



A Markov Chain Analysis of Synthetic Film Genre Transitions on Streaming Platforms

Juan Prihanda Nainggolan^{1*}

¹Department of Mathematics, Universitas Sumatera Utara, Medan, 20155, Indonesia

Prihandajuan@gmail.com

ABSTRACT

This study examines a first-order Markov chain as an interpretable model of transitions between film genres. The analysis uses a supplied synthetic frequency table containing 340 transitions among Action, Comedy, Horror, Romance, and Fantasy; it does not use observed streaming-platform records. Transition probabilities were obtained by normalizing each row of the count matrix, and the stationary distribution was calculated from the full-precision probabilities under a time-homogeneous model. The resulting long-run probabilities were 0.231016 for Action, 0.216088 for Comedy, 0.211556 for Fantasy, 0.175860 for Horror, and 0.165480 for Romance. Action had the largest stationary share and Romance the smallest. All states communicated, and the existence of return paths of lengths two and three established aperiodicity despite zero diagonal entries. The example therefore had a unique stationary distribution to which the modeled genre distribution converged. These results describe the supplied synthetic construction and do not establish actual audience preferences or recommendation accuracy. The model offers a transparent baseline for explaining genre transitions, but empirical use would require documented event sequences, validation of modeling assumptions, and comparison with temporal and personalized recommendation methods.

Keywords: Markov chain; synthetic data; film genre; stationary distribution

Received: 2 January 2025

Revised: 4 February 2025

Published: 30 March 2025

INTRODUCTION

Streaming recommendations concern both the relevance of individual titles and the order in which viewers consume them. General recommender-system research describes a range of approaches and limitations (Adomavicius & Tuzhilin, 2005), while the Netflix case illustrates how recommendation algorithms operate within a broader service and evaluation process (Gomez-Urbe & Hunt, 2015). A transition between two genre labels is consequently a limited representation of viewing behavior: it captures sequence order but omits title attributes, exposure, and reasons for choosing a film.

Temporal and sequential modeling are established areas of recommendation research. Factorized personalized Markov chains combine sequential behavior with user-specific preferences (Rendle et al., 2010). Self-attentive models (Kang & McAuley, 2018), bidirectional transformer models (Sun et al., 2019), and convolutional sequence models (Tang & Wang, 2018) provide alternatives to a single-step transition model. The contribution of the present study is therefore a transparent worked example, rather than the introduction of Markov chains to recommendation research.

This study uses a supplied synthetic table to construct and verify a five-state transition matrix and its stationary distribution. The objectives are to explain the calculation, interpret the ranking implied by the matrix, and distinguish mathematical properties of the example from claims requiring real viewing data. The analysis does not evaluate a deployed recommendation system or a community intervention.

METHODS

The analysis uses synthetic aggregate counts across five genres: Action (A), Comedy (C), Horror (H), Romance (R), and Fantasy (F). The source manuscript describes a scenario of 120 users and 340 transitions. However, individual histories, a generation algorithm, random seed, and sequence lengths were not supplied. Accordingly, the unit analyzed here is the aggregate transition table. Its total of 340 transitions is reproducible, whereas the stated number of users cannot be verified from that table alone.

Each modeled event has one genre label. Time denotes the index of successive events, rather than elapsed calendar time. The first-order assumption makes the next genre depend only on the current genre; time homogeneity makes the transition matrix constant over steps. Neither assumption was empirically tested. Films with multiple genre labels would require an explicit assignment rule before applying this state representation to real data.

$$Pr(X_{t+1} = j | X_t = i, \dots, X_0) = Pr(X_{t+1} = j | X_t = i) = p_{ij} \quad (1)$$

In Equation (1), X_t is the genre at event t , and p_{ij} is the probability of a transition from i to j . The matrix was estimated by row normalization of the supplied counts, using Equation (2). Rows represent origin genres and columns destination genres. A dash on the source table's diagonal was interpreted as zero, consistently with the zero diagonal in its supplied probability matrix. This convention restricts the example to switches between different genres; it is not evidence that real users never watch consecutive films in the same genre.

$$p_{ij} = n_{ij} / n_i, \quad n_i = \sum_k n_{ik} \quad (2)$$

Here, n_{ij} is the count from genre i to j and n_i is the origin-row total. All row totals are positive, so normalization is defined. Displayed probabilities are rounded, but all calculations use the exact count ratios. For a row-vector distribution q_t , the next-step distribution is $q_{t+1} = q_t P$. The stationary row vector satisfies Equation (3).

$$\pi = \pi P, \quad \sum_i \pi_i = 1, \quad \pi_i \geq 0 \quad (3)$$

The stationary solution was computed by solving the linear system formed from $P^T - I$, replacing one redundant equation with the unit-sum constraint. Verification comprised row sums, nonnegativity, normalization of the stationary vector, and the residual $\pi P - \pi$. Convergence was assessed from the transition graph: every off-diagonal transition is

positive, and return paths of lengths two and three exist. Their greatest common divisor is one, establishing aperiodicity. Together with irreducibility in a finite state space, this guarantees a unique stationary distribution and convergence from any initial distribution. These standard finite-chain results follow the framework described by the Markov chain literature (Norris, 1997).

No prediction benchmark was run because event-level sequences and a temporal test set were unavailable. This restriction is material: empirical comparisons of session-based algorithms show the need to assess methods under a defined evaluation protocol (Ludewig & Jannach, 2018; Ludewig et al., 2021). The present calculations establish internal consistency of the synthetic example, not predictive performance.

RESULT AND DISCUSSION

Table 1 gives the complete synthetic count matrix. Four origin genres have 70 transitions each, while Horror has 60, giving a total of 340. The unequal row totals are accounted for through row-wise normalization; using a common denominator of 340 would instead estimate joint transition proportions.

Table 1. Synthetic transition counts by origin and destination genre

From / To	Action	Comedy	Horror	Romance	Fantasy	Total
Action	0	22	14	10	24	70
Comedy	25	0	18	12	15	70
Horror	16	20	0	14	10	60
Romance	12	18	16	0	24	70
Fantasy	26	14	12	18	0	70

Note. Synthetic counts supplied in the source manuscript.

Table 2 reports the resulting transition probabilities. Conditional on Action, Fantasy has the largest next-step probability, $24/70 = 0.342857$. Conditional on Fantasy, Action has the largest probability, $26/70 = 0.371429$. These are properties of the constructed table. They are not estimates of population behavior on a named streaming service.

Table 2. Row-normalized transition probabilities

sdsdsssd	Action	Comedy	Horror	Romance	Fantasy
Action	0.000000	0.314286	0.200000	0.142857	0.342857
Comedy	0.357143	0.000000	0.257143	0.171429	0.214286
Horror	0.266667	0.333333	0.000000	0.233333	0.166667
Romance	0.171429	0.257143	0.228571	0.000000	0.342857
Fantasy	0.371429	0.200000	0.171429	0.257143	0.000000

Note. Values are rounded to six decimal places.

The exact rows sum to one. Small deviations in sums of displayed entries arise only from rounding. Zero diagonal entries do not make the chain periodic: Action can return through Comedy in two steps or through Comedy and Horror in three. Since every state communicates with every other state, all five states are recurrent in this finite chain. In particular, a low stationary share cannot justify calling Horror a transient state.

Table 3. Stationary probabilities and genre ranking

Genre	Stationary probability	Rank
Action	0.231016	1
Comedy	0.216088	2
Horror	0.175860	4
Romance	0.165480	5
Fantasy	0.211556	3

Note. Values are rounded to six decimal places.

The stationary distribution in Table 3 ranks Action first, followed by Comedy, Fantasy, Horror, and Romance. In the state order A, C, H, R, F, the full-precision solution is approximately (0.23101620, 0.21608816, 0.17585953, 0.16548003, 0.21155609). The maximum absolute stationarity residual is below 10^{-12} . The sum of the displayed six-decimal entries is one. Romance has the lowest stationary probability; Horror is fourth. The result follows from the complete network of transitions rather than from a single row or column alone.

A stationary share is the limiting proportion of modeled events in a state under repeated application of the same matrix. It is neither a measure of time spent watching films nor a prediction of market share. It also differs from the next-genre distribution for a user currently in a particular state: that conditional distribution is the corresponding row of P . Using the same stationary ranking for every viewer would discard the immediate-state information that motivates the Markov formulation.

The model also pools all users into one matrix, so it cannot distinguish different viewing profiles. Personalized Markov models address a different task by incorporating user variation (Rendle et al., 2010). Session-based graph models (Wu et al., 2019) and attentive session models (Li et al., 2017) offer richer representations of within-session behavior. Their existence demonstrates that the simple model should be presented as an interpretable baseline rather than as an untested replacement for established methods.

Recent model development includes self-supervised objectives based on mutual information (Zhou et al., 2020), contrastive learning (Xie et al., 2022), filter-enhanced multilayer perceptrons (Zhou et al., 2022), and disentangling long- and short-term interests (Zheng et al., 2022). These approaches motivate possible comparisons once suitable sequences are available. They were not trained or tested in this study, and no superiority over them is claimed. A replication study of BERT4Rec further illustrates why implementation and tuning details matter in such comparisons (Petrov & Macdonald, 2022).

Any empirical extension should retain event order, document session boundaries, and separate model selection from testing. The MovieLens dataset description is relevant when considering public movie-recommendation benchmarks, although ratings and their timestamps must not automatically be interpreted as verified viewing events (Harper & Konstan, 2015). A unified experimental framework such as RecBole can support implementation consistency (Zhao et al., 2021), but it cannot correct an inappropriate interpretation of the underlying data.

Evaluation should compare the Markov baseline with simple popularity and other relevant baselines, using a stated temporal protocol and identical candidate sets. Research on progress in recommendation algorithms emphasizes the importance of strong baselines (Ferrari Dacrema et al., 2019). Work on the robustness and discrimination of ranking metrics (Valcarce et al., 2020) and evaluation bias (Jeunen, 2019) also cautions against treating a single offline metric as a complete measure of system quality. Neither recommendation accuracy, diversity, user satisfaction, nor engagement was measured here.

The principal limitation is the undocumented construction of the synthetic table. Without the original sequences or generation mechanism, the example cannot support uncertainty estimates for a real population or tests of the first-order assumption. Zero self-transitions further constrain the interpretation. Future work should specify whether consecutive same-genre events are retained, provide a reproducible construction or empirical data protocol, and examine the stability of the resulting stationary ranking under changes in transition counts.

CONCLUSION

Row normalization of the supplied 340 synthetic transitions produces a valid five-state Markov chain with a unique stationary distribution. Action has the largest long-run share at approximately 23.10%, followed by Comedy at 21.61%, Fantasy at 21.16%, Horror at 17.59%, and Romance at 16.55%. Every state is recurrent; the matrix provides no basis for classifying Horror as transient. The analysis supplies a reproducible numerical example of genre transitions under fixed assumptions. Application to streaming recommendations requires documented event-level data, assessment of the modeling assumptions, and comparative evaluation before claims about audience preferences or recommendation benefits can be made.

BIBLIOGRAPHY

- Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734–749. <https://doi.org/10.1109/tkde.2005.99>
- Ferrari Dacrema, M., Cremonesi, P., & Jannach, D. (2019). Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems* (pp. 101–109). <https://doi.org/10.1145/3298689.3347058>
- Gomez-Uribe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), 1–19. <https://doi.org/10.1145/2843948>
- Harper, F. M., & Konstan, J. A. (2015). The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4), 1–19. <https://doi.org/10.1145/2827872>
- Jeunen, O. (2019). Revisiting offline evaluation for implicit-feedback recommender systems. In *Proceedings of the 13th ACM Conference on Recommender Systems* (pp. 596–600). <https://doi.org/10.1145/3298689.3347069>

- Kang, W.-C., & McAuley, J. (2018). Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)* (pp. 197–206).
<https://doi.org/10.1109/icdm.2018.00035>
- Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., & Ma, J. (2017). Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (pp. 1419–1428).
<https://doi.org/10.1145/3132847.3132926>
- Ludewig, M., & Jannach, D. (2018). Evaluation of session-based recommendation algorithms. *User Modeling and User-Adapted Interaction*, 28(4-5), 331–390.
<https://doi.org/10.1007/s11257-018-9209-6>
- Ludewig, M., Mauro, N., Latifi, S., & Jannach, D. (2021). Empirical analysis of session-based recommendation algorithms: A comparison of neural and non-neural approaches. *User Modeling and User-Adapted Interaction*, 31(1), 149–181.
<https://doi.org/10.1007/s11257-020-09277-1>
- Norris, J. R. (1997). *Markov chains*. Cambridge University Press.
<https://doi.org/10.1017/cbo9780511810633>
- Petrov, A., & Macdonald, C. (2022). A systematic review and replicability study of BERT4Rec for sequential recommendation. In *Proceedings of the 16th ACM Conference on Recommender Systems* (pp. 436–447).
<https://doi.org/10.1145/3523227.3548487>
- Rendle, S., Freudenthaler, C., & Schmidt-Thieme, L. (2010). Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on World wide web* (pp. 811–820).
<https://doi.org/10.1145/1772690.1772773>
- Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., & Jiang, P. (2019). BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (pp. 1441–1450).
<https://doi.org/10.1145/3357384.3357895>
- Tang, J., & Wang, K. (2018). Personalized top-N sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (pp. 565–573).
<https://doi.org/10.1145/3159652.3159656>
- Valcarce, D., Bellogín, A., Parapar, J., & Castells, P. (2020). Assessing ranking metrics in top-N recommendation. *Information Retrieval Journal*, 23(4), 411–448.
<https://doi.org/10.1007/s10791-020-09377-x>
- Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., & Tan, T. (2019). Session-based recommendation with graph neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 346–353.
<https://doi.org/10.1609/aaai.v33i01.3301346>

- Xie, X., Sun, F., Liu, Z., Wu, S., Gao, J., Zhang, J., Ding, B., & Cui, B. (2022). Contrastive learning for sequential recommendation. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)* (pp. 1259–1273). <https://doi.org/10.1109/icde53745.2022.00099>
- Zhao, W. X., Mu, S., Hou, Y., Lin, Z., Chen, Y., Pan, X., Li, K., Lu, Y., Wang, H., Tian, C., Min, Y., Feng, Z., Fan, X., Chen, X., Wang, P., Ji, W., Li, Y., Wang, X., & Wen, J.-R. (2021). RecBole: Towards a unified, comprehensive and efficient framework for recommendation algorithms. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (pp. 4653–4664). <https://doi.org/10.1145/3459637.3482016>
- Zheng, Y., Gao, C., Chang, J., Niu, Y., Song, Y., Jin, D., & Li, Y. (2022). Disentangling long and short-term interests for recommendation. In *Proceedings of the ACM Web Conference 2022* (pp. 2256–2267). <https://doi.org/10.1145/3485447.3512098>
- Zhou, K., Wang, H., Zhao, W. X., Zhu, Y., Wang, S., Zhang, F., Wang, Z., & Wen, J.-R. (2020). S3-Rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (pp. 1893–1902). <https://doi.org/10.1145/3340531.3411954>
- Zhou, K., Yu, H., Zhao, W. X., & Wen, J.-R. (2022). Filter-enhanced MLP is all you need for sequential recommendation. In *Proceedings of the ACM Web Conference 2022* (pp. 2388–2399). <https://doi.org/10.1145/3485447.3512111>