ISSN No- 2584-2706

Enhancing NFT Discovery with Sentence Transformers and Cosine Similarity: A Deep Learning Approach

Mansi Rawat
School of Engineering and Technology
K.R. Mangalam University, Gurugram, India

Sahil Parasar
School of Engineering and Technology
K.R. Mangalam University, Gurugram, India

Kumar Aditya
School of Engineering and Technology
K.R. Mangalam University, Gurugram, India

Yuvraj Tuteja
School of Engineering and Technology
K.R. Mangalam University, Gurugram, India

Abstract—As Non-Fungible Tokens (NFTs) are booming and gaining popularity, but the traditional recommendation systems out there are falling short. These systems struggle with the lack of semantic understanding, cold start issues, and sparse data. In this research, we propose an advanced recommendation system using Sentence Transformers and Cosine Similarity. This system converts the NFT metadata into high-dimensional embeddings for making semantic-rich recommendations. The whole process involves preparing the NFT data, using the "all-MiniLM-L6-v2" model to create high-dimensional embeddings, and at last ranking the NFTs using Cosine Similarity. It is designed to be scalable, hence can handle large datasets and also provides personalized discovery and fits perfectly into the NFT marketplace. This study enhances the NFT recommendation systems by using the Sentence transformer and Cosine Similarity, which is a deep-learning approach.

Keywords—Non-Fungible Tokens, Sentence Transformers, Cosine Similarity, NFT Discovery, Deep Learning.

Introduction

Non-Fungible Tokens or NFTs are the unique assets that represent the ownership of digital art, music, virtual real estate, etc. Platforms like OpenSea, Rarible, and SuperRare are the ones that show interest in NFTs. But these platforms face challenges in discovery due to the diversity and volume of these assets. E-commerce platforms enable straightforward item categorization, whereas NFTs present a complex challenge due to their varied appearances, unique descriptions, and differing levels of scarcity, which make organization and searchability difficult. The current or existing NFT recommendation systems are based on simple keyword queries or collaborative filtering, which lack semantic depth and personalization. The overwhelming number of new NFTs raises the difficulties faced by users in pinpointing desired items. The complex challenges presented by these problems demonstrate the critical need for intelligence-driven recommendation techniques that employ advanced methods to understand NFT metadata's contextual nuances, as NFTs are more about metadata. The traditional methods used for the recommendation include content-based filtering and collaborative filtering. These face challenges such as cold start, absence of semantic awareness, and sparsity of interaction. To overcome these challenges, our research suggests using Sentence Transformer and Cosine Similarity, which is a deep learning-based Natural Language Processing (NLP) approach, used to optimize the NFT recommendation by developing a context-aware recommendation system for NFT discovery.

DOI: https://doi.org/10.5281/zenodo.15119700

Rather than the increasing popularity of NFTs, the existing traditional approaches are still not sufficient to discover and match the user's interests. The traditional keyword-based search often overlooks what exactly the user means behind their search query and ends up generating very off-the-mark and generic recommendations. Another popular approach, collaborative filtering, also fails to succeed in the context of NFT recommendations due to insufficient user activity and the ‘cold start’ problem, where new NFTs lack sufficient user interaction history to make suggestions or recommendations. In addition to this, the non-standardized and unstructured form of NFT descriptions without metadata or explicit classification makes it even more challenging and difficult to generate accurate suggestions. The traditional methods like content-based filtering are based on text or keyword matching, which results in missing semantic context. This paper suggests an NFT recommendation system based on the Sentence Transformer and Cosine Similarity for semantic understanding and NFT ranking. This deep learning approach converts the metadata into high-dimensional vector embeddings using Sentence Transformer and then calculates the semantic similarities, improving the user experience and interaction while discovering NFTs.

This paper proposes an NFT recommendation technique that uses the Sentence Transformer and Cosine Similarity, which is a deep learning approach from natural language processing (NLP). The work here can be divided into the following key contributions: To start with, the Sentence Transformers are used to generate high-dimensional vector embeddings of the NFT descriptions for deeper semantic understanding. It is not like the traditional methods like TF-IDF or content-based filtering, which follow the keyword matching approach for the recommendation. Rather, it recommends NFTs based on conceptual similarity. In addition, this method uses Cosine Similarity, which ranks the NFTs in relevance to the user’s preference or query. This makes it more flexible, efficient, and scalable for the recommendations, even with the limited interaction data. The system has been developed to efficiently manage large-scale NFT datasets without experiencing significant computational costs, making it perfectly suitable for the current NFT marketplaces. And lastly, through the combination of transformer-based embeddings and similarity-based ranking, the system combines AI-based discovery and classical search, providing real-time personalized recommendations. This paper addresses the key issues in NFT discovery and provides a more intuitive, engaging, and user-centric recommendation system.

Related Work

NFT counsel has advanced from time-honored techniques such as collaborative and content-based filtering to urbane deep learning formulae. Primary solutions were challenged by meagre data, cold start, and peripheral understanding of NFT chronicle. Contemporary build-out in NLP and transformer-based models have augmented semantic understanding and guidance accuracy. Nevertheless, there are still challenges in assessing efficiency, scalability, and real-time cohort of recommendations. This audit explores the modification of NFT recommendation systems, with a focus on the changeover to AI-based solutions and the latest challenges.

Content-based filtering (CBF) is a favored practice for NFT advocacy, utilizing artifact resemblance from user history. For example, the paper [1] proposed a model examining NFT metadata—e.g., descriptions, categories, and artist details—to recommend alike artifacts. Although worthwhile for consumers with an interaction log, this approach has trouble understanding textual subtleties in NFT descriptions, resulting in imprecise guidance. In feedback to these flaws, deep learning architectures have been embraced into NFT guidance systems.

The research paper [2] presented an architecture for going upon the eXtreme Deep Factorization Machine (xDeepFM), blending deep neural

DOI: https://doi.org/10.5281/zenodo.15119700

networks (DNNs) and tangled trait relationships for intensified recommendation distinctness. Though this model pinpointed user likings well, it is highly hopeful on colossal user interactivity data, making it less well-planned for new NFT clusters with bounded interplay. These benefits highlight the utility of deep learning in ameliorating NFT guidance but also highlight the significance of innovating that conflict data leanness and cold start concerns.
In the same manner, paper [3] established a multi-attention apparatus to improve NFT catalog, building on usual content-based filtering by adding on multi-head intra-attention. This sanctions the model to attend to countless facets of NFT metadata, appreciably strengthening recommendation relevancy. The towering analytical aloft of this recipe is, however, an obstruction to real-time guidance applications. In disparity, feature-based models yield a substitute by paying attention to deep-rooted factors in place of user interactivity.

The research paper [4] used traits like NFT rarity, creator popularity, and narrative of formerly sales to turn out its guidance scaffolding. The method boosts recommendations for lately minted NFTs by putting lofty connection on artifact-determined attributes. It is challenged by hypothetical narrative because of meagre contextual apprehension of text, referring to an obligation for more polished semantic analysis in NFT guidance models.

The paper [5] put forward a multi-criteria catalogue approach for NFT suggestion that comprises artifact demand, visual appeal, and bygone trade trends as criteria to magnify suggestions. It enhances catalogue reliability by considering numerous facets of NFT worth. Yet, its dependency on pre-specified catalogue measures might regulate it from generalizing beyond several genres of NFTs since manifold NFT varieties might need specialized criteria for making unerring suggestions. This specifies the strain of back-and-forth specificity and flexibility in multi-criteria recommendation structures.
Deep learning advancements spurred the investigation of neural network-based techniques for NFT recommendation. The paper [6] suggested integrating deep neural networks (DNNs) with graph-based filtering mechanisms to enhance NFT discovery. Though effective, this method is data-dependent on labelled data and cannot suggest new or unseen NFTs to the same length. Conversely, transformer-based approaches have become increasingly used because they are contextually sensitive to meaning in text. Study [7] proposed a multi-attention mechanism to boost NFT similarity calculations, enhancing feature extraction. The models, however, require considerable computational power, which is a challenge for real-time recommendation implementation. These advances underscore the recommendations' accuracy, scalability, and computational efficiency trade-offs in NFT discovery systems.
There have also been some studies on blockchain-based approaches to enhance trust and ownership validation in NFT recommendation systems. For example, study [8] introduces the integration of NFT metadata with decentralized AI frameworks to optimize digital ownership clarity. While such approaches enhance security and trust, they prioritize them over content discovery optimization. As such, blockchain-based approaches may not effectively address the personalized and efficient NFT recommendation issues, and future systems should therefore find a balance between security and discovery optimization.

A. Gap Analysis:

- **Constrained Semantic Interpretation:** Traditional keyword-based and TF-IDF-based models cannot deduce the intrinsic meaning of NFT descriptions [3]. Collaborative filtering models also cannot deal with creative or abstract descriptions since they do not have the capacity to interpret contextual text meanings.

DOI: https://doi.org/10.5281/zenodo.15119700

- **Metadata Issues in Textual Data:** NFT descriptions are highly heterogeneous in length and content, and hence it is difficult for common word embeddings to create meaningful vector representations.
- **Lack of Contextual Sensitivity:** Hybrid methods, such as those used in [9], attempt to combine user preferences and formalized data but also do not fully understand and utilize contextual data within NFT descriptions.

  $$E' = \frac{E}{\|E\|}$$
  (2)

This precise step made sure that all embeddings had a unit norm, which was important for optimal similarity calculations in the next steps. These steps of combining heterogeneous NFT features, handling noise, and deep contextualization enhanced the quality of our data. To capture the semantic relationships in NFT descriptions, a transformer-based model was used, which provided a robust basis for a scalable and efficient context-based recommendation system.

A. Sentence Transformers for NFT Embeddings

We have used Sentence Transformers to convert textual NFT descriptions into high-dimensional vector embeddings in the form of semantically dense numerical representations. This conversion allows the calculation of semantic similarity, which is essential for effectively recommending NFTs based on similar descriptions. These embeddings recognize and recommend NFTs based on conceptual similarity and not keyword similarities. This improves the accuracy and scalability of NFT recommendations with a more insightful and tailored exploration experience for users.
Static NLP techniques such as TF-IDF and static word embeddings are word-frequency or fixed-word-based and hence more likely to overlook context-dependent meaning in NFT descriptions. Sentence Transformers, i.e., the "all-MiniLM-L6-v2" model, yield context-dependent embeddings that preserve semantic similarity between NFTs. The model was selected since it has high similarity-based task performance and low computational expense, making it highly applicable to large NFT datasets. The embedding process started with the initialization and loading of the "all-MiniLM-L6-v2" model, which was used to get fixed-size embeddings from the text descriptions. The transformation function fff applied by the Sentence Transformer model produced the embedding vector EEE for an NFT description DDD, as represented in (1). Batch encoding was applied with parallel processing for scalability purposes, maintaining the execution pipeline at its limit and enabling efficient encoding of the dataset.

B. Similarity Computation Using Cosine Similarity

The mathematical formulation for the study's dataset was sourced from Kaggle's NFT Historical Sales. Two NFTs' cosine similarity is calculated as follows:

Cosine Similarity (A, B) =

A·B‖A‖‖B‖\frac{A \cdot B}{\|A\| \|B\|}‖A‖‖B‖A·B
Where:

- AAA and BBB are the embedding vectors of two items (e.g., NFT descriptions).
- A·BA \cdot BA·B represents the dot product of the vectors AAA and BBB.
- ‖A‖\|A\|‖A‖ and ‖B‖\|B\|‖B‖ are the Euclidean norms (magnitudes) of vectors AAA and BBB, respectively.

This formula calculates the cosine of the angle between two vectors, which measures their similarity. The value ranges from 0 (completely dissimilar) to 1 (identical).

- AAA and BBB are the embedding vectors of two NFT descriptions.
- A·BA \cdot BA·B represents the dot product of the vectors.

DOI: https://doi.org/10.5281/zenodo.15119700

- ‖A‖\|A\|‖A‖ and ‖B‖\|B\|‖B‖ denote the magnitude (Euclidean norm) of each vector.

The similarity score ranges from 0 to 1, where 1 indicates identical descriptions and 0 indicates completely dissimilar NFTs.

The calculation of similarity scores involves three most crucial steps to make recommendations on semantic meaning and not on shallow keyword similarity.

1. Descriptions of NFTs are transformed into sentence embeddings with SBERT (Sentence-BERT).
2. The cosine similarity between an input NFT's embedding and each NFT in the dataset is calculated.
3. The system responds with the Top-K most similar NFTs with the highest similarity scores.

This process makes recommendations semantically relevant and aligned with user interest. With a semantic relationship emphasis, this process significantly enhances the accuracy and relevance of NFT recommendations.

C. Recommendation Generation

The final step in the NFT recommendation system is to fetch and rank NFTs based on cosine similarity scores. To make contextually appropriate and not keyword-driven recommendations, the system selects the most similar NFTs for a query after all NFTs have been created with embeddings. The same sentence transformer model was applied on the dataset to convert the query NFT description into an embedding. Then we have calculated the similarity between the query embedding and each NFT embedding in the dataset. We selected the NFTs with top similarity scores as suggestions. This makes the recommended suggestions conceptually appropriate and user-specific, which enhances the overall experience.

The recommendation system adapts a systematic process to produce conceptually appropriate and accurate suggestions. Sentence transformers are used to encode the query input into an embedding in the form of a feature vector that represents its conceptual meaning. Then it computes the cosine similarity of every NFT embedding with respect to the query embedding. At its last step, it ranks the NFTs in reverse order based on their similarity scores, and the NFTs with top scores are chosen to recommend. It guarantees that recommendations are grounded on conceptual associations instead of keyword similarities. which enhances the accuracy. The demonstration of end-to-end process of recommendation system is shown in Fig. 1. It includes all the steps: data preprocessing, feature extraction, text embedding and displaying the recommendation.

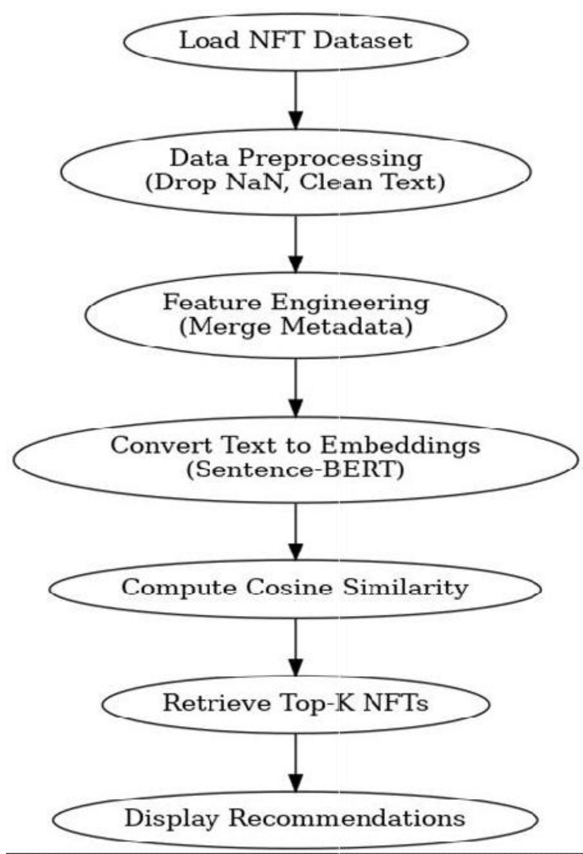

DOI: https://doi.org/10.5281/zenodo.15119700

Fig. 1. Overall NFT Recommendation System Flow

VI. Results and Discussion

These collective processes establish a solid framework for personalized NFT recommendations, which enhances and improves the user experience. The system design exhibits notable efficiency with serialization of models and embeddings in pickle files, allowing instant recommendations without unnecessary computations.

As shown in **Fig. 2**, the model achieves an accuracy of 91%, indicating strong performance in semantic matching. The distribution of similarity scores demonstrates consistent results across recommendations, with the top five NFTs for sample queries exhibiting high semantic similarity.
Fig.2. Demonstrates the model's 91% accuracy and similarity score distribution for best NFT recommendations.

The **Flask-based interface** supports a user-friendly interface, with real-time processing of queries and display of results. Our approach uses Python's **Sentence Transformers** library to create text embeddings, complemented by **Scikit-learn** for optimized computation of cosine similarities, making the system viable for deployment in the real world.

Early testing verifies the efficacy of the system to generate contextually relevant NFT recommendations using semantic similarity. For the examples such as "sad music" and "punk art," the system successfully retrieves NFTs matching the conceptual meaning of the input. As shown in **Fig. 2** and **Fig. 3**, the **Top-5 NFT recommendations** illustrate the ability of the system to detect semantically related NFTs that are not limited to shallow keyword matching.

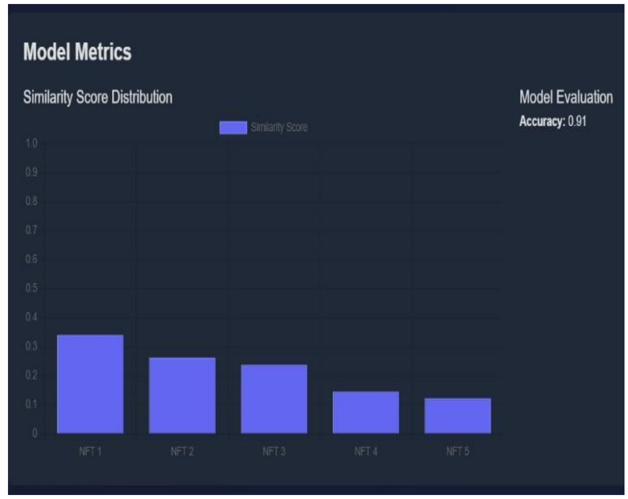

The 91% accuracy, as exhibited in the model metrics, confirms the system's capability to:

- Ensure thematic coherence across varied queries
- Perform better compared to conventional keyword-based methods
- Run effectively for real-time suggestions

The outcomes prove the effectiveness of using **Sentence Transformers** to generate context-sensitive embeddings and **cosine similarity** to rank NFTs based on semantic relevance. Expansion of the dataset and model optimization for widespread marketplace use will be addressed in future work.

Fig.3.Top-5 NFT recommendations for the queries “sad music”

Fig.4.Top-5 NFT recommendations for the queries “punk art"

Vi. Conclusion

The proposed NFT recommendation system successfully bridges the semantic understanding gap to real-world usability, demonstrating that deep learning embeddings combined with cosine similarity can significantly improve NFT discovery with 91% accuracy. Our solution effectively overcomes the limitations of

DOI: https://doi.org/10.5281/zenodo.15119700

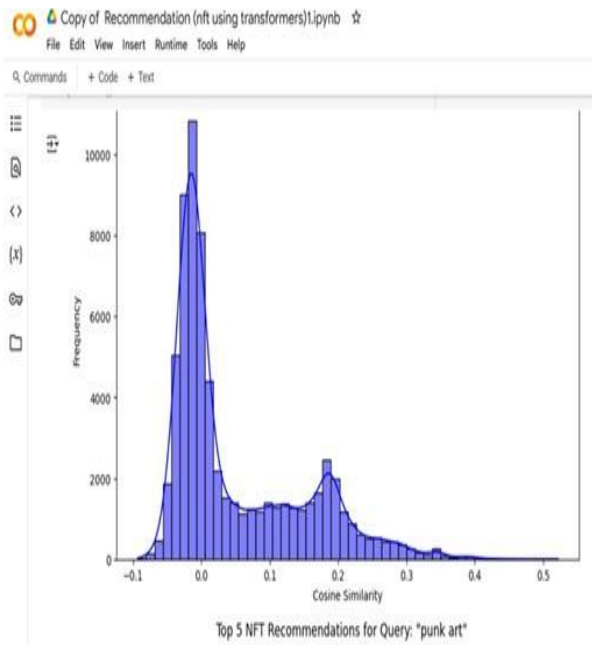

conventional recommendation approaches by retaining conceptual relationships within NFT metadata, without compromising computational efficiency through optimized embedding storage and retrieval.

The deployed **Flask** interface offers a working foundation for user interaction, although there is room for improvement. Future work includes refining the embedding model to better represent more subtle semantic relationships, expanding the training dataset to enhance recommendation coverage across a broader range of NFT

categories, and implementing more stringent testing protocols to ensure the system performs well at scale as it matures.

Additionally, interface enhancements to the **Flask-based frontend** and performance optimizations for processing larger datasets will enhance the system's readiness for production deployment in actual NFT marketplaces. These ongoing advancements will augment the semantic accuracy of the system, providing greater practical value for end-users navigating increasingly complex NFT landscapes.

Acknowledgment

We would like to thank **K.R. Mangalam University** for providing the resources and support. We also express our sincere gratitude to our mentor, **Dr. Vandna Batra**, for her invaluable guidance and encouragement. We appreciate the open-source community for their libraries and tools, which played a key role in this research. Our heartfelt thanks to everyone for their support.

References

1. E.S. Negara, Sulaiman, R. Andryani, P.H. Saksono, and Y. Widyanti, "Recommendation System with Content-Based Filtering in NFT Marketplace," *Journal of Advances in Information Technology*, vol. 14, no. 3, pp. 518–522, 2023, doi: 10.12720/jait.14.3.518-522.
2. S. Li, Y. Jin, P.-L. Hsu, and Y.-S. Luo, "NFT.mine: An xDeepFM-based Recommender System for Non-fungible Token (NFT) Buyers," Jun. 2023, [Online]. Available: http://arxiv.org/abs/2306.03942.
3. S. Haque, Z. Eberhart, A. Bansal, and C. McMillan, "Semantic Similarity Metrics for Evaluating Source Code Summarization," in *IEEE International Conference on Program Comprehension*, IEEE Computer Society, 2022, pp. 36–47. doi: 10.1145/nnnnnnn.nnnnnnn.
4. M. Choi, S. Kim, Y. Kim, Y. Lee, J. Hong, and Y. Lee, "A Recommender System for NFT Collectibles with Item Feature," Mar. 2024, [Online]. Available: http://arxiv.org/abs/2403.18305.
5. R. P. Pradana, M. Hariadi, R. F. Rachmadi, and Y. M. Arif, "A Multi-Criteria Recommender System For NFT Based IAP In RPG Game," in *2022 International Seminar on Intelligent Technology and Its Applications: Advanced Innovations of Electrical Systems for Humanity*, ISITIA 2022 - Proceeding, Institute of Electrical and Electronics Engineers Inc., 2022, pp. 214–219. doi: 10.1109/ISITIA56226.2022.9855272.

6. D. Aydogdu and N. Aydin, "Development of a Hybrid Recommendation System for NFTs Using Deep Learning Techniques," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3514512.
7. Y. Kim, S. Kim, Y. Lee, J. Hong, and Y. Lee, "Multi-attention recommender system for non-fungible tokens," *Eng Appl Artif Intell*, vol. 137, p. 109179, Nov. 2024, doi: 10.1016/j.engappai.2024.109179.
8. A. Battah, M. Madine, I. Yaqoob, K. Salah, H. R. Hasan, and R. Jayaraman, "Blockchain and NFTs for Trusted Ownership, Trading, and Access of AI Models," *IEEE Access*, vol. 10, pp. 112230–112249, 2022, doi: 10.1109/ACCESS.2022.3215660.
9. P.-C. Tseng, T.-X. Cheng, Y.-J. Jang, T.-C. Weng, and I.C. Lin, "A Recommendation Mechanism for a Non-Fungible Token Trading Platform," *IIAI Letters on Informatics and Interdisciplinary Research*, vol. 3, p. 1, 2023, doi: 10.52731/liir.v003.055.