# How Convolutional Neural Networks Can Help Predict NFT Prices by Analyzing Images

By Shunsuke Nakatani

## Abstract

This research explores the use of Convolutional Neural Networks (CNNs) to predict the market prices of Non-Fungible Tokens (NFTs) by analyzing their visual features. While traditional price prediction models rely heavily on historical price data and financial indicators, this study focuses on the unique artistic and visual aspects of NFTs, which often play a crucial role in determining their value. Using a dataset from the Moonbirds NFT collection, a CNN model was developed to extract image-based features such as colors, textures, and shapes, and predict corresponding market prices. The results reveal that while the CNN model could identify significant visual patterns, it often underestimated prices and produced a wide range of errors. This indicates the need for further refinement of the model and the incorporation of additional factors, such as metadata and market trends, to improve accuracy. Despite its limitations, this research provides a novel approach to NFT price prediction and offers insights for future work in blending visual data with traditional financial models to enhance market forecasting tools for investors.

## 1. Introduction

Non-fungible tokens (NFTs) represent a revolutionary way to prove ownership of digital assets (Sharma 2023). Unlike cryptocurrencies (The Investopedia Team 2024), which are fungible and identical to one another, NFTs are unique and indivisible, making them akin to digital collectibles or trading cards. Each NFT can represent anything from digital art to virtual real estate, and their value is often driven by their rarity, demand, and the perceived utility or access they provide, such as exclusive memberships or privileges. The NFT market has seen explosive growth, with some projects fetching extraordinary prices due to their unique characteristics and the benefits they offer to owners.

However, predicting the future value of NFTs remains a significant challenge. The rarity of an NFT and the specific attributes of its associated image play crucial roles in determining its market value. The sheer diversity and volume of NFT images make it nearly impossible for human evaluators to accurately assess the potential market price of each individual token. This is where machine learning, and more specifically, Convolutional Neural Networks (CNNs), comes into play (IBM 2024).

CNNs, a class of deep learning models particularly adept at image analysis, offer a promising solution to this problem. By analyzing the visual features of NFT images, CNNs can identify patterns and characteristics that correlate with higher market values. While previous research has explored the prediction of NFT prices based on historical price data, these approaches often overlook the significant influence of the visual appeal and artistic quality of the NFTs themselves (Ghosh et al. 2023).

This research aims to bridge this gap by developing models that predict NFT prices through image analysis using CNNs. By focusing on the visual elements of NFTs, we hope to provide a more comprehensive and accurate prediction model. Ultimately, our goal is to reduce the number of investment failures in the NFT market by equipping investors with better tools to evaluate the potential value of NFTs based on their images.

## 2. Background

Research on forecasting NFT prices is still in its early stages. One significant approach is presented by Ghosh, I., Alfaro-Cortés, E., Gámez, M., and García-Rubio, N (Ghosh et al. 2023), which predicts NFT and DeFi prices using machine learning models. This study combines technical indicators, macroeconomic data, and sentiment analysis to create accurate price predictions, particularly during periods of financial instability. The models used, such as Gradient Boosting and Random Forest, have demonstrated strong performance in capturing trends based on past price movements and external factors. While this approach effectively incorporates financial and economic indicators, it does not consider the unique artistic and visual aspects of NFTs, which often play a significant role in their valuation.

Additionally, Baals, L. J., Lucey, B. M., Vigne, S. A., and Long's research offers a systematic projection of future research directions for the NFT market (Baals et al. 2022). This work emphasizes the need for further investigation into unexplored areas, highlighting the importance of integrating more diverse data sources and methodologies to fully understand the NFT market dynamics. While insightful, this research primarily focuses on market trends and does not delve into the potential of image-based analysis for NFTs.

Despite these advances, one key limitation remains: neither of these approaches fully addresses the influence of the art itself, which is often a major factor in determining an NFT's price. Unlike traditional financial assets, NFTs are largely valued based on the perceived quality and uniqueness of the artwork, and predicting their price without accounting for visual elements can lead to significant gaps in accuracy.

This gap in existing research motivates the need to explore how the art aspect of NFTs can be incorporated into price-prediction models. Specifically, convolutional neural networks (CNNs) have shown potential in analyzing visual data and could play a crucial role in forecasting NFT prices by interpreting the artistic features that drive value. By combining CNNs with traditional financial models, my research aims to provide a more comprehensive and accurate approach to NFT price prediction, addressing the limitations of prior studies.

## 3. Dataset

For this research, the dataset used includes both numerical and vision data. Specifically, the dataset provides the price at which NFTs were sold, along with links to the corresponding

NFT images. Initially, this dataset can be classified as a language dataset, as it includes text-based information such as the image URLs. However, after downloading the images from these links, the dataset also becomes a vision dataset.

The dataset consists of NFT data from five different projects: Bored Ape Yacht Club (Amure 2024), Mutant Ape Yacht Club (Daly 2024), The Otherdeeds ("Otherdeed NFT Traits | Otherside Wiki"), Azuki (Daly 2024), and Moonbirds (Georgiev 2022). Given that each project varies significantly in terms of price range, utility, and popularity, the research focuses on one project as a proof of concept—Moonbirds. This decision helps isolate the impact of image-based features on NFT pricing without introducing variability from differing utilities across projects.

The dataset contained over 10 columns of information, such as block number, transaction hash, and marketplace, which were unnecessary for predicting NFT prices. Therefore, the dataset was simplified by retaining only the essential features: usd_price (the price of the NFT in USD), name (the name of the NFT), image_url (the link to the NFT's image), and collection_name (the name of the NFT project). These features were relevant to understanding the relationship between visual aspects of the NFTs and their market price.

In total, the dataset consists of 2,000 samples from the Moonbirds collection. The data was split into two sets: 80% (1,600 samples) for training and 20% (400 samples) for evaluation. This split allowed us to train our model on a large portion of the data while keeping a separate set for evaluating its performance.

## 4. Methodology

### 4.1. Modeling

A Convolutional Neural Network (CNN) was built to predict NFT prices by analyzing image features. The architecture consists of two convolutional layers with 32 and 64 filters, respectively, followed by max-pooling layers to reduce the spatial dimensions. After extracting key image features such as colors, textures, and shapes, the output was flattened into a one-dimensional vector. This was then passed through fully connected layers, leading to a final prediction for the NFT price.

For training, 80% of the 2,000 NFT images were used. The model trained for 10 epochs, processing mini-batches of data to refine its predictions. Once trained, the model was saved for future use in predicting NFT prices based on image data.

### 4.2. Evaluation

In the evaluation stage, the model's performance was tested using 20% of the dataset, consisting of 400 NFT images. These images were not seen by the model during training, making them ideal for assessing how well the model generalizes to new data. The model's predictions were compared against the actual NFT prices, providing a measure of its accuracy.

During evaluation, the model was set to inference mode (model.eval()) to prevent further training. Predictions were made without adjusting the model’s parameters. The predicted prices were then collected and compared with the true prices from the dataset. This comparison allowed for an assessment of how close the model's predictions were to the actual NFT prices, giving insight into the model’s overall accuracy and performance.

## 5. Results

The performance of the Convolutional Neural Network (CNN) model in predicting NFT prices was evaluated by comparing its predicted values with the actual prices. Below are the key visual analyses from the evaluation phase.

### 5.1. Histogram of Differences Between Predictions and Actual Prices

The histogram (Fig. 1) plots the frequency of differences between predicted and actual NFT prices. The x-axis represents the differences between the model’s predictions and the true values, while the y-axis shows how often these differences occurred. From the histogram, it’s evident that there are substantial discrepancies between the predicted and actual prices, with many predictions significantly lower than the actual values. This pattern indicates that the model frequently underestimates NFT prices, which suggests a potential bias in the model's predictions that needs to be addressed. Additionally, the broad distribution of prediction errors highlights the model’s inconsistency in predicting NFT prices accurately.

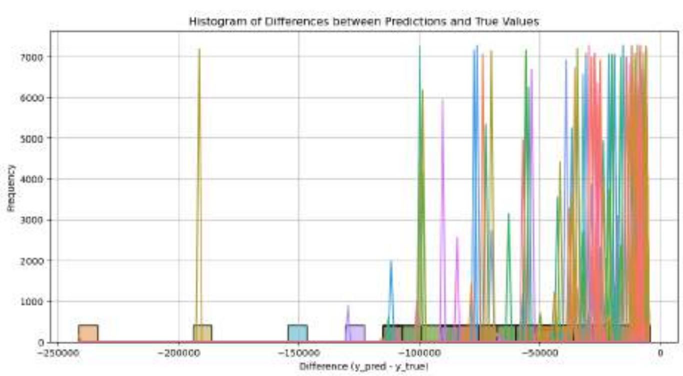


Figure1: Histogram Plot of Differences Between Predictions and Actual Prices

### 5.2. Box Plot of Prediction Differences

A more detailed analysis is provided by the box plot (Fig. 2), which summarizes the distribution of prediction errors. The box represents the interquartile range (middle 50%) of the errors, while the orange line inside the box denotes the median error. The whiskers extend to show the overall range of errors, with several data points lying far from the whiskers, representing extreme outliers. The presence of these large prediction errors and outliers further underscores the model's lack of precision, as well as a significant variance in its predictions. The clustering of errors below the median line supports the earlier observation that the model tends to underestimate NFT prices. These results suggest that while the model captures some patterns in the data, its predictions are not sufficiently reliable for practical use without further refinement.

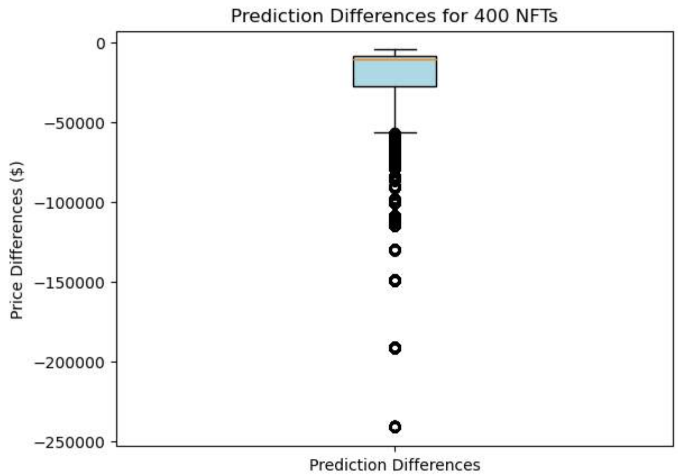


Figure 2: Box Plot of Differences Between Predictions and Actual Prices

### 5.3. Q-Q Plot of Prediction Differences

Lastly, the Q-Q plot (Fig. 3) compares the distribution of prediction errors to a normal distribution. Ideally, the points should align along the red diagonal line if the errors were normally distributed. However, in our case, many points deviate significantly from this line, particularly at the tails, indicating that the prediction errors are not normally distributed. This further demonstrates that the model has difficulty generalizing across the entire price range, leading to high variance in its predictions, especially for extreme cases.

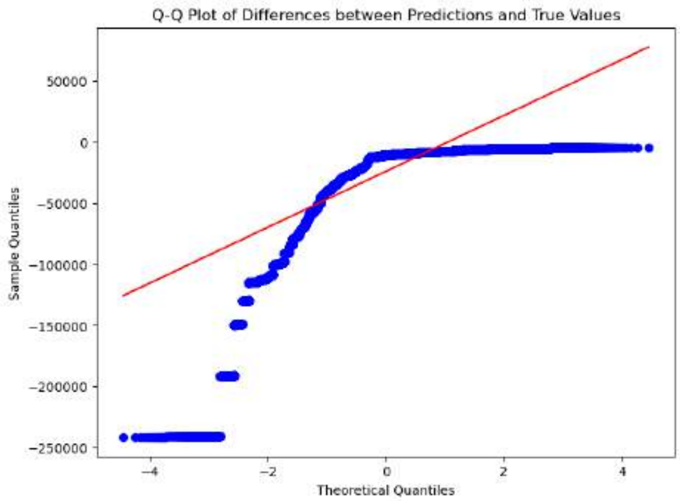


Figure 3: Q-Q Plot of Differences Between Predictions and Actual Prices

**6. Discussion**

The evaluation results show that while the CNN model can extract meaningful features from NFT images, it struggles with accurately predicting their market prices. The model frequently underestimates prices, as demonstrated by both the histogram and box plot analyses, and exhibits large prediction errors with several outliers. The Q-Q plot confirms that the model's prediction errors deviate from a normal distribution, particularly at the extremes.

These issues could stem from various factors, including limitations in the model architecture, insufficient training data, or complexity in the NFT market that the current model does not account for. To improve model performance, further optimization steps could include:

- Data Augmentation: Expanding the dataset or augmenting the existing data with variations of the images could help the model learn more robust features.
- Advanced Architectures: Implementing deeper architectures like ResNet or experimenting with transfer learning might improve the model's ability to capture more complex patterns in the data.
- Fine-Tuning Hyperparameters: Adjusting hyperparameters such as learning rate, batch size, and the number of epochs could lead to better convergence and accuracy.
- Price-Related Features: Incorporating non-image features, such as metadata about NFTs (e.g., artist reputation, rarity) could provide additional context for the price prediction.

Ultimately, while the CNN model shows some promise in predicting NFT prices based on image data, it requires significant improvements to achieve consistent and reliable results.

## 7. Conclusion

In conclusion, this research highlights the potential of using Convolutional Neural Networks (CNNs) to predict the prices of NFTs by analyzing the visual characteristics of their images. The model successfully demonstrated that image-based features such as colors, textures, and shapes play a significant role in price determination, offering an innovative approach to understanding the NFT market. However, the results also exposed several challenges. The model tended to underestimate prices and produced a wide range of prediction errors, including notable outliers. This inconsistency suggests that while visual features are important, they alone may not provide a complete picture of an NFT's market value.

To enhance the model's accuracy, future research should consider integrating additional factors, such as NFT metadata (e.g., rarity, artist reputation, or utility) and exploring more advanced architectures like ResNet or transfer learning to capture complex patterns more effectively. Data augmentation could also help improve generalization by allowing the model to learn from a broader set of image variations. Lastly, fine-tuning hyperparameters such as learning rates and batch sizes could lead to better model performance. Despite the challenges, this research provides valuable insights into the role of image analysis in predicting NFT prices and opens avenues for further exploration in combining visual and non-visual features for more accurate market forecasting.

## 8. Acknowledgement

I would like to express my sincere gratitude to Joe Xiao for his invaluable guidance and support throughout my research journey. His extensive expertise in AI and machine learning, combined with his passion for mentoring students, has greatly enriched my understanding of the field. I am particularly thankful for his encouragement and insight, which have inspired me to pursue independent research and strive for excellence in my work.

## Works Cited

Sharma, Rakesh. "Non-Fungible Token Definition: Understanding NFTs." *Investopedia*, 6 Apr. 2023, www.investopedia.com/non-fungible-tokens-nft-5115211. Accessed 12 Aug. 2024.

The Investopedia Team. "Cryptocurrency Explained with Pros and Cons for Investment." Investopedia, 15 June 2024, www.investopedia.com/terms/c/cryptocurrency.asp. Accessed 12 Aug. 2024.

Ghosh, Indranil, et al. "Prediction and Interpretation of Daily NFT and DeFi Prices Dynamics: Inspection through Ensemble Machine Learning & XAI." *International Review of Financial Analysis*, Feb. 2023, p. 102558, https://doi.org/10.1016/j.irfa.2023.102558. Accessed 12 Aug. 2024.

IBM. "What Are Convolutional Neural Networks? | IBM." Www.ibm.com, IBM, 2024, www.ibm.com/topics/convolutional-neural-networks. Accessed 13 Aug. 2024.

Baals, Lennart John, et al. "Towards a Research Agenda on the Financial Economics of NFT's." SSRN Electronic Journal, Mar. 2022, https://doi.org/10.2139/ssrn.4070710.

Amure, Tobi. "Bored Ape Yacht Club." Investopedia, 25 July 2024, www.investopedia.com/bored-ape-yacht-club-6747275. Accessed 13 Aug. 2024.

Daly, Lyle. "What Is Mutant Ape Yacht Club? | the Motley Fool." The Motley Fool, 26 Mar. 2024, www.fool.com/terms/m/mutant-ape-yacht-club. Accessed 13 Aug. 2024.

"Otherdeed NFT Traits | Otherside Wiki." Otherside-Wiki.xyz, www.otherside-wiki.xyz/otherdeed. Accessed 14 Aug. 2024.

Daly, Lyle. "What Is Azuki NFT? | the Motley Fool." The Motley Fool, 20 Mar. 2024, www.fool.com/terms/a/azuki/. Accessed 14 Aug. 2024.

Georgiev, George. "What Are Moonbirds? Guide to the NFT Collection." CryptoPotato, 22 Apr. 2022, cryptopotato.com/what-is-moonbirds-nft. Accessed 14 Aug. 20