NFT SALES CHARACTERISTICS AND PRICE PREDICTION BY TRANSFER LEARNING VISUAL ATTRIBUTES

A Thesis

by

Mustafa Pala

Submitted to the
Graduate School of Sciences and Engineering
in Partial Fulfillment of the Requirements for
the Degree of

Master of Science

in the
Department of Computer Science

Özyeğin University
February 2024

NFT SALES CHARACTERISTICS AND PRICE PREDICTION BY TRANSFER LEARNING VISUAL ATTRIBUTES

Approved by:

Asst. Prof. Dr. Emre Sefer, Advisor
Department of Computer Science
Özyeğin University

Prof. Dr. Hasan Fehmi Ateş
Department of Electrical & Electronics Engineering
Özyeğin University

Prof. Dr. Vahap Burak Saltoğlu
Department of Economics
Boğaziçi University

Date Approved: December 29, 2023

To my family, my friends, and all of my educators who guided me...

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. NFTs are almost always traded by using cryptocurrencies such as Ethereum and blockchains are utilized to encode them. There was a recent surge of interest in trading them which makes them another type of alternative investment. While the existing research predominantly focuses on the technical aspects of NFTs, limited attention has been given to predicting their prices. The inherent volatility of NFT prices, attributed to factors such as over-speculation, liquidity constraints, rarity, subjectivity, 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 two related prediction problems over NFTs: Predicting NFTs' sale price, and inferring whether a given NFT will participate in a secondary sale. 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, trader behavior, seller's and buyer's centralities in the trading network, and collection's resale probability, and we assess the reliability of these features. We categorize input NFTs into six categories based on their characteristic features: Art, Collectibles, Games, Metaverse, Utility, and Others. We train several different machine learning methods on these attributes to answer these two questions. Across detailed experiments, we found visual attributes obtained from deep pre-trained models to increase the prediction performance in all cases, even though the pre-trained model giving the optimal result may change depending on

the problem type. In general, none of the learning algorithms consistently outperformed the rest of them across all categories. Our code is publicly available at `https://github.com/seferlab/deep_nft`.

ÖZETÇE

Nitelikli Fikri Tapu'lar (NFT'ler); temel olarak resimler, sanat eserleri, hareketli görseller ve mizahi tasarımlar gibi görselleştirilebilecek herhangi bir nesneyi temsil eden, mülkiyeti blok zinciri üzerinde tanımlanmış dijital varlıklardır. Genellikle Ethereum gibi kripto para birimleri ile çevrimiçi alınıp satılabilirler.

NFT'lere ve ticaretine olan ilginin artışı, NFT'leri alternatif bir yatırım aracı haline getirmektedir. Ancak mevcut araştırma ve çalışmalar genellikle NFT'lerin teknik yönlerine odaklanmıştır ve fiyat tahminine yönelik çalışmalar sınırlı kalmıştır. NFT piyasasının karmaşık yapısı, spekülasyon ve manipülasyona açık olması, likidite eksikliği, nadirlik, öznellik ve volatilite gibi faktörler fiyat tahminini oldukça zor hale getirmektedir. Bu bağlamda, analiz ve tahminler için makine öğrenimi güçlü bir çözüm sunmaktadır.

Bu çalışmada iki tahmin problemine odaklanılmıştır: NFT'lerin satış fiyatı ve ikincil satış olasılığı. Önceden eğitilmiş derin öğrenme algoritmaları kullanılarak NFT'lerin görsel özellikleri analiz edilmiş, bu görsel bilgi ile satış geçmişi, tüccar davranışı ve ikincil satışların fiyat üzerindeki etkisi de analiz ve öğrenme sürecine dahil edilmiş ve bu özelliklerin güvenilirlikleri değerlendirilmiştir. NFT'leri sanat, koleksiyon, oyun, metaverse, işlevsellik ve diğer olmak üzere 6 kategoriye ayırarak makine öğrenimi algoritmalarıyla satışlar modellenmiş ve bu problemlere cevap aranmıştır.

Denemeler sonucunda, görsel özelliklerin tahmin performansını tüm durumlarda artırdığı, optimum performansı gösteren modelin ise problem tipine göre değiştiği görülmüştür. Genel olarak, öğrenme algoritmalarının hiçbiri tüm kategorilerde sürekli olarak diğerlerinden daha iyi performans gösterememiştir. Kodlarımız, `https://github.com/seferlab/deep_nft` adresinde herkese açık olarak sunulmuştur.

ACKNOWLEDGEMENTS

Firstly, I would like to extend my thankfulness to my thesis advisor Assistant Prof. Emre Sefer for his precious orchestration, guidance, and support. I am sincerely grateful to my educators.

I would also like to thank my thesis defense jury members Prof. Hasan Fehmi Ateş, and Prof. Burak Saltoğlu for their valuable time and effort in evaluating my thesis.

Finally, special thanks to my family for their unwavering love and belief.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

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, music, images, video recordings, and anything else that can be visualized as a representation of objects [1]. As an asset, using NFT may depend on obtaining the proper license [2]. 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 [3].

In terms of functionality, NFTs are similar to cryptographic tokens which can be defined as access rights or assets that can be programmed [4]. 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 [5]. Different than cryptographic tokens and cryptocurrencies, people generally cannot interchange Non-Fungible Tokens reciprocally where the non-fungible name comes from [6]. A non-fungible token 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 [7].

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 [8]. In general, NFTs have been seen and traded as speculative asset classes. Some investors still consider NFT trading dynamics similar to a Ponzi scheme [9]. Additionally, NFTs have been criticized for their higher energy usage, increased carbon footprint over several types of blockchain, and increased appearance in art frauds [10]. 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 [11]. 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 [12]. Such decentralized intellectual property rights are an alternative to centralized copyright control which is handled by formal government institutes, brokers, and intermediaries [13].

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 [14]. 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, 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 [15]. 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.

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 [16, 17]. Another class of studies that analyze NFT market dynamics is limited to some well-known collections such as CryptoKitties [18, 19]. Moreover, those studies may also focus on more isolated NFT markets such as Decentraland [20, 21]. Lastly, more recent studies have started to interpret and analyze NFTs as tradable assets [21, 22, 23]. Once analyzed as tradable assets, predicting NFTs price 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 [22, 24]. 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-16 [25], EfficientNet [26], and ResNet50 [27] via transfer learning. The existing techniques do not use these pre-trained models, do not fully utilize image attributes of NFTs which are publicly available.

In this thesis, we tackle the problem of predicting NFT primary sale price, and its participation in a secondary sale over time by various machine and 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.

More formally, we mainly apply eXtreme Gradient Boosting (XGBoost) [28], AdaBoost [29], Random Forest [30], and CNN [31] to the following regression and classification problems respectively: 1- Primary sale price prediction, and 2- 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 prices, especially when we train an independent model for each NFT category. On the other hand, CNN and Random Forest perform better than the rest of the methods in predicting secondary sale participation. Once hyperparameters are tuned accordingly, data imbalance is addressed via Random Oversampling and SMOTE [32], 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 ResNet50 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.

CHAPTER II

METHODOLODY

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

2.1 Logistic Regression

The simplest classifier we used in our secondary sale prediction problem is Logistic regression [33]. 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.

2.2 AdaBoost

Boosting method was first introduced in [34]. As described in [35], 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 [29], 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 weighted error
   $$\mathrm{err}_m = \frac{\sum_{i=1}^{N} w_i I(y_i \neq G_m(x_i))}{\sum_{i=1}^{N} w_i}.$$
   (c) Compute classifier weight $\alpha_m = \log((1 - \mathrm{err}_m)/\mathrm{err}_m)$.
   (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 $G(x) = \mathrm{sign}\left[\sum_{m=1}^{M} \alpha_m G_m(x)\right]$.

2.3 XGBoost

Gradient Boosting [36] 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 [28], is a powerful and efficient ensemble learning algorithm that falls under the gradient boosting framework. It is renowned for its speed and accuracy. Let N be the number of training points, and M be the number of base learners, its algorithm is summarized in Algorithm 2.

2.4 Random Forest

A Random Forest [30] 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 Figure 1.

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

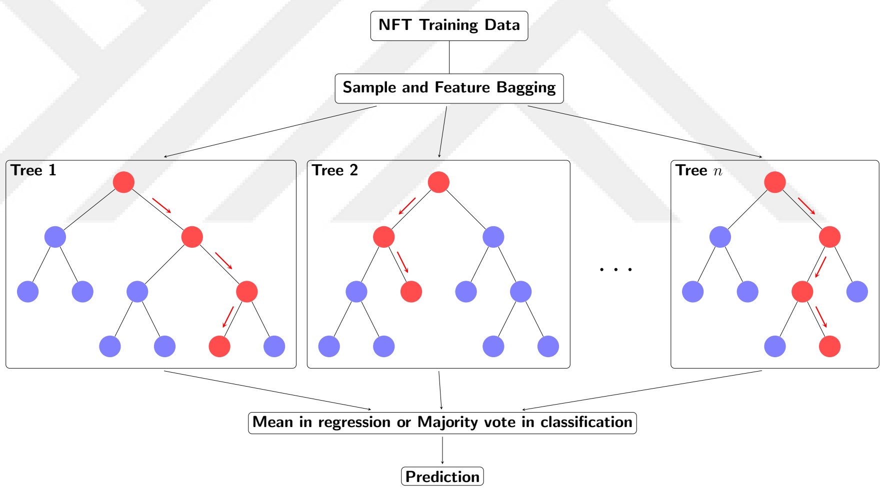

Figure 1: Random Forest Visualization

2.5 CNN (Convolutional Neural Network)

CNNs [31] are a class of deep neural networks particularly effective for image classification and visual recognition tasks [37]. They have been used for time-series prediction tasks in finance as well [38]. Consisting of 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 pre-trained convolutional neural network architectures that are pre-trained for image detection and segmentation tasks. VGG-16, ResNet50, and EfficientNet are popular architectures that are used to extract the features from the images. In this process, the dense layer values just before the classification layer were used to extract the latent vector for each distinct image in the dataset. The extracted features are high-dimensional representations that contain the various aspects of the image. The outputs of the pre-trained architecture are connected to another CNN structure.

This CNN structure contains a convolution 2D layer that have kernel size $(3, 3)$, stride $(1, 1)$ with ReLU activation, max pooling layer with pooling $(2, 2)$, stride $(2, 2)$ which are connected by batch normalization. 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. We use Adam optimizer [39] with learning rate set to $1e^{-4}$. We train CNN for 10 epochs, which is reasonable given that image processing steps are already from pretrained models. If the problem is a regression problem, we use the linear activation function. If the problem is a classification problem, we use the sigmoid activation function in the end for binary classification. Figure 2 summarizes our CNN structure used in this problem.

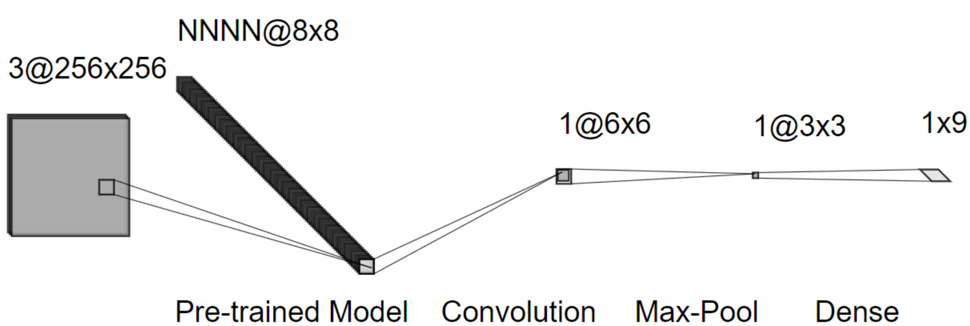

Figure 2: Convolutional Neural Network integrated with Pretrained Models for NFT Predictions

2.6 Pretrained Visual Models

2.6.1 VGG-16

VGG-16 [25] is a convolutional neural network that is 16 layers deep in total. It consists of 13 convolutional layers, 5 max-pooling layers, and 3 fully connected layers, where the number of layers having tunable parameters is 16 (13 convolutional layers and 3 fully connected layers). It is pretrained over more than a million images from the ImageNet database. The pretrained network can classify images into 1000 object categories, such as keyboard, mouse, pencil, and many animals. As a result, the network has learned rich feature representations for a wide range of images. The network has an image input size of 224×224. VGG-16 architecture is summarized in Figure 3. In our case, we get visual embeddings after applying convolutional layers, just before fully connected layers.

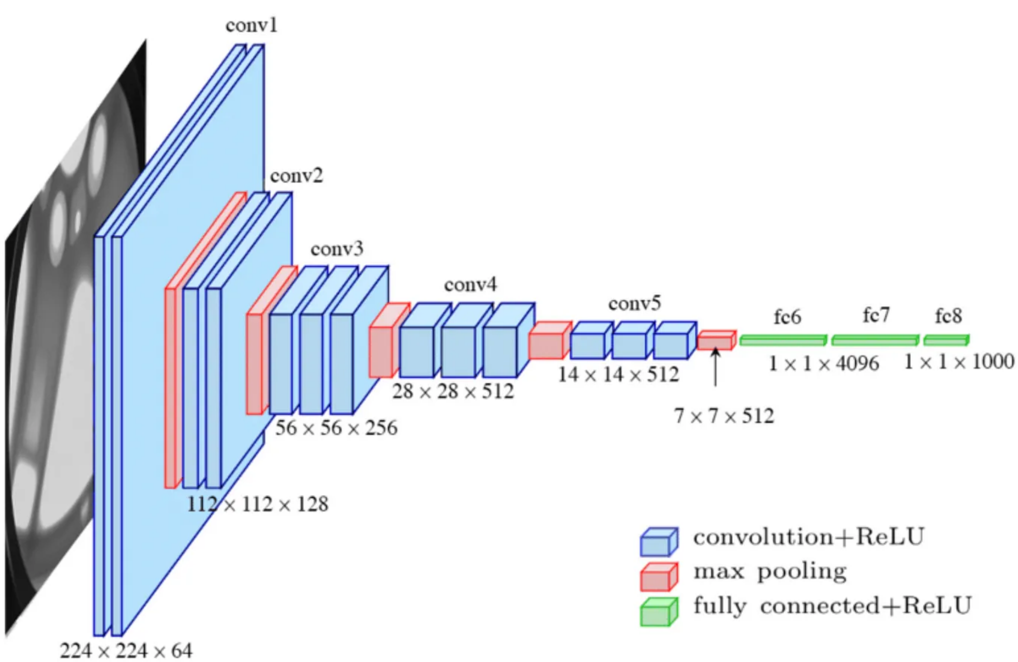

Figure 3: VGG-16

2.6.2 ResNet50

ResNet50 [27] consists of 50 layers and achieved state-of-the-art performance on the ImageNet dataset. ResNet50 consists of 16 residual blocks, with each block consisting of several convolutional layers with residual connections. The architecture also includes pooling layers, fully connected layers, and a softmax output layer for classification. In detail, the input layer of ResNet50 takes an image of size $224 \times 224 \times 3$ as input. This outputs a feature map with dimensions $(112, 112, 64)$. This layer uses 64 filters, each with a kernel size of (7, 7) and a stride of (2, 2) to downsample the input image by a factor of 2 in both the width and height dimensions. After the convolutional layer, a batch normalization layer is applied to normalize the activations, followed by a ReLU activation function to introduce nonlinearity. Finally, a max pooling layer with a pool size of (3, 3) and a stride of (2, 2) is applied to further downsample the feature map by a factor of 2 in both dimensions. This produces a feature map with dimensions (56, 56, 64), which serves as the input to the subsequent convolutional blocks in the ResNet50 architecture. ResNet50 architecture is summarized in Figure 4. In our case, we get visual embeddings after applying convolutional layers, just before fully connected layers.

2.6.3 EfficientNet

EfficientNet [26] is a convolutional neural network that is based on the notion of compound scaling. This compounding scaling idea tackles the age-old trade-off between model size, accuracy, and computing efficiency. Compound scaling is intended to scale three critical aspects of a neural network: breadth, depth, and resolution which are defined as below. EfficientNet architecture is summarized in Figure 5. In our case, we get visual embeddings after applying convolutional layers, just before fully connected layers.

- Width: The number of channels in each layer of the neural network is referred

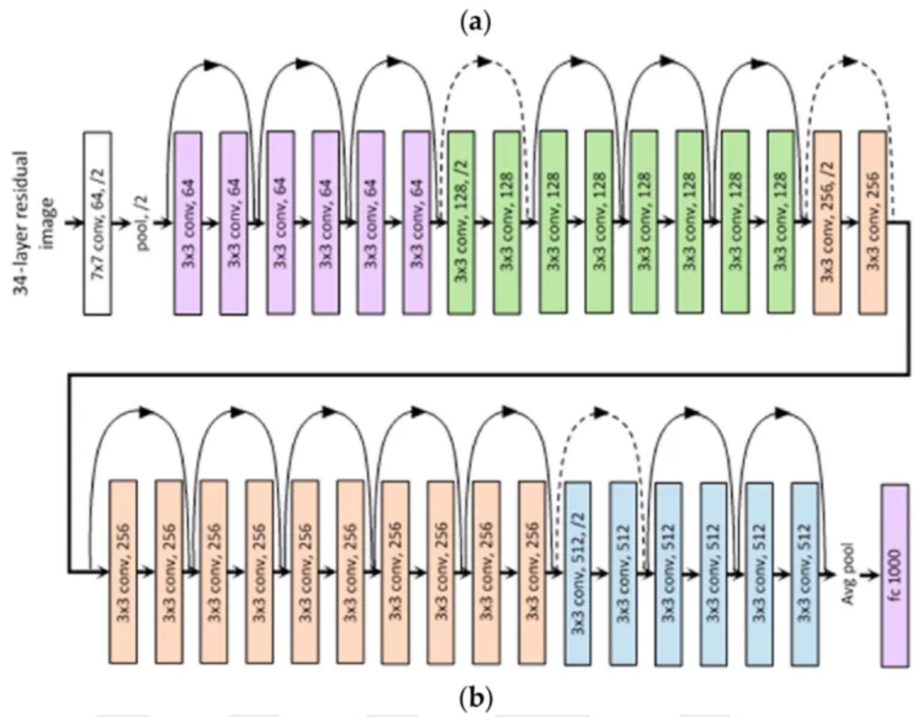

Figure 4: ResNet50

to as width scaling. By enlarging the width, the model may capture more complicated patterns and characteristics, resulting in higher accuracy. Reducing the width, on the other hand, results in a more lightweight AI model that is ideal for low-resource applications.

- Depth: The total number of layers in the network is referred to as depth scaling. Deeper models can capture more complex data representations, but they also need more computing resources. Shallower models, on the other hand, are more computationally efficient but may forfeit accuracy.

- Resolution: Resolution scaling entails changing the size of the input image. Higher-resolution images give more detailed information, which may result in improved performance. They do, however, need larger memory and processing power. Lower-resolution images, on the other hand, use fewer resources but may lose fine-grained information.

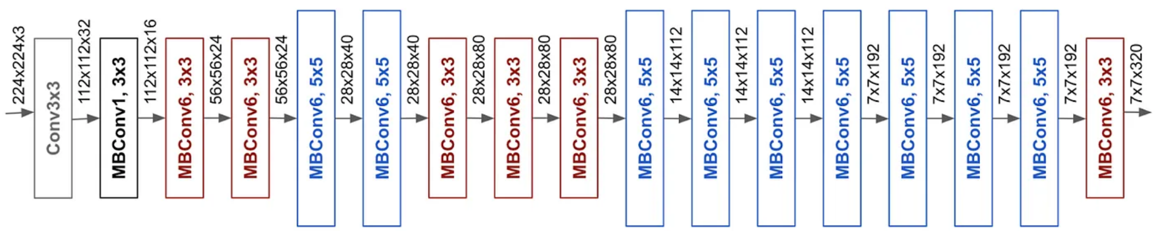

Figure 5: EfficientNet

CHAPTER III

EXPERIMENTS

3.1 Datasets

Our NFT sales dataset [24] comprises only NFT purchase transactions, where NFT's possession changes after the transaction is executed. In this case, transactions that represent bids as trained of an auction or minted NFTs will be removed from our analysis. The dataset is composed and tracked over multiple cryptocurrencies. CryptoKitties sales [40], Gods-Unchained [41], Decentraland [42], and OpenSea [43] 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 [44], which builds NFT valuations by tracking NFT sales information. Other than Ethereum, WAX blockchain transactions are also tracked via Atomic API [45]. 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 256×256 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 256×256 resolution. The SVGs were rasterized into 256×256 PNG images.

All conversion and processing methods were implemented in ImageMagick [46].

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-16 [25], EfficientNet [26], and ResNet50 [27] pre-trained models. All 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) [47] 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 10 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.

3.2 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.

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	0x76481caa104b5f6bccb540dae4cefaf1c398ebea	Seller's address in the blockchain
Buyer_address	0xe0fb7622091e3d9ef9b438471b10b9ea88c7cf6b	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

3.3 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 secondary sale participation 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:

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

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

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

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

We also evaluated the secondary sale prediction performance in terms Area Under the ROC Curve (AUC) score which shows the classifier's performance at distinct threshold values. The ROC curve plots recall against the False Positive Rate (FPR) at different thresholds which is defined as below. Afterwards, area under this curve is reported. The higher area corresponds to a better performance.

$$FPR = \frac{TP}{FP + TN} \tag{6}$$

3.4 Handling Imbalanced Datasets and Outlier Removal

Random oversampling [48] 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) [32] is another 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 - 1.5IQR and above Q_3 + 1.5IQR 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.

3.5 Implementation

Our approach is implemented in Python by using Pytorch [49] and Sklearn [50] libraries. Our machine learning models Random Forest, XGBoost, and CNN optimal hyperparameters are obtained via 5-fold cross-validation. For CNN, we set the size of the convolution kernel to 3 with stride 1. We use Adam optimizer [39] 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. Our code is publicly available at `https://github.com/seferlab/deep_nft`.

CHAPTER IV

RESULTS AND DISCUSSION

4.1 Analyzing NFT Dataset for Feature Engineering

4.1.1 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, ResNet50 [27], VGG-16 [25], and EfficientNet [26] 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 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) [47] 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 Figure 6 and Figure 7 respectively for the topmost 3 principal components.

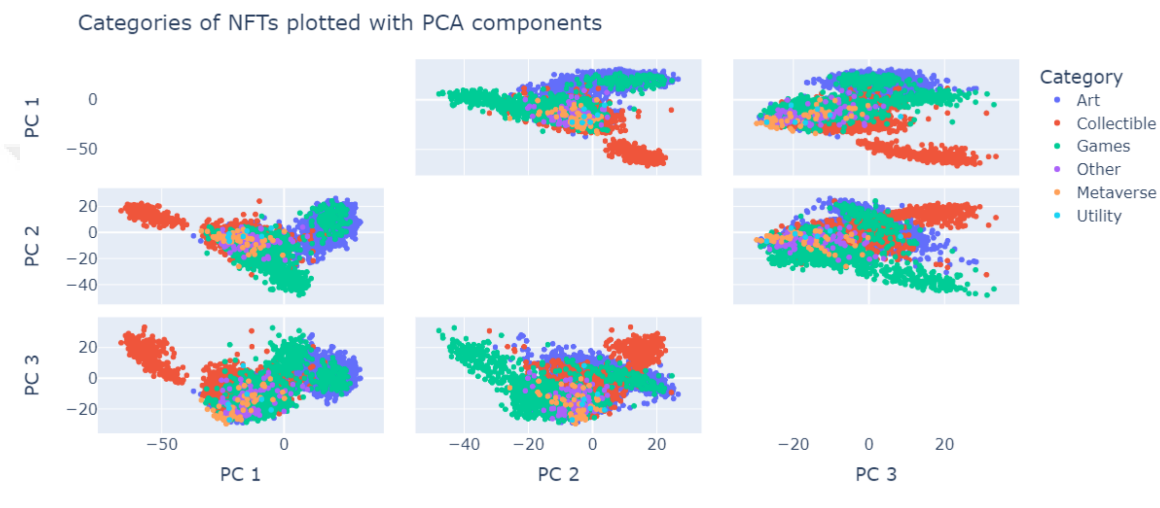

Figure 6: PCA Components Category Relations

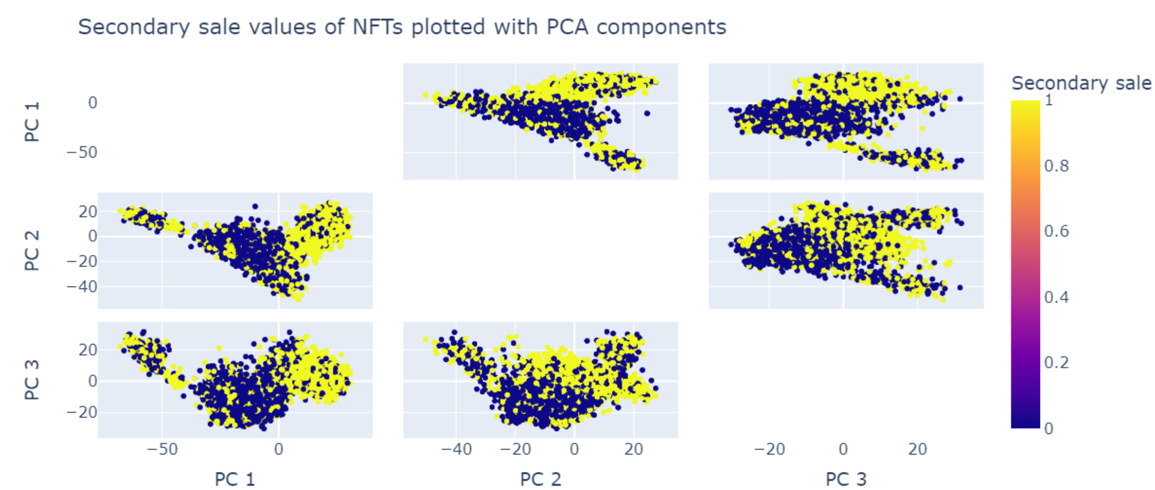

Figure 7: PCA Components Secondary Sale Relations

For both predicting the primary sale prices and secondary sale participation with CNN, we only used the image features that are extracted from the pre-trained models without applying PCA.

4.1.2 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 Figure 8.

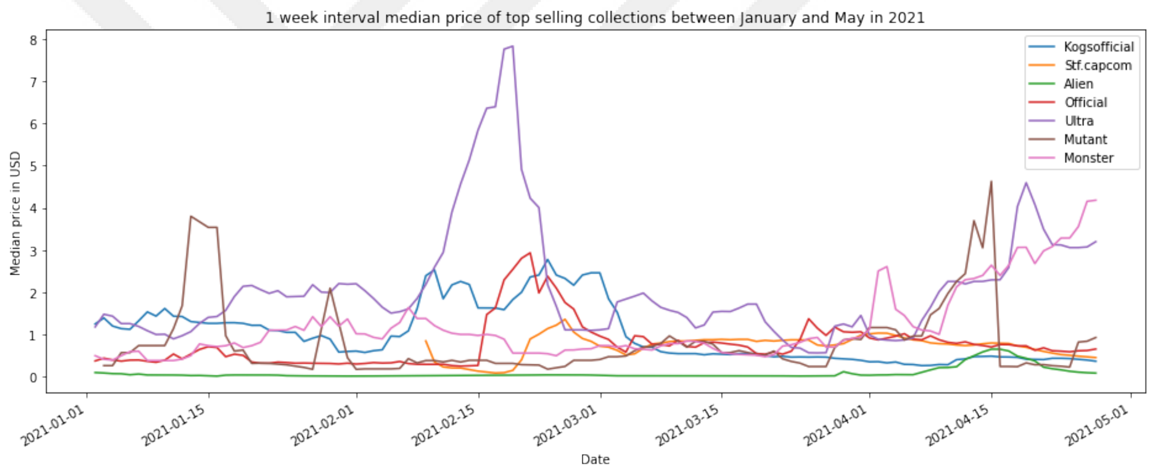

Figure 8: Median Prices of Weekly Intervals of NFTs

4.1.3 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.

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$_3$

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 Figure 9.

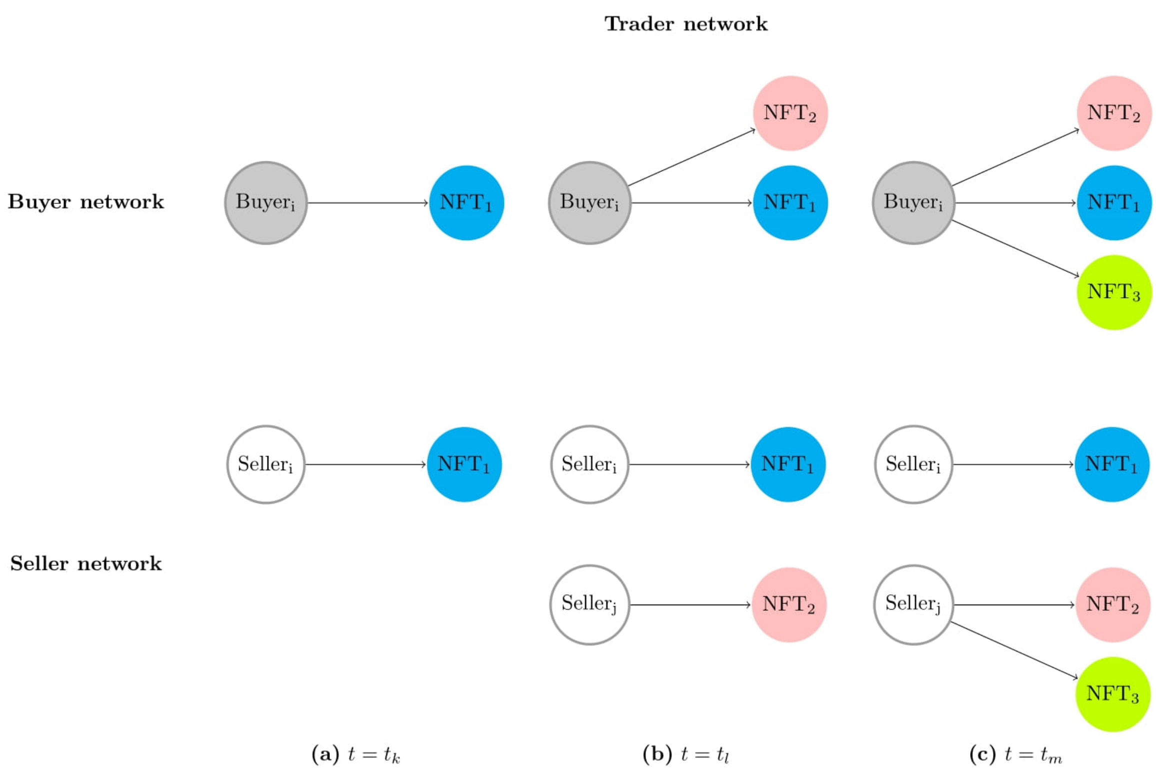

Figure 9: Trader Networks Over Time Steps

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.

4.1.4 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:

$$Resale = \frac{0.5}{n+1} + \frac{s}{n+1} \tag{7}$$

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.

4.2 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 $500,000$ samples being dropped, with almost 2.9 million samples left for further analysis. Figure 10 includes approximately 3.4 million NFT samples without missing features. The green area shows the 95% confidence interval area of the price distribution.

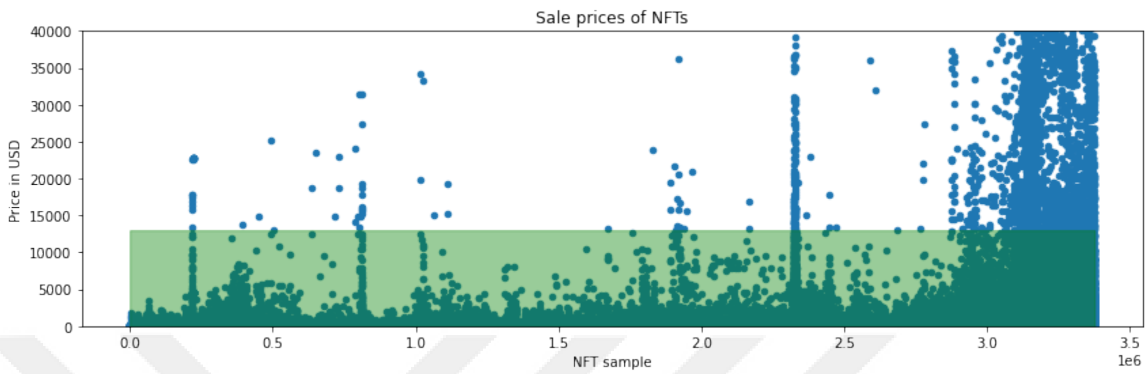

Figure 10: NFT Prices Scatter Plot

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, Utility, 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 ResNet50 which is in line with EfficientNet's better performance on various vision tasks [26]. 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 [51]. 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.

Table 3: NFT primary sale price prediction performance in terms of R^2 score for the best performing two methods XGBoost and CNN over ResNet50 and EfficientNet pre-trained visual models.

	XGBoost		CNN	
	ResNet50	EfficientNet	ResNet50	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**

Table 4 shows the performance comparison of all methods including the simplest logistic regression when we use EfficientNet as a pre-trained model. According to the table, we obtain the best NFT price prediction performance in CNN across 6 out of 7 categories. We found Random Forest and XGBoost to perform similarly across distinct categories. In general, visual attributes are informative even for logistic regression when the primary sale price is predicted.

Table 4: NFT primary sale price prediction performance comparison in terms of R^2 score over EfficientNet pre-trained visual model.

	Logistic Regression	Random Forest	XGBoost	CNN
Games	0.632	0.724	0.711	**0.761**
Art	0.415	0.509	0.505	**0.541**
Other	0.534	**0.647**	0.589	0.644
Collectible	0.650	0.655	0.687	**0.715**
Utility	0.184	0.342	0.351	**0.388**
Metaverse	0.699	0.741	0.752	**0.781**
All	0.591	0.615	0.625	**0.655**

Additionally, training and validation loss curves in Figure 11 show that the models were stopped before overfitting as they were monitored with early stopping algorithms.

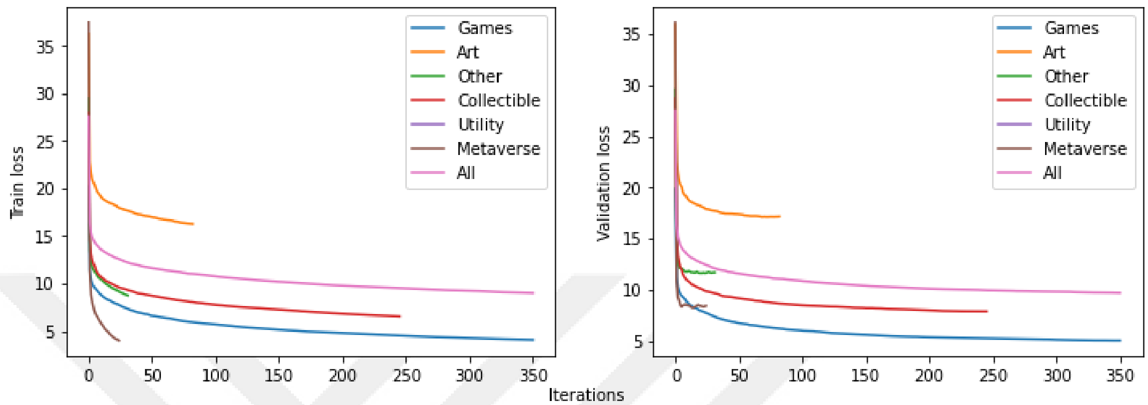

Figure 11: Train and Validation Loss

4.2.1 Ablation Studies

We also evaluate the performance of various learning algorithms by removing certain subsets of features as seen in Table 5. As seen in the Table, using all attributes performs the best. However, the performance drops the most when we remove the visual embedding attributes.

Table 5: Ablation studies.

	LR	RF	XGBoost	CNN
W/o Embedding	0.502	0.495	0.491	**0.512**
W/o Collection Median Price	0.523	0.514	0.588	**0.592**
W/o Trader Centrality	0.552	0.555	0.604	**0.621**
W/o Resale Probability	0.582	0.595	0.613	**0.643**
All Attributes	0.591	0.615	0.625	**0.655**

4.3 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, Utility, and all of the categories at once. In this case, the dataset is imbalanced as seen in Figure 12, 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.

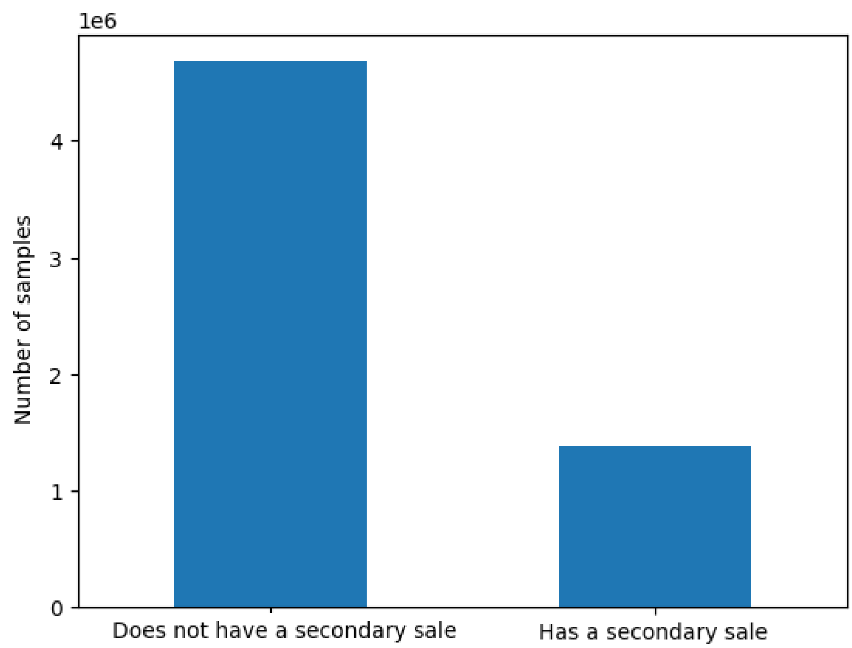

Figure 12: Distribution of Secondary Sale Participation for NFTs.

Once the data is balanced, Tables 6-7 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 mostly in CNN when we use ResNet50, and in a 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 [26]. 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 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.

Table 6: 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
ResNet50	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

Once the data is balanced via SMOTE, Table 8 shows the performance comparison of all methods including the simplest logistic regression when we use EfficientNet as a pre-trained model. According to the table, we obtain the best classification performance mostly in CNN and Random Forest. We found Random Forest to consistently outperform XGBoost across many categories. In general, visual attributes are informative even for logistic regression when they participate in secondary sale prediction.

When we repeat the same analysis over Random Oversampling technique, results

Table 7: 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
ResNet50	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

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 9 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.

4.3.1 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.

4.3.2 Ablation Studies

We also evaluate the performance of various learning algorithms on secondary sale participation prediction by removing certain subsets of features as seen in Table 10. As

Table 8: Comparison of multiple methods for NFT secondary sale prediction over EfficientNet when the data is balanced via SMOTE.

Method	Category	Accuracy	Recall	F1 Score	Precision	AUC
Logistic Regression	Games	0.636	0.544	0.400	0.315	0.721
	Art	0.537	0.478	0.613	0.855	0.807
	Other	0.580	0.625	0.282	0.182	0.708
	Collectible	0.668	0.497	0.462	0.432	0.745
	Utility	0.493	0.705	0.369	0.250	0.730
	Metaverse	0.629	0.506	0.571	0.655	0.751
	All	0.645	0.588	0.576	0.566	0.768
Random Forest	Games	0.826	0.661	0.628	0.597	0.875
	Art	0.836	0.863	0.890	0.918	0.888
	Other	0.823	0.496	0.425	0.371	0.845
	Collectible	0.780	0.654	0.652	0.650	0.823
	Utility	0.770	0.363	0.400	0.444	0.835
	Metaverse	0.775	0.805	0.777	0.751	0.865
	All	0.818	0.795	0.782	0.770	0.852
XGBoost	Games	0.701	0.588	0.467	0.387	0.844
	Art	0.766	0.770	0.835	0.912	0.884
	Other	0.656	0.561	0.301	0.205	0.851
	Collectible	0.740	0.541	0.544	0.546	0.855
	Utility	0.722	0.546	0.453	0.387	0.852
	Metaverse	0.740	0.731	0.733	0.734	0.843
	All	0.739	0.713	0.692	0.672	0.833
CNN	Games	0.841	0.588	0.622	0.660	0.865
	Art	0.852	0.905	0.903	0.901	0.915
	Other	0.888	0.357	0.456	0.631	0.886
	Collectible	0.812	0.621	0.654	0.691	0.891
	Utility	0.785	0.296	0.366	0.481	0.876
	Metaverse	0.769	0.794	0.770	0.748	0.843
	All	0.822	0.776	0.782	0.788	0.854

seen in the Table, similar to the primary sale price prediction, using all the attributes performs the best. However, the performance drops the most when we remove the visual embedding attributes.

Table 9: 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	AUC
XGBoost	Games	0.638	0.543	0.400	0.317	0.865
	Art	0.537	0.478	0.613	0.841	0.887
	Other	0.596	0.593	0.279	0.183	0.844
	Collectible	0.669	0.496	0.462	0.432	0.854
	Utility	0.512	0.727	0.386	0.262	0.871
	Metaverse	0.643	0.513	0.578	0.662	0.821
	All	0.645	0.588	0.576	0.566	0.798
CNN	Games	0.695	0.610	0.471	0.386	0.856
	Art	0.759	0.756	0.828	0.915	0.899
	Other	0.678	0.574	0.319	0.222	0.876
	Collectible	0.752	0.518	0.545	0.575	0.867
	Utility	0.708	0.614	0.470	0.380	0.856
	Metaverse	0.739	0.740	0.743	0.730	0.866
	All	0.741	0.713	0.693	0.675	0.854

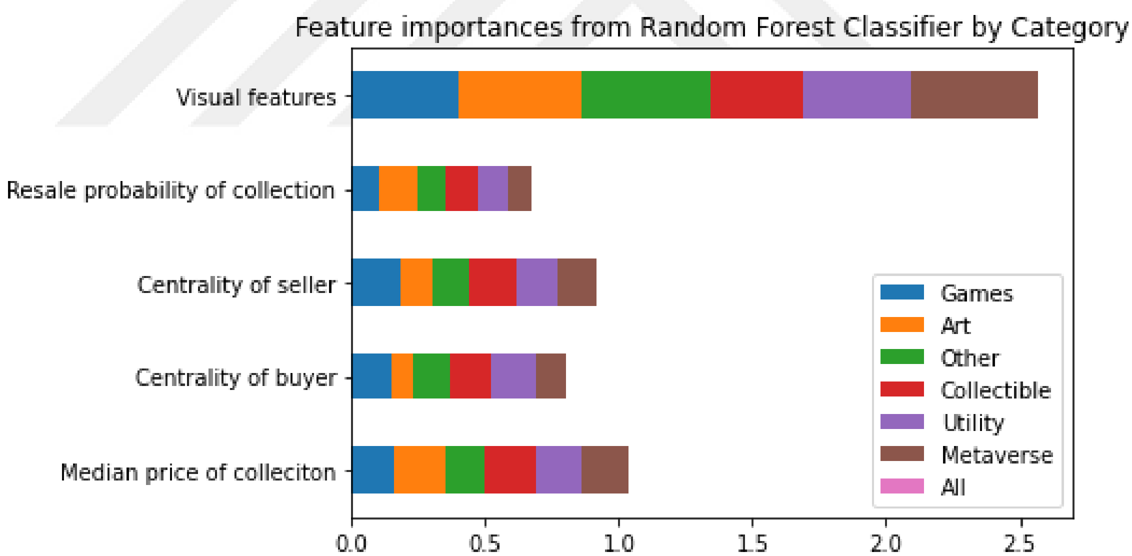

Figure 13: Feature Importance Levels

Table 10: Ablation studies on secondary sale participation prediction when we use random oversampling.

	LR	RF	XGBoost	CNN
W/o Embedding	0.634	0.633	0.651	**0.687**
W/o Collection Median Price	0.677	0.656	0.644	**0.777**
W/o Trader Centrality	0.701	0.653	0.632	**0.804**
W/o Resale Probability	0.713	0.666	0.647	**0.821**
All Attributes	0.741	0.693	0.675	**0.854**

CHAPTER V

RELATED WORK

5.1 *Machine Learning and Deep Learning Techniques for Cryptocurrency Predictions*

In [52], 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 [53] 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 [54], 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.

5.2 *NFTs and Their Financial Relationship with Bitcoin and Ethereum*

The authors in [55] 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.

5.3 NFT Dynamics and Price Prediction

Innovative and novel NFTs main characteristics and underlying protocol standards have been analyzed in [16]. Up to now, the most seminal NFT work has been carried out in [24] where a large collection of NFT images and associated attributes have been made publicly available in addition to their large-scale NFT analysis. In [24], 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, [56] 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 [21] 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 [23] 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 [22]. 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 [57, 23, 22], 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 [58, 59] have forecasted NFTs financial performance, such as NFTs primary sale price and secondary sale participation, via considering raw dataset of images and textual descriptions. [58] predicts NFT primary sale prices by deep learning, where the NFT category is an important factor in the model. On the other hand, [59] 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.

CHAPTER VI

CONCLUSION AND FUTURE WORK

6.1 Conclusion

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.

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.

REFERENCES

[1] T. Yilmaz, S. Sagfossen, and C. Velasco, "What makes nfts valuable to consumers? perceived value drivers associated with nfts liking, purchasing, and holding," *Journal of Business Research*, vol. 165, p. 114056, 2023.

[2] I. E. Okonkwo, "NFT, copyright and intellectual property commercialization," *International Journal of Law and Information Technology*, vol. 29, pp. 296–304, 11 2021.

[3] H. Taherdoost, "Non-fungible tokens (nft): A systematic review," *Information*, vol. 14, no. 1, 2023.

[4] C. Cachin and N. Chandran, "A secure cryptographic token interface," in *2009 22nd IEEE Computer Security Foundations Symposium*, pp. 141–153, 2009.

[5] B. White, A. Mahanti, and K. Passi, "Characterizing the opensea nft marketplace," in *Companion Proceedings of the Web Conference 2022*, WWW '22, (New York, NY, USA), p. 488–496, Association for Computing Machinery, 2022.

[6] Z. Wang, J. Gao, and X. Wei, "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*, WWW '23, (New York, NY, USA), p. 2099–2109, Association for Computing Machinery, 2023.

[7] K. Ito, K. Shibano, and G. Mogi, "Predicting the bubble of non-fungible tokens (nfts): An empirical investigation," 03 2022.

[8] N. H. Park, H. Kim, C. Lee, C. Yoon, S. Lee, Y. jin, and S. Shin, "A deep dive into nft whales: A longitudinal study of the nft trading ecosystem," 2023.

[9] N. Kshetri, "Scams, frauds, and crimes in the nonfungible token market," *Computer*, vol. 55, no. 4, pp. 60–64, 2022.

[10] J. Huang, N. He, K. Ma, J. Xiao, and H. Wang, "A deep dive into NFT rug pulls," *CoRR*, vol. abs/2305.06108, 2023.

[11] A. Yakovenko, "Solana: A new architecture for a high performance blockchain v0. 8.13," *Whitepaper*, 2018.

[12] M. Madine, K. Salah, R. Jayaraman, and J. Zemerly, "Nfts for open-source and commercial software licensing and royalties," *IEEE Access*, vol. 11, pp. 8734–8746, 2023.

[13] M. K. Hari, A. Agrawal, R. Bhatia, A. Bhatia, and K. Tiwari, "T-pass: A blockchain-based nft enabled property management and exchange system," in *2023 International Conference on Information Networking (ICOIN)*, pp. 140–145, 2023.

[14] G. Colavizza, "Seller-buyer networks in nft art are driven by preferential ties," *Frontiers in Blockchain*, vol. 5, p. 1073499, 01 2023.

[15] L. Du, M. Kim, and J. Lee, "The art nfts and their marketplaces," 2022.

[16] Q. Wang, R. Li, Q. Wang, and S. Chen, "Non-fungible token (nft): Overview, evaluation, opportunities and challenges," 05 2021.

[17] M. Westerkamp, F. Victor, and A. Küpper, "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, 2018.

[18] K. Sako, S. Matsuo, and S. Meier, "Fairness in erc token markets: A case study of cryptokitties," in *Financial Cryptography and Data Security. FC 2021 International Workshops* (M. Bernhard, A. Bracciali, L. Gudgeon, T. Haines, A. Klages-Mundt, S. Matsuo, D. Perez, M. Sala, and S. Werner, eds.), (Berlin, Heidelberg), pp. 595–610, Springer Berlin Heidelberg, 2021.

[19] A. Serada, T. Sihvonen, and J. T. Harviainen, "Cryptokitties and the new ludic economy: How blockchain introduces value, ownership, and scarcity in digital gaming," *Games and Culture*, vol. 16, no. 4, pp. 457–480, 2021.

[20] M. Dowling, "Fertile land: Pricing non-fungible tokens," *Finance Research Letters*, vol. 44, p. 102096, 2022.

[21] M. Dowling, "Is non-fungible token pricing driven by cryptocurrencies?," *Finance Research Letters*, vol. 44, p. 102097, 2022.

[22] A. Kapoor, D. Guhathakurta, M. Mathur, R. Yadav, M. Gupta, and P. Kumaraguru, "Tweetboost: Influence of social media on nft valuation," in *Companion Proceedings of the Web Conference 2022*, WWW '22, (New York, NY, USA), p. 621–629, Association for Computing Machinery, 2022.

[23] 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.

[24] 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.

[25] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (Y. Bengio and Y. LeCun, eds.), 2015.

[26] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," 2020.

[27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

[28] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, (New York, NY, USA), p. 785–794, Association for Computing Machinery, 2016.

[29] R. E. Schapire, "Explaining adaboost," in *Empirical inference*, pp. 37–52, Springer, 2013.

[30] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[31] K. O'Shea and R. Nash, "An introduction to convolutional neural networks," 2015.

[32] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: Synthetic minority over-sampling technique," *J. Artif. Int. Res.*, vol. 16, p. 321–357, jun 2002.

[33] D. R. Cox, "The regression analysis of binary sequences," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 20, no. 2, pp. 215–232, 1958.

[34] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.

[35] R. E. Schapire, "A brief introduction to boosting," in *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'99, (San Francisco, CA, USA), p. 1401–1406, Morgan Kaufmann Publishers Inc., 1999.

[36] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer, 2 ed., 2009.

[37] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

[38] T. Tuncer, U. Kaya, E. Sefer, O. Alacam, and T. Hoser, "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, (New York, NY, USA), p. 79–86, Association for Computing Machinery, 2022.

[39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (Y. Bengio and Y. LeCun, eds.), 2015.

[40] N. De La Rouviere, "A subgraph to index & explore cryptokitties auctions. (2021)." `https://thegraph.com/explorer/subgraph/nieldlr/cryptokitties-sales`, Accessed 4 May 2021.

[41] D. Rosenbaum, "Gods-unchained marketplace. (2021)." `https://github.com/djrosenbaum/unchained-transactions`, Accessed 4 May 2021.

[42] D. Team, "Decentraland marketplace. (2021)." `https://thegraph.com/explorer/subgraph/decentraland/marketplace`, Accessed 4 May 2021.

[43] O. Team, "Api overview. (2021)." `https://docs.opensea.io/reference`, Accessed 4 May 2021.

[44] N. Team, "The best place to analyze, track, and discover nfts. (2021)." `https://nonfungible.com/`, Accessed 4 May 2021.

[45] A. M. Team, "Atomic market api. (2021)." `https://wax.api.atomicassets.io/atomicmarket/docs/swagger/`, Accessed 4 May 2021.

[46] The ImageMagick Development Team, *Imagemagick*, 2021.

[47] K. P. F.R.S., "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*, vol. 2, no. 11, pp. 559–572, 1901.

[48] T. Wongvorachan, S. He, and O. Bulut, "A comparison of undersampling, oversampling, and smote methods for dealing with imbalanced classification in educational data mining," *Information*, vol. 14, p. 54, 01 2023.

[49] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, Curran Associates, Inc., 2019.

[50] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[51] H. Axelsen, J. R. Jensen, and O. Ross, "When is a dao decentralized?," *Complex Systems Informatics and Modeling Quarterly*, p. 51–75, July 2022.

[52] S. McNally, J. Roche, and S. Caton, "Predicting the price of bitcoin using machine learning," in *2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, p. 339–343, 2018.

[53] J. Hamilton, *Time series analysis.* Princeton, NJ: Princeton Univ. Press, 1994.

[54] J. Chevallier, D. Gu´egan, and S. Goutte, "Is it possible to forecast the price of bitcoin?," *Forecasting*, vol. 3, no. 2, p. 377–420, 2021.

[55] L. Ante, "The non-fungible token (nft) market and its relationship with bitcoin and ethereum," *FinTech*, vol. 1, no. 3, p. 216–224, 2022.

[56] K. Vasan, M. Janosov, and A.-L. Barabási, "Quantifying nft-driven networks in crypto art," *Scientific Reports*, vol. 12, no. 1, p. 2769, 2022.

[57] K.-H. Ho, Y. Hou, T.-T. Chan, and H. Pan, "Analysis of non-fungible token pricing factors with machine learning," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1161–1166, 2022.

[58] 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*, WWW '23, (New York, NY, USA), p. 1875–1885, Association for Computing Machinery, 2023.

[59] 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*, ICAIF '23, (New York, NY, USA), p. 116–123, Association for Computing Machinery, 2023.

VITA

Mustafa Pala graduated from Ege University in Electrical and Electronics Engineering in 2019. He began working as a certification and test engineer. He has been working as a backend developer for more than 2 years in the field of EV charging station management systems. Besides software development, he is interested in economics and he is a former NFT trader.