQLite design is described. The pipeline has six steps: (1) build the learner profile; (2) extract syllabus text from PDF, DOCX, or TXT; (3) filter candidates cumulatively by level, then language, certificate, source type, and downloadability; (4) score candidates against the syllabus; (5) re-weight by modality and prior knowledge; and (6) rank within the seven-module curriculum template, return up to five resources per module, and persist the roadmap and progress.




Fig. 2. Freeversity pipeline (rules only; no trained neural model).

D. Level Detection

The quiz score is the fraction of correct answers: below 0.40 maps to beginner, below 0.75 to intermediate, and otherwise advanced; learners may also select a level manually (Fig. 3). These thresholds are fixed engineering rules and are not psychometrically validated. Calibrated alternatives, such as knowledge-tracing or item-response-based learner models, exist in the literature but are not implemented [8], [48].




Fig. 3. Level detection from the placement quiz and the cumulative filter cascade.

E. Scoring and Re-ranking

Let x_s be the extracted syllabus text and x_r a resource text built from its title, description, and tags. TF-IDF vectors are fitted on the syllabus plus the candidate resource texts, giving a cosine similarity $C(s,r)$ in $[0,1]$ [41], [43]. Let K_s be the set of lowercase syllabus tokens of at least three letters after stop-word removal and T_r the set of semicolon-separated, lowercased resource tags. The tag-overlap term is $O(s,r) = |K_s \cap T_r| / \max(|T_r|, 1)$, and the displayed score is



$$\text{Match}(s,r) = 100 \times [0.6 C(s,r) + 0.4 O(s,r)] \quad (1)$$

The score is interpretable in that both components and the overlapping tags can be inspected, a property valued in explainable recommendation [19], but it is not a calibrated probability. Multiword or hyphenated tags (e.g., “linear-regression”) are kept as whole strings while syllabus keywords are single tokens, so such tags rarely overlap; this under-counts the tag term and needs a canonicalization step.

With a syllabus, the score is multiplied by a modality weight: the preferred modality receives 1.00, and the other two receive 0.45 and 0.35 as specified by the implementation; “Any” gives 1.00 to all. Without a syllabus the base score is 60 before weighting. Resources whose tags overlap the learner’s prior-knowledge list are multiplied by 0.60. If a module has no candidate after filtering, up to three resources are returned with the level restriction relaxed and a baseline score of 50, labelled as previews above the current level (Fig. 4). All factors are hand-set and not optimized on validation data. Modality is used only as a soft preference [38].




Fig. 4. Scoring and re-ranking flowchart, including the no-syllabus branch and the fallback for empty modules.

F. Evaluation Procedure

Five descriptive checks were run on the source archive: (i) a field-completeness check of the 39 catalogue records; (ii) execution of the automated tests (27 methods, of which the 18 non-API tests could be run); (iii) whole-syllabus ranking of all 39 resources with Eq. (1) and no learner filters; (iv) a component-sensitivity comparison of TF-IDF alone, tag overlap alone, and the hybrid; and (v) a per-module roadmap for an illustrative intermediate, watching-preference, English-language profile. Because no relevance labels exist, none of these checks is a measure of recommendation accuracy [20].

G. Planned Validation Protocol

A publication-ready study should (a) collect several Data Science syllabi and segment them into topics or learning outcomes, as section-level alignment requires [30]; (b) have two subject-matter reviewers rate each topic–resource pair on an ordinal 0–3 scale, report Cohen’s kappa, and adjudicate disagreements; (c) compare tag-overlap-only, TF-IDF-only, the 0.6/0.4 hybrid, and a Sentence-BERT retriever [31] on identical catalogues, topics, and filters, with ablations removing modality re-ranking and level filtering and with strong lexical baselines kept in view [14], [35]; (d) compute Precision@5, Recall@5, MRR, and NDCG@5 with per-topic and macro averages and bootstrap confidence intervals, without tuning weights on evaluation topics [20], [29]; (e) check constraint satisfaction (language, source, certificate, level, download) separately and mark fallback items; and (f) conduct any usability study only after consent and ethics review [12]. None of these labels or metrics has been collected. The protocol is summarised in Fig. 5.




Fig. 5. Planned validation protocol (steps a–f); no labels or metrics have yet been collected.

4. RESULTS AND DISCUSSION

A. Catalogue and Test Results

All 39 records contain the expected fields. Of the 27 test methods, the 18 non-API tests (quiz thresholds, modality weights, language and preference filters, deterministic TF-IDF scores, hybrid ordering, and roadmap assembly) all passed. The remaining nine exercise Flask routes and were not executed because Flask was not installed; the result is therefore 18/18 for core logic, not 27/27 for the full system. A clean end-to-end run is needed before external submission, in line with reproducibility recommendations [42].

B. Whole-Syllabus Ranking

Thirty-eight of 39 resources received a non-zero score; the minimum was 0.0, the maximum 46.5, the mean 18.57, and the median 15.6 on the 0–100 heuristic scale. Table III and Fig. 6(a) show the top five. The list is dominated by statistics and Python resources, consistent with the sample syllabus, but this reflects the code's behavior on one input, not instructor- or student-validated relevance. Because each resource is compared with the whole syllabus, a resource aligned with a large share of the document can outrank one that is highly relevant to a smaller topic, which motivates section-level matching [28], [30].

TABLE III. TOP FIVE RESOURCES FOR THE SAMPLE SYLLABUS (HYBRID SCORE)

Rank	Resource	Module	Score
1	Khan Academy: Statistics and Probability	statistics	46.5
2	CS50's Introduction to Programming with Python	python-basics	41.6
3	Think Stats (free book)	statistics	40.4
4	StatQuest with Josh Starmer	statistics	38.7
5	CodeWithHarry: Python tutorial series (Hindi)	python-basics	38.1

C. Component Sensitivity and Per-Module Output

TF-IDF alone ranked StatQuest first (22.8/100); tag overlap alone ranked Khan Academy first (85.7/100); the hybrid ranked Khan Academy first (46.5/100) (Fig. 6(b)). The blend therefore changes the top item relative to the lexical-only view, but the component scales differ and no labels exist, so this does not show that the hybrid is more accurate than either part; the same caution applies to claims about weak or strong baselines [14], [35]. For the illustrative intermediate profile, all seven modules returned resources; the highest final scores were Python Basics 42.0, Statistics 46.6, Data Cleaning 8.6, Data Visualization 5.5, Linear Regression 15.3, Classification 17.2, and Applied ML Practice 11.8 (Fig. 6(c)). The low Data Cleaning and Data Visualization values indicate where section segmentation and tag normalization would help [17], [30].





Fig. 6. (a) Top-five hybrid scores, (b) top score under each scoring component, (c) best final score per module for the illustrative profile.

D. Discussion

The prototype shows an end-to-end design that goes beyond a static list of links, and a reviewer can trace each recommendation to resource text, tags, preferences, and module, which supports inspectability [19]. A running parser, API route, or progress control, however, shows only that a component works, not that recommendations are educationally appropriate [20], [33]. A score shown as a percentage remains a rescaled heuristic until it is evaluated against independent judgments. The findings therefore concern the behavior of the baseline and the evidence still required (RQ2, RQ3), while RQ1 is addressed by the pipeline in Section 3.C.

5. PROPOSED WORK

The proposed work is organised in five phases (Fig. 7). Phase 1 covers consistent tag normalization and URL and license verification [39]. Phase 2 adds syllabus segmentation into outcome-level units and per-module topic-to-resource scoring [17], [30]. Phase 3 adds a sentence-embedding model [10], [31] and multilingual encoders [25] as measured baselines rather than assumed improvements [35]. Phase 4 introduces prerequisite graphs, calibrated learner models and knowledge tracing [6], [8], [48], and open learner models [22]. Phase 5 adds knowledge-graph and LLM components with grounded explanations [1], [3], [9], and weight optimization from labelled data. Each extension should be justified by measured gains over the transparent baseline [33].



Each phase is justified only by measured gains over the transparent baseline

Fig. 7. Proposed work roadmap, from data hygiene to learned and LLM-based components.

6. CONCLUSION

Freeversity is a syllabus-aware, preference-guided prototype that assembles free educational resources into a module-organized roadmap using a fixed quiz, metadata filters, TF-IDF cosine similarity, tag overlap, hand-set modality weights, and



SQLite persistence. On the bundled catalogue and sample syllabus, 38 of 39 resources scored above zero and 18 core-logic tests passed. These results describe implementation behavior, not recommendation accuracy or improved learning. The practical value lies in a transparent baseline, in line with calls for rigorous and reproducible evaluation in learning-path personalization [18], [33], [42]. Publication-level claims require independent relevance judgments, comparative ranking metrics, a broader corpus, and full backend testing [20].

ARTIFACT AVAILABILITY

A browser-side showcase is available at <https://stirringsquirrel-efdcf6.netlify.app/>. The deployed page and archive use the spelling “Freevarsity” while this paper uses “Freeversity”; the name should be harmonized before submission. The canonical implementation analyzed here is the Flask/SQLite archive. For public submission, the exact source version, seed CSV files, sample syllabus, test instructions, and any annotation data should be archived with a stable identifier, and resource URLs and licenses re-checked [39], [42].

REFERENCES

- [1] W. Yun, B. Akbayan, K. Aray, and T. Azamat, “Generative LLMs for educational recommendation,” *Front. Artif. Intell.*, vol. 9, 2026, doi: 10.3389/frai.2026.1827553.
- [2] W. Yao and X. Hu, “Large language model-empowered course recommendation with learning interest-goal contrastive learning,” *J. Syst. Sci. Syst. Eng.*, 2026, doi: 10.1007/s11518-025-5702-8.
- [3] J. Chen and T. Zhu, “Personalized learning resource recommendation framework based on knowledge graph and large language model,” in *Knowledge Science, Engineering and Management (KSEM 2025)*, *Lecture Notes in Computer Science*, vol. 15919. Springer, 2026.
- [4] S. C. Chang and N. D. Dao, “An intelligent recommender system based on K-nearest neighbors to foster self-regulated learning and reduce cognitive load in online higher education,” *Comput. Educ.*, vol. 240, Art. no. 105470, 2026.
- [5] F. Sadeghi, A. Degbelo, C. Kessler, and R. Zolnouri, “Applying the FAIR principles to open educational resources: A semantic similarity approach to improve resource discovery,” *IEEE Access*, vol. 13, pp. 136626–136642, 2025, doi: 10.1109/ACCESS.2025.3594963.
- [6] Z. Liu et al., “Deep learning based knowledge tracing: A review, a tool and empirical studies,” *IEEE Trans. Knowl. Data Eng.*, 2025, doi: 10.1109/TKDE.2025.3552759.
- [7] S. Wang et al., “Large language models for education: A survey and outlook,” *arXiv:2403.18105*, 2024.
- [8] G. Abdelrahman, Q. Wang, and B. Nunes, “Knowledge tracing: A survey,” *ACM Comput. Surv.*, vol. 55, no. 11, Art. no. 224, 2023.
- [9] E. Kasneci et al., “ChatGPT for good? On opportunities and challenges of large language models for education,” *Learn. Individ. Differ.*, vol. 103, Art. no. 102274, 2023.
- [10] N. Muennighoff, N. Tazi, L. Magne, and N. Reimers, “MTEB: Massive text embedding benchmark,” in *Proc. EACL*, 2023, pp. 2014–2037.
- [11] S. Bulathwela, M. Pérez-Ortiz, E. Yilmaz, and J. Shawe-Taylor, “Power to the learner: Towards human-intuitive and integrative recommendations with open educational resources,” *Sustainability*, vol. 14, no. 18, Art. no. 11682, 2022, doi: 10.3390/su141811682.
- [12] W. Holmes et al., “Ethics of AI in education: Towards a community-wide framework,” *Int. J. Artif. Intell. Educ.*, vol. 32, no. 3, pp. 504–526, 2022.
- [13] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang, “Language-agnostic BERT sentence embedding,” in *Proc. ACL*, 2022, pp. 878–891.
- [14] N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, and I. Gurevych, “BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models,” in *Proc. NeurIPS Datasets and Benchmarks Track*, 2021.
- [15] M. C. Urdaneta-Ponte, A. Méndez-Zorrilla, and I. Oleagordia-Ruiz, “Recommendation systems for education: Systematic review,” *Electronics*, vol. 10, no. 14, Art. no. 1611, 2021.
- [16] M. L. Bernacki, M. J. Greene, and N. G. Lobczowski, “A systematic review of research on personalized learning: Personalized by whom, to what, how, and for what purpose(s)?” *Educ. Psychol. Rev.*, vol. 33, no. 4, pp. 1675–1715, 2021.
- [17] H. Chau, I. Labutov, K. Thaker, D. He, and P. Brusilovsky, “Automatic concept extraction for domain and student modeling in adaptive textbooks,” *Int. J. Artif. Intell. Educ.*, vol. 31, pp. 820–846, 2021.
- [18] A. H. Nabizadeh, J. P. Leal, H. N. Rafsanjani, and R. R. Shah, “Learning path personalization and recommendation methods: A survey of the state-of-the-art,” *Expert Syst. Appl.*, vol. 159, Art. no. 113596, 2020, doi: 10.1016/j.eswa.2020.113596.
- [19] Y. Zhang and X. Chen, “Explainable recommendation: A survey and new perspectives,” *Found. Trends Inf. Retr.*, vol. 14, no. 1, pp. 1–101, 2020.
- [20] R. Cañameres, P. Castells, and A. Moffat, “Offline evaluation options for recommender systems,” *Inf. Retr. J.*, vol. 23, no. 4, pp. 387–410, 2020.
- [21] S. Bulathwela, M. Pérez-Ortiz, E. Yilmaz, and J. Shawe-Taylor, “TrueLearn: A family of Bayesian algorithms to match lifelong learners to open educational resources,” in *Proc. AAAI Conf. Artif. Intell.*, 2020.
- [22] S. Abdi, H. Khosravi, S. Sadiq, and D. Gasevic, “Complementing educational recommender systems with open learner models,” in *Proc. 10th Int. Conf. Learn. Analytics Knowl. (LAK)*, 2020.
- [23] S. S. Khanal, P. W. C. Prasad, A. Alsadoon, and A. Maag, “A systematic review: Machine learning based recommendation systems for e-learning,” *Educ. Inf. Technol.*, vol. 25, no. 4, pp. 2635–2664, 2020.
- [24] V. Karpukhin et al., “Dense passage retrieval for open-domain question answering,” in *Proc. EMNLP*, 2020, pp. 6769–6781.
- [25] A. Conneau et al., “Unsupervised cross-lingual representation learning at scale,” in *Proc. ACL*, 2020, pp. 8440–8451.
- [26] S. Raschka, J. Patterson, and C. Nolet, “Machine learning in Python: Main developments and technology trends in data science, machine learning, and artificial intelligence,” *Information*, vol. 11, no. 4, Art. no. 193, 2020.
- [27] A. Shemshack and J. M. Spector, “A systematic literature review of personalized learning terms,” *Smart Learn. Environ.*, vol. 7, Art. no. 33, 2020.
- [28] Z. A. Pardos and A. Nam, “A university map of course knowledge,” *PLoS ONE*, vol. 15, no. 9, Art. no. e0233207, 2020.
- [29] W. Krichene and S. Rendle, “On sampled metrics for item recommendation,” in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2020, pp. 1748–1757.



- [30] S. Saha et al., "Aligning learning outcomes to learning resources: A lexico-semantic spatial approach," in Proc. 28th Int. Joint Conf. Artif. Intell. (IJCAI), 2019, pp. 5168–5174, doi: 10.24963/ijcai.2019/718.
- [31] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in Proc. EMNLP-IJCNLP, 2019, pp. 3982–3992, doi: 10.18653/v1/D19-1410.
- [32] UNESCO, Recommendation on Open Educational Resources (OER). Paris, France: UNESCO, 2019.
- [33] M. Ferrari Dacrema, P. Cremonesi, and D. Jannach, "Are we really making much progress? A worrying analysis of recent neural recommendation approaches," in Proc. 13th ACM Conf. Recommender Systems (RecSys), 2019, pp. 101–109.
- [34] J. Reich and J. A. Ruipérez-Valiente, "The MOOC pivot," *Science*, vol. 363, no. 6423, pp. 130–131, 2019.
- [35] J. Lin, "The neural hype and comparisons against weak baselines," *ACM SIGIR Forum*, vol. 52, no. 2, pp. 40–51, 2019.
- [36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, 2019, pp. 4171–4186.
- [37] D. Wiley and J. L. Hilton III, "Defining OER-enabled pedagogy," *Int. Rev. Res. Open Distrib. Learn.*, vol. 19, no. 4, 2018.
- [38] P. M. Newton and M. Miah, "Evidence-based higher education – Is the learning styles 'myth' important?" *Front. Psychol.*, vol. 8, Art. no. 444, 2017.
- [39] M. D. Wilkinson et al., "The FAIR Guiding Principles for scientific data management and stewardship," *Sci. Data*, vol. 3, Art. no. 160018, 2016.
- [40] J. Hilton III, "Open educational resources and college textbook choices: A review of research on efficacy and perceptions," *Educ. Technol. Res. Dev.*, vol. 64, no. 4, pp. 573–590, 2016.
- [41] J. Beel, B. Gipp, S. Langer, and C. Breiteringer, "Research-paper recommender systems: A literature survey," *Int. J. Digit. Libr.*, vol. 17, no. 4, pp. 305–338, 2016.
- [42] J. Beel, C. Breiteringer, S. Langer, A. Lommatzsch, and B. Gipp, "Towards reproducibility in recommender-systems research," *User Model. User-Adapt. Interact.*, vol. 26, no. 1, pp. 69–101, 2016.
- [43] C. C. Aggarwal, *Recommender Systems: The Textbook*. Cham, Switzerland: Springer, 2016.
- [44] P. M. Newton, "The learning styles myth is thriving in higher education," *Front. Psychol.*, vol. 6, Art. no. 1908, 2015.
- [45] B. A. Rogowsky, B. M. Calhoun, and P. Tallal, "Matching learning style to instructional method: Effects on comprehension," *J. Educ. Psychol.*, vol. 107, no. 1, pp. 64–78, 2015.
- [46] H. Drachsler, K. Verbert, O. C. Santos, and N. Manouselis, "Panorama of recommender systems to support learning," in *Recommender Systems Handbook*, 2nd ed., F. Ricci, L. Rokach, and B. Shapira, Eds. Boston, MA, USA: Springer, 2015, pp. 421–451.
- [47] A. Klačnja-Milićević, M. Ivanović, and A. Nanopoulos, "Recommender systems in e-learning environments: A survey of the state-of-the-art and possible extensions," *Artif. Intell. Rev.*, vol. 44, no. 4, pp. 571–604, 2015.
- [48] C. Piech et al., "Deep knowledge tracing," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 28, 2015.
- [49] K. Jordan, "Massive open online course completion rates revisited: Assessment, length and attrition," *Int. Rev. Res. Open Distrib. Learn.*, vol. 16, no. 3, pp. 341–358, 2015.
- [50] D. Orr, M. Rimini, and D. Van Damme, *Open Educational Resources: A Catalyst for Innovation*. Paris, France: OECD Publishing, 2015.