

# MultiNFT: A Multimodal Dataset for Non-Fungible Tokens Market Analysis

Shuying Liu  
Adam Smith Business School  
University of Glasgow  
Glasgow, United Kingdom  
2595459L@student.gla.ac.uk

Bowei Chen  
Adam Smith Business School  
University of Glasgow  
Glasgow, United Kingdom  
Bowei.Chen@glasgow.ac.uk

Cathy Yi-Hsuan Chen  
Adam Smith Business School  
University of Glasgow  
Glasgow, United Kingdom  
CathyYi-Hsuan.Chen@glasgow.ac.uk

Yu Zheng  
School of Finance  
Southwestern University of Finance and Economics  
Chengdu, China  
zhengyu@swufe.edu.cn

Jun Wang  
Department of Computer Science  
University College London  
London, United Kingdom  
jun.wang@cs.ucl.ac.uk

**Abstract**—Non-fungible tokens (NFTs) are unique digital assets that play an increasingly important role in decentralized markets, supporting new forms of ownership, valuation, and exchange. Their inherently multimodal structure, which encompasses visual content, metadata, and trading history, has led to a growing academic interest in modeling NFT pricing and market behavior. However, existing research is limited by the lack of comprehensive datasets that unify these modalities with consistent formatting and longitudinal coverage. To address this gap, we introduce MultiNFT, a large-scale multimodal dataset comprising 50 curated profile picture (PFP) NFT collections, including 523,020 unique assets and 2.38 million transaction records from April 2021 to September 2025. MultiNFT integrates standardized images, structured metadata, and time-series trading data, along with rarity scores and aesthetic features, offering a unified foundation for multimodal learning and NFT analytics. Unlike prior datasets that focus on visual similarity or static snapshots, MultiNFT captures evolving valuation dynamics across market cycles and connects them to trait-level characteristics. We demonstrate the utility of the dataset through three case studies, including within-collection rarity-price analysis, visual feature clustering across collections, and quantifying feature contributions in a comprehensive pricing model. By bridging computer vision, behavioral modeling, and financial forecasting, MultiNFT supports a wide range of interdisciplinary research and practical use cases. The dataset is publicly available and is intended to promote reproducible experimentation and further exploration of the mechanisms driving value in digital asset ecosystems.

**Index Terms**—Non-fungible token, multimodal dataset, asset pricing, visual analytics, big data

## I. INTRODUCTION

Non-fungible tokens (NFTs) are unique digital assets recorded on the blockchain that serve as certificates of ownership and authenticity for various digital content such as artworks, collectibles, and gaming items. Unlike interchangeable cryptocurrencies, NFTs are inherently distinct and irreplaceable. These characteristics have led to widespread applications of NFTs in digital ownership, crypto art, and metaverse ecosystems [1]–[3]. Today, NFTs have become an important

component of digital markets. Understanding the underlying value drivers, pricing mechanisms, and marketing strategies has become a shared focus of both academia and industry.

NFT data integrates on-chain transaction records, metadata descriptions, and rich visual and multimedia content, naturally possessing multimodal data characteristics. This diversity in data types makes them ideal subjects for multimodal machine learning, data fusion analysis, and information systems research. In recent years, with the rapid development of NFT markets, both academic and industrial communities have shown growing demands for comprehensive multimodal NFT datasets. The integrated datasets can provide a foundation for various economic research, including NFT price prediction, user behavior analysis, and market dynamic modeling. Moreover, they also provide deep insights into NFT visual aesthetic characteristics, offering data-driven guidance for digital art design optimization and innovative marketing strategies.

However, existing NFT data resources still face numerous challenges, including data fragmentation, missing modalities, inconsistent formatting, and limited coverage. Current datasets typically address only single aspects of NFT data, focusing exclusively on transaction histories or metadata while lacking systematic cross-modal integration. The visual components often remain unprocessed and inconsistent in format, limiting the applicability for rigorous quantitative analysis and machine learning applications. Additionally, differences in data collection and processing methods result in limited compatibility and comparability among datasets. These limitations constrain the development of large-scale data-driven automated analysis, machine learning model optimization, and end-to-end intelligent system development.

Motivated by these limitations, we construct and release MultiNFT, a large-scale, multimodal, and carefully curated NFT dataset. MultiNFT encompasses complete transaction histories for 50 profile-picture (PFP) NFT projects over a 54-month period, together with asset metadata and standardized

NFT images obtained and processed through a unified pipeline. All data undergo quality validation, attribute alignment, and cross-modal association, thereby supporting a wide range of analytical applications and modeling tasks.

Our work makes three primary contributions. First, as a data resource, MultiNFT brings together complete metadata and standardized visual content for 523,020 NFT avatars, along with 2.38 million transaction records from April 2021 to September 2025, offering coherent coverage across temporal, structural, and visual dimensions. In doing so, it introduces a large-scale cross-modal open resource that has so far been largely absent in the NFT domain. Second, on the methodological side, we propose a standardized pipeline for multimodal NFT data collection, processing, and integration, which can serve as a practical template for constructing similar datasets and help lower technical barriers for researchers and practitioners working with NFT data. Third, at the application level, MultiNFT enables in-depth studies of pricing dynamics and market behavior, and supports advanced machine learning tasks—such as classification, representation learning, and forecasting—by jointly leveraging visual, textual, and transactional information. Analyses based on MultiNFT can shed light on the functioning of digital asset markets, with implications for decision-making, rarity assessment, and risk evaluation, and can also contribute to broader conversations around industry practice and regulatory perspectives.

## II. REVIEW OF NFT RESEARCH AND DATASETS

NFT research has rapidly evolved to address the unique challenges of valuing and analyzing digital assets that combine multiple data modalities. NFTs, particularly in digital art and PFP collections, exhibit inherently multimodal characteristics by integrating structured transaction records, textual trait metadata, and visual content [4], [5]. Early studies primarily focused on market dynamics and pricing mechanisms using traditional econometric approaches, but recent work has expanded significantly in scope and methodology. Contemporary research explores diverse aspects, including price prediction and market forecasting [5], [6], fraud detection and security analysis [7], and behavioral phenomena such as herding effects, social influence, and collector preferences [8]–[10].

A notable shift has occurred toward investigating content-driven value determinants, with researchers examining visual aesthetics, trait-based rarity metrics, and descriptive attributes as key pricing factors. This evolution has spurred innovative approaches in visual analytics [11], the development of rarity-based pricing models [9], [12], and the creation of multimodal recommendation systems [13], [14]. These developments have driven methodological advances from traditional unimodal baselines toward sophisticated integrated frameworks that combine behavioral, visual, and semantic signals. However, these advanced multimodal approaches are increasingly data-dependent, requiring comprehensive and well-structured datasets to achieve robust performance across different collections and market conditions.

TABLE I  
COMPARISON OF VARIABLE COVERAGE ACROSS NFT DATASETS.

Variables/features	Related datasets					MultiNFT
	[11]	[20]	[9]	[15]	[14]	
Sale/transaction price	✓	✓	✓	✓		✓
Buyer/seller address	✓	✓				✓
NFT images (raw)					✓	✓
Visual features	✓	✓				✓
Trait types and values			✓		✓	✓
Rarity scores			✓			✓
Cryptocurrency price				✓		✓
Gas fees				✓		✓

Note: ✓ indicates that the variable is included in the dataset.

Despite these methodological advances, further progress is significantly constrained by the limitations of existing NFT datasets. Available resources typically suffer from fragmentation, single-modality focus, inconsistent formatting, or inadequate coverage across collections and temporal periods [14], [15]. These limitations present significant challenges for developing interpretable models and building end-to-end multimodal learning pipelines. While traditional methods struggle to unify heterogeneous data types for comprehensive market modeling [16], emerging AI tools show promising potential [17]–[19]. However, the absence of standardized, large-scale multimodal datasets remains a critical bottleneck for advancing the field. Table I presents a systematic comparison of representative NFT datasets, highlighting the gaps in multimodal coverage that our work addresses.

## III. MULTINFT CONSTRUCTION AND DESCRIPTION

In this section, we outline the construction of the MultiNFT dataset, including data sources, preprocessing procedures, and the structure of the final multimodal resource.

### A. Data Source and Coverage

Our data is sourced from OpenSea,<sup>1</sup> the largest digital marketplace for NFT creation and trading. Using its official API,<sup>2</sup> we systematically collect multiple data types: historical sales records, structured metadata, and image assets.

To inform the selection of NFT collections included in MultiNFT, we conducted a systematic analysis of the top 100 collections on OpenSea, ranked by cumulative trading volume by the end of 2024. As shown in Fig. 1, PFP collections dominate the landscape, accounting for 54 of the top 100 entries, followed by Art collections with 21 entries. PFP NFTs are particularly well-suited for multimodal research due to their generative nature, structured and interpretable trait schemas, and highly engaged user communities. These characteristics facilitate cross-modal learning and comparative analysis at both the asset and collection levels. Based on this insight, we curated a subset of 50 representative PFP collections launched between April 2021 and March 2023, which is a critical period encompassing the peak and subsequent stabilization of the NFT market. This sampling strategy ensures a heterogeneous

<sup>1</sup><https://opensea.io>

<sup>2</sup><https://docs.opensea.io/reference/api-overview>

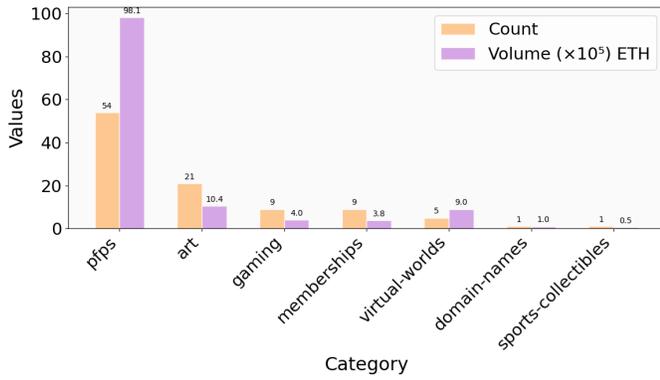


Fig. 1. Summary statistics of the top 100 collections.

and thematically coherent dataset, capturing variation in visual aesthetics, economic valuation, and community interaction patterns within the NFT ecosystem.

### B. Data Collection Pipeline

The data collection pipeline comprises four main stages, designed to acquire three data modalities and ensure data integrity through systematic cleaning and validation:

- **Sales data collection:** We systematically collect historical sales records spanning from April 2021 to September 2025, providing 54 months of comprehensive transaction history. This includes sale prices, buyer and seller addresses, transaction timestamps, payment tokens, and gas fees, enabling detailed market dynamics analysis.
- **Metadata extraction:** Complete structured metadata is retrieved for all tokens within selected collections, including trait types and values, token descriptions, creation dates, and external URL references. This structured information enables trait-level analysis, rarity assessment, and semantic understanding of asset characteristics.
- **Image download and standardization:** High-resolution images are systematically downloaded via metadata-linked URLs and standardized to  $224 \times 224$  RGB PNG format. This preprocessing ensures compatibility with standard deep learning architectures while maintaining the visual fidelity necessary for aesthetic analysis.
- **Data cleaning and quality control:** To ensure dataset integrity, we implement comprehensive quality control measures, including removal of entries with missing or corrupted URLs, incomplete metadata fields, or inaccessible resources. Content flagged as inappropriate or disabled is excluded, and invalid transactions such as zero-payment transfers are filtered out. All transaction values are standardized into USD, where ETH and WETH are converted using historical ETH-USD exchange rates from Yahoo Finance,<sup>3</sup> while tokens DAI and USDC are treated as direct 1:1 USD equivalents, enabling consistent economic analysis across payment tokens and time periods.

<sup>3</sup><https://finance.yahoo.com>



Fig. 2. Examples of NFT trait types and values from different collections.

### C. Rarity Scoring Methods

PFP NFTs are characterized by structured trait compositions, where each asset possesses a combination of trait types and values that define its visual and conceptual properties. As illustrated in Fig. 2, different collections employ varying trait schemas. For instance, Bored Ape Yacht Club utilizes traits such as "Background", "Hat", "Eyes", and "Clothes". In contrast, Azuki uses traits including "Type", "Background", "Eyes", and "Hair". These trait combinations enable quantitative assessment of rarity within each collection, as assets with uncommon trait values or rare trait combinations are typically perceived as more valuable by collectors.

To provide a comprehensive rarity assessment, we implement two widely used computational methods that offer complementary perspectives on asset scarcity: inverse frequency-based rarity [21] and information content-based rarity [22].

#### Inverse frequency-based rarity (IF):

$$S_{IF} = \sum_{j=1}^M s_{i,j} + s_c, \quad (1)$$

where  $s_{i,j} = \frac{N}{n_{i,j}} \times \frac{1}{u_j}$  is the score for trait value  $i$  under type  $j$ ,  $s_c = \frac{N}{n_c} \times \frac{1}{u_c}$  is the score for non-null trait counts,  $N$  is collection size,  $n_{i,j}$  and  $n_c$  are frequencies of trait value  $i$  of type  $j$  and non-null trait count  $c$  respectively,  $u_j$  and  $u_c$  are unique counts of trait values for type  $j$  and possible trait counts,  $M$  is the maximum trait count, and  $c$  is the actual trait count. When  $c < M$ , missing trait positions are scored as null values for trait equalization.

#### Information content-based rarity (IC):

$$S_{IC} = \frac{I(x)}{E[I(x)]}, \quad (2)$$

where  $I(x) = \sum_{j=1}^M I_j(x)$  and  $I_j(x) = -\log_2(\mathbb{P}(\text{trait}_{i,j}))$ . Here  $\mathbb{P}(\text{trait}_{i,j}) = n_{i,j}/N$  represents the probability of specific trait value  $i$  of type  $j$ , and  $M$  is the equilibrated trait count including implicit null values.

The IF assigns higher rarity scores to infrequent traits, while the IC uses Shannon entropy to quantify the "surprisal" of an NFT's trait configuration. Applied across all assets, these

methods produce complementary rarity features to facilitate critical analyses such as price forecasting and trait-based clustering.

#### D. Dataset Structure and Accessibility

To support open and reproducible research, the MultiNFT dataset is released under a CC BY-NC license at:

<https://multinft-dataset.github.io>

The dataset consists of four main components, with a total size of 25.56 GB, and is tailored to support diverse multimodal research applications. In particular:

- **collections.csv** documents collection-level metadata for 50 carefully curated PFP projects, including unique identifiers, API endpoints, social media links, and smart contract addresses, serving as the structural backbone for linking assets across modalities;
- **assets.csv** provides detailed token-level metadata for 523,020 NFTs, encompassing structured trait information and precomputed rarity scores based on both inverse frequency and information-theoretic methods, enabling granular trait-based analysis and ranking;
- **sales.csv** records comprehensive transaction histories, including timestamps, payment tokens, normalized USD prices, and log-transformed values, providing a robust foundation for econometric modeling, market behavior analysis, and predictive tasks;
- **images** directory contains standardized NFT images in  $224 \times 224$  PNG format, organized by collection slug to streamline retrieval, image-based modeling, and cross-modal linkage with textual and transactional data.

Together, these components offer a unified, high-quality resource for studying the economic, visual, and semantic dimensions of NFTs at scale. Detailed attribute descriptions for each component are provided in Table II. All files are relationally structured and can be seamlessly joined using token identifiers and collection slugs, enabling flexible and coherent integration across visual, textual, and transactional modalities. Each record has undergone rigorous filtering and validation procedures to ensure completeness, consistency, and research-grade quality. This structured interoperability supports a broad spectrum of downstream tasks, including multimodal representation learning, trait-based market analysis, predictive modeling, and generative applications.

### IV. CASE STUDIES

This section presents brief case studies showing how MultiNFT supports analyses of pricing dynamics, trading behavior, and rarity assessment by combining visual, textual, and transactional information. They highlight the practical utility of the dataset and illustrate its potential for future research.

#### A. Understanding How Rarity Shapes NFT Pricing Behavior

To examine how rarity contributes to NFT valuation, the distributional characteristics of the IF and IC rarity scores are analyzed across all 50 collections in MultiNFT. Fig. 3

presents violin plots of the normalized score distributions for each collection, with both rarity metrics min–max scaled to the standardized range 0–100 to facilitate cross-collection comparison.

The IF scores (blue) exhibit relatively uniform distributions across collections, with the majority of NFTs concentrated in the lower rarity ranges. This pattern reflects the metric’s focus on individual trait frequency, under which highly uncommon combinations become exponentially scarce and thus occupy only the upper tail of the distribution. By contrast, the IC scores (orange) display a wider variety of distributional shapes, ranging from bimodal to uniform, indicating greater sensitivity to differences in trait schema complexity. Overall, the IF method yields more stable and comparable results across collections, whereas the IC approach captures finer variations in trait structure.

Fig. 4 examines how rarity relates to market value using Pearson and Spearman correlation coefficients. The relationships are weak to moderate, with most collections showing correlations below 0.3. IF-based rarity exhibits more stable and generally higher correlations, with several collections reaching 0.4–0.5, while IC-based correlations vary more widely across collections, suggesting that the information-theoretic measure is more collection-specific and context-dependent.

Across both metrics, Spearman correlations consistently exceed Pearson correlations, implying non-linear associations between rarity and price. However, the overall strength of these relationships remains modest, underscoring the heterogeneous nature of NFT markets. Different communities appear to value rarity through distinct mechanisms, shaped by collection identity, aesthetic preferences, and social signaling. These findings suggest that while rarity contributes meaningfully to NFT pricing behavior, it alone cannot fully explain valuation differences across diverse collections.

#### B. Visual Feature Analysis and Market Valuation

To investigate the visual characteristics of NFT collections, we extract deep visual features using a pre-trained ResNet-50 backbone, yielding 2,048-dimensional embeddings that are reduced to 50 dimensions via principal component analysis (PCA). Fig. 5 shows a t-SNE visualization of this feature space, where each collection is indicated by a distinct color. The clustering patterns reveal that inter-collection differences substantially exceed intra-collection variation, with each collection occupying a distinct region. Collections with similar artistic styles lie close to one another, while visually distinctive projects form clearly separated clusters, indicating that deep features effectively capture the visual signatures that differentiate NFT collections.

In addition to the PCA-processed visual features, we extract interpretable visual features using traditional computer vision techniques, including color properties, complexity measures (image entropy, edge density), structural characteristics (contrast, symmetry), and geometric properties. We also employ the Neural Image Assessment (NIMA) framework [23] to evaluate aesthetic quality using a model pre-trained on the

TABLE II  
OVERVIEW OF MULTINFT DATA STRUCTURE AND KEY VARIABLES.

Table	Attribute/variable	Description
Collections	collection	Collection slug. A unique string to identify a collection on OpenSea
	name	Name of the collection
	description	Description of the collection
	image_url	Link to the square image representing the collection
	banner_image_url	Link to the banner image representing the collection
	owner	The public blockchain address of the owner
	category	Category of the collection (e.g., PFPs, Memberships, Art)
	trait_offers_enabled	If trait offers are currently accepted for the collection
	opensea_url	OpenSea link to collection
	project_url	External URL for the collection's website
	wiki_url	External URL for the collection's wiki
	discord_url	External URL for the collection's Discord server
	telegram_url	External URL for the collection's Telegram group
	twitter_username	Username for the collection's Twitter account
	instagram_username	Username for the collection's Instagram account
	contracts	Nested attribute including contract address and blockchain chain
	editors	List of editor addresses for the collection
fees	Nested attribute including fee percentage, recipient address, and whether the fee is required	
payment_tokens	Nested attribute including symbol, address, chain, image, name, decimals, ETH price, and USD price	
total_supply	Total supply of the collection (minted minus burned)	
created_date	Date the collection was created	
Assets	identifier	NFT's unique identifier within the smart contract (token_id)
	collection	Collection slug (unique identifier on OpenSea)
	contract	The public blockchain address of the contract
	token_standard	ERC standard of the token (erc721, erc1155)
	name	Name of the NFT
	description	Description of the NFT
	image_url	Link to the image associated with the NFT
	display_image_url	Link to the image displayed on OpenSea
	display_animation_url	Link to the animation displayed on OpenSea
	metadata_url	Link to the offchain metadata store
	opensea_url	Link to the NFT on OpenSea
	updated_at	Last time the NFT's metadata was updated by OpenSea
	is_disabled	If the item can be bought or sold on OpenSea
is_nsfw	If the item is classified as Not Safe for Work by OpenSea	
Sales	event_type	Defaults to sale
	order_hash	Order hash for the fulfilled order
	chain	OpenSea supported chains
	protocol_address	Exchange contract address which fulfilled the order
	closing_date	Posix timestamp converted to datetime
	nft_identifier	NFT's unique identifier (token_id)
	nft_collection	Collection slug
	nft_contract	Public blockchain address of the contract
	nft_token_standard	ERC standard (erc721, erc1155)
	nft_name	Name of the NFT
	nft_description	Description of the NFT
	nft_image_url	Link to the image
	nft_metadata_url	Link to offchain metadata
	nft_opensea_url	Link to NFT on OpenSea
	nft_updated_at	Last metadata update time
	quantity	Number of assets transferred
	seller	Seller of the NFT
	buyer	Buyer of the NFT
	payment_quantity	Amount of tokens in the order
	payment_token_address	Contract address for ERC20 token
	payment_decimals	Number of decimals the token uses
	payment_symbol	Symbol of the token (ETH, WETH, etc.)
	transaction	Transaction hash
event_timestamp	Posix timestamp converted to datetime	
payment_usd_price	Payment amount converted to USD	
log_usd_price	Natural logarithm of USD price	

**Note:** This table summarizes the key variables and their descriptions across the three datasets: Collections, Assets, and Sales. For simplicity, some less important variables are not shown here, while the full data description is available on the dataset webpage: <https://multinft-dataset.github.io>.

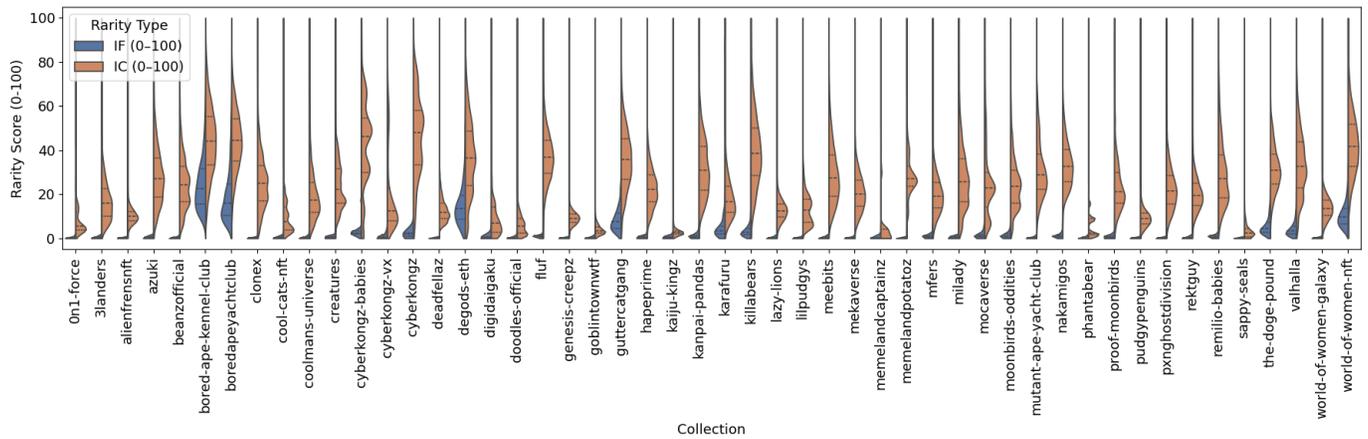


Fig. 3. Distribution of normalized IF and IC rarity scores across 50 PFP collections.

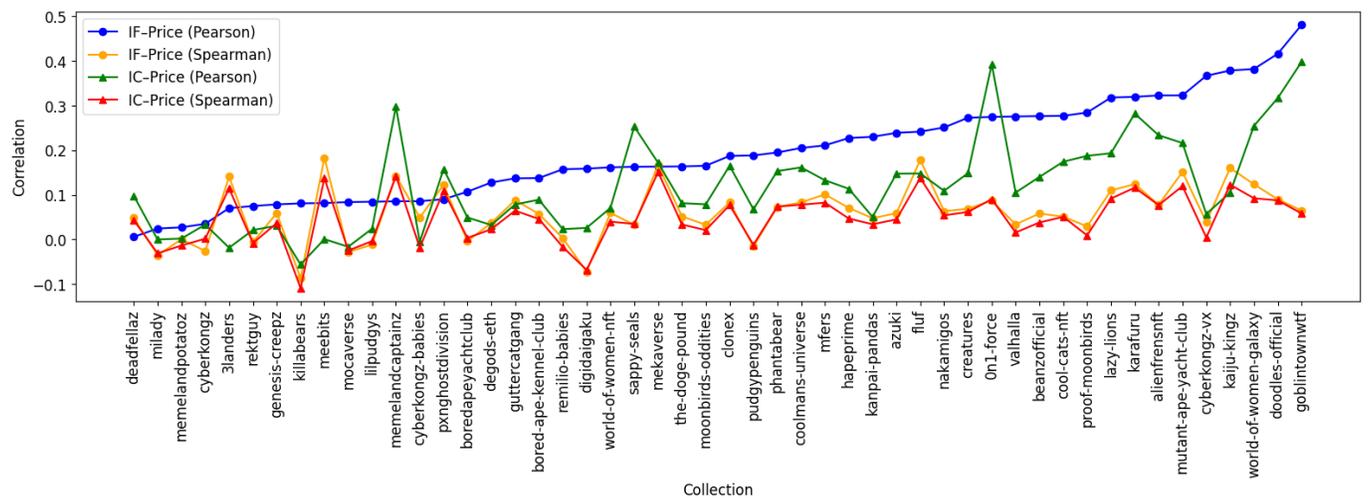


Fig. 4. Pearson and Spearman correlations between rarity scores and sale prices across collections.

AVA dataset [24]. NIMA outputs aesthetic scores on a 1-10 scale representing predicted human aesthetic preferences.

We evaluate different visual feature representations using XGBoost regression models for NFT price prediction. Three configurations are tested: explainable visual features alone ( $R^2 = 0.406$ ,  $RMSE = 1.202$ ), PCA embeddings ( $R^2 = 0.446$ ,  $RMSE = 1.161$ ), and a combined approach ( $R^2 = 0.462$ ,  $RMSE = 1.144$ ). The results demonstrate progressive improvements, with the combined model achieving the best performance. The result of SHAP (Shapley Additive Explanations) [25] analysis, as shown in Fig. 6, reveals that  $pca\_4$  dominates feature importance with significantly higher predictive weight than other variables. Among interpretable features, image entropy and symmetry show the strongest contributions. Traditional computer vision features contribute meaningfully but with lower impact than PCA components. The dominance of learned representations suggests promising potential for end-to-end deep learning approaches that could directly learn price-relevant visual features from raw images.

### C. NFT Price Prediction and Feature Importance Analysis

To demonstrate the full potential of MultiNFT, we develop an integrated XGBoost regression model that combines multiple feature types: normalized rarity scores (IF and IC), visual features (50-dimension PCA components), aesthetic scores (NIMA), market context variables (ETH/USD price, gas fees), user experience metrics (seller/buyer experience), temporal features (sale month), and collection identifiers. This comprehensive approach achieves strong predictive performance with  $R^2 = 0.947$ ,  $RMSE = 0.361$ , and  $MAE = 0.243$ , demonstrating the effectiveness of multimodal feature integration for NFT price prediction.

The SHAP analysis reveals a clear hierarchy of predictive importance across different model configurations. In the full model including collection identifiers, as shown in Fig. 7, average gas price emerges as the most important factor, reflecting the strong influence of network transaction costs on pricing decisions. Collection-specific effects dominate the feature ranking, with collections like Mutant Ape Yacht Club

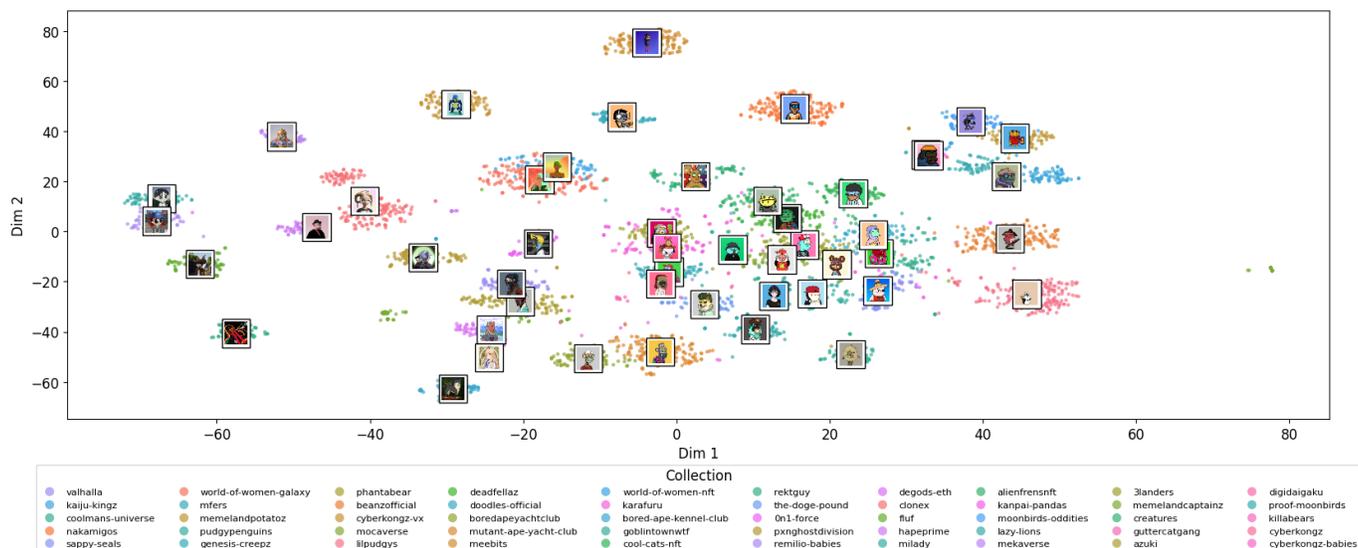


Fig. 5. T-SNE visualization of NFT collections in PCA-reduced visual feature space.

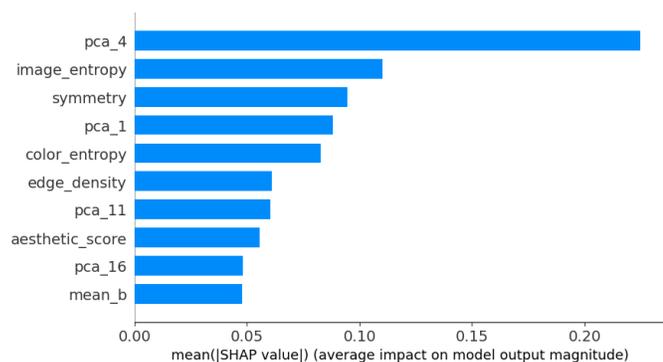


Fig. 6. SHAP feature importance for visual-based price prediction model.

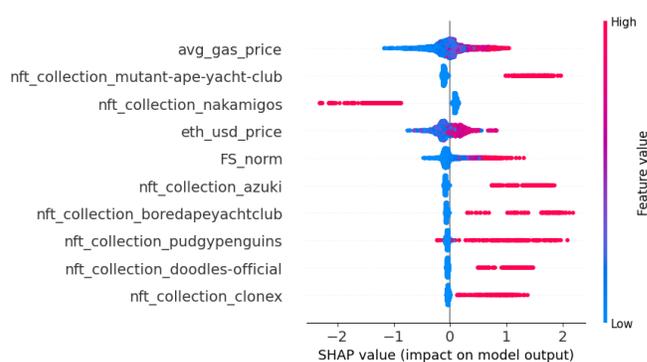


Fig. 7. Full model SHAP value distributions.

and Nakamigos showing substantial positive price premiums. Among content-based features, the frequency-based rarity score (FS\_norm) ranks prominently, validating the importance of scarcity in valuation, while visual features like *pca\_4* contribute meaningfully but with lower individual impact.

To assess the role of collection identity versus intrinsic asset characteristics, a second model is estimated without collection identifiers. This restricted specification achieves  $R^2 = 0.867$ ,  $RMSE = 0.568$ , and  $MAE = 0.388$ . The reduction in explanatory power relative to the full model underscores the substantial influence of collection reputation and brand effects in NFT markets. As shown in Fig. 8, once collection identifiers are removed, visual features gain markedly in importance, with several PCA components (*pca\_4*, *pca\_1*, *pca\_12*) emerging among the top predictors. Rarity metrics (IS\_norm and FS\_norm) and aesthetic scores also move up in the importance ranking, indicating that intrinsic asset characteristics partially substitute for collection identity in the pricing process.

The grouped SHAP analysis in Fig. 9 quantifies this substitution effect across the two model configurations. In the

full model (left), collection effects dominate the prediction, accounting for 53.1% of the total contribution, while temporal and market variables contribute 19.1% and 14.7%, respectively. Visual features and rarity play a comparatively smaller role, explaining 7.4% and 5.7% of the overall importance. When collection identifiers are removed (right), visual features become the primary driver at 55.9%. Temporal and market factors remain relatively stable (18.5% and 18.1%), and the contribution of rarity metrics changes only marginally to 7.5%.

This shift in contributions demonstrates that visual characteristics capture meaningful pricing information that partially overlaps with collection-specific effects, though the overall predictive performance remains superior when collection identity is included through the temporal and market context features. The findings reveal that while intrinsic asset characteristics (visual appeal, rarity, and aesthetics) influence pricing, market timing and transaction context play crucial roles in NFT valuation. The substantial increase in visual feature importance when collections are removed indicates that visual attributes serve as important proxies for collection identities.

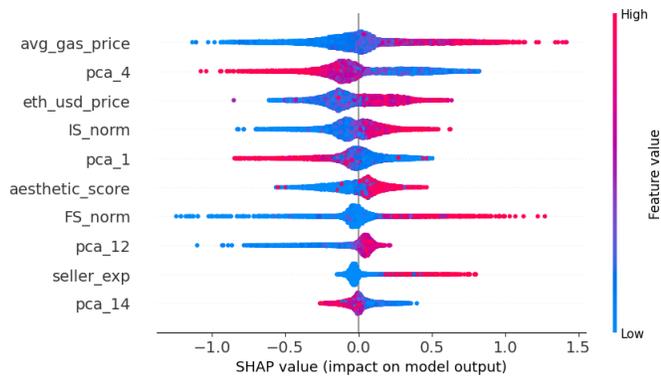


Fig. 8. No-collection model SHAP value distributions.

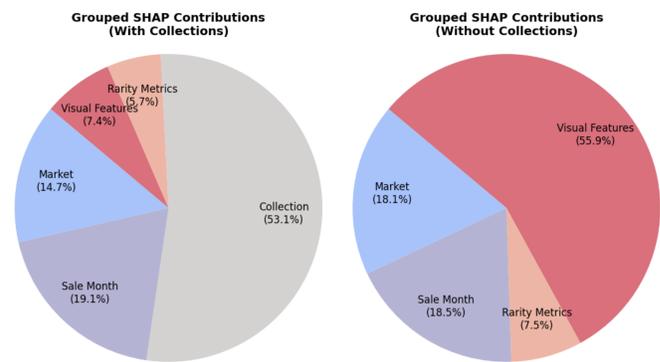


Fig. 9. Grouped SHAP contributions with and without collection identifiers.

## V. CONCLUSION

This study introduces MultiNFT, a comprehensive multi-modal dataset for examining NFT valuation dynamics across diverse collections and market conditions. Through case studies on aesthetic evaluation and predictive modeling, we show the complex interplay between visual traits, market indicators, and pricing, demonstrating that while collection-level factors serve as the strongest determinants, aesthetic and rarity features also meaningfully inform valuation models. MultiNFT provides a unified foundation for digital asset analytics (e.g., price determinants and liquidity trends), market prediction (e.g., temporal patterns and macro-event impacts), and multimodal learning (e.g., combining image, metadata, and transaction records for classification or generative tasks). We invite researchers across machine learning, economics, HCI, and digital media to use MultiNFT as a benchmark for advancing multimodal digital asset research.

## REFERENCES

- [1] K. B. Wilson, A. Karg, and H. Ghaderi, "Prospecting non-fungible tokens in the digital economy: Stakeholders and ecosystem, risk and opportunity," *Business Horizons*, vol. 65, no. 5, pp. 657–670, 2022.
- [2] E. Sung, O. Kwon, and K. Sohn, "NFT luxury brand marketing in the metaverse: Leveraging blockchain-certified nfts to drive consumer behavior," *Psychology & Marketing*, vol. 40, no. 11, pp. 2306–2325, 2023.
- [3] R. Belk, M. Humayun, and M. Brouard, "Money, possessions, and ownership in the metaverse: NFTs, cryptocurrencies, Web3 and wild markets," *Journal of Business Research*, vol. 153, pp. 198–205, 2022.
- [4] S. Alsultan, A. Kourtis, and R. N. Markellos, "Can we price beauty? aesthetics and digital art markets," *Economics Letters*, vol. 235, p. 111572, 2024.
- [5] D. Costa, L. La Cava, and A. Tagarelli, "Show me your NFT and I tell you how it will perform: Multimodal representation learning for NFT selling price prediction," in *Proceedings of the ACM Web Conference 2023*. ACM, 2023, pp. 1875–1885.
- [6] T. Wang, L. Deng, C. Wang, J. Lian, Y. Yan, N. J. Yuan, Q. Zhang, and H. Xiong, "COMET: NFT price prediction with wallet profiling," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2024, pp. 5893–5904.
- [7] Y. Niu, X. Li, H. Peng, and W. Li, "Unveiling wash trading in popular NFT markets," in *Companion Proceedings of the ACM Web Conference 2024*, 2024, pp. 730–733.
- [8] G. Fridgen, R. Kräussl, O. Papageorgiou, and A. Tugnetti, "Pricing dynamics and herding behaviour of NFTs," *European Financial Management*, vol. 31, no. 2, pp. 670–710, 2025.
- [9] A. Mekacher, A. Bracci, M. Nadini, M. Martino, L. Alessandretti, L. M. Aiello, and A. Baronchelli, "Heterogeneous rarity patterns drive price dynamics in NFT collections," *Scientific Reports*, vol. 12, no. 1, p. 13890, 2022.
- [10] K. Brahmstaedt, "Community and consumer dynamics in NFTs: Understanding digital asset value through social engagement," *Journal of Consumer Behaviour*, 2025.
- [11] M. Nadini, L. Alessandretti, F. Di Giacinto, M. Martino, L. M. Aiello, and A. Baronchelli, "Mapping the NFT revolution: Market trends, trade networks, and visual features," *Scientific Reports*, vol. 11, no. 1, p. 20902, 2021.
- [12] Y. Liu, B. Huangfu, Z. Qiao, and X. Zhao, "Rarity, the WTA-WTP disparity, and price adjustments in the NFT market," *International Journal of Research in Marketing*, 2024.
- [13] Y. Kim, S. Kim, Y. Lee, J. Hong, and Y. Lee, "Multi-attention recommender system for non-fungible tokens," *Engineering Applications of Artificial Intelligence*, vol. 137, p. 109179, 2024.
- [14] S. Wang, Y. Lei, Z. Zhang, W. Liu, H. Liu, L. Yang, B. Li, W. Li, J. Gao, and W. Hu, "NFT1000: A cross-modal dataset for non-fungible token retrieval," in *Proceedings of the 32nd ACM International Conference on Multimedia*. ACM, 2024, pp. 2214–2222.
- [15] F. Horvay, C. Rachel, and J. Fidrmuc, "Price determinants of non-fungible tokens in the digital art market," *Finance Research Letters*, vol. 48, p. 103007, 2022.
- [16] M. Dowling, "Is non-fungible token pricing driven by cryptocurrencies?" *Finance Research Letters*, vol. 44, p. 102097, 2022.
- [17] E. Iosif and L. Katelaris, "Integrating machine learning with non-fungible tokens," *Computers*, vol. 13, no. 6, p. 147, 2024.
- [18] B. Seyhan and E. Sefer, "NFT primary sale price and secondary sale prediction via deep learning," in *Proceedings of the Fourth ACM International Conference on AI in Finance*. ACM, 2023, pp. 116–123.
- [19] M. Pala and E. Sefer, "NFT price and sales characteristics prediction by transfer learning of visual attributes," *The Journal of Finance and Data Science*, vol. 10, p. 100148, 2024.
- [20] D. Costa, L. La Cava, and A. Tagarelli, "Unraveling the NFT economy: A comprehensive collection of non-fungible token transactions and metadata," *Data in Brief*, vol. 51, p. 109749, 2023.
- [21] PopRank, "Poprank: NFT trait-based rarity rankings," <https://github.com/poprank/rankings>, 2023, accessed: 2025-10-10.
- [22] OpenRarity, "Openrarity methodology documentation," 2023, accessed: 2025-10-10. [Online]. Available: <https://openrarity.gitbook.io/developers/fundamentals/methodology>
- [23] H. Talebi and P. Milanfar, "NIMA: Neural image assessment," *IEEE Transactions on Image Processing*, vol. 27, no. 8, pp. 3998–4011, 2018.
- [24] N. Murray, L. Marchesotti, and F. Perronnin, "AVA: A large-scale database for aesthetic visual analysis," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2012, pp. 2408–2415.
- [25] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, "From local explanations to global understanding with explainable AI for trees," *Nature Machine Intelligence*, vol. 2, no. 1, pp. 56–67, 2020.