The Journal of Finance and Data Science 10 (2024) 100148

Contents lists available at ScienceDirect

The Journal of Finance and Data Science

journal homepage: www.keaipublishing.com/en/journals/jfds



Research Article

NFT price and sales characteristics prediction by transfer learning of visual attributes

Mustafa Pala, Emre Sefer *

Artificial Intelligence and Data Engineering Department, Ozyegin University, Nisantepe, Orman Sk. No:13, 34794, Istanbul, Turkey

ARTICLE INFO

Keywords:
NFTs
Blockchain
Deep learning
Transfer learning
Temporal price prediction

ABSTRACT

Non-fungible tokens (NFTs) are unique digital assets whose possession is defined over a blockchain. NFTs can represent multiple distinct objects such as art, images, videos, etc. There was a recent surge of interest in trading them which makes them another type of alternative investment. The inherent volatility of NFT prices, attributed to factors such as over-speculation, liquidity constraints, rarity, and market volatility, presents challenges for accurate price predictions. For such analysis and forecasting, machine learning methods offer a robust solution framework.

Here, we focus on three related prediction problems over NFTs: Predicting NFTs sale price, inferring whether a given NFT will participate in a secondary sale, and predicting NFT's sale price change over time. We analyze and learn the visual characteristics of NFTs by deep pre-trained models and combine such visual knowledge with additional important non-visual attributes such as the sale history, seller's and buyer's centralities in the trading network, and collection's resale probability. We categorize input NFTs into six categories based on their characteristics. Across detailed experiments, we found visual attributes obtained from deep pre-trained models to increase the prediction performance in all cases, and EfficientNet seems to perform the best. In general, CNN and XGBoost consistently outperformed the rest of them across all categories. We also publish our novel NFT dataset with temporal price knowledge, which is the first dataset to have NFT prices over time rather than at a single time point. Our code and NFT datasets are publicly available at https://github.com/seferlab/deep_nft.

1. Introduction

A Non-Fungible Token (NFT) is a data file, which is kept on blockchain. In this case, blockchain behaves as a kind of digital ledger and NFTs can be traded. NFTs are distinctive digital identifiers that are utilized for ownership and genuineness verification. NFTs could not be divided into multiple parts, replaced, or duplicated. Blockchain records NFT's right of possession, but such right of possession could be moved to another person via the permission of the owner by selling and trading. NFTs can be used to represent various types of physical or digital assets, such as art, memes, GIFs, photos, musics, images, video recordings, and anything else that can be visualized as a representation of objects (Yilmaz et al., 2023). As an asset, using NFT may depend on obtaining the proper license (Okonkwo, 2021). Nevertheless, NFT trading is still extralegal to a certain extent. So, buying and selling NFTs may end up in unofficial possession modification for the underlying asset without any legal imposition (Taherdoost, 2023).

* Corresponding author.
E-mail address: emre.sefer@ozyegin.edu.tr (E. Sefer).
Peer review under the responsibility of KeAi Communications Co., Ltd.

https://doi.org/10.1016/j.jfds.2024.100148
Received 19 December 2023; Received in revised form 28 August 2024; Accepted 11 November 2024
Available online 16 November 2024
2405-9188/© 2024 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

NFTs are uniquely identifiable by the address of the smart contract that governs their collection and a distinct identifier. Each token is linked to specific metadata, which provides details about the asset it represents. This metadata is usually formatted in JavaScript Object Notation (JSON) and typically includes a name, description, a link to the media file's storage location, and various traits that describe the asset. The media file itself can be anything from an image or video to a website, piece of code, or any type of data (White et al., 2022a). Metadata play a crucial role in NFTs, as they are used by marketplaces to showcase each NFT's characteristics and determine the token's rarity (Wang et al., 2021). The initial generation of NFTs, often referred to as static NFTs or NFT 1.0, featured metadata that was fixed at the time of creation and could not be altered afterward. A typical example is the Bored Ape NFT, where the metadata, including the image and its traits stored as a JSON object, are immutable (for instance, the metadata for a Bored Ape NFT can be found here: https://ipfs.io/ipfs/QmeSjSinHpPnmXmspMjwiXyN6zS4E9zccariGR3jxcaWtq/0, accessed on 23 May 2023). While static metadata work well for art and collectible NFTs, a more dynamic approach is necessary for building a versatile and expansive NFT ecosystem that extends beyond these applications. For instance, in gaming, a player's character level could change based on their in-game performance, necessitating updates to the NFT's metadata. Moving beyond static images and basic collectibles is essential to fully enhance user interactions with NFTs. The potential applications of NFTs are vast, ranging from healthcare and digital passports to real estate and art authentication. To unlock this potential, developers must rethink the role of NFTs, focusing on practical, real-world uses that address user needs, ensuring that people can fully benefit from this evolving technology.

In terms of functionality, NFTs are similar to cryptographic tokens which can be defined as access rights or assets that can be programmed (Cachin and Chandran, 2009). Cryptographic tokens are controlled by smart contracts and associated decentralized ledgers. Cryptographic tokens can only be accessed by the owner having the private key for the corresponding address as well as could be signed solely via that private key (White et al., 2022b). Different than cryptographic tokens and cryptocurrencies, people generally cannot interchange Non-Fungible Tokens reciprocally where the non-fungible name comes from (Wang et al., 2023a). An NFT includes links associated with the data storage. For instance, such links may give detailed information on the associated art's storage. However, such data links could be easily impacted by the deterioration of hyperlinks over time (Ito et al., 2022).

Although NFTs have started to be produced in 2014, they left their mark in 2021. The public's interest in NFTs skyrocketed in 2021, a year in which their market saw record sales, but little is known about the industry's general makeup and course of development. NFT trading has jumped to $17 billion in 2021, which was just $82 million in 2020 (Park et al., 2023). In general, NFTs have been seen and traded as speculative asset classes. Some investors still consider NFT trading dynamics similar to a Ponzi scheme (Kshetri, 2022). Additionally, NFTs have been criticized for their higher energy usage, increased carbon footprint over several types of blockchain, and increased appearance in art frauds (Huang et al., 2023). The first appearance of NFTs has been on the Ethereum blockchain. Following such an introduction, multiple blockchains have also integrated them. In such blockchains, smart contracts are used to define NFTs where their owner history and minting could be tracked at all. As of now, Ethereum, Solana, and Cardano blockchains compose the largest NFT platforms (Yakovenko, 2018). Starting from 2022, the trading volume of NFTs has dropped almost 90 % relative to its peak in 2021.

Exchanging digital tokens that are connected to a digital file asset has been NFTs one main usage. NFT's possession has frequently been connected to a license to utilize such an associated digital asset. However, that does not always transfer copyright to the buyer. For instance, depending on the agreement type, the license may be granted personal or non-commercial usage, as well as commercial usage of the associated digital asset (Madine et al., 2023). Such decentralized intellectual property rights are an alternative to centralized copyright control which is handled by formal government institutes, brokers, and intermediaries (Hari et al., 2023).

As an example, one main usage of Non-Fungible Tokens is digital artworks. Such digital art NFTs distinguished public auctions have been a major milestone in the NFTs' life cycle. For instance, the artwork named Merge by artist Pak has been sold at $91.8 million, and artist Mike Winkelmann's Everydays: the First 5000 Day artwork has been sold at $69.3 million in 2021 (Colavizza, 2023). On the other hand, certain types of digital art NFTs which could be seen as generative art examples have been popular. Those NFT collections include pixel art characters, where Bored Apes, CryptoKitties, CryptoPunks (Wang et al., 2023b), and EtherRocks collections could be considered generative art examples. In these collections, collecting and combining a choice of simple picture components across multiple distinct combinations ends up in multiple distinct image generation (Du et al., 2022). Among these collections, one of the most popular ones is CryptoKitties which is composed of unique virtual cat images and they can be traded over Ethereum blockchain. On the other hand, metaverse NFTs have become quite popular in their social ecosystem (Guidi and Michienzi, 2023).

Even though public attention has increased towards NFTs, the number of academic studies on NFTs is quite limited. A subset of these studies tries to understand NFTs from a technological perspective, such as unique and efficient protocols to track them, blockchain protocols, standardized protocols, etc (Wang et al., 2021; Westerkamp et al., 2018). Another study explores the relationship between rarity and price of profile picture NFT over BAYC NFT collection (Lee et al., 2024). Another class of studies that analyze NFT market dynamics is limited to some well-known collections such as CryptoKitties (Sako et al., 2021; Serada et al., 2021). Moreover, those studies may also focus on more isolated NFT markets such as Decentraland (Dowling, 2022a, 2022b). Lastly, more recent studies have started to interpret and analyze NFTs as tradable assets (Dowling, 2022b; Kapoor et al., 2022; Mekacher et al., 2022). Once analyzed as tradable assets, predicting NFTs price over time by deep learning approaches has been one of the most interesting and important questions researchers focused on. However, these limited number of existing techniques have several limitations: 1- Most of them do not benefit from deeper techniques even though the NFT dataset is large. Instead, they make predictions only by using more traditional Machine Learning approaches such as regression, and logistic regression (Kapoor et al., 2022; Nadini et al., 2021). 2- NFT's image content and visual attributes carry major knowledge about its possible price, and we can extract such embedded knowledge over pre-trained large vision models such as deep convolutional VGG-19 (Simonyan et al., 2015), InceptionNet (Szegedy et al., 2015), EfficientNet (Tan and Le, 2020), and RestNet50 (He et al., 2016) via transfer learning. The existing techniques that do not use these pre-trained models do not fully utilize image attributes of NFTs which are publicly available.

Here, we tackle the problem of predicting NFT primary sale price, sale price direction change over time, and its participation in a secondary sale over time by various deep learning techniques. The highly volatile nature of NFT markets as well as the lower trading volume makes price forecasting a challenging work. Numerous factors contribute to price fluctuations including speculation and manipulation, lack of liquidity, rarity, subjectivity, etc. Additionally, the image content and associated visual attributes also indirectly represent the knowledge about rarity. To accurately predict the price of NFTs, good predictors are required. In our work, in addition to the publicly available NFT dataset of images, we also crawl and publish our novel NFT dataset with temporal price knowledge, which is, to our best knowledge, the first dataset to have NFT prices over time rather than at a single time point. We use the dataset to predict NFT price and sale dynamics over time. We combine such an image dataset with NFTs temporal price change dataset.

More formally, we apply eXtreme Gradient Boosting (XGBoost) (Chen and Guestrin, 2016), AdaBoost (Schapire, 2013), CNN, and LSTM to the following three problems where the first one is a regression problem whereas the remaining two are classification problems: 1- Primary sale price prediction, 2- Primary sale price direction change, and 3- Secondary sale participation which corresponds to the binary classification of NFTs are based on whether they have more than one sale (1) or only one sale (0). The data is categorized into Art, Collectibles, Games, Metaverse, Utility, and Others. We found both XGBoost and CNN to be more effective in predicting NFT price dynamics, especially when we train an independent model for each NFT category. Moreover, we found temporal deep learning methods such as LSTM to be more effective in predicting NFT price dynamics, especially when we train an independent model for each NFT category. Once hyperparameters are tuned accordingly, data imbalance is addressed via Random Oversampling and SMOTE (Chawla et al., 2002), and the outlier data caused by extremely high and low sales is filtered, we obtain the best performance for all considered problems. Mostly the rarest NFTs' sales and mistakenly listed NFTs at a wrong price are the reasons for these extreme values. In this case, outliers are found with confidence intervals and interquartile ranges before filtering them. When integrating visual attributes of NFTs by pre-trained models, EfficientNet performs better than ResNet 50 for many categories in terms of price prediction. In summary, our research offers a comprehensive comparison of different machine and deep learning algorithms in predicting NFT prices. In summary, our research does not only enhance the previous static NFT price prediction by offering a comprehensive comparison of different machine and deep learning algorithms. But, to our best knowledge, it also predicts time-changing prices for NFTs.

2. Related work

2.1. Machine learning and deep learning techniques for cryptocurrency predictions

In (McNally et al., 2018), the authors predict Bitcoin USD price by deep learning. Implementing a Bayesian-optimized Recurrent Neural Network (RNN) and a Long Short Term Memory (LSTM) network allows for variable degrees of success in accomplishing the challenge. With 52 % classification accuracy and an RMSE of 8 %, LSTM performs best according to their results. In contrast to these deep learning algorithms, ARIMA (Hamilton, 1994) was also used to forecast future Bitcoin prices from historical prices. When compared to nonlinear deep learning techniques, ARIMA has performed worse. Despite employing Bayesian optimization to enhance the dropout selection, LSTM is better at spotting longer-term dependencies than RNNs. LSTM slightly outperformed RNN, while requiring a lot more training time. In (Chevallier et al., 2021), the authors compared several machine learning methods such as AdaBoost, and Random Forest to forecast Bitcoin price under market growth. In general, they found AdaBoost and Random Forest to outperform the rest of the techniques. They found that the prediction outcomes vary greatly depending on the prediction time frame and training input data.

2.2. NFTs and their relationship with Bitcoin and Ethereum

The authors in (Ante, 2022) examined the link between NFT sales, NFT users, Ethereum price, and Bitcoin price. Over daily data, they found that a spike in Bitcoin price causes a rise in NFT sales and that a spike in Ethereum price causes a drop in the number of active NFT wallets. The bigger cryptocurrency markets influence the relatively smaller NFT markets expansion and growth, but there is no apparent opposite relationship.

2.3. NFT dynamics and price prediction

Innovative and novel NFTs main characteristics and underlying protocol standards have been analyzed in (Wang et al., 2021). Up to now, the most seminal NFT work has been carried out in (Nadini et al., 2021) where a large collection of NFT images and associated attributes have been made publicly available in addition to their large-scale NFT analysis. In (Nadini et al., 2021), the authors first identified the same collection of NFTs to be visually homogeneous. Moreover, even though NFT data is limited to becoming static over time, they show that NFT price predictability could be enhanced by integrating NFT visual attributes compared to the case of only using the transaction history in forecasting. On the other hand, from a graph analysis perspective (Vasan et al., 2022), has come up with a graph analysis of a limited subset of NFTs.

In addition to the research on analyzing NFTs as mentioned above, another line of research has analyzed NFTs financial dimension by treating them as financial assets. For instance, several previous studies (Dowling, 2022b) could not identify a statistically significant correlation between NFTs prices and large market cap cryptocurrency prices, although the financial performance of cryptocurrencies can be insightful in NFT pricing. Additional work in (Mekacher et al., 2022) has analyzed the impact of a subset of important factors in NFT price formation over a massive dataset. In this case, several effective factors that are associated with NFTs financial performance are a visual rarity where NFTs with a high rarity score have a lower potential of negative return so these NFTs could be traded at a higher

price. In contrast, another study has shown the price prediction performance of NFTs increases if they integrate external factors such as the Twitter social media dataset into the original dataset (Kapoor et al., 2022). They do not only show the efficiency of social media data in NFT price prediction, but also they discovered that OpenSea characteristics like bids withdrawal or presale turns are crucial indicators as well. They also intended to address various security concerns regarding the environment in case of fraud and create mechanisms to quickly identify such fraudulent transactions that try to create artificial market signals. Among these earlier works on NFTs financial performance prediction (Ho et al., 2022; Mekacher et al., 2022; Kapoor et al., 2022), visual features incorporated in these techniques prediction models are directly from publicly available metadata, without being the result of training. Moreover, a subset of these methods predictions are based on a predefined set of classes.

More recent work on NFTs (Costa et al., 2023; Seyhan and Sefer, 2023) have forecasted NFTs financial performance, such as NFTs primary sale price and secondary sale participation, via considering raw dataset of images and textual descriptions (Costa et al., 2023). predicts NFT primary sale prices by deep learning, where the NFT category is an important factor in the model. On the other hand (Seyhan and Sefer, 2023), focuses on integrating multimodal image and textual datasets and enhances the prediction performance by using novel Differences Similarities-based deep learning architecture.

In general, our proposed approaches have several advantages while inferring NFTs distinct financial characteristics: 1- We integrate NFT image content and detailed visual attributes via transfer learning over pre-trained deep large vision models which can better utilize the image's representation. Those deep representation learning models do not focus on identifying the optimal subset of attributes for the prediction tasks by feature engineering. 2- The proposed approaches will be perfect solution candidates for people without NFT trading experience since no prior knowledge is required by them regarding NFT domain. 3- We make dynamic price predictions over time rather than predicting NFT prices at a static snapshot. To our best knowledge, none of the existing approaches consider the temporal price dynamics of NFTs, instead they simply assume NFT price is static over time and make price predictions accordingly.

3. Methods

In this study, we focus on the tasks of 1- Predicting the price at which NFT will be sold, 2- Predicting the price at which NFT will be sold at a given future time point, and 3- Detecting whether NFT will be sold in a secondary sale. The first two problems are regression problems, whereas the third problem is a classification problem. We tried several algorithms for all these problems.

3.1. Logistic regression

The simplest classifier we used in our secondary sale prediction and price direction change problems is Logistic regression (Cox, 1958). It models the probability of a binary outcome using the logistic function. This algorithm is suitable for scenarios where the dependent variable is categorical. However, its performance is quite low compared to other more advanced techniques.

3.2. AdaBoost

Boosting method was first introduced in (Freund and Schapire, 1997). As described in (Schapire, 1999), it works by combining several weak learning algorithms with weights into a strong learning algorithm which yields higher accuracy than each base learning algorithm. AdaBoost (Schapire, 2013), short for Adaptive Boosting, is an ensemble learning method designed for both classification and regression. It combines the outputs of multiple weak classifiers (learners with modest predictive power) to create a strong, accurate model. AdaBoost assigns weights to instances in the dataset, focusing on misclassified samples to improve overall performance iteratively. Let x_i, y_i represent training input and output data for training sample i respectively, Adaboost algorithm is summarized in Algorithm 1.

Algorithm 1. AdaBoost Algorithm

1. Initialize the observation weights $w_i = 1/N$, $i = 1, 2, \ldots, N$
2. For $m = 1$ to M:
   (a) Fit a classifier $G_m(x)$ to the training data using weights w_i.
   (b) Compute
   $$\mathrm{err}_m = \frac{\sum_{i=1}^{N} w_i I(y_i \neq G_m(x_i))}{\sum_{i=1}^{N} w_i}.$$
   (c) Compute $\alpha_m = \log((1 - \mathrm{err}_m)/\mathrm{err}_m)$
   (d) Set $w_i \leftarrow w_i \cdot \exp[\alpha_m \cdot I(y_i \neq G_m(x_i))]$, $i = 1, 2, \ldots, N$
3. Output $G(x) = \mathrm{sgn}\left[\sum_{m=1}^{M} \alpha_m G_m(x)\right]$

3.3. XGBoost

Gradient Boosting (Hastie et al., 2009) is a powerful ensemble learning technique employed for both classification and regression problems. It builds decision trees sequentially, each correcting the errors of the previous ones. The model minimizes a loss function and adjusts weights to improve predictions iteratively, making it robust and effective for complex datasets. XGBoost, or eXtreme Gradient Boosting (Chen and Guestrin, 2016), is a powerful and efficient ensemble learning algorithm that falls under the gradient boosting framework. It is designed for classification and regression tasks and is renowned for its speed and accuracy. Its algorithm is summarized in Algorithm 2. Let *N* be the number of training points, and *M* be the number of base learners, its algorithm is summarized in Algorithm 2.

Algorithm 2. XGBoost Algorithm

Data: Observations D $= \{(X_1, y_1, u_1, v_1), \ldots, (X_N, y_N, u_N, v_N)\}$

Result: Linear model set $\mathcal{F}^{\mathcal{M}} = \{F_1^M, \ldots, F_N^M\}$

for $n = 1$ to N do
  $\hat{\beta}_n^1 = \arg\min_{\beta_n^1} \frac{1}{2} \sum_{i=1}^{N} w_i (y_i - f_{\beta_n^1}(x_i))^2$
  $F_i^1 = f_{\hat{\beta}_n^1}$
end
for $m = 2$ to M do
  for $n = 1$ to N do
    $r_n = \lambda \cdot [y_n - F_n^{m-1}(x_n)]$
    $\hat{\beta}_n^m = \arg\min_{\beta_n^m} \frac{1}{2} \sum_{i=1}^{N} w_i (r_n - f_{\beta_n^m}(x_i))^2$
    $F_n^m = F_n^{m-1} + f_{\hat{\beta}_n^m}$
  end
end

3.4. Random Forest

A Random Forest (Breiman, 2001) is an ensemble learning method that constructs a multitude of decision trees during training and outputs the class that is the mode of the classes (classification) or mean prediction (regression) of the individual trees. It excels in handling complex datasets and reducing the risk of overfitting. Its algorithm is visualized in Fig. 1.

3.5. CNN (convolutional neural network)

CNNs (O' et al., 2015) are a class of deep neural networks particularly effective for image classification and visual recognition tasks (Lecun et al., 1998). They have been used for time-series prediction tasks in finance as well (Tuncer et al., 2022). Consisting of

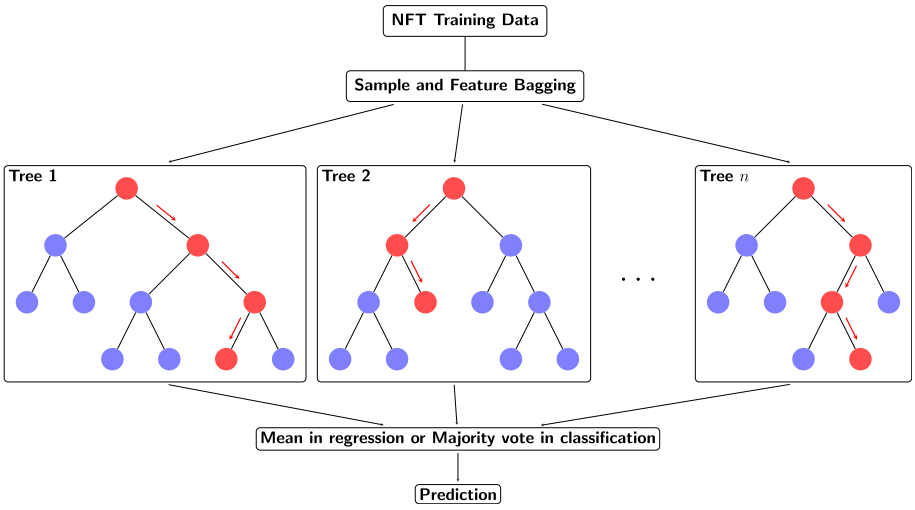

Fig. 1. Random forest visualization.

convolutional layers, pooling layers, and fully connected layers, CNNs automatically learn hierarchical representations of features from input data. They have revolutionized image processing and are widely used in computer vision applications.

In our case, we use a 2 convolutional layer neural network, in which convolutional layers are connected by batch normalization and max pooling layers. These layers have kernel size (3, 3), stride (1, 1) with ReLU activation for the convolution 2D layer and pooling (2, 2), stride (2, 2) for the max pooling layer. After two convolutional operations and max pooling, the dimension of processed data becomes $32 \times 63 \times 63$. After flattening this matrix, it is randomized to prevent overfitting using a dropout layer with a rate of 0.5. We apply two fully connected layers and generate a single prediction output. If the problem is a classification problem, we use the sigmoid activation function in the end for binary classification. Fig. 2 summarizes our CNN structure used in this problem.

3.6. LSTM (long short-term memory)

LSTMs (Hochreiter and Schmidhuber, 1997) are a class of deep recurrent neural networks particularly effective for time-series prediction problems such as stock price prediction, electricity prediction, demand forecasting, etc (Tuncer et al., 2022). LSTMs solve the problem of vanishing gradients better than the earlier deep recurrent learning variants since it has a higher memory power. LSTM's main difference from the more traditional feedforward neural networks is that their feedback connections retain useful information from the previous data in the sequence without having to treat each point independently. In our case, we use a single LSTM layer with 64 units. We use a historical window of size 5 to generate the corresponding train and test data. Among the different layer counts and unit combinations we tried, a single LSTM layer with 64 units gave the best performance. In addition to these steps, we increased the learning ability of the model by decreasing the learning rate and started using the “tanh” activation function.

4. Experiments

4.1. Datasets

Our NFT sales dataset (Nadini et al., 2021) comprises only NFT purchase transactions, where NFT's possession changes after the transaction is executed. In this case, transactions that represent bids as part of an auction or minted NFTs will be removed from our analysis. The dataset is composed and tracked over multiple cryptocurrencies. CryptoKitties sales (De La Rouviere, 2021), Gods-Unchained (Rosenbaum, 2021), Decentraland (Team, 2021a), and OpenSea (Team, 2021b) are 4 open source APIs where a subset of Ethereum datasets are collected. The remaining Ethereum blockchain data for SuperRare, Makersplace, Knownorigin, Cryptopunks, and Asyncart are obtained from NonFungible Corporation (Team, 2021c), which builds NFT valuations by tracking NFT sales information. Other than Ethereum, WAX blockchain transactions are also tracked via Atomic API (Team, 2021d). The whole set of NFTs is partitioned into 6 categories: Arts, Collectible, Games, Metaverse, Utility, and Other. In this case, the Metaverse category consists of virtual world items. The utility category represents elements with a certain function. Art consists of digital items such as GIFs, videos, and images. Games include competitive game objects. The collectible category includes elements that will be focused on by collectors. Lastly, the Other category consists of all remaining collections. We downloaded image objects that represent each NFT's URL and then resized images to 256x256 resolution. Almost all NFTs in the dataset were in typical image formats such as PNG and JPG. Other formats such as GIFs (Graphics Interchange Format) and SVGs (Scalable Vector Graphics) were also present. For GIFs, the first frame of the animation was converted into a 256x256 resolution. The SVGs were rasterized into 256x256 PNG images. All conversion and processing methods were implemented in ImageMagick (The ImageMagick Development Team, 2021).

While analyzing images for NFTs and extracting visual attributes over them, we use multiple large-scale pre-trained models for image characterization which generates hidden embeddings from the image via deep neural network. We focus on VGG-19 (Simonyan et al., 2015), InceptionNet (Szegedy et al., 2015), EfficientNet (Tan and Le, 2020), and RestNet50 (He et al., 2016) pre-trained models. All

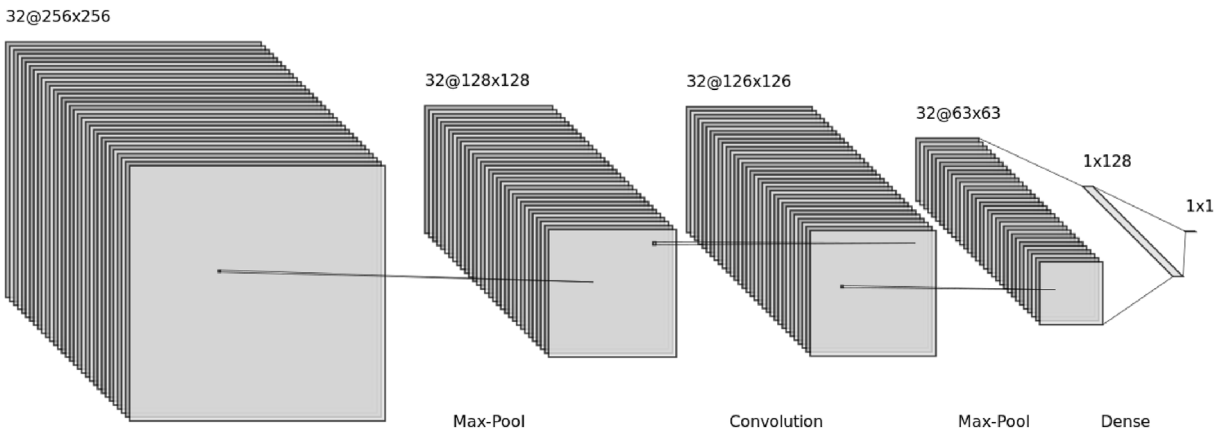

Fig. 2. Convolutional neural network for nft predictions.

these pre-trained models pass a given input image through more than one layer of transformation. The last layer vectors could be employed for multiple distinct tasks such as classification, clustering, similarity ranking, etc. We decrease the dimensionality of those pre-trained models embedding vectors, we run Principal Component Analysis (PCA) (Pearson, 1901) to select the topmost relevant components. Principal Component Analysis is a well-studied statistical method that preserves the variation in the data as high as possible while projecting each high-dimensional point to a smaller number of dimensions.

In summary, our dataset is made up of NFTs extracted from mostly Ethereum and WAX cryptocurrencies and covering dates between June 23, 2017, and April 27, 2021. For every sample in the dataset, many attributes including knowledge such as the NFT price, traders, and the category of the NFT are present. Out of the 24 attributes, we have utilized 12 of them to process and analyze to create features for individual NFT transactions. A brief description of these attributes as well as an example sample is given below with Table 1.

4.2. Dataset labeling

The prices for the primary sale prediction were already present in the dataset. Labeling the data for secondary sale prediction was done by labeling NFTs with more than one sale within a period after the primary sale with the value 1, and labeling the rest with 0. The time interval we used in our labeling process was 2 years.

4.3. Evaluation

For price prediction performance evaluation, let y_i be the true NFT price whereas we use $\hat{y}_i$ to represent the predicted NFT price. R^2 score to measure the amount of variance explained by a regression-based machine learning model as defined in:

$$R^2 = 1 - \frac{\sum_i y_i - \hat{y}_i^2}{\sum_i y_i - \bar{y}_i^2} \tag{1}$$

We compare the performance of our method with the existing techniques in terms of price direction and secondary sale participation predictions by using the common classification metrics. Let *TP*, *TN*, *FP*, and *FN* stand for True Positive, True Negative, False Positive, and False Negative respectively. Among them, Accuracy indicates how often a machine learning model can correctly predict an outcome from the total number of predictions. Recall metric measures what fraction of positive instances the machine learning model correctly predicts as a positive. Precision metric measures the performance of a model by considering the fraction of positively predicted labels that are correct. Lastly, F1 score uses the harmonic mean of the precision and recall. These metrics are defined as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{2}$$

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

$$Recall = \frac{TP}{TP + FN} \tag{4}$$

$$F1 = \frac{2{*}Precision{*}Recall}{Precision + Recall} = \frac{2{*}TP}{2{*}TP + FP + FN} \tag{5}$$

4.4. Handling imbalanced datasets and outlier removal

Random oversampling (Wongvorachan et al., 2023) is an oversampling method that randomly takes samples from the minority class and adds duplicates of them to the training dataset. SMOTE (Synthetic Minority Oversampling Method) (Chawla et al., 2002) is another

Table 1
Explanation of dataset attributes.

Attribute	Sample value	Description
Unique_id_collection	"('Godsunchained', '87094722.0')"	Unique identifier for a given NFT
Price_USD	0.030318	Price of the sold NFT converted into USD with a daily resolution
Seller_address	0x76481caa104b5f6bccb540da e4cefaf1c398ebea	Seller's address in the blockchain
Buyer_address	0xe0fb7622091e3d9ef9b438471 b10b9ea88c7cf6b	Buyer's address in the blockchain
Image_url_1	card.godsunchained.com/?id = 33&q = 4	URL to the digital object of the given NFT with more URLs as fallbacks
Image_url_2	…	URL fallback 1
Image_url_3	…	URL fallback 2
Datetime_updated	2019-11-30 00:00:00	Time of the transaction with a day resolution
Collection_cleaned	Godsunchained	Collection of a given NFT cleaned and grouped up
Category	Games	Category of a given NFT such as Games and art

oversampling method that creates synthetic examples for the minority class. We use these two oversampling techniques to handle imbalanced datasets to increase the performance of our models.

Interquartile range (*IQR*) measures the spread of the data and describes the middle group of a population. A method for finding outliers using *IQR* declares observations below Q_1 - 1.5*IQR* and above Q_3+1.5*IQR* outliers. Q_3 and Q_1 stand for the upper and lower quartiles of the data, which correspond to the 25th and the 75th percentile.

4.5. Implementation

Our approach is implemented in Python by using Pytorch (Paszke et al., 2019) and Sklearn (Pedregosa et al., 2011) libraries. Our machine learning models Adaboost, Random Forest, and XGBoost optimal hyperparameters are obtained via 5-fold cross-validation. For CNN, we set the size of the convolution kernel to 1 with stride 1. We use Adam optimizer (Kingma et al., 2015) with learning rate set to $1e^{-4}$. We train CNN for 10 epochs, which is reasonable given that image processing steps are already from pre-trained models.

5. Results

5.1. Analyzing NFT dataset for feature engineering

5.1.1. Visual feature extraction

Visual features of NFTs play an essential role in predicting NFT prices. This is mainly because image features could be tied to the rarity. NFT producers can provide rarity sheets with the launch of the NFT collections. Rarity gives information about the image features. Rarity sheets show the image attributes and represent how common these attributes are with the probabilities. It is possible to extract the ranking of all NFTs in the collection from the possibilities of the visual attributes. However, these rankings are not sufficient for assessing the NFTs visually as well. Let's consider 2 NFTs with similar rankings, where one has medium rare visual attributes and another has one ultra rare, common other visual features. The second NFT will be sold for a greater amount.

To extract features from the images, convolutional neural network models that are pre-trained on the ImageNet dataset were used. Three models, ResNet 50 (He et al., 2016), VGG-19 (Simonyan et al., 2015), and EfficientNet (Tan and Le, 2020) were utilized and EfficientNet was chosen as it proved to be significantly better than the remaining pre-trained models. In this process, the dense layer values just before the classification layer were used to extract the latent vector of size 4096 for each distinct image in the dataset.

To use these vectors for prediction and model training, we mapped these vectors into smaller dimension vectors. To decrease the size of these vectors, Principal Component Analysis (PCA) (Pearson, 1901) was performed on the whole dataset. By lowering the dimensionality of datasets with higher dimensionality, PCA minimizes the information loss while improving interpretability. We have reduced the dimensionality of these vectors to 7, which corresponds to the 7 components of PCA. The separation and grouping of categories and secondary sales are shown in Figs. 3 and 4 respectively for the topmost 3 principal components.

5.1.2. Collection median price

Another feature that accurately explains the variance in the dependent variable while predicting the NFT primary sale prices is the median price of the collection the NFT belongs to. For each NFT transaction in the dataset, the median price of all the NFT transactions made in the same collection until the resolution day of that transaction within a time interval is calculated. The median price of the NFTs sold in a week's interval served the best results for predictions. Its fluctuations can be seen in Fig. 5.

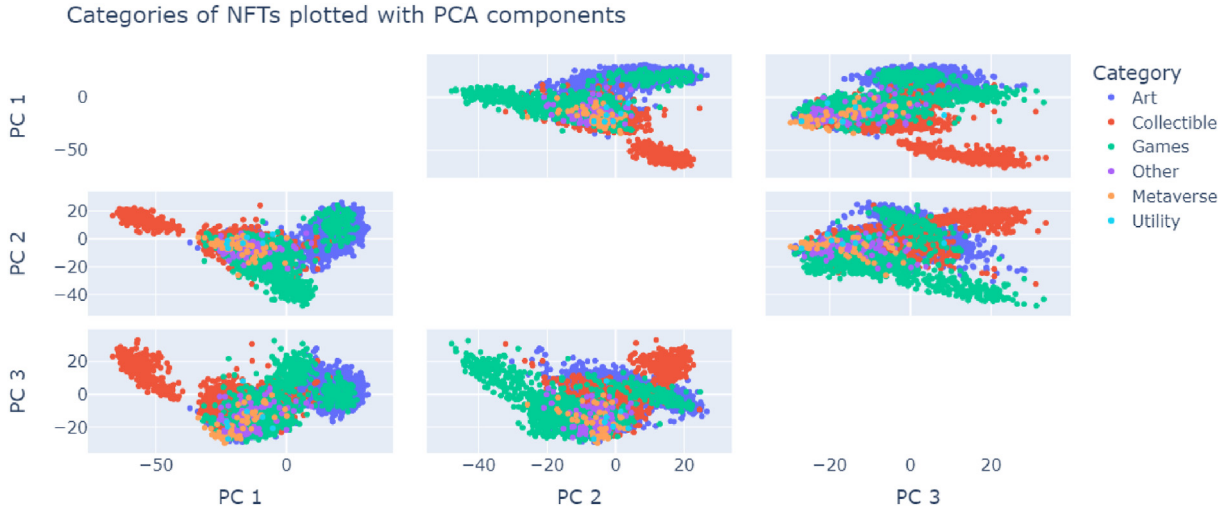

Fig. 3. PCA components category relations.

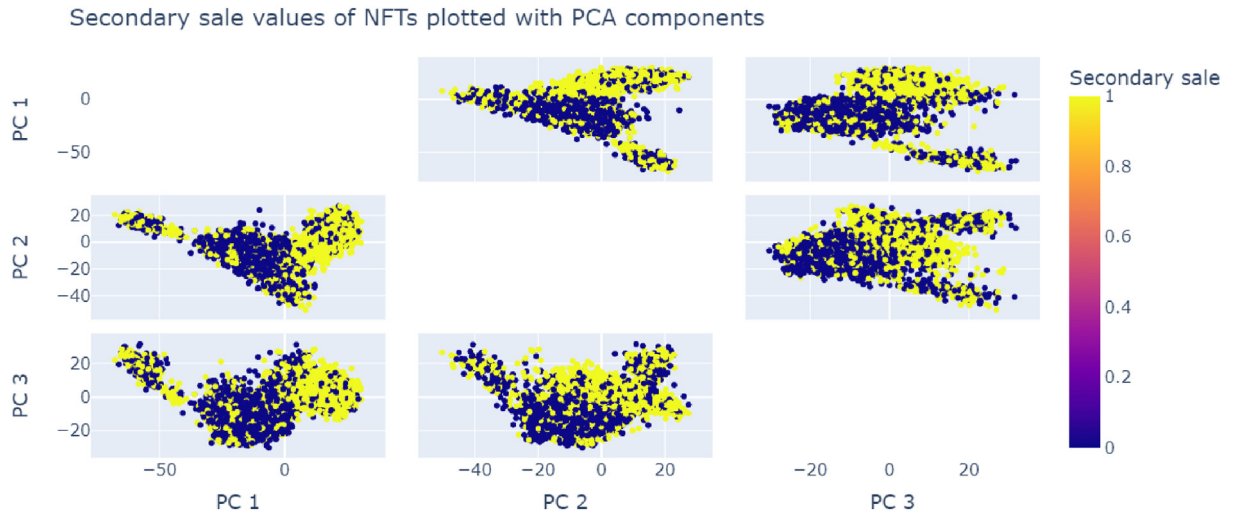

Fig. 4. PCA components secondary sale relations.

5.1.3. *Trader centrality*

Trader behaviors and tendencies will also be important when analyzing NFTs. To account for such trader dynamics, we construct two networks consisting of sellers and buyers. These networks are directed graphs that are constructed by creating directed edges between traders and the NFT that is being traded. Through this network, the degree centrality of the buyer and seller of the NFT transactions is used as a feature. The degree centrality of a node in a directed graph is the number of edges it is connected to, normalized by the maximum degree of any node in the graph.

These two networks are constructed by using every transaction in the dataset. To explain the mechanism, let us consider Table 2, with three transactions made in times t_k, t_l, and t_m. When NFT_1 is sold from trader Seller_i to Buyer_i at t_k, a directed link from Seller_i to NFT_1 is formed in the Seller network. Simultaneously, another directed link from Buyer_i to NFT_1 is created in the Buyer network. The seller and buyer network creation of each step can be seen in Fig. 6.

When calculating the centralities for each NFT transaction, transactions that are only resolved the day before are considered, which demands the calculation of networks separately for each day.

5.1.4. *Resale probability*

Another feature that is calculated on a collection basis is the resale probability of prior sales in that collection. A resale is defined as a binary attribute of an NFT derived from the number of that NFT resale. Only one sale indicates no resale and two or more sales indicate resale. The frequency of resales in a collection may likely be an indicator for resales of subsequent NFTs. For example, an NFT belonging to a trading card collection that has many prior resales in that collection is more likely to be resold than other collections with lower resales. The resale probability is defined as:

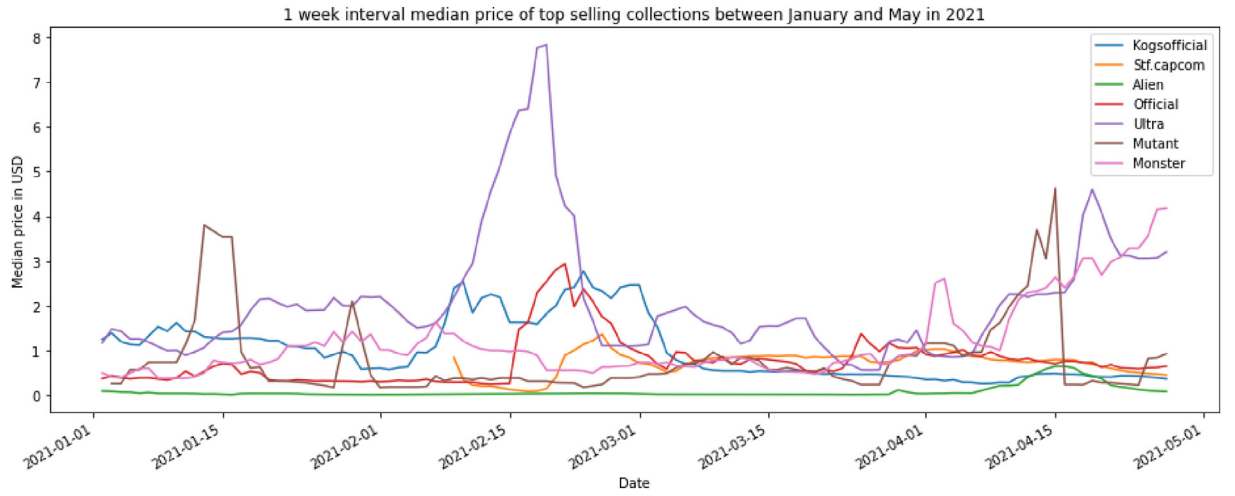

Fig. 5. Median prices of weekly intervals of NFTs.

Table 2
Example transactions.

Transaction time	Seller	Buyer	NFT
t_k	Seller_i	Buyer_i	NFT_1
t_l	Seller_j	Buyer_i	NFT_2
t_m	Seller_j	Buyer_i	NFT_3

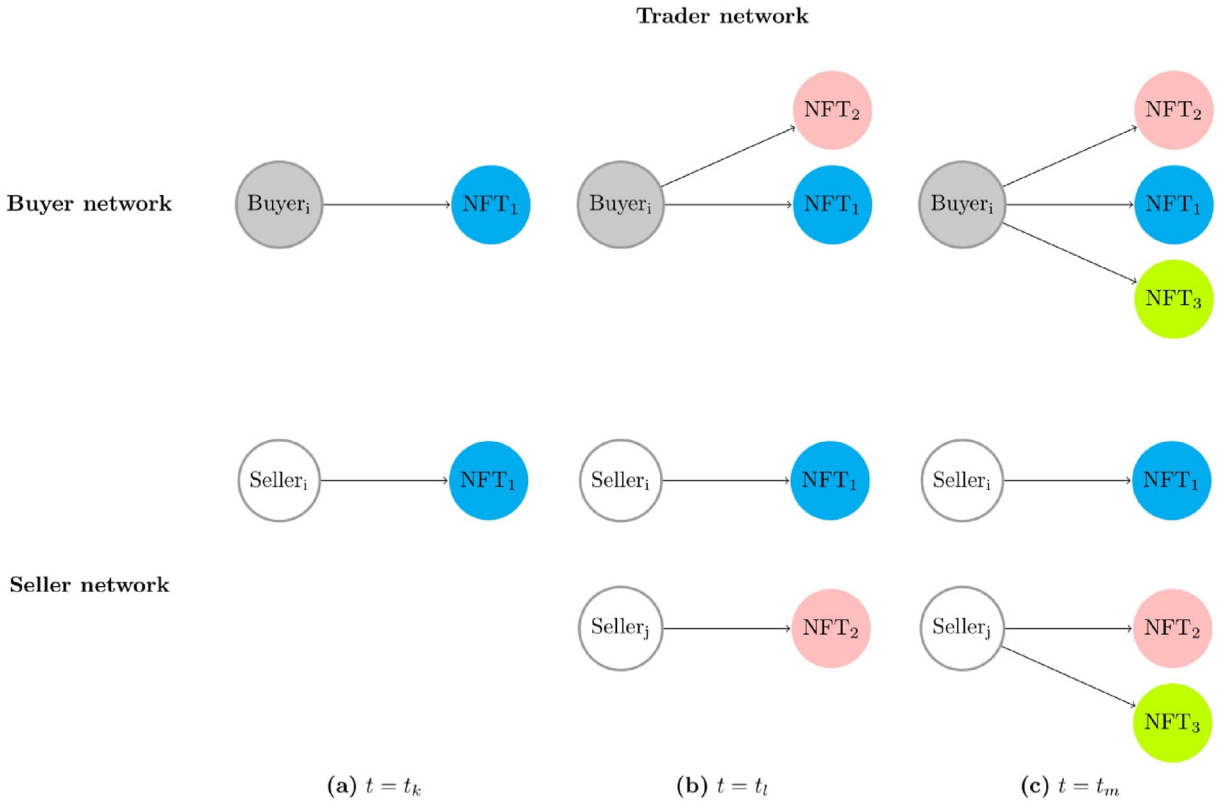

Fig. 6. Trader networks over time steps.

$$Presale = \frac{0.5}{n+1} + \frac{s}{n+1} \tag{6}$$

where n is the number of unique NFTs sold in that collection and s is the number of unique NFTs that were sold twice or more. This value models the resale frequency of a collection with 1 being highly resold and 0 being no resales. The NFTs that are sold first in a collection are assigned 0.5.

5.2. NFT primary sale price prediction

We combine all these features mentioned in the previous step to predict NFT prices in terms of USD. Once trained with the complete dataset, we obtained lower performance mainly due to outliers. To enhance the performance, we have applied outlier detection and removal. For the whole dataset, the standard deviation of the NFT price was 6430.50 with a mean of 131.62 and the max value of 7501893.00. We find the outliers by comparing two methods: Confidence intervals of the prices and the interquartile ranges. In this case, we have decided to drop samples with values less than $Q1$ - 1.5IQR or greater than $Q3 + 1.5IQR$. This led to almost 50,000 samples being dropped, with almost 3,25 million samples left for further analysis. Fig. 7 includes approximately 3.4 million NFT samples without missing features. The green area shows the 95 % confidence interval area of the price distribution.

To better understand the given dataset and compare separate models, instead of only training models on every sample, separate models were trained per NFT category. A total of 7 models for each problem were trained which are separate models considering only Games, Art, Other, Collectible, Metaverse, Utilility, and all of the categories at once. Table 3 shows the performance of the best-performing two learning algorithms results for NFT primary sale price prediction over various pre-trained visual models. In most of the categories, using the EfficientNet pre-trained model to preprocess visual attributes provided better performance than ResNet 50

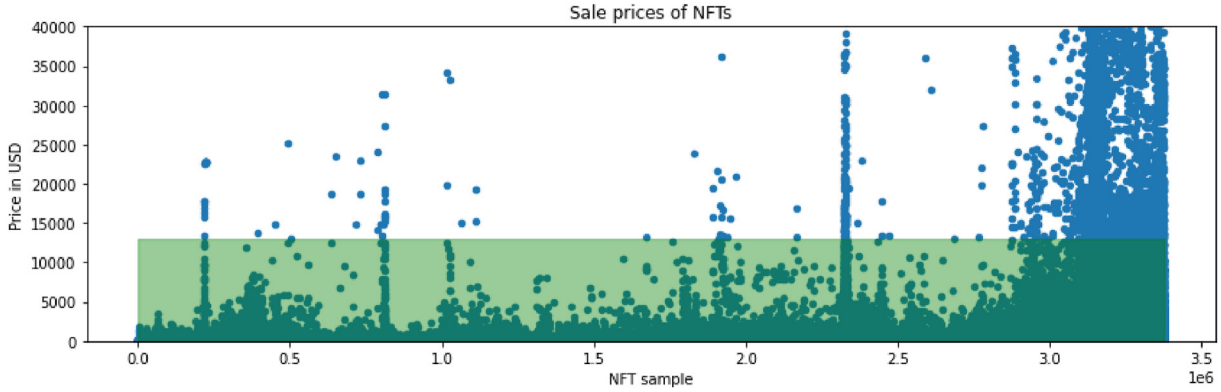

Fig. 7. Nft prices scatter plot.

Table 3
NFT primary sale price prediction performance in terms of R^2 score for the best performing two methods XGBoost and CNN over ResNet 50 and EfficientNet pre-trained visual models.

	XGBoost		CNN	
	ResNet 50	EfficientNet	ResNet 50	EfficientNet
Games	0.695	**0.711**	0.732	**0.761**
Art	0.474	**0.505**	0.495	**0.541**
Other	**0.601**	0.589	0.631	**0.644**
Collectible	**0.692**	0.687	0.705	**0.715**
Utility	0.196	**0.351**	0.298	**0.388**
Metaverse	0.727	**0.752**	0.757	**0.781**
All	0.611	**0.625**	0.633	**0.655**

which is in line with EfficientNet's better performance on various vision tasks (Tan and Le, 2020). CNN remarkably outperforms XGBoost in all categories when we preprocess the image content for both ResNet50 and EfficientNet pre-trained models. In all cases, the significantly lower R^2 score of the Utility category compared to the other categories suggests that our designed visual and non-visual features are not sufficient to explain the high variance in this category quite well. One of the main reasons for such lower performance is that Utility NFTs are targeted to make something more useful such as tools, services, royalties, and DAOs (Axelsen et al., 2022). Utility NFTs provide authorization of access, and these NFTs mostly have the same usuals or have tier-based visuals that define the features that users can reach according to their NFT tier. Therefore, visual attributes become less effective for Utility NFTs. Results summarized in Table 3 show the importance of investigation of NFTs on a categorical basis, which might lead to better insight and better predictions.

Additionally, training and validation loss curves in Fig. 8 show that the models were stopped before overfitting as they were monitored with early stopping algorithms.

5.3. NFT primary sale price prediction over time

We combine all these features mentioned in the previous step to predict NFT prices in terms of USD over time. Similar to Section 5.2, once trained with the complete dataset, we obtained lower performance mainly due to outliers. To enhance the performance, we have applied outlier detection and removal. Similar to Section 5.2, we find the outliers by comparing two methods: Confidence intervals of the prices and the interquartile ranges. In this case, we have decided to drop samples with values less than $Q1 - 1.5IQR$ or greater than $Q3 + 1.5IQR$.

To better understand the given dataset and compare separate models, instead of only training models on every sample, separate models were trained per NFT category. A total of 7 models for each problem were trained which are separate models considering only Games, Art, Other, Collectible, Metaverse, Utilility, and all of the categories at once. Table 4 shows the performance of the best-performing two learning algorithms (LSTM and CNN) for NFT primary sale price prediction over time, where images are first passed through various pre-trained visual models. Similar to the results without the temporal dynamics considered, utilizing the EfficientNet pre-trained model to preprocess visual attributes provided better performance than ResNet 50. LSTM significantly outperforms CNN in all categories when we preprocess the image content for both ResNet50 and EfficientNet pre-trained models. Similar to the results in Section 5.2, we observe quite low R^2 score of the Utility category. According to Table 4 results, analysis of NFTs on a categorical basis is important since it leads to better predictions.

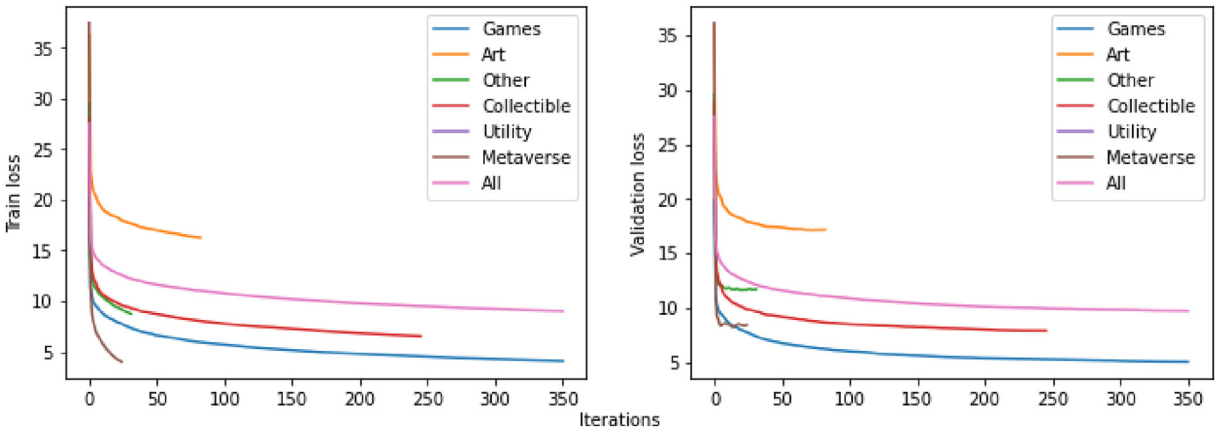

Fig. 8. Train and validation loss.

Table 4
NFT primary sale price prediction performance over time in terms of R^2 score for the best performing two methods: LSTM and CNN. We obtain visual embeddings for NFT images over ResNet 50 and EfficientNet pre-trained visual models.

	LSTM		CNN	
	ResNet 50	EfficientNet	ResNet 50	EfficientNet
Games	0.681	**0.702**	0.622	0.643
Art	0.435	**0.496**	0.468	0.478
Other	0.550	**0.579**	0.522	0.548
Collectible	0.634	**0.672**	0.614	0.635
Utility	0.192	**0.210**	0.205	0.207
Metaverse	0.698	**0.725**	0.720	0.710
All	0.605	**0.619**	0.625	0.626

5.4. NFT secondary sale participation prediction

Secondary sale refers to the binary class of NFTs the NFT belongs to, 1 being an NFT with more than one sale and 0 being an NFT with only one sale. In our dataset, a secondary sale is defined as the sale 2 years after the primary sale, which results in a classification problem. To better understand the given dataset and compare separate models, instead of only training models on every sample, separate models were trained per NFT category. A total of 7 models for each problem were trained which are separate models considering only Games, Art, Other, Collectible, Metaverse, Utilility, and all of the categories at once. In this case, the dataset is imbalanced as seen in Fig. 9), since approximately 20 % of the NFTs have no resale. To overcome this issue, we use two over-sampling methods: Random Oversampling and SMOTE. After rearranging the minority class according to these balancing techniques, we use 90 % of the dataset for training the model. In our analysis, we have employed 10-fold cross-validation for hyperparameter selection (see Fig. 10).

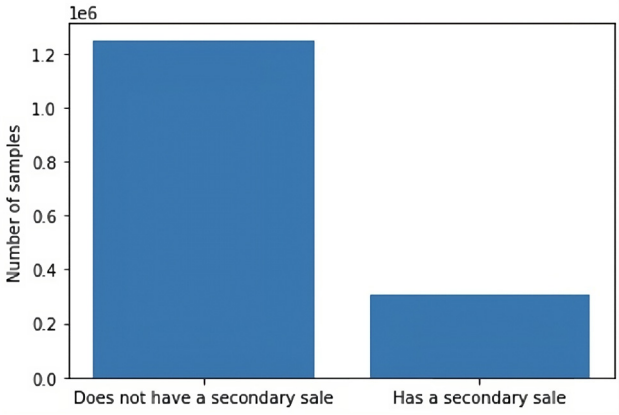

Fig. 9. Distribution of secondary sale participation for NFTs.

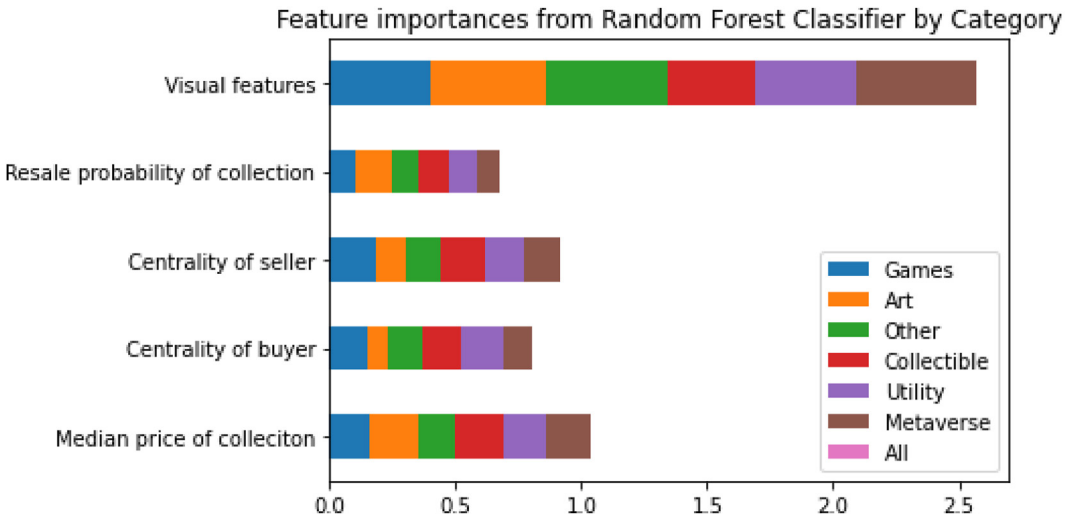

Fig. 10. Feature importance levels.

Once the data is balanced, Tables 5 and 6 show the performance of the CNN and XGBoost classifiers for NFT secondary sale prediction over various pre-trained visual models, when data is rebalanced by SMOTE. According to the table, we obtain the best classification performance always in CNN when we use ResNet 50, and in few categories by XGBoost. The model where the visual attributes are obtained via EfficientNet performs better than the one where they are obtained via ResNet50 for both methods. This is in line with the better performance of EfficientNet on multiple different computer vision tasks (Tan and Le, 2020). The performance of the Other category is worse than the rest of the category-specific models since Other category NFTs do not depend on image attributes as much as

Table 5
XGBoost Results for NFT secondary sale prediction over various pre-trained visual representation models when data is balanced via SMOTE.

Pre-trained Model	Category	Accuracy	Recall	F1 Score	Precision
ResNet 50	Games	0.691	0.573	0.451	0.373
	Art	0.726	0.729	0.803	0.894
	Other	0.637	0.561	0.290	0.195
	Collectible	0.744	0.487	0.522	0.563
	Utility	0.679	0.682	0.472	0.362
	Metaverse	0.712	0.726	0.711	0.696
	All	0.712	0.698	0.663	0.632
EfficientNet	Games	0.702	0.582	0.466	0.389
	Art	0.726	0.732	0.807	0.899
	Other	0.634	0.622	0.306	0.203
	Collectible	0.751	0.507	0.549	0.599
	Utility	0.681	0.646	0.466	0.364
	Metaverse	0.715	0.744	0.726	0.709
	All	0.714	0.699	0.682	0.666

Table 6
CNN Results for NFT secondary sale prediction over various pre-trained visual representation models when the data is balanced via SMOTE.

Pre-trained Model	Category	Accuracy	Recall	F1 Score	Precision
ResNet 50	Games	0.701	0.588	0.467	0.387
	Art	0.766	0.770	0.835	0.912
	Other	0.656	0.561	0.301	0.205
	Collectible	0.740	0.541	0.544	0.546
	Utility	0.722	0.546	0.453	0.387
	Metaverse	0.740	0.731	0.733	0.734
	All	0.739	0.713	0.692	0.672
EfficientNet	Games	0.841	0.588	0.622	0.660
	Art	0.852	0.905	0.903	0.901
	Other	0.888	0.357	0.456	0.631
	Collectible	0.812	0.621	0.654	0.691
	Utility	0.785	0.296	0.366	0.481
	Metaverse	0.769	0.794	0.770	0.748
	All	0.822	0.776	0.782	0.788

Table 7
Comparison of CNN and XGBoost results for NFT secondary sale prediction over EfficientNet when the data is balanced via Random Oversampling.

Method	Category	Accuracy	Recall	F1 Score	Precision
XGBoost	Games	0.638	0.543	0.400	0.317
	Art	0.537	0.478	0.613	0.841
	Other	0.596	0.593	0.279	0.183
	Collectible	0.669	0.496	0.462	0.432
	Utility	0.512	0.727	0.386	0.262
	Metaverse	0.643	0.513	0.578	0.662
	All	0.645	0.588	0.576	0.566
CNN	Games	0.695	0.610	0.471	0.386
	Art	0.759	0.756	0.828	0.915
	Other	0.678	0.574	0.319	0.222
	Collectible	0.752	0.518	0.545	0.575
	Utility	0.708	0.614	0.470	0.380
	Metaverse	0.739	0.740	0.743	0.730
	All	0.741	0.713	0.693	0.675

the categories. Although visual attributes may not be so informative for price prediction for the Utility category as seen in Table 3 mainly due to variability, the same is not true for secondary sale participation prediction.

When we repeat the same analysis over Random Oversampling technique, results are slightly lower than the results above for SMOTE. However, CNN and EfficientNet combination also appears to be the best among different combinations for Random Oversampling as well. Table 7 shows that CNN again outperforms XGBoost in many categories. We can easily observe that for all four methods the highest F1 scores generally belong to the Art category.

5.4.1. Feature importance

When we analyzed the importance of the considered features in each category over Random Forest, we found visual features to play a huge role in secondary sale classification for each category, building up our path to focus more on visual features in the follow-up studies. In terms of the lowest important features, our approach to integrating resale probabilities was not as significant as we anticipated.

6. Conclusions

This study proposes a comparison of multiple machine learning and deep learning algorithms to predict the primary sale price of NFTs as well as the existence of a secondary sale. We demonstrate that the NFTs visual features, previous collection sale prices, and trader's activity significantly impact the quality of prediction results. Our results reveal that NFT primary sale prices and secondary sales can be predicted moderately. Investors and companies can use our proposed models to formulate their trading strategies. Nevertheless, many other factors affect the price of NFTs and secondary sales, such as the price of the respective cryptocurrency, social media, and the creators of NFTs, besides the features we take into consideration. Additionally, there is always a margin of uncertainty due to the market's volatile nature. The NFT market remains a highly risky market for investors irrespective of any methods used in determining investment strategies. Developing models specific to NFT categories leads to increased prediction performance. We also publish our novel NFT dataset with temporal price knowledge, which is, to our best knowledge, the first dataset to have NFT prices over time rather than at a single time point.

In future work, we plan to take into account social media trends which can be another source of data regarding NFT market behaviors. In this case, Google Trends and Twitter can be two main sources. In the analyzed dataset, we have no information on the creator of a given NFT. Collection and analysis of creators might be considered for a better prediction performance. Moreover, we are likely to have missed any independent NFT producers, since our dataset consists of data mainly from NFT marketplaces and directly from the WAX or Ethereum blockchains. Incorporating such data into the prediction framework may also increase the adaptability of our prediction method. Additionally, data from other active blockchains such as Solana could be used. Lastly, our novel NFT dataset will help develop more complex NFT price prediction methods which will take a time-series dataset and predict the future NFT price.

Declaration of competing interest

This manuscript has not been published and is not under consideration for publication elsewhere. We have no conflicts of interest to disclose.

References

Ante, L., 2022. The non-fungible token (nft) market and its relationship with bitcoin and ethereum. FinTech 1 (3), 216–224.
Axelsen, H., Jensen, J.R., Ross, O., 2022. When is a dao decentralized? Complex Systems Informatics and Modeling Quarterly 31, 51–75. https://doi.org/10.7250/csimq.2022-31.04.
Breiman, L., 2001. Random forests. Machine Learning 45 (1), 5–32. https://doi.org/10.1023/A:1010933404324.

Cachin, C., Chandran, N., 2009. A secure cryptographic token interface. In: 2009 22nd IEEE Computer Security Foundations Symposium, pp. 141–153. https://doi.org/10.1109/CSF.2009.7.
Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P., 2002. Smote: synthetic minority over-sampling technique. J. Artif. Int. Res. 16 (1), 321–357.
Chen, T., Guestrin, C., 2016. Xgboost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16. Association for Computing Machinery, New York, NY, USA, pp. 785–794. https://doi.org/10.1145/2939672.2939785.
Chevallier, J., Guégan, D., Goutte, S., 2021. Is it possible to forecast the price of bitcoin? Forecasting 3 (2), 377–420. https://doi.org/10.3390/forecast3020024.
Colavizza, G., 2023. Seller-buyer networks in nft art are driven by preferential ties. Frontiers in Blockchain 5. https://doi.org/10.3389/fbloc.2022.1073499.
Costa, D., La Cava, L., Tagarelli, A., 2023. 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. Association for Computing Machinery, New York, NY, USA, pp. 1875–1885. WWW '23.
Cox, D.R., 1958. The regression analysis of binary sequences. Journal of the Royal Statistical Society: Series B (Methodological) 20 (2), 215–232.
De La Rouviere, N., 2021. A Subgraph to Index & Explore Cryptokitties Auctions. https://thegraph.com/explorer/subgraph/nieldlr/cryptokitties-sales. (Accessed 4 May 2021).
Dowling, M., 2022a. Fertile land: pricing non-fungible tokens. Finance Research Letters 44, 102096.
Dowling, M., 2022b. Is non-fungible token pricing driven by cryptocurrencies? Finance Research Letters 44, 102097.
Du, L., Kim, M., Lee, J., 2022. The art nfts and their marketplaces. arXiv:2210.14942.
Freund, Y., Schapire, R.E., 1997. A decision-theoretic generalization of on-line learning and an application to boosting. Journal of Computer and System Sciences 55 (1), 119–139. https://doi.org/10.1006/jcss.1997.1504.
Guidi, B., Michienzi, A., 2023. The social impact of nfts in the metaverse economy. In: Proceedings of the 2023 ACM Conference on Information Technology for Social Good, GoodIT '23. Association for Computing Machinery, New York, NY, USA, pp. 428–436. https://doi.org/10.1145/3582515.3609564.
Hamilton, J., 1994. Time Series Analysis. Princeton Univ. Press, Princeton, NJ. URL. http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+126800421&sourceid=fbw_bibsonomy.
Hari, M.K., Agrawal, A., Bhatia, R., Bhatia, A., Tiwari, K., 2023. T-pass: a blockchain-based nft enabled property management and exchange system. In: 2023 International Conference on Information Networking (ICOIN), pp. 140–145. https://doi.org/10.1109/ICOIN56518.2023.10048973.
Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning, second ed. Springer http://www-stat.stanford.edu/~tibs/ElemStatLearn/,/bib/hastie/Hastie2001/weatherwax_epstein_hastie_solutions_manual.pdf.
He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. https://doi.org/10.1109/CVPR.2016.90.
Ho, K.-H., Hou, Y., Chan, T.-T., Pan, H., 2022. Analysis of non-fungible token pricing factors with machine learning. In: 2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 1161–1166.
Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Computation 9 (8), 1735–1780.
Huang, J., He, N., Ma, K., Xiao, J., Wang, H., 2023. A deep dive into nft rug pulls. arXiv:2305.06108.
Ito, K., Shibano, K., Mogi, G., 2022. Bubble prediction of non-fungible tokens (nfts): an empirical investigation. arXiv:2203.12587.
Kapoor, A., Guhathakurta, D., Mathur, M., Yadav, R., Gupta, M., Kumaraguru, P., 2022. Tweetboost: influence of social media on nft valuation. In: Companion Proceedings of the Web Conference 2022. Association for Computing Machinery, New York, NY, USA, pp. 621–629. WWW '22.
Kingma, D.P., Ba, J., 2015. Adam: a method for stochastic optimization. In: Bengio, Y., LeCun, Y. (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.
Kshetri, N., 2022. Scams, frauds, and crimes in the nonfungible token market. Computer 55 (4), 60–64. https://doi.org/10.1109/MC.2022.3144763.
Lecun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE 86 (11), 2278–2324. https://doi.org/10.1109/5.726791.
Lee, H., Lee, G.-C., Koo, H.-Y., 2024. Exploring the relationship between rarity and price of profile picture nft: a formal concept analysis on the bayc nft collection. Blockchain: Research and Applications 5 (2), 100191. https://doi.org/10.1016/j.bcra.2024.100191.
Madine, M., Salah, K., Jayaraman, R., Zemerly, J., 2023. Nfts for open-source and commercial software licensing and royalties. IEEE Access 11, 8734–8746. https://doi.org/10.1109/ACCESS.2023.3239403.
McNally, S., Roche, J., Caton, S., 2018. Predicting the price of bitcoin using machine learning. In: 2018 26th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), pp. 339–343. https://doi.org/10.1109/PDP2018.2018.00060.
Mekacher, A., Bracci, A., Nadini, M., Martino, M., Alessandretti, L., Aiello, L.M., Baronchelli, A., 2022. Heterogeneous rarity patterns drive price dynamics in nft collections. Scientific Reports 12 (1), 13890.
Nadini, M., Alessandretti, L., Di Giacinto, F., Martino, M., Aiello, L.M., Baronchelli, A., 2021. Mapping the nft revolution: market trends, trade networks, and visual features. Scientific Reports 11 (1), 20902.
Okonkwo, I.E., 2021. NFT, copyright and intellectual property commercialization. International Journal of Law and Information Technology 29 (4), 296–304. https://doi.org/10.1093/ijlit/eaab010.
O'Shea, K., Nash, R., 2015. An introduction to convolutional neural networks. arXiv:1511.08458.
Park, N.H., Kim, H., Lee, C., Yoon, C., Lee, S., jin, Y., Shin, S., 2023. A deep dive into nft whales: a longitudinal study of the nft trading ecosystem. arXiv:2303.09393.
Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: an imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems, 32, pp. 8024–8035.
Pearson, K., 1901. LIII. On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 2 (11), 559–572. https://doi.org/10.1080/14786440109462720.
Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. Journal of Machine Learning Research 12, 2825–2830.
Rosenbaum, D., 2021. Gods-unchained Marketplace. https://github.com/djrosenbaum/unchained-transactions. (Accessed 4 May 2021).
Sako, K., Matsuo, S., Meier, S., 2021. Fairness in erc token markets: a case study of cryptokitties. In: Bernhard, M., Bracciali, A., Gudgeon, L., Haines, T., Klages-Mundt, A., Matsuo, S., Perez, D., Sala, M., Werner, S. (Eds.), Financial Cryptography and Data Security. FC 2021 International Workshops. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 595–610.
Schapire, R.E., 1999. A brief introduction to boosting. In: Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'99. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 1401–1406.
Schapire, R.E., 2013. Explaining adaboost. In: Empirical Inference. Springer, pp. 37–52.
Serada, A., Sihvonen, T., Harviainen, J.T., 2021. Cryptokitties and the new ludic economy: how blockchain introduces value, ownership, and scarcity in digital gaming. Games and Culture 16 (4), 457–480.
Seyhan, B., Sefer, E., 2023. Nft primary sale price and secondary sale prediction via deep learning. In: Proceedings of the Fourth ACM International Conference on AI in Finance, ICAIF '23. Association for Computing Machinery, New York, NY, USA, pp. 116–123. https://doi.org/10.1145/3604237.3626896.
Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: Bengio, Y., LeCun, Y. (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.
Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2015. Rethinking the inception architecture for computer vision. CoRR abs/1512.00567 arXiv:1512.00567.
Taherdoost, H., 2023. Non-fungible tokens (nft): a systematic review. Information 14 (1). https://doi.org/10.3390/info14010026.
Tan, M., Le, Q.V., 2020. Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv:1905.11946.
Team, D., 2021a. Decentraland Marketplace. https://thegraph.com/explorer/subgraph/decentraland/marketplace. (Accessed 4 May 2021).
Team, O., 2021b. Api Overview. https://docs.opensea.io/reference. (Accessed 4 May 2021).
Team, N., 2021c. The Best Place to Analyze, Track, and Discover Nfts. https://nonfungible.com/. (Accessed 4 May 2021).

Team, A.M., 2021d. Atomic Market Api. https://wax.api.atomicassets.io/atomicmarket/docs/swagger/. (Accessed 4 May 2021).
The ImageMagick Development Team, Imagemagick, 2021. https://imagemagick.org.
Tuncer, T., Kaya, U., Sefer, E., Alacam, O., Hoser, T., 2022. Asset price and direction prediction via deep 2d transformer and convolutional neural networks. In: Proceedings of the Third ACM International Conference on AI in Finance, ICAIF '22. Association for Computing Machinery, New York, NY, USA, pp. 79–86. https://doi.org/10.1145/3533271.3561738.
Vasan, K., Janosov, M., Barabási, A.-L., 2022. Quantifying nft-driven networks in crypto art. Scientific Reports 12 (1), 2769.
Wang, Q., Li, R., Wang, Q., Chen, S., 2021. Non-fungible token (nft): overview, evaluation, opportunities and challenges. arXiv:2105.07447.
Wang, Z., Gao, J., Wei, X., 2023a. Do nfts' owners really possess their assets? a first look at the nft-to-asset connection fragility. In: Proceedings of the ACM Web Conference 2023. Association for Computing Machinery, New York, NY, USA, pp. 2099–2109. WWW '23 https://doi.org/10.1145/3543507.3583281.
Wang, J.-N., Lee, Y.-H., Liu, H.-C., Hsu, Y.-T., 2023b. Dissecting returns of non-fungible tokens (nfts): evidence from cryptopunks. The North American Journal of Economics and Finance 65, 101892. https://doi.org/10.1016/j.najef.2023.101892.
Westerkamp, M., Victor, F., Küpper, A., 2018. Blockchain-based supply chain traceability: token recipes model manufacturing processes. In: 2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data. SmartData, pp. 1595–1602.
White, B., Mahanti, A., Passi, K., 2022a. Characterizing the opensea nft marketplace. In: Companion Proceedings of the Web Conference 2022, pp. 488–496.
White, B., Mahanti, A., Passi, K., 2022b. Characterizing the opensea nft marketplace. In: Companion Proceedings of the Web Conference 2022. Association for Computing Machinery, New York, NY, USA, pp. 488–496. WWW '22 https://doi.org/10.1145/3487553.3524629.
Wongvorachan, T., He, S., Bulut, O., 2023. A comparison of undersampling, oversampling, and smote methods for dealing with imbalanced classification in educational data mining. Information 14, 54. https://doi.org/10.3390/info14010054.
Yakovenko, A., 2018. Solana: a new architecture for a high performance blockchain v0. 8.13. Whitepaper.
Yilmaz, T., Sagfossen, S., Velasco, C., 2023. What makes nfts valuable to consumers? perceived value drivers associated with nfts liking, purchasing, and holding. Journal of Business Research 165, 114056. https://doi.org/10.1016/j.jbusres.2023.114056.