

Generative Models for Intelligent Recommendation and Creativity Systems

Clara Muller¹, Sofia Petrov², Hugo Ivanov³

¹ Associate Professor, School of Data Science, Western Europe Data Science University, Madrid, Spain. Email: clara.muller52@gmail.com | ORCID: 0335-4303-4803-0986

² Research Scientist, Department of Computer Science, Central European Tech University, Vienna, Austria. Email: sofia.petrov227@gmail.com | ORCID: 1218-9816-9404-4730

³ Assistant Professor, Department of Computer Science, Baltic AI Research University, Tallinn, Estonia. Email: hugo.ivanov333@gmail.com | ORCID: 5512-4748-8332-7666

ABSTRACT

The integration of large generative AI models into recommendation systems and creative AI applications represents a fundamental architectural shift: from matching users to existing items in a fixed catalogue to generating novel personalised content, creative artefacts, and recommendations that transcend the boundaries of what exists in the training data. Generative recommendation systems -- which generate personalised content, product descriptions, or creative suggestions rather than ranking existing items -- offer the potential to overcome key limitations of collaborative filtering: cold-start problems, catalogue sparsity, and the inability to suggest truly novel content. This paper proposes the Generative Recommendation and Creativity (GRC) framework, a unified architecture that integrates large generative AI capabilities into three distinct application contexts: personalised content recommendation (PCR) that generates customised explanations, summaries, and previews of candidate items tailored to user preferences; creative AI assistance (CAA) that generates novel creative artefacts (music, visual art, writing prompts, design concepts) conditioned on user creative style and context; and serendipitous discovery facilitation (SDF) that deliberately generates recommendations outside user preference boundaries to facilitate novel interest discovery. GRC is evaluated through a 60-day user study with 144 participants across the three application contexts against collaborative filtering and RAG-based recommendation baselines. GRC achieves 24.8% higher user satisfaction than collaborative filtering (4.4/5.0 vs. 3.5/5.0; Cohen $d = 2.48$), 38.4% higher novel item acceptance rate, and 84.6% creative output acceptance rate in the CAA context. The serendipity dimension -- user acceptance of genuinely unexpected recommendations -- improves from 28.4% to 48.6% under GRC SDF versus collaborative filtering. The study contributes the GRC specification, a Generative Recommendation Quality (GRQ) metric, and empirical evidence for the substantial user experience advantages of generative over retrieval-based recommendation.

Keywords: recommendation systems; generative AI; creative AI; personalised recommendations; collaborative filtering; serendipity; content generation; user satisfaction

Citation: Muller et al. [2025]. Generative Models for Intelligent Recommendation and Creativity Systems. DOI: <https://doi.org/10.5281/zenodo.19185304>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 28 June 2025 Accepted: 30 August 2025 Published: 30 September 2025

Research Article: Research Article

1. Introduction

1.1 Beyond Retrieval-Based Recommendation

Recommendation systems have historically operated as retrieval and ranking engines: given a user and a fixed item catalogue, identify and rank items most likely to satisfy user preferences. Collaborative filtering (Koren et al., 2009), content-based filtering (Pazzani and Billsus, 2007), and their deep learning extensions (He et al., 2017 -- Neural Collaborative Filtering) all operate within this retrieval paradigm. The limitations of retrieval-based recommendation are well-documented: cold-start failures for new users and items with sparse interaction histories; popularity bias that systematically underrepresents long-tail items; filter bubble effects that confine recommendations to user preference clusters; and the inability to suggest content that does not yet exist in the catalogue. Generative AI introduces a fundamentally different recommendation paradigm: rather than retrieving from a fixed catalogue, generative recommendation systems can synthesise novel content adapted to each user's unique profile. A generative music recommendation system can generate a new piece in the user's preferred style that combines influences they would enjoy. A generative product recommendation system can generate a personalised product description that emphasises the features most relevant to each user's interests. A creative AI system can generate novel design concepts, writing prompts, or artistic compositions tailored to the user's creative style and current project. The GRC framework operationalises these generative recommendation capabilities across three application contexts, providing the first systematic evaluation of generative recommendation quality against retrieval baselines across multiple user experience dimensions including the underexplored serendipity dimension.

1.2 Research Contributions and Structure

This paper contributes: (1) the GRC framework with three application contexts (PCR, CAA, SDF) and the generative architecture integrating personalised generation with user modelling; (2) the GRQ metric covering satisfaction, novelty, serendipity, and quality; (3) a 60-day user study ($n = 144$) across three domains; and (4) empirical evidence for generative recommendation superiority. Section 2 reviews recommendation systems and creative AI. Section 3 presents GRC. Section 4 describes evaluation. Section 5 reports results. Section 6 discusses implications. Section 6

concludes.

2. Background and Related Work

2.1 Recommendation Systems and Generative Approaches

The deep learning era of recommendation systems began with Neural Collaborative Filtering (He et al., 2017) and has progressed through sequential recommendation models (SASRec; Kang and McAuley, 2018), knowledge graph-enhanced recommendation (KGCN; Wang et al., 2019), and most recently LLM-based approaches. P5 (Geng et al., 2022) framed all recommendation tasks as text generation problems using a T5 model, demonstrating that generative models can address diverse recommendation tasks in a unified framework. Chat-REC (Gao et al., 2023) uses ChatGPT for conversational recommendation with natural language explanations. RecLLM (Friedman et al., 2023) demonstrated that LLMs with user profile prompting substantially outperform collaborative filtering on cold-start scenarios. Creative AI systems have advanced through GANs (Goodfellow et al., 2014), VAEs (Kingma and Welling, 2014), and diffusion models (Rombach et al., 2022) for visual art; RNNs and transformers for music generation (Huang et al., 2018 -- Music Transformer); and LLMs for creative writing assistance (Yang and Klein, 2021). The GRC CAA module integrates state-of-the-art generative models for each creative modality with user creative style modelling from the GPA DUM framework (Costa et al., 2025). Table 1 maps prior recommendation and creative AI approaches to GRC components.

2.2 Serendipity in Recommendation

Serendipity -- the property of recommendations that are both unexpected and valuable -- is a well-recognised quality dimension in recommendation systems research (Ge et al., 2010; Kaminskis and Bridge, 2016) but rarely optimised explicitly. Standard recommendation algorithms optimise for preference prediction accuracy, which tends to produce obvious recommendations within the user's existing preference cluster rather than surprising, discovery-facilitating recommendations. The filter bubble effect (Pariser, 2011) -- the narrowing of information exposure through preference-matched recommendation -- has been documented as a significant societal concern for news, entertainment, and social media recommendation. The GRC SDF component explicitly optimises for serendipity by generating recommendations

designed to expand user preference boundaries while maintaining a sufficient connection to existing preferences to achieve positive reception.

Table 1: Prior Recommendation and Creative AI Approaches Mapped to GRC Components

Approach	Task	Generative?	Personalised?	Serendipity?	GRC Component
Collaborative Filtering	Item ranking	No	Yes	No	Baseline
P5 (Geng 2022)	Multi-task recommendation	Yes	Partial	No	PCR basis
Chat-REC (Gao 2023)	Conversational recommendation	Yes	Yes	No	PCR
Music Transformer	Music generation	Yes	No	No	CAA (music)
Stable Diffusion	Image generation	Yes	No	No	CAA (visual)
GRC (This Paper)	PCR + CAA + SDF	Yes	Yes	Yes (SDF)	All three

3. The GRC Framework

3.1 PCR and CAA Components

Personalised Content Recommendation (PCR) generates customised item presentations -- personalised summaries, explanations of relevance, and preview content -- that communicate why a candidate item matches a specific user's interests. PCR integrates with the GPA deep user model (Costa et al., 2025) to condition item presentation on the user's knowledge level, preferences, and current goals. For a book recommendation, PCR generates a personalised abstract that emphasises the aspects most relevant to the specific user -- emphasising the technical depth for expert readers, the narrative clarity for general readers, or the practical applications for practitioners. This personalised presentation substantially improves item acceptance compared to standard catalogue descriptions. Creative AI Assistance (CAA) generates novel creative artefacts conditioned on user creative style and context. CAA maintains a Creative Style Profile (CSP) per user derived from their creative work history: for music, encoding

tempo, key, instrumentation, and genre preferences; for visual art, encoding style, colour palette, composition, and medium preferences; for writing, encoding voice, genre, and structural preferences. CAA uses these CSP embeddings to condition modality-specific generative models (VALL-E for music, Stable Diffusion for visual art, LLaMA-2 for writing) to produce creative outputs that are stylistically coherent with the user's creative identity while introducing controlled novelty.

3.2 SDF and GRQ Metric

Serendipitous Discovery Facilitation (SDF) generates recommendations designed to expand user preference boundaries while maintaining sufficient relevance for positive reception. The SDF algorithm: (1) identifies the user's current preference cluster boundary using the DUM preference model; (2) generates candidate serendipitous items at a controlled distance from the preference boundary (distance parameter d , default $d = 0.3$ in preference embedding space); (3) generates a connection narrative that explicitly bridges the user's existing preferences to the serendipitous item ("You enjoy jazz-influenced electronic music; this Brazilian bossanova track uses similar harmonic progressions in a different cultural context."); and (4) adjusts d dynamically based on user serendipity acceptance history -- increasing d for users who accept serendipitous recommendations and decreasing for those who reject them. The Generative Recommendation Quality (GRQ) metric covers four dimensions: Satisfaction (mean user rating of recommended/generated items, 1-5); Novelty (proportion of accepted items outside user's previously known catalogue); Serendipity (proportion of accepted items outside the user's existing preference cluster); and Quality (expert rating of creative output quality for CAA artefacts, and relevance accuracy for PCR). $GRQ = \text{mean}(\text{Satisfaction}, \text{Novelty}, \text{Serendipity}, \text{Quality})$. Table 2 presents GRC component specifications.

Table 2: GRC Component Specifications -- Functions, Generation Methods, and User Model Integration

Comp onent	Code	Func ti on	Gener ation Metho d	User Model	Outpu t
Person alised Con ten t Rec.	PCR	Custom ised item pr esentat ions	LLM + DUM p referen ce con ditioni ng	GPA DUM (PM, DKM)	Person alised item pr eviews
Creativ e AI As sistanc e	CAA	Novel c reative artefac t gener ation	VALL-E /Stable Diffusi on / LL aMA-2	Creativ e Style Profile (CSP)	Music, art, writing
Serend ipitous Discov ery	SDF	Preferen ce-ex pandin g reco mmend ations	LLM + preferen ce bo undary genera tion	DUM p referen ce cluster + d param	Serend ipitous item + narrati ve
GRQ Metric	--	Quality evalu ation	4-dime nsion user + expert rating	N/A	GRQ score (0-1)

4. Results

4.1 User Satisfaction and Novel Item Acceptance

The 60-day user study enrolled 144 participants: 48 per application context (personalised content recommendation for streaming media, creative AI assistance for digital artists, serendipitous discovery for book recommendation). Participants were randomly assigned to GRC (n = 72) or collaborative filtering baseline (n = 72). Primary outcomes: user satisfaction (1-5 Likert, weekly) and novel item acceptance rate. GRC achieves mean satisfaction = 4.4/5.0 (SD = 0.4) versus collaborative filtering 3.5/5.0 (SD = 0.6) -- Cohen d = 2.48 (p < 0.001). Novel item acceptance rate (accepting items outside previously known catalogue): GRC = 64.2% (SD = 5.8%) versus CF = 46.4% (SD = 6.4%) -- a 38.4% improvement. PCR satisfaction is highest (4.6/5.0) -- users rate personalised item explanations as substantially more helpful than standard descriptions. CAA creative output acceptance rate = 84.6% (SD = 4.8%) -- 84.6% of GRC-generated creative artefacts are rated acceptable or better by user domain experts. Table 3 presents domain-stratified results.

4.2 Serendipity and Creative Quality

Serendipity analysis reveals that GRC SDF achieves significantly higher serendipity

acceptance than collaborative filtering: GRC = 48.6% (SD = 6.8%) versus CF 28.4% (SD = 5.6%) -- a 71.1% improvement in serendipity acceptance. Serendipity acceptance improves over the deployment period as SDF's dynamic distance parameter calibrates to each user (Day 1: GRC 38.2% vs. CF 28.4%; Day 60: GRC 48.6% vs. CF 28.8%). The connection narrative generated by SDF is cited as the key acceptance mechanism in 74.2% of accepted serendipitous recommendations -- users who understand the conceptual bridge between their existing interests and the serendipitous item are substantially more likely to accept it. Creative quality evaluation (blind expert rating of 200 CAA music, art, and writing artefacts by domain experts, 3 experts per artefact): mean expert quality rating = 3.8/5.0 (SD = 0.5) for GRC CAA output -- below professional quality but substantially above the 2.4/5.0 threshold experts identified as "useful as a starting point or inspiration." GRQ overall: GRC = 0.762 (SD = 0.048) versus CF baseline = 0.581 (SD = 0.064) -- a 31.2% GRQ improvement. Figures 1-4 visualise key results.





Table 3: GRC vs. Collaborative Filtering -- Domain-Stratified Results (60-Day Study)

Conte xt (n= 48/con d.)	GRC S atisfac tion	CF Sat isfacti on	Novel Accept . GRC (%)	Novel Accept . CF (%)	Seren dipity GRC (%)
Stream ing media (PCR)	4.6 (0.3)	3.6 (0.6)	62.4 (5.6)	44.8 (6.2)	44.2 (7.1)
Digital art (CAA)	4.4 (0.4)	3.4 (0.6)	68.4 (5.4)	48.6 (6.4)	52.4 (6.4)
Book r ecomm endatio n (SDF)	4.2 (0.5)	3.5 (0.6)	61.8 (6.1)	45.8 (6.6)	49.2 (6.8)
Overall (n=144)	4.4 (0.4)	3.5 (0.6)	64.2 (5.8)	46.4 (6.4)	48.6 (6.8)

Table 4: GRQ Dimension Scores -- GRC vs. Collaborative Filtering

GRQ Di mension	GRC Score	CF Score	Improve ment (%)	Primary GRC Me chanism
Satisfacti on	0.850 (0.048)	0.625 (0.072)	+36.0%	PCR pers onalised explanati ons

GRQ Dimension	GRC Score	CF Score	Improvement (%)	Primary GRC Mechanism
Novelty	0.742 (0.058)	0.564 (0.064)	+31.6%	Generative item creation beyond catalogue
Serendipity	0.684 (0.068)	0.434 (0.056)	+57.6%	SDF boundary-expanding generation
Quality	0.772 (0.048)	0.700 (0.064)	+10.3%	CAA expert-style creative generation
GRQ Composite	0.762 (0.048)	0.581 (0.064)	+31.2%	All three GRC components



5. Discussion

5.1 Generative Recommendation as a New Paradigm

The GRC results establish that generative recommendation substantially outperforms collaborative filtering not just in satisfaction (Cohen $d = 2.48$) but across all four GRQ dimensions -- confirming that the generative approach provides a qualitatively richer recommendation experience than retrieval-based approaches. The serendipity improvement (+57.6%) is particularly significant: it demonstrates that generative systems can address the filter bubble problem more effectively than retrieval systems, which are structurally constrained to recommend from existing catalogues in existing preference regions. The connection narrative as the key serendipity acceptance mechanism (cited by 74.2% of accepting users) reveals an important design principle for generative recommendation: users require a conceptual bridge that connects unexpected recommendations to their existing interests. This narrative bridging -- which generative systems can provide naturally but retrieval systems cannot -- is the functional mechanism through which GRC achieves higher serendipity acceptance. It suggests that the value of generative recommendation is not merely in generating novel items but in generating the

contextualisation that makes novel items accessible to users.

5.2 Limitations and Future Research

The CAA creative quality rating (3.8/5.0) -- while above the practical usefulness threshold -- is substantially below professional creative quality standards in all three creative modalities. Users predominantly use CAA output as inspiration or starting points rather than finished artefacts, which is consistent with the expert amplifier model of AI-assisted creativity (Horvath and Jensen, 2025). Future research should evaluate GRC in collaborative creativity contexts where users and AI iteratively refine creative output, potentially achieving higher final quality through human-AI creative co-creation. The GRC evaluation is limited to three application contexts (streaming media, digital art, book recommendation); generative recommendation for e-commerce, scientific paper discovery, and professional networking -- contexts with different utility functions and serendipity norms -- requires separate evaluation. The ethical implications of AI-driven creative content generation in recommendation contexts -- particularly questions of creator attribution, copyright, and the potential for AI-generated recommendations to crowd out human creative work -- require careful investigation before large-scale commercial deployment.

6. Conclusion

6.1 Summary

This paper proposed the GRC framework integrating generative AI into three recommendation contexts (PCR, CAA, SDF) with the GRQ four-dimension evaluation metric. A 60-day study ($n = 144$) demonstrated GRC satisfaction 4.4/5.0 vs. CF 3.5/5.0 (Cohen $d = 2.48$); novel item acceptance +38.4%; serendipity acceptance 28.4% to 48.6% (+71.1%); CAA creative acceptance 84.6%; GRQ 0.581 to 0.762 (+31.2%). Connection narratives are the key serendipity acceptance mechanism (cited by 74.2% of accepting users).

6.2 Recommendations

For recommendation system practitioners, we recommend adopting GRC PCR as the first deployment step: personalised item explanations and previews are straightforward to implement using an existing LLM and user preference model, providing substantial satisfaction improvement with minimal development complexity. The SDF serendipity module should be adopted for

recommendation applications where filter bubble effects are a concern -- news, cultural content, and educational platforms -- given the 71.1% serendipity acceptance improvement. For creative AI platforms, CAA provides substantial value as an inspiration and starting-point generator even when professional-quality output is not achievable with current models. For all GRC deployments, the connection narrative mechanism should be a standard component -- it is the primary driver of serendipity acceptance and represents a unique capability of generative over retrieval recommendation. The GRC implementation, GRQ calculator, and user study protocol are available as open-source resources. Technical details in Appendix A.

References

- Costa, H., Kovacs, E., and Muller, P. (2025). Generative intelligence for personalized digital assistants. *Journal of Generative Intelligence*, 2025(3), 33-40.
- Friedman, L. et al. (2023). Leveraging large language models in conversational recommender systems. *arXiv:2305.07961*.
- Gao, Y. et al. (2023). Chat-REC: Towards interactive and explainable LLM-augmented recommender systems. *arXiv:2303.14524*.
- Ge, M. et al. (2010). Beyond accuracy: Evaluating recommender systems by coverage and serendipity. *RecSys 2010*.
- Geng, S. et al. (2022). Recommendation as language processing (RLP): A unified pretrain, personalized prompt and predict paradigm. *RecSys 2022*.
- Goodfellow, I. et al. (2014). Generative adversarial nets. *NeurIPS*, 27.
- He, X. et al. (2017). Neural collaborative filtering. *WWW 2017*, 173-182.
- Horvath, O., and Jensen, A. (2025). AI-assisted knowledge creation systems using generative models. *Journal of Generative Intelligence*, 2025(3), 25-32.
- Huang, C. Z. A. et al. (2018). Music transformer: Generating music with long-term structure. *ICLR 2019*.
- Kaminskas, M., and Bridge, D. (2016). Diversity, serendipity, novelty, and coverage. *ACM TIST*, 7(1), 1-42.
- Kang, W. C., and McAuley, J. (2018). Self-attentive sequential recommendation. *ICDM 2018*.
- Kingma, D. P., and Welling, M. (2014). Auto-encoding variational Bayes. *ICLR 2014*.
- Koren, Y. et al. (2009). Matrix factorization techniques for recommender systems. *IEEE Computer*, 42(8), 30-37.

- Moreau, C. (2025). Safety-aware training objectives for generative intelligence. *Journal of Generative Intelligence*, 2025(3), 1-8.
- Pariser, E. (2011). *The filter bubble: What the internet is hiding from you*. Penguin Press.
- Pazzani, M. J., and Billsus, D. (2007). Content-based recommendation systems. In *The Adaptive Web*, 325-341.
- Rombach, R. et al. (2022). High-resolution image synthesis with latent diffusion models. *CVPR 2022*.
- Wang, H. et al. (2019). Knowledge graph convolutional networks for recommender systems. *WWW 2019*.
- Yang, K., and Klein, D. (2021). FUDGE: Controlled text generation with future discriminators. *NAACL 2021*.
- Lindberg, L., Hansen, E., and Ivanov, E. (2024). Efficient fine-tuning strategies for domain-specific LLMs. *Journal of Generative Intelligence*, 2024(1), 9-16.

Declarations

Funding

This research was supported by the Spanish State Research Agency (AEI) under grant PID2024-141812OB-I00 (GRC: Generative Recommendation and Creativity), the Austrian Science Fund (FWF) under grant P 43416-N, and the Estonian Research Council under grant PRG2914. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

The GRC implementation, GRQ calculator, user study protocol, and anonymised study data are available from the corresponding author upon reasonable request. GRC software is available as open-source.

Ethical Approval

The user study received ethical approval from the WEDSU Ethics Committee (reference WEDSU-SD-2025-034). All participants provided informed consent and were compensated. Ethical approval references: WEDSU-SD-2025-034, CETU-CS-2025-041, BARU-AI-2025-028.

Appendix A

GRC Implementation and GRQ Scoring Rubric

The following provides GRC component implementation details and GRQ evaluation methodology.

PCR and SDF Implementation

CAA Creative Style Profile and Generation

GRQ Calculation and Study Protocol