POLITECNICO DI TORINO

Master of Science's Degree in Data Science and Engineering

Master of Science's Degree Thesis

Supervised Contrastive Learning for Classification of Market Stock Series

Supervisors

Prof. Luca CAGLIERO

Prof. Jacopo FIOR

Candidate

Arcangelo FRIGIOLA

Academic Year 2023/24

Abstract

In the complex world of financial markets, the quest for innovative methods to analyze and interpret market behavior is ongoing. This thesis explores the potential of Supervised Contrastive Learning (SCL) as a novel approach to classifying stock market states, aiming to offer a more nuanced understanding of market dynamics. By constructing a dataset from NASDAQ 100 index prices, the study employs a deep neural network model to examine how SCL performs in comparison to traditional machine learning and deep learning techniques. This thesis attempts to contribute to the broader discourse on financial technology innovation, underlining the importance of continued exploration and experimentation in the development of financial analysis tools.

The focus is on the process and methodology of applying SCL to financial time series analysis, emphasizing the exploratory nature of this research in seeking new pathways for financial analysis. The research provides valuable insights into the applicability of SCL in financial markets, suggesting directions for future work in enhancing the accuracy and efficiency of market state classification. The task involves categorizing financial time series into three main trends: buy, hold, and sell, over different future time frames—specifically 3, 5, and 7 days following the targeted period. In addition to the empirical evaluation of SCL against traditional and deep learning models, this thesis embarks on a qualitative analysis of the latent representations generated by the SCL model, compared with other models, such as TS2Vec. This analysis seeks to uncover whether these representations can illuminate significant patterns within the financial time series data, offering insights into the underlying mechanisms of market behavior. The ultimate aim is to showcase how the Supervised Contrastive Learning (SCL) approach can be effectively applied to forecast financial time series.

Table of Contents

List of Tables

List of Figures

Acronyms

AI
: Artificial Intelligence

ANN
: Artificial Neural Network

CNN
: Convolutional Neural Network

DL
: Deep Learning

GNN
: Graph Neural Networks

GRU
: Gated Recurrent Units

LSTN
: Long Short-Term Memory

ML
: Machine Learning

NN
: Neural Network

NLP
: Natural Language Processing

RL

Reinforcement Learning

ResNet

Residual Networks

SupCon

Supervised Contrastive

SCL

Supervised Contrastive Learning

Chapter 1

Introduction

In today's digital age, the term "Big Data" has become synonymous with the vast expanse of information generated every second across various sectors. One of the most common forms of this data is time series data, which is represented sequence of information measured at successive steps in time. Time series data is prevalent in many fields including finance, healthcare, engineering, and more. The analysis of time series data is crucial for uncovering underlying patterns and trends which, in turn, are instrumental in making informed decisions.

Machine Learning (ML), a subset of artificial intelligence (AI), has proven to be a powerful tool in extracting meaningful insights from big data. Coined by Arthur Samuel while working for IBM in 1959, the ML term describe the pattern recognition tasks that delivered the "learning" component on the pioneering systems.[1] ML algorithms can learn from this vast amount of information and improve over time, making them highly effective for various tasks, like forecasting, clustering, or classification. However, traditional ML approaches often require manual feature engineering and may fall short in capturing temporal dependencies present in time series data.

This challenge has led to the rise of Deep Learning (DL), an advanced subset of ML that excels in identifying complex patterns in large datasets without explicit programming. DL is particularly adept at processing time series data, thanks to neural networks and, more recently, transformer models, which can learn directly from raw data.

One of the most impacted fields by DL is certainly finance. In the financial domain, the analysis of time series data is quintessential. Financial time series data, such as stock prices, trading volumes, and other related metrics, are crucial for a myriad of financial tasks including but not limited to forecasting, risk management, and algorithmic trading. The ability to accurately classify or predict financial time series data can lead to more informed and timely decision-making, which is a critical asset in the fast-paced and often volatile financial markets.

Before the advent of artificial intelligence (AI) and machine learning (ML), classification problems in financial markets were often approached using statistical methods. One such method is logistic regression, which has been widely used in finance for binary classification problems, like predicting whether a stock price will go up or down based on historical data. With the advent of AI, the approach towards classification problems in financial markets has evolved significantly. The emerging use of deep algorithms within financial systems is said to be disrupting and transforming industries and societies [2]. The financial services industry, in particular, has entered what's described as the AI phase of the digital marathon, where companies have transitioned from core systems modernization to intelligent automation enabled by AI.

The evolution of AI techniques in finance, propelled by the increasing computational power and abundance of available data, has allowed for more sophisticated analysis and decision-making processes. This transition has also brought about a shift from traditional statistical methods to more advanced ML techniques, including deep learning, which is particularly adept at handling time series data. As you delve deeper into the domain of finance, you could highlight the critical importance of being able to classify financial time series data accurately, especially when it comes to predicting market states like bearish or bullish trends which are essential for informed decision-making in the stock market. Deep Learning has shown promise in tackling financial time series data. By employing sophisticated neural network architectures, DL can help in deriving meaningful patterns from financial time series data, aiding in the classification or prediction of stock market states, which is essential for guiding investment decisions.

The problem of classifying stock market states into bearish (downtrending) or bullish (uptrending) states based on time series data poses a significant challenge yet holds immense value. Accurate classification can provide actionable insights for investors, traders, and financial analysts, aiding in the formulation of investment strategies and risk management practices. This thesis aims to explore the application of supervised contrastive learning techniques for time series classification, with a primary focus on classifying market stock states to aid in financial decision-making.

This thesis explores the application of supervised contrastive learning, a novel technique in ML, for the classification of financial time series data, specifically aiming to differentiate between uptrending (bullish) and downtrending (bearish) market states. Such classifications are invaluable for investors, traders, and analysts, offering insights that can guide strategic investment decisions and risk management. Focus of this work, supervised contrastive learning enhances the ability of models to understand and differentiate between data points by learning from contrasts or comparisons. Unlike traditional supervised learning, which focuses on matching inputs directly to labels, supervised contrastive learning works by comparing how similar or different data points are to each other, even within the same class. This approach is particularly useful in tasks where understanding the nuanced differences and similarities between data points can lead to more accurate classifications. In extension to the focus on SCL and deep learning techniques, this thesis incorporates a comprehensive discussion on the broader field of representation learning from time series data. It delves into various techniques developed for extracting meaningful information from time series, emphasizing their relevance in financial analysis. This section aims to bridge the gap between traditional time series analysis methods and the cutting-edge approaches facilitated by machine learning.

By integrating these aspects, the thesis aims to offer a holistic view of the current state and future potential of machine learning applications in financial time series analysis. It contributes to the academic discourse by providing a detailed examination of SCL's role in financial decision-making, alongside a critical comparison with unsupervised learning models.

Structured to provide a comprehensive overview of deep learning and supervised contrastive learning techniques, this thesis delves into their application in financial time series analysis. It begins with the development of a model based on deep neural networks, designed to efficiently encode time series data. The efficacy of this model is initially tested against a benchmark dataset to validate the approach. Subsequent analysis focuses on real-world data, employing a novel dataset compiled from NASDAQ 100 index stock prices [3]. The final goal is to demonstrate the practical value of the proposed model through rigorous experimentation and analysis.

By contributing to the field of financial time series classification through advanced ML techniques, this work aims to enhance financial decision-making processes, offering a valuable resource for both academic research and practical applications in the financial industry. Through rigorous experimentation and analysis, this work aspires to contribute to the growing body of knowledge in applying advanced machine learning techniques to financial time series classification, with a vision to enhance decision-making processes in the financial domain. Furthermore, this thesis extends its exploration by conducting a qualitative analysis of the latent representations generated through supervised contrastive learning (SCL). This analysis aims to discern whether these representations can unveil significant patterns within the analyzed financial time series data, thereby contributing to a deeper understanding of market dynamics. By comparing these representations with those derived from unsupervised models, such as TS2Vec, this work seeks to highlight the nuances and potential advantages of SCL in capturing the intricate details of financial time series data.

In the course of this research, the model developed demonstrated promising performance on a controlled dataset. This achievement underscores the potential of supervised contrastive learning as a method for distinguishing between different representations under simplified conditions.

However, when transitioning from the controlled environment to the complexity and unpredictability of real-world financial data, the model faced significant challenges. The results on real-world data did not meet the expectations set by its previous performance. This discrepancy highlights the intricate and often unpredictable nature of financial markets, where numerous variables and external factors can influence outcomes in ways that are difficult to replicate in a controlled setting.

It's important to view these findings within the larger narrative of scientific progress, where every outcome, regardless of its immediate applicability, contributes to our understanding and sparks further inquiry. The exploration of supervised contrastive learning in this context has shed light on the limitations and challenges of applying new machine learning techniques to financial time series data. This insight is invaluable, as it delineates areas for future research, such as refining the model, exploring additional data preprocessing techniques, or integrating domain-specific knowledge into the learning process.

In conclusion, while the model did not achieve the anticipated results in real-world data application, this research contributes to the evolving dialogue on machine learning applications in finance. It lays the groundwork for future studies to build upon, guiding them toward areas ripe for discovery and improvement.

Chapter 2

AI, Machine Learning and Deep Learning Foundations

2.1 Artificial Intelligence

Artificial Intelligence (AI) is a branch of computer science that aims to create machines capable of intelligent behavior. It seeks to develop algorithms, models, and techniques that enable computers to perform tasks that typically require human intelligence. These tasks include problem-solving, speech recognition, planning, learning, perception, reasoning, and the ability to move and manipulate objects.

The concept of AI has been around for centuries, but it was formally introduced in the 1950s. As defined by John McCarty, AI is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.[4]

Contrary to what people might think, artificial intelligence (AI) is hardly a new topic. It has been around since 1956 when the seminal summer workshop was organized at Dartmouth College, New Hampshire, US.

For long time, AI remained a simple field of research in universities, or an inspiration for science-fiction writers. Nevertheless, thanks to the acceleration in new technologies, both hardware and software, AI began to cover a pivotal role in real-world applications, including business context. [5] The raising of AI, and even more of DL, is due to three important aspects:

- Data availability: the new era of so called "Big Data" made the rising of these technologies possible. Everything around us produces data, from our smartwatch, to out television, from our domotic-house devices, to our smartphones.
- Computing power: recently, GPUs have significantly improved cost-efficiency

in computing, prompting a shift in knowledge acquisition. As the demand for advanced GPUs rises, they are expected to address computational challenges in their own design. While Moore's Law continues its exponential trajectory, advancements in GPU technology will likely represent the supporting data points for this trend.

- Algorithm efficiency: the emergence of generative AI and tools like Midjourney or ChatGPT has showcased the remarkable capabilities of these tools to a wider audience, making unprecedented abilities readily accessible to everyone.

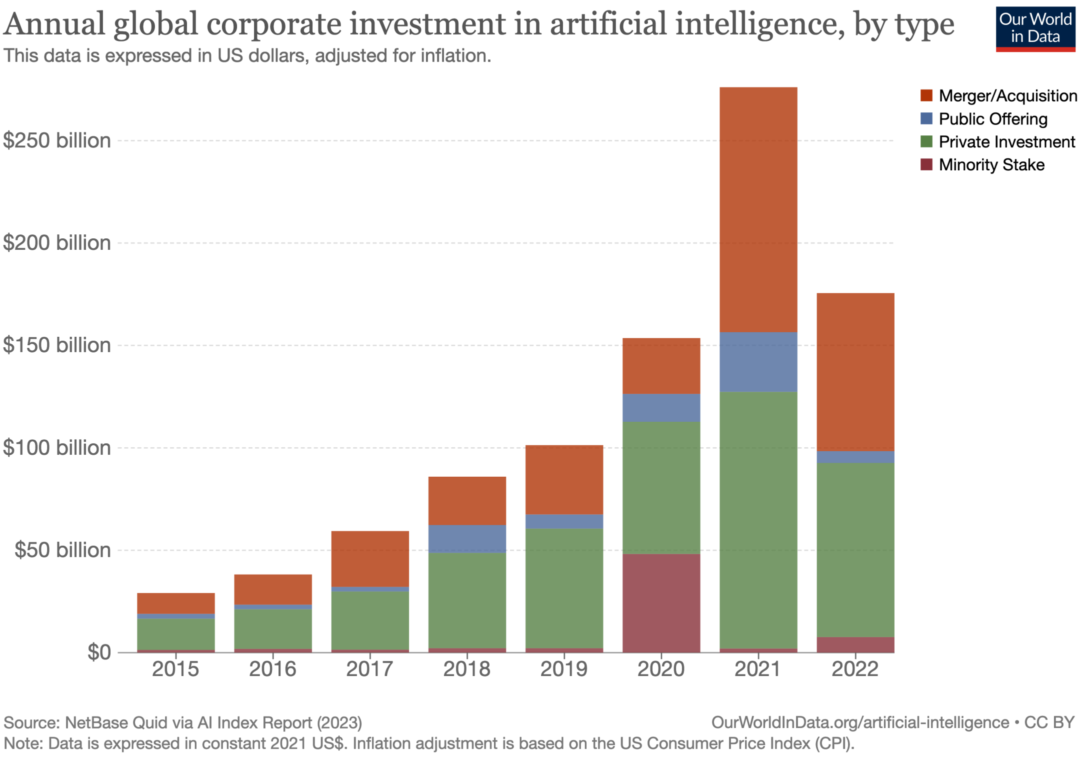

Figure 2.1: Annual global corporate investment in AI, by type.[6]

2.1.1 Artificial Intelligence in Finance

Consequences of this rapid growth, venture-capital (VC) investments in artificial-intelligence startups and companies have increased sharply in recent years, from less than \$50 billions to over \$250 billions in 2021, as shown in Fig.2.1.

In the realm of finance, the application of artificial intelligence (AI) spans across five primary sectors[5]:

- **Investment and Asset Management**: Algorithms can sift through a plethora of data to identify correlations between global events and their repercussions on asset values. They can also harness information from public social media channels to predict market trends.

- **Credit Analysis and Underwriting**: Machine learning can bolster the accuracy of credit underwriting decisions by lenders. Additionally, advanced computer vision technology, leveraging geospatial and aerial imagery, proves instrumental in insurance or property underwriting processes.

- **Regulatory Adherence and Fraud Detection**: By employing advanced pattern-recognition analytics, various data channels and types can be scrutinized for any fraudulent activities. In contemporary anti-money laundering procedures, automated scans of both inbound and outbound transactions are conducted based on preset criteria like country of origin/destination, customer name, etc. The prevailing systems tend to generate numerous false positives, which are then individually controlled by middle-office or compliance personnel. Machine learning can refine this process by identifying users for whitelisting, discerning patterns for rule engine inclusion, and significantly reducing false positive occurrences. This not only curtails costs but also enhances the quality of the screening process.

- **Market Research and Reporting**: Intelligent agents are capable of curating and semantically indexing financial market research content. They can also automate the creation of reports, personalized websites, emails, articles, and more utilizing natural language generation software. This work aims at producing a tools which can guide decision making, through classification of market states. This goals fits the model produced both in this sectors as much as in the first cited above.

- **Customer Support and Assistance**: Intelligent agents can analyze incoming communications, streamline case routing, provide precise suggestions to customer service representatives, or aid in optimizing personal finance management.

2.2 Machine Learning

We now further explore more in detail how AI works, and what is beneath the surface. Key concepts of AI include Machine Learning (ML), a subset field that focuses on the development of algorithms that can learn from and perform predictive or other kinds of analytics based on data. ML provides a robust framework to delve into complex data-driven tasks, paving the way for deeper insights and real-world solutions in various domains, including finance. Machine learning models fall into three primary categories: Supervised, Unsupervised, and Semi-Supervised learning.[7] We can further include an extension of supervised learning type, which is Reinforcement Learning.

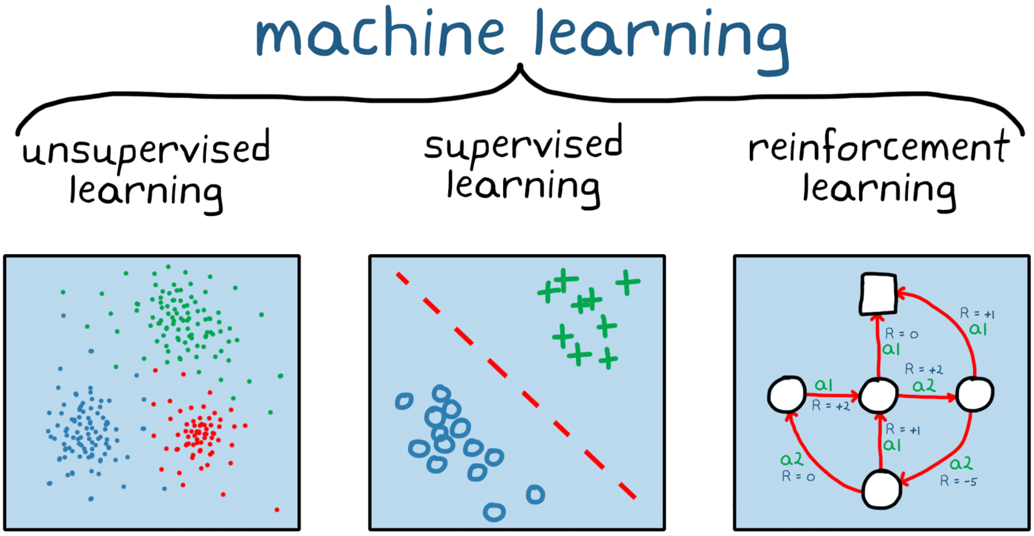

Figure 2.2: Supervised, Unsupervised and Reinforcement Learning representation.[8]

2.2.1 Supervised Learning

Supervised learning is a type of learning paradigm where the model is trained on labeled data. The training dataset includes input data along with corresponding correct labels. Through iterative learning, the algorithm analyzes the training data and learns a function f such that $y = f(X)$, where X is the data input and y is the output label. It aims to minimize the error in predicting the labels of the training data and eventually generalizes well to unseen data. Common supervised learning algorithms include Artificial Neural Networks (ANNs), Logistic Regression, Linear Regression, Support Vector Machine (SVM), and Random Forest. Applications span across various domains including image recognition, fraud detection, financial forecasting, predictive maintenance, and medical diagnosis.

Self-Supervised Learning

Self-supervised Learning is a hybrid approach aiming to mitigate the limitations of supervised learning, particularly the need for extensive labeled data. In this paradigm, the algorithm generates its own supervisory signal from the data itself, thus creating a pretext task that it can learn from. This learning paradigm consists of two phases:

- *Feature Extraction*: Initially, task-independent features are extracted from the data in an unsupervised manner.
- *Task-Specific Learning*: The extracted features are then utilized for the main task, transferring the knowledge gained from the first phase.

Self-supervised learning finds its applications predominantly in natural language processing (e.g., Word2Vec, GloVE, fastText, BERT), image processing, and time series forecasting. The methods within self-supervised learning can be broadly categorized into:

- *Generative Methods*: Aimed at learning the data distribution to generate new data points.
- *Adversarial Methods*: Employ game theory to train models in a competitive manner.
- *Contrastive Methods*: Learn representations by contrasting positive and negative examples.

Reinforcement Learning

Reinforcement Learning (RL) is a type of machine learning, similar to supervised learning, where an agent learns how to behave in an environment by performing actions and observing the consequences of those actions. Unlike supervised learning, where the correct answers are explicitly provided, in reinforcement learning, the agent receives feedback in the form of rewards or penalties, which guide the learning process.

The learning process in RL consists of the following components:

- Agent: The decision-maker in the system.
- Environment: The external system with which the agent interacts.
- State: A specific situation or configuration the environment can be in.
- Action: Operations that the agent can perform.
- Reward: Immediate feedback received post-action to indicate the benefit of the action.

The agent interacts with the environment by executing actions, transitioning between states, and receiving rewards. The goal is to find a policy, a strategy for choosing actions over time, that maximizes the cumulative reward, often referred to as the reward signal.

The process is iterative, with the agent continually updating its understanding of the environment and improving its policy as it gathers more experience. The exploration vs exploitation dilemma is a core challenge in RL, where the agent needs to balance the exploration of unknown, potentially rewarding actions, and the exploitation of known, rewarding actions.

Reinforcement Learning is used in various fields including robotics, game playing, natural language processing, and finance among others.

Reinforcement Learning's strength lies in solving complex, interactive, and uncertain problems, making it a crucial paradigm within machine learning and AI.

2.2.2 Unsupervised Learning

Unsupervised Learning, on the other hand, works with datasets without labels. The objective here is to discover inherent patterns and structures within the data. The primary tasks within unsupervised learning include:

- *Clustering*: Grouping data points based on similarity.
- *Association*: Identifying relationships and rules among data objects.
- *Anomaly Detection*: Recognizing outliers or unusual data points.
- *Dimensionality Reduction*: Like autoencoders, reducing the number of variables under consideration to discover a simpler structure in the data.

Unsupervised learning facilitates the analysis of data in scenarios where labeled data is scarce or unavailable. Semi-supervised Learning

2.2.3 Semi-Supervised Learning

Semi-supervised Learning strikes a balance between supervised and unsupervised learning by utilizing a small amount of labeled data to guide the learning process on a larger set of unlabeled data. This paradigm is beneficial in situations where obtaining labeled data is expensive or time-consuming, like in medical image analysis.

2.3 Deep Learning

Deep Learning (DL) is a further subset of ML, involving algorithms inspired by the structure and function of the brain called artificial neural networks. Deep learning algorithms are excellent for handling large amounts of data and identifying patterns, making them crucial for various AI applications. Natural Language Processing (NLP) is the ability of a computer program to understand human language as it is spoken, and is used in chatbots, translators, and personal assistants like Siri and Alexa. Robotics is a field of engineering focused on the design and manufacturing of robots, aiming to develop machines that can substitute for humans, especially in hazardous or repetitive tasks.

2.3.1 Neural Networks

Deep learning utilizes multi-layer neural networks to analyze data. The architecture comprises input, hidden (one or more), and output layers. Each layer consists of nodes (neurons) connected by weighted pathways. As data propagates through the network, each layer processes an aspect of the data, refining the input for the subsequent layer. One of the key aspects of NN architectures and algorithms is *backpropagation*. This is a vital step in DL which adjusts the weights of the connections in the network to minimize the error between the predicted and actual outcomes. Through multiple iterations, the network learns and improves its performance. [9]

Deep learning can be applied in supervised, unsupervised, and self-supervised learning contexts, adapting to the data availability and task requirements. It can handle labeled, partially labeled, or unlabeled data, making it a versatile tool for various domains.

Some DL application domains are:

- *Computer Vision*: Deep learning excels in image and video recognition tasks, enabling applications like facial recognition, object detection, and autonomous vehicles.
- *Natural Language Processing (NLP)*: It powers language translation, chatbots, and sentiment analysis by understanding and processing human language.
- *Audio Recognition*: It's used in voice-activated assistants, speech-to-text systems, and other audio recognition tasks.

TensorFlow and PyTorch: These are among the most popular frameworks for developing deep learning models due to their robustness and community support. For the sake of this work, the framework use is the latter.

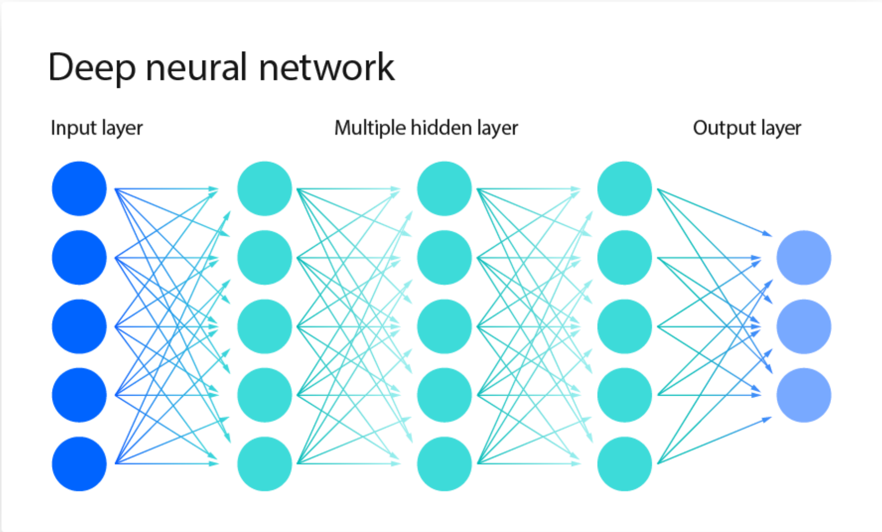

Figure 2.3: Neural Network Design. [10]

2.3.2 Residual Networks

The work presented here utilizes a special type of NN, born to encode images information, but soon applied for wider contexts. Residual Networks, or ResNet, paved the way for better performance in deep learning tasks by tackling the vanishing gradient problem which hampers the training of deep networks. The notable introduction of "skip connections" allows certain layers to be skipped during the forward pass of the network, which facilitates the training of very deep networks by learning residual functions with respect to the layer inputs. [11]

In the realm of finance, where time series data is a fundamental piece, ResNet has found its footing. Analyzing financial time series data, like stock prices or trading volumes, requires a keen understanding of the temporal dependencies present in the data. Deep learning models like ResNet provide a way to capture these dependencies and unveil the underlying patterns that govern market dynamics.

The architecture of ResNet, with its various configurations, has been instrumental in tackling challenges tied to financial time series analysis. The skip connections in ResNet allow for the efficient backpropagation of gradients even in very deep networks, making it a suitable architecture for handling the sequential nature of financial data.

Moreover, the ResNet architecture has spurred further innovation. For instance, its principles have been adapted to create new architectures tailored for time series data, enhancing the capability of models to better understand financial markets and provide more accurate forecasts.

The foray of ResNet into the financial domain highlights the flexibility and power of deep learning architectures in tackling domain-specific challenges. By delving into the temporal intricacies of financial data, ResNet and its derivatives are aiding financial analysts, traders, and decision-makers in making more informed and data-driven decisions.

Chapter 3

Contrastive Learning Techniques

3.1 Introduction

The domain of machine learning has continually evolved with the underlying goal of enabling machines to learn from data and make predictions or decisions without being explicitly programmed. One significant avenue through which this objective has been pursued is the development of algorithms capable of learning meaningful representations from data, which forms the crux of representation learning. Among the various approaches toward representation learning, *contrastive learning* has emerged as a compelling paradigm, especially in the realm of unsupervised and self-supervised learning.

Historically, the roots of contrastive learning can be traced back to the idea of learning by comparison. The seminal work by Hadsell et al. in 2006 laid down the foundation of this concept by introducing a contrastive loss function aimed at learning a metric space where "neighbors are pulled together and non-neighbors are pushed apart.[12] This early endeavor demonstrated the potential of learning representations by harnessing the power of contrast, setting the stage for subsequent advancements in this domain. As the field matured, numerous variations and enhancements to the original idea of contrastive learning were proposed. The introduction of triplet loss by Schroff et al. in 2015 [13] further propelled the development of contrastive learning techniques by extending the notion of pairwise comparisons to triplets, comprising an anchor, a positive sample, and a negative sample.

In recent times, the principles of contrastive learning have been extended to the supervised setting, giving rise to Supervised Contrastive Learning (SCL). Unlike traditional contrastive learning, which operates in an unsupervised or self-supervised manner, Supervised Contrastive Learning leverages label information to guide the learning of representations.

3.2 Contrastive Learning principles

Contrastive learning emerges as a beacon within the representation learning landscape, particularly under unsupervised or self-supervised settings. Its foundational idea lies in learning robust representations by contrasting positive pairs (instances that are similar or related) against negative pairs (instances that are dissimilar or unrelated). This paradigm facilitates the learning of a feature space where similar instances are mapped close together, while dissimilar instances are mapped far apart.

3.2.1 Loss Function

Central to the mechanism of contrastive learning are contrastive loss functions, which quantify the similarity and dissimilarity among instances. The seminal work by Hadsell et al.[12] introduced a contrastive loss function that aims to minimize the distance between similar instances while maximizing the distance between dissimilar instances. This loss function is defined as:

$$\mathcal{L}(\mathbf{x}_i, \mathbf{x}_j, y) = (1 - y) \cdot \frac{1}{2} D(\mathbf{x}_i, \mathbf{x}_j)^2 + y \cdot \frac{1}{2} \max(0, m - D(\mathbf{x}_i, \mathbf{x}_j))^2 \tag{3.1}$$

where:

- $x_i \cdot x_i$ and $x_i \cdot x_j$ are instances,
- y is a binary label indicating whether the instances are similar (0) or dissimilar (1),
- D is a distance metric,
- and m is a margin that specifies a radius within which similar instances are pulled together.

Triplet-loss

Following this foundational work, the concept of triplet loss was introduced by Schroff et al.[13], which extends the pairwise comparison to triplets consisting of an anchor, a positive sample, and a negative sample. The triplet loss aims to ensure that the anchor is closer to the positive sample than to the negative sample by at least a margin m, and is defined as:

$$\mathcal{L}(\mathbf{x}_a, \mathbf{x}_p, \mathbf{x}_n) = \max(0, D(\mathbf{x}_a, \mathbf{x}_p) - D(\mathbf{x}_a, \mathbf{x}_n) + m) \tag{3.2}$$

where:

- x_a, x_p, and x_n are, respectively, the anchor, the positive, and the negative instances,
- D is a distance metric,
- and m is the margin.

As shown in Fig.3.1, the goal is to bring the anchor and the positive, which share the same label or are augmentations of each other, closer in the feature space, and push the anchor and the negative farther apart. This way, over time, the algorithm learns to place similar items close together and dissimilar items far apart in the feature space, aiding in tasks like classification or retrieval.

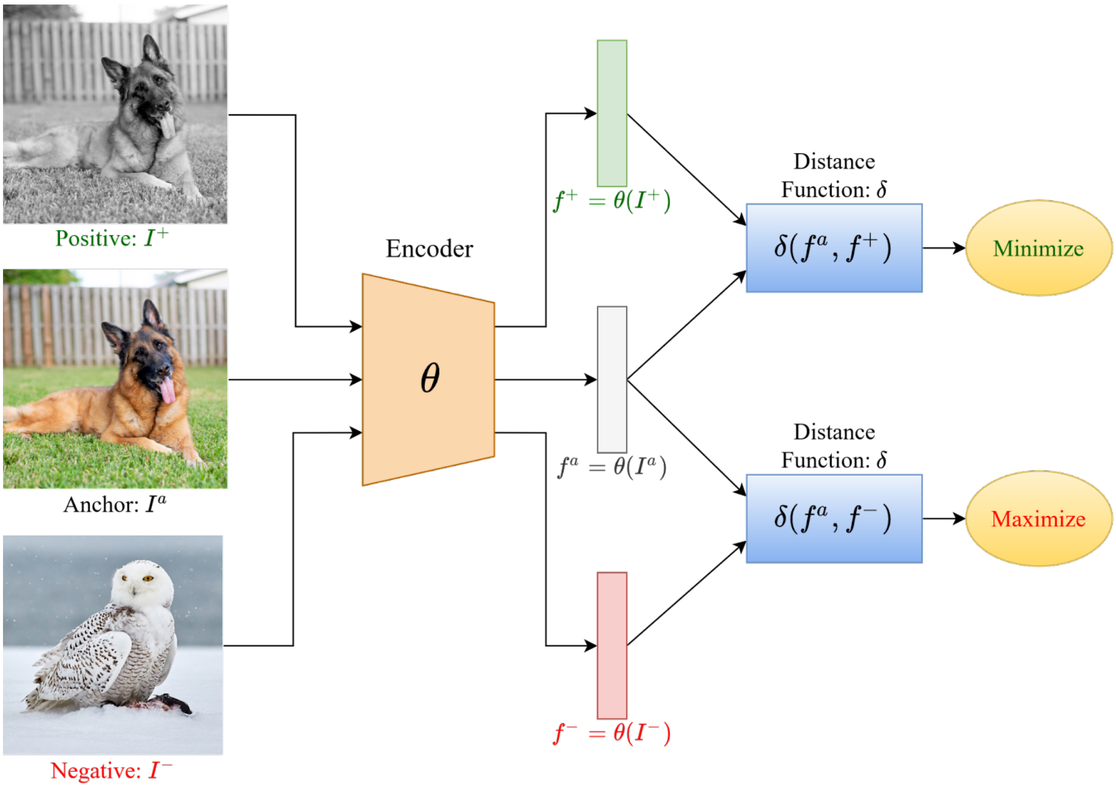

Figure 3.1: Representation of Contrastive Learining [14]

Embedding Space and Similarity Metrics

The effectiveness of contrastive learning is heavily contingent on the choice of the embedding space and the similarity metric. The objective is to learn an embedding space where the contrastive loss is minimized, leading to a meaningful separation of similar and dissimilar instances. Common similarity metrics include Euclidean distance and cosine similarity, each with its own set of advantages and considerations.

The construction of *positive* and *negative* pairs is a crucial aspect of contrastive learning. Various strategies can be employed to form these pairs, such as random sampling, hard negative mining, or data augmentation. Data augmentation, in particular, has shown to be effective in generating varied yet semantically consistent views of the data, thereby enriching the set of positive pairs and facilitating better representation learning.

3.2.2 Data Augmentation

Data augmentation techniques encompass a diverse array and must be judiciously selected to introduce varying perspectives of a particular instance without sacrificing or distorting the inherent information. The efficacy and appropriateness of these techniques can be contingent on the domain of application. For instance, in the realm of instance discrimination-based contrastive learning for image data, prevalent augmentation methods include:

- **Colour Jittering**: This technique entails random alterations in the brightness, contrast, and saturation of an RGB image. It aids in ensuring that a model does not overly rely on color attributes for object recognition. While the resultant image colors might appear aberrant to human observers, such augmentations prompt the model to discern object edges and shapes beyond mere color characteristics.

- **Image Rotation**: Images are subjected to random rotations within a specified range, typically 0-90 degrees. Given that rotation preserves the fundamental information within an image (e.g., a depicted dog remains identifiable as a dog), this technique trains models to exhibit rotation invariance, thereby enhancing prediction robustness.

- **Image Noising**: Random noise is integrated into the images on a pixel-wise basis, challenging the model to distinguish signal from noise, and hence bolstering its robustness against image alterations during testing phases. For instance, the introduction of salt-and-pepper noise by randomly toggling some pixels to white or black exemplifies this technique.

Figure 3.2: Image data augmentation [15]

In the context of time series data, the intricacies of augmentation escalate, necessitating domain-aware techniques to prevent information loss or distortion within the time series stream. Several time series data augmentation strategies exist, such as:

- **Jittering**: Adding small random noise to the time series data to enhance the model's robustness.
- **Scaling**: Multiplying the time series data by a random scaling factor to learn scale-invariant representations.
- **Slicing**: Extracting sub-sequences from the original time series to expand the dataset and enable the model to learn from shorter sequences.
- **Permutation**: Rearranging chunks of the time series data to create new sequences while retaining the inherent dynamics.
- **Time Warping**: Altering the speed of certain sections of the time series to simulate variations in the temporal dynamics.
- **Window Slicing**: Generating multiple shorter sub-sequences from a longer time series to increase the diversity of the training data.

These augmentation techniques strive to diversify the training dataset and furnish the model with a broader understanding of the underlying patterns and variations within the time series data, all while preserving the essential temporal characteristics intrinsic to the domain.

3.3 Supervised and Self-Supervised Contrastive Learning

3.3.1 Supervised Constrastive Learning

Supervised Contrastive Learning (SCL) is a learning paradigm that leverages label information to guide the learning of representations. Unlike traditional supervised learning which directly optimizes for a discriminative task, Supervised Contrastive Learning seeks to learn representations where similar instances (instances of the same class) are pulled together, and dissimilar instances (instances of different classes) are pushed apart in the embedding space. This is accomplished by employing a contrastive loss function which is minimized when similar instances are close and dissimilar instances are far apart in the learned feature space.

The core advantage of Supervised Contrastive Learning is that it can learn more informative representations by utilizing the available label information, which can subsequently be used for a variety of tasks. Moreover, it can lead to better generalization and performance on downstream tasks.

3.3.2 Self-Supervised Constrastive Learning

Self-Supervised Contrastive Learning, on the other hand, does not rely on external labels but rather constructs its own supervisory signal from the data. A typical approach in self-supervised contrastive learning is to create two augmented views of each data instance and treat them as a positive pair, while all other instances in the dataset are treated as negative examples. The objective then is to learn representations such that the augmented views of the same instance are brought closer together, while being pushed away from representations of other instances.

The primary allure of Self-Supervised Contrastive Learning is its ability to learn useful representations from unlabeled data, thereby reducing the dependency on large labeled datasets. This is particularly beneficial in domains where labeled data is scarce or expensive to obtain.

Both supervised and self-supervised contrastive learning approaches leverage the idea of learning by comparing similar and dissimilar instances. However, while supervised contrastive learning benefits from the explicit guidance of labels, self-supervised contrastive learning exploits the data itself to generate supervisory signals. In the context of this work, we will overcome a typical lack of the financial domaain, which is missing labels, taking advantage of SCL and auto-labeling the time series windows from the stock market. The labeling method will be further explained in the Methodology section.

Chapter 4

Analysis and Classification of Financial Time Series

4.1 Introduction

Financial Time Series analysis is a domain that unveils a plethora of techniques to scrutinize financial market data, inherently temporal, aiming to prognosticate future market behaviors. This domain is a linchpin for a multitude of stakeholders like investors, financial analysts, traders, and policymakers, offering a prism through which the dynamism and complexities of financial markets can be dissected and comprehended. The meticulous analysis of financial data over time not only paves the way for informed decision-making but also unveils underlying market trends and potential investment opportunities.

One of the seminal approaches within this domain is *Technical Analysis*, defined as *the study of market action, primarily through the use of charts, with the objective of forecasting future price trends* [16]. The term "market action" encapsulates three principal sources of information available to the technician: price, volume, and open interest, the latter being pertinent only in futures and options. Although the term "price action" is often utilized, it seems too narrow as most technicians integrate volume and open interest as pivotal parts of their market analysis. With this distinction delineated, the terms "price action" and "market action" are used interchangeably throughout the discourse.

The philosophy of Technical Analysis, as reported by John J. Murphy in his assay "Technical Analysis of the Financial Markets" [16], is based on three fundamental premises:

- Market action discounts everything,
- Prices move in trends,
- History repeats itself.

The cornerstone of Technical Analysis is the axiom, "market action discounts everything." This axiom posits that every conceivable factor that could impact the price - be it fundamental, political, psychological, or otherwise - is invariably reflected in the market price. This implies that a thorough study of price action is quintessential. The logic behind this premise grows more compelling with accruing market experience. It's stated that if all factors affecting market price are ultimately mirrored in the market price, then a study of market price is all that's requisite. By analyzing price charts and a plethora of supporting technical indicators, the chartist essentially allows the market to elucidate the most probable direction it's likely to take.

The concept of trend is indispensable to the technical approach. The primary objective of charting the price action of a market is to identify trends in their nascent stages for the purpose of trading in the direction of those trends. A significant segment of technical analysis and the study of market action is dedicated to understanding human psychology. Chart patterns, which have been recognized and categorized over the past century, manifest certain pictorial representations on price charts, reflecting the bullish or bearish psychology of the market. Given that these patterns have demonstrated efficacy in the past, it's conjectured that they will continue to be efficacious in the future. They are predicated on the study of human psychology, which tends to remain constant over time. This last premise, that history repeats itself, underscores the importance of studying the past, postulating that the future is but a repetition of the past.

4.2 Financial Time Series Analysis

Financial Time Series Analysis endeavors to unravel the complex dynamics of financial markets through a temporal lens. The primary objective is to forecast future price trends based on historical market data. This domain is broadly bifurcated into technical and fundamental analysis, each with its unique set of principles and methodologies.

4.2.1 Technical Analysis

Technical analysis is a methodology employed to evaluate investments and identify trading opportunities by analyzing statistical trends gathered from trading activity, such as price movement and volume. Unlike fundamental analysts who attempt to evaluate a security's intrinsic value, technical analysts focus on charts of price movement and various analytical tools to evaluate a security's strength or weakness and forecast future price changes.

John J. Murphy, in [16], delineates technical analysis as the study of market action, primarily through the use of charts, with the aim of forecasting future price trends. The term "market action" encompasses three principal sources of information available to the technician: price, volume, and open interest (pertinent only in futures and options). The belief is that all possible factors affecting prices, whether they are fundamental, political, psychological, or otherwise, are already reflected in the market price, hence, the focus on price action.

4.2.2 Fundamental Analysis

On the other end of the spectrum, fundamental analysis delves into the intrinsic value of financial instruments by meticulously examining a multitude of economic, financial, and other qualitative and quantitative factors. This analysis is rooted in the examination of the economic health, performance metrics of companies, and broader macroeconomic indicators. Fundamental analysts strive to ascertain whether a security is overvalued or undervalued, providing insights into potential investment opportunities.

Key components of fundamental analysis include the scrutiny of financial statements, management and industry analysis, and the macroeconomic environment. The objective is to garner a profound understanding of the underlying factors driving the market, thereby aiding investors in making informed decisions.

4.2.3 Technical Versus Fundamental Analysis

While technical analysis zeroes in on the study of market action, fundamental analysis is anchored on the economic forces of supply and demand that drive prices to ascend, descend, or remain static. The fundamental approach delves into all relevant factors affecting the price of a market to discern its intrinsic value based on the law of supply and demand. If the intrinsic value is below the current market price, the market is deemed overpriced and should be sold; if above, it's undervalued and should be bought.

Both these paradigms of market forecasting endeavor to resolve the quintessential problem of determining the likely direction of price movements, albeit from divergent vantage points. The fundamentalist delves into the causes of market movement, while the technician studies the effects. The technician contends that the effect is all that is requisite, rendering the causes or reasons as superfluous, whereas the fundamentalist is always driven to know why.

Traders often identify themselves as either technicians or fundamentalists. However, in reality, there exists a significant overlap. Many fundamentalists possess a working knowledge of basic chart analysis tenets, while many technicians have at least a passing awareness of fundamental principles. The best strategy would involve fundamental and technical analyses tailored to the user's investment goals and risk tolerance.[17]

Market price acts as a precursor to the fundamentals or the prevailing conventional wisdom. While known fundamentals have already been discounted and are "in the market," prices are now reacting to unknown fundamentals. Some of the most dramatic bull and bear markets in history have been inaugurated with little or no perceived change in the fundamentals. By the time these changes became conspicuous, the new trend was well on its course. In accepting the premises of technical analysis, it becomes clear why technicians deem their approach superior to fundamentalists. If the fundamentals are indeed reflected in market prices, then studying those fundamentals becomes redundant. Chart reading emerges as a streamlined form of fundamental analysis. On the flip side, fundamental analysis does not encompass a study of price action. It's conceivable to trade financial markets using just the technical approach, whereas trading based solely on fundamentals, without any consideration of the technical side of the market, seems dubious at best.

4.3 Other Instruments for Technical Analysis

Technical indicators are pivotal tools utilized in technical analysis, providing a lens through which market trends and dynamics can be analyzed. These indicators are mathematical calculations based on price, volume, or open interest of a security. By providing a graphical or numerical representation of market trends and patterns, technical indicators aid traders and analysts in the formulation of trading strategies and decision-making. Herein, we delve into some of the most well-regarded technical indicators in the financial market analysis realm.

4.3.1 Basic Market Attributes

Technical analysis commences with a scrutiny of basic market attributes, which encapsulate the essence of market activity within a specific timeframe:

- *Open Price*: The inaugural price at which a security trades upon the market's opening on a given trading day.
- *Close Price*: The terminal price at which a security trades before the market's closure on a given trading day.

- *High Price*: The zenith price at which a security trades during the course of the trading day.
- *Low Price*: The nadir price at which a security trades during the course of the trading day.
- *Volume*: The aggregate number of shares or contracts traded in a security or an entire market during a stipulated period.

These fundamental attributes serve as the underpinning for deriving more sophisticated technical indicators, which furnish deeper insights into market dynamics.

4.3.2 Moving Average

Moving Averages (MAs) are quintessential trend-smoothing tools, employed to average a security's price over a specified span of periods, thereby ameliorating random fluctuations and rendering a clearer depiction of the overall trend direction.

Simple Moving Average (SMA)

$$\text{SMA} = \frac{1}{n}\sum_{i=1}^{n} P_i \tag{4.1}$$

where n is the number of periods, and P_i is the price of the security at period i.

Exponential Moving Average (EMA)

$$\text{EMA} = \left(\frac{C - P}{n}\right) + P \tag{4.2}$$

where C is the current price, P is the previous period's EMA, and n is the smoothing factor.

4.3.3 On-Balance Volume - OBV

On-Balance Volume (OBV) is a cumulative indicator, leveraging volume and price to discern whether a security is being accumulated or distributed.

$$\text{OBV} = \sum(V \cdot D) \tag{4.3}$$

where V is the volume of the security, and D is the direction of the price (1 if the price increased, -1 if the price decreased).

4.3.4 Bollinger Bands

Bollinger Bands are composed of a middle band being an NN-period simple moving average (SMA), an upper band, and a lower band.

Middle Band

$$\text{Middle Band} = \text{SMA}(N) \tag{4.4}$$

Upper Band

$$\text{Upper Band} = \text{SMA}(N) + (K \cdot \text{SD}) \tag{4.5}$$

Lower Band

$$\text{Lower Band} = \text{SMA}(N) - (K \cdot \text{SD}) \tag{4.6}$$

where N is the number of periods, K is the number of standard deviations, and SD is the standard deviation of the price over N periods.

4.3.5 Further Technical Indicators

Transitioning beyond these indicators, the technical analysis landscape is replete with a myriad of other indicators like the Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), Stochastic Oscillator, among others, each providing unique lenses through which to analyze market behavior and formulate trading strategies.

4.4 Market State Classification

Time series data is a cornerstone in financial analysis, providing a chronological trail of market variables such as prices, trading volumes, and other financial metrics. Each data point in a time series is time-stamped, thus preserving the temporal order of observations. This temporal order is crucial as it encapsulates the dynamics of market evolution.

Time series data in financial markets can be broadly categorized into two types: univariate and multivariate.

- *Univariate Time Series*: A univariate time series consists of single observations recorded sequentially over time. For instance, the daily closing prices of a particular stock form a univariate time series. Each data point in a univariate time series represents a single variable's value at a specific time.

- *Multivariate Time Series*: Contrarily, a multivariate time series comprises multiple variables recorded at each time step. For example, the daily opening price, closing price, high, low, and trading volume of a stock constitute a multivariate time series. Each data point in a multivariate time series encapsulates the values of multiple variables at a particular time.

The distinction between univariate and multivariate time series is crucial as it influences the choice of analytical methods and the complexity of the analysis. Market State Classification (MSC) is a robust analytical framework that ventures to delineate different market states or regimes. The objective is to dissect the underlying market dynamics that significantly transmute from one state to another. Here's a more detailed exposition:

4.4.1 Base Models for Financial Time Series Analysis

In the rapidly evolving landscape of financial markets, the ability to accurately classify and predict market states has become a cornerstone for effective trading and investment strategies. With the advent of machine learning and deep learning techniques, the analytical capabilities available to financial analysts and traders have significantly expanded. These technologies, leveraging algorithms like Support Vector Machines (SVM) and Random Forests (RF), along with advanced deep learning architectures, have redefined the approach to understanding and navigating the complex dynamics of financial markets. This section delves into the state-of-the-art methodologies employed in market state classification, highlighting the role of clustering, classification, and complex system analysis, and how they pave the way for constructing robust trading strategy ensembles. The incorporation of frameworks like FinRL further exemplifies the shift towards automation, leveraging the predictive power of deep learning models to optimize trading strategies dynamically.

Machine Learning Algorithms

Machine learning algorithms, particularly Support Vector Machines (SVM) and Random Forests (RF), have shown high accuracy in classification tasks, including market state classification (MSC). These algorithms, by learning from historical data, have become vital tools for predicting or making decisions, especially in classifying market states through both historical and real-time market data analysis.

The techniques of clustering and classification are pivotal in distinguishing different market states, utilizing the historical and real-time data to group similar data points. This enables the identification and understanding of market states and their transitions, enhancing the analysis of complex systems that underpin market dynamics. Understanding these complex interactions is crucial for classifying financial market states more accurately.

Furthermore, the classification of market states facilitates the construction of trading strategy ensembles, improving navigation through diverse market conditions. This approach not only allows for the comparison of results across different studies but also highlights the effectiveness of various trading strategies under different market conditions. These methodologies represent the forefront of market state classification, offering sophisticated tools to financial analysts and traders for navigating the complex landscape of financial markets. The growing prominence of machine learning and statistical techniques promises to further refine the precision and effectiveness of market state classification.

Advanced Deep Learning Architectures

Deep learning, with its advanced architectures, has revolutionized financial forecasting by identifying complex patterns in market data without manual feature extraction. These models, categorized into individual and ensemble models, enhance prediction accuracy by combining multiple models' strengths. Among these, Deep Neural Networks (DNNs), with their multilayered architecture, have been instrumental in capturing nonlinear relationships in financial data. They process input through successive layers of neurons, each layer transforming the data to higher abstraction levels, making DNNs particularly effective for complex prediction tasks. However, they might not inherently capture temporal dependencies in time series data, a limitation overcome by other specialized models.

1D CNNs excel in identifying local patterns within time series data, utilizing convolutional filters to extract features across the temporal dimension. This thesis work employs ResNet, a variant of 1D CNNs known for its deep architecture facilitated by skip connections, allowing it to learn from data effectively while avoiding the vanishing gradient problem. This characteristic makes ResNet particularly suited for financial time series analysis, where capturing long-term dependencies is crucial.

RNNs and their variants, such as LSTM (Long Short-Term Memory) and GRU (Gated Recurrent Unit), are designed to handle sequential data explicitly. By maintaining a memory of previous inputs, they can capture temporal dependencies within the series, essential for accurate financial forecasting. LSTMs, with their unique gating mechanism, are adept at learning from long sequences without the risk of gradient vanishing or explosion, making them a popular choice for financial time series modeling. Beyond individual models, ensemble approaches combine multiple predictive models to improve forecast accuracy. Hybrid models, like

Convolutional-Recurrent Neural Networks (CRNNs), merge the spatial feature extraction capabilities of CNNs with the temporal modeling strengths of RNNs, offering a powerful tool for analyzing financial time series data that exhibits both spatial and temporal dynamics.

In the realm of deep learning architectures, FinRL, short for Financial Reinforcement Learning, represents a cutting-edge intersection between finance and reinforcement learning, offering a robust framework tailored for the development of sophisticated automated trading strategies. It harnesses the predictive prowess of advanced deep learning models, such as ResNet, to facilitate decision-making in trading activities. The core functionality of FinRL lies in its ability to optimize trading strategies dynamically over time by interacting directly with market data. This is achieved through a process where the framework learns from the market's behavior, making adjustments to the trading strategies to maximize performance.

The integration of FinRL into the landscape of financial forecasting and trading strategy development marks a significant shift towards the automation of trading systems and the utilization of intelligent systems in decision-making processes. The framework's ability to leverage deep learning models for real-time decision-making underscores the potential of AI in transforming financial strategy formulation and execution. FinRL's innovative approach to automated trading strategy development not only streamlines the trading process but also enhances the capability of financial models to adapt to changing market conditions, promising to revolutionize financial market analysis and strategy development in the years to come.

Chapter 5

Time Series Representation Learning

5.1 Introduction

Time series data permeates every facet of the real world, from environmental monitoring to financial markets. The quest to derive meaningful and actionable insights from time series data has led to the evolution of various representation learning methods. These methods aim to extract and infer valuable information from time series data, facilitating a deeper understanding of complex dynamics and informed decision-making processes. Among various approaches, deep learning stands out for its exceptional ability to uncover hidden patterns and features without the necessity for manual feature engineering. This chapter delves into the state-of-the-art in time series representation learning, spotlighting significant advancements and methodologies that have shaped the field.

5.2 State of the Art in Time Series Representation Learning

A time series comprises a sequence of data points collected in time order, capturing the intricate behaviors of specific variables or events as they unfold. This type of data encapsulates valuable insights across different fields at various moments, facilitating strategic decisions and forecasts. Examples include sensor outputs in the Internet of Things (IoT), data from cyber-physical systems, changes in stock market prices, and physical activity tracked by wearable technology. Nonetheless, deciphering the rich information embedded within such complex data sequences necessitates a method for effectively representing time series. This necessity gave rise to research focused on time series representation. By developing new ways to represent time series data, it becomes possible to adeptly conduct a range of subsequent analyses, including but not limited to forecasting, classification, regression, and identifying anomalies.

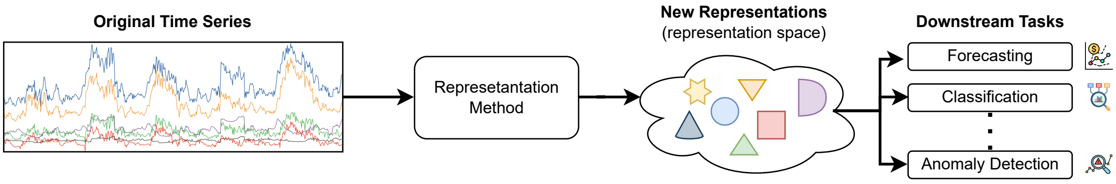

Figure 5.1: Basic concept of time-series representation methods.[18]

The landscape of time series representation learning is rich and varied, with deep learning methods at the forefront of extracting intricate patterns from time series data. A survey by Trirat et al.[18] introduces a novel taxonomy for universal time series representation learning, categorizing methods based on neural architecture, learning objectives, and training data utilization. This comprehensive review illuminates how these components enhance the quality of learned representations and sets the stage for future research directions in the field.

5.2.1 Time Series Properties

In this section, we delve into the distinct attributes of time series data identified by previous research in the field of time series representation learning [19]. These attributes highlight the intricate world of time series representation learning.

Temporal Dependency

Time series data is inherently dependent on chronological order, meaning the value at a certain point is related to its preceding values. For a given input x_t at time t, the model might predict y_t, but this input could lead to a different prediction at another time. To capture this temporal dependency, models often use windows or subsequences of previous observations. The challenge lies in determining the optimal window length to effectively capture these dependencies, which might be variable. Moreover, temporal dependencies can be local, relating to sudden changes or noise, or global, relating to overarching trends or patterns.

Noise and Dimensionality

Time series data, especially from real-world sources, is typically noisy and high-dimensional. Noise can stem from measurement inaccuracies or uncertainties.

Techniques like dimensionality reduction and wavelet transforms help mitigate these issues by filtering out noise and compressing the data. However, this process can lead to the loss of crucial information and often requires domain-specific expertise to choose the appropriate methods.

Relationships among Variables

This aspect is especially significant in multivariate time series data, where the interaction between variables can be complex and not always apparent. Analyzing a limited set of variables without considering inter-relationships may not fully capture the dynamics of the system. For example, an array of sensors detecting various gases to identify a specific smell or monitoring a single stock in a complex financial system might not provide comprehensive insights into the overall state or future trends.

Variability and Nonstationarity

Time series data exhibits variability and nonstationarity, with statistical properties such as mean, variance, and frequency changing over time. These variations often manifest as seasonal patterns, long-term trends, or fluctuations. Seasonality involves repetitive cycles at regular intervals, while trends indicate directional shifts over a longer period. Sometimes, changes in frequency are crucial to the analysis, making frequency domain methods more advantageous than time-domain approaches.

5.3 State-of-the-Art Models for Time Series Representation Learning

In the rapidly evolving field of time series analysis, the development of sophisticated representation models has been pivotal in unlocking new insights and applications across various domains such as finance, healthcare, and environmental monitoring. We will now explore some of the state of the art models in this field, which contributed to the foundations of this work.

5.3.1 TS2Vec: Towards Universal Representation of Time Series

TS2Vec (Time Series to Vector) is a novel framework designed for learning representations of time series data across various lengths and domains. This framework aims to address the challenges of time series representation learning by providing a flexible, efficient, and powerful method that captures the inherent temporal dynamics and patterns within time series data.

Methodology and Architecture

TS2Vec operates on the principle of self-supervised learning, where the model learns rich representations by maximizing the agreement between representations of different segments within the same time series under various augmentations. This approach leverages contrastive learning, a technique that learns to distinguish between similar (positive) and dissimilar (negative) pairs of data points.

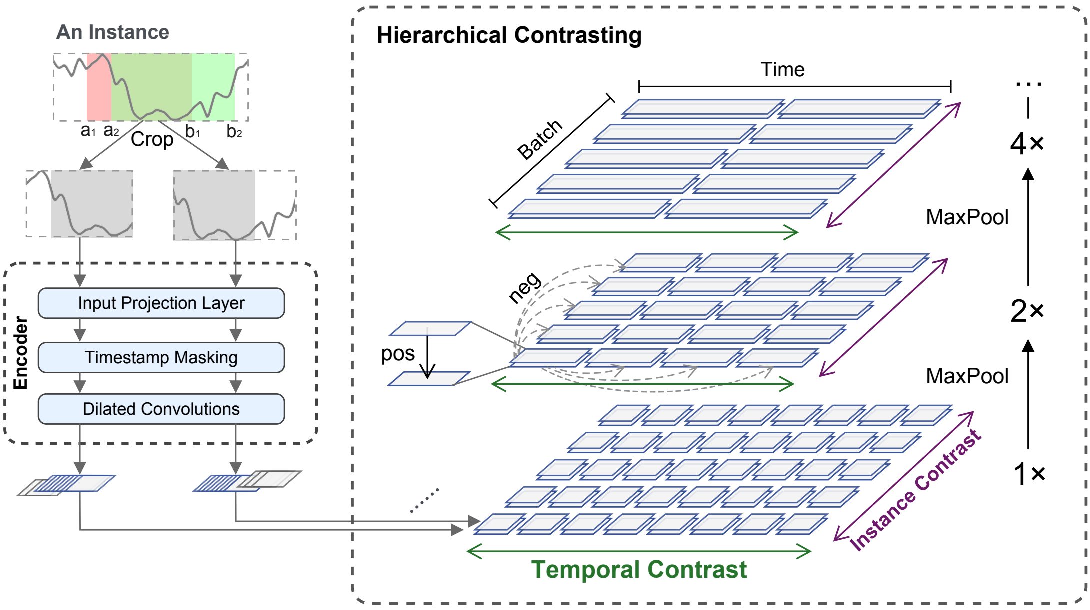

Figure 5.2: Proposed TS2Vec architecture, shown with a univariate time series example, adaptable to multivariate inputs. In the diagram, each parallelogram represents a timestamp's representation vector for an instance.[20]

As shown in Figure 5.2, the TS2Vec model employs a hierarchical contrastive learning objective that captures temporal dependencies at multiple scales. It uses a convolutional neural network (CNN) architecture to process time series data. The CNN extracts features from raw time series inputs, generating representations at multiple layers of the network. These representations are then used in a contrastive learning setting, where the model is trained to bring closer the representations of augmented versions of the same time series segment while pushing apart representations of different segments.

To achieve this, TS2Vec utilizes a variety of data augmentation techniques that preserve the temporal dynamics of the series, such as time warping, magnitude scaling, and jittering. These augmentations ensure that the model learns invariant and robust features across different transformations of the input data. One of the key advantages of TS2Vec is its flexibility in handling time series of varying lengths and domains without the need for manual feature engineering or domain-specific knowledge. This makes it a powerful tool for a wide range of applications, from financial time series analysis to healthcare monitoring.

5.3.2 Time Series Classification from Scratch with Deep Neural Networks: A Strong Baseline

In the paper ”Time Series Classification from Scratch with Deep Neural Networks: A Strong Baseline” [21], the authors explore the effectiveness of deep neural networks (DNNs) in classifying time series data without the need for extensive data preprocessing or feature engineering. Among the architectures evaluated, the Residual Network (ResNet) model is of particular interest for its ability to handle very deep architectures through the use of shortcut connections.

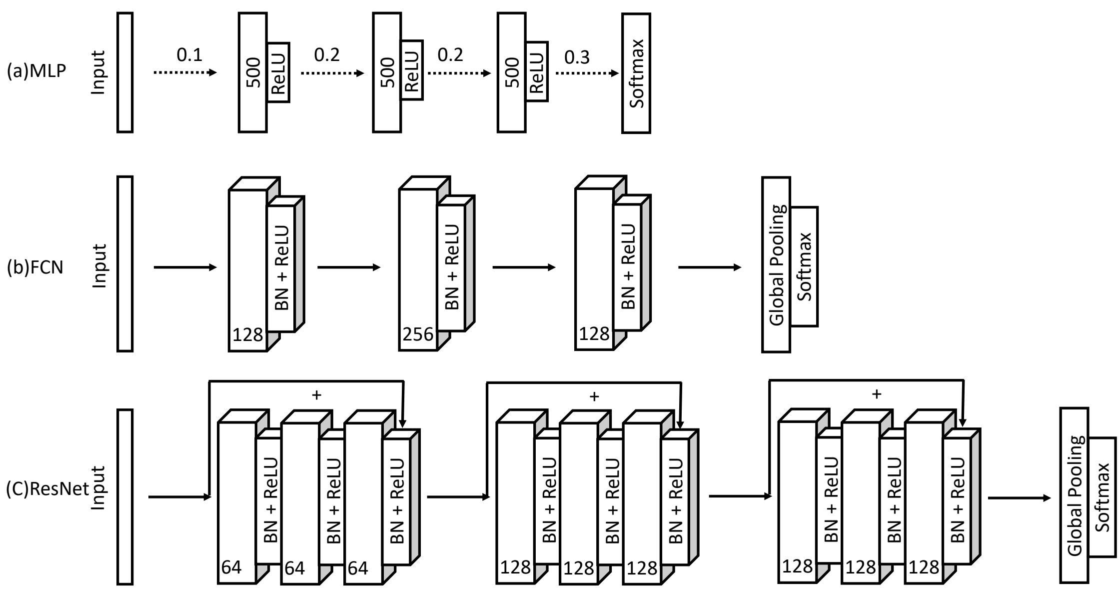

Figure 5.3: The network structure of three tested neural networks in [21].

The ResNet model is designed to push the boundaries of how deep neural networks can go. It introduces shortcut connections that skip one or more layers to prevent the vanishing gradient problem and facilitate easier learning by allowing the gradient to flow through the network more efficiently. This is particularly beneficial for time series data, which may require capturing long-term dependencies that span across many time steps.

In the context of time series classification, each residual block within the ResNet architecture uses a series of convolutional layers followed by batch normalization and ReLU activation. The model employs a stack of three residual blocks with varying numbers of filters, specifically designed to capture the temporal dynamics and dependencies within time series data at different scales.

The paper highlights that while ResNet is capable of achieving state-of-the-art performance on time series classification tasks, it also tends to overfit more easily compared to other models, such as the Fully Convolutional Network (FCN). This is attributed to the relatively small size and lack of diversity in the UCR time series datasets used for evaluation. However, despite this tendency, ResNet's inclusion in the evaluation demonstrates its potential in handling complex patterns in time series data when adequately regularized and tuned.

5.3.3 T-Loss and TNC (Temporal Neighborhood Coding)

This subsection delves into two significant contributions to time series representation learning: Temporal Neighborhood Coding (TNC)[22] and a Triplet Loss-based approach (T-Loss)[23], highlighting their methodologies, innovations, and applications.

Temporal Neighborhood Coding (TNC)

Temporal Neighborhood Coding (TNC) emerges as a powerful self-supervised framework tailored for complex, multivariate, and non-stationary time series data. At its core, TNC exploits the concept of temporal neighborhoods—segments of time series that share similar characteristics due to the local smoothness in their generative processes. The uniqueness of TNC lies in its ability to learn representations by differentiating between the distributions of neighboring and non-neighboring signals within these temporal segments. This distinction is achieved through a debiased contrastive learning objective, which ensures that representations maintain the local stationary properties inherent to each neighborhood.

TNC's applicability is particularly pronounced in medical settings, where understanding and tracking the dynamic nature of physiological signals can aid in diagnosing, monitoring, and predicting patient states. The framework's self-supervised nature, requiring no explicit labels for learning, makes it an ideal candidate for handling sparsely labeled or unlabeled time series data prevalent in healthcare.

Triplet Loss-Based Approach (T-Loss)

The Triplet Loss-based approach (T-Loss) for time series representation learning offers a different perspective on capturing the nuances of time series data. By employing a triplet loss function, T-Loss aims to learn representations that bring closer the embeddings of similar time series segments (positive pairs) while pushing apart those of dissimilar segments (negative pairs). This method focuses on maximizing the margin between positive and negative examples in the embedding space, thereby enhancing the separability and interpretability of the learned representations.

T-Loss shines in scenarios where precise modeling of time series dynamics is crucial for tasks like classification, clustering, or anomaly detection. Its ability to handle high-dimensional, multivariate time series makes it suitable for a wide range of applications, from financial time series analysis to environmental monitoring.

Both TNC and T-Loss introduce innovative ways to tackle the challenges of time series representation learning. While TNC leverages the local stationarity within temporal neighborhoods to encode time series data, T-Loss employs a triplet loss function to differentiate between similar and dissimilar segments. These methodologies offer new avenues for extracting meaningful and generalizable features from time series, crucial for downstream tasks like classification, forecasting, and anomaly detection.

Chapter 6

Methodology

This chapter dives into the detailed methodologies adopted in our research, aiming to explore the use of supervised contrastive learning techniques in time series classification, specifically focusing on stock market data. Our main goal is to improve the precision and efficiency of predicting financial market movements, which requires leveraging cutting-edge deep learning models. Given the complexity and unpredictability of financial markets, accurately classifying and forecasting market trends is crucial for both theoretical research and practical applications in financial analysis and trading.

6.1 Problem Definition

The cornerstone of our research involves understanding and optimizing the way time series data $\mathcal{X} = x_1, x_2, \ldots, x_N$, which includes N instances, is processed and interpreted for the purpose of accurate classification. At the heart of this endeavor lies the development of a sophisticated embedding function f, designed through the capabilities of a deep neural network (DNN). The mathematical representation of this embedding function is expressed as follows:

$$f : \mathcal{X} \to \mathbb{R}^d \tag{6.1}$$

Here, d symbolizes the dimensionality of the space into which the time series data is embedded. This embedded space is conceived to encapsulate the intrinsic patterns and characteristics of the data in a more compact and discriminative form, thereby facilitating more nuanced and precise classification tasks.

The process begins with the organization of input data into batches, which are then fed into a Residual Network (ResNet). ResNet is chosen for its remarkable ability to mitigate the vanishing gradient problem common in deep networks, thus enabling the learning of complex patterns without compromising network depth.

The architecture of ResNet, featuring skip connections that allow gradients to flow through the network more effectively, proves to be particularly adept at handling the intricate patterns present in time series data. Once the data has been processed through the ResNet, a projection network takes over to compute a supervised contrastive loss on the outputs. This computation is pivotal, as it refines the embeddings by enforcing a clear margin of separability between classes (inter-class separability) and ensuring that instances of the same class are pulled closer together in the embedded space (intra-class compactness). This step is critical for enhancing the model's ability to discern between different classes, a feature especially valuable in the context of financial time series, where minute distinctions can significantly impact classification accuracy.

To capitalize on the refined representations obtained through the supervised contrastive loss, a dense network layer is subsequently introduced. This layer is trained atop the embeddings, employing cross-entropy loss as its objective function. The use of cross-entropy loss in this phase is instrumental in fine-tuning the embeddings for optimal classification performance, effectively leveraging the nuanced representations learned in the earlier phase for the direct task of classifying the time series data.

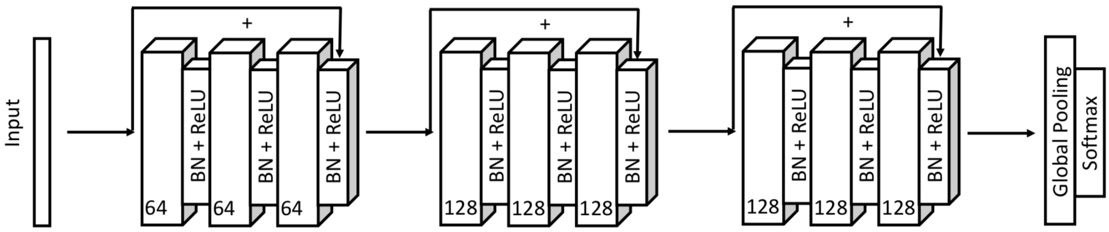

Figure 6.1: Residual Network structure from [21]

6.2 Approach

6.2.1 Data pre-processing

One of the core tenets of this research is the emphasis on minimal preprocessing, inspired by the approach outlined in [21]. This approach pivots away from traditional heavy preprocessing techniques, aligning with the broader objective of harnessing Deep Neural Networks (DNNs) for end-to-end time series classification. By doing so, the research aims to demonstrate the capability of DNNs to directly process and recognize patterns in raw time series data, eliminating the need for intricate preprocessing steps or manual feature extraction.

The only preprocessing steps applied are ment to bring the dataset to the same level of the benchmark of reference in TS2Vec [20]:

- **Handling Missing Values**: Given that real-world time series data can often have gaps or missing values, a crucial preprocessing step is ensuring these missing values are addressed. This is essential to maintain the consistency and quality of the data, ensuring that DNNs can process them without disruptions.

- **Centering Varying Length Series**: Time series data, especially from diverse sources or domains, can have varying lengths. To ensure uniformity and facilitate efficient batching during training, the time series are centered using a specialized function. This operation aligns the meaningful (non-NaN) parts of the series, ensuring that any gaps or missing values are uniformly distributed around the actual data.

By maintaining a near-zero preprocessing approach, this research underscores the power and versatility of DNNs in handling raw time series data, pushing the boundaries of what's possible in time series classification.

6.2.2 Model Architecture

Deep neural networks, especially deep convolutional neural networks (CNNs), have demonstrated impressive performance in various domains. One significant challenge with deep networks is the degradation problem: as the network grows deeper, its accuracy can saturate and then rapidly degrade. ResNet, or Residual Network, was introduced to address this problem.

The core idea behind ResNet is the introduction of "skip connections" or "shortcuts" that bypass one or more layers. These connections allow the network to learn the "residual" functions, which, in essence, means that these layers can learn to identify and forward the essential features without any modification if needed. This mechanism helps in alleviating the degradation problem. For time series data, these residual connections can capture both short-term and long-term dependencies in the data, making ResNet architectures particularly effective.

The ResNet architecture employed in this research is specifically tailored for time series classification. Taking inspiration by the model used in [21] and shown in Fig.6.1, the network starts with an initial convolutional layer designed to capture basic patterns in the time series data.

$$h_0 = Conv1D_d$$

where d is the dimensionality of the input time series data. This is followed by batch normalization and ReLU activation. Each residual block in the network is composed of two convolutional layers with batch normalization and ReLU activations. Additionally, if there's a change in the number of channels or if a stride other than 1 is used, a shortcut connection with a 1x1 convolution and batch normalization is introduced to match the dimensions. In the current architecture, the blocks and their configurations are:

$$
\begin{aligned}
h_1 &= Block_{k_1}(x) \\
h_2 &= Block_{k_2}(x) \\
y &= h_2 + x \\
\hat{h} &= \mathrm{ReLU}(y)
\end{aligned}
$$

Each block ensures that the network learns and captures both short-term and long-term patterns in the data. After the residual blocks, the feature maps are pooled to produce a fixed-size vector, effectively summarizing the features learned from the entire time series through an AdapriveAvgPoll with output dimension equals to 1.

The ResNet architecture described above is adept at processing time series data, extracting essential features, and producing embeddings that can be further used for classification or other downstream tasks. The residual connections, characteristic of ResNet architectures, ensure stable and effective training even as the network depth increases.

6.2.3 Supervised Contrastive Loss

Supervised Contrastive Learning is a recent learning approach that merges the strengths of supervised learning with the principles of contrastive learning. While traditional supervised learning focuses on predicting correct labels, contrastive learning is concerned with the quality of the learned representations. Its primary aim is to ensure that representations of samples from the same class (positive pairs) are closer to each other in the embedding space, while those from different classes (negative pairs) are farther apart.

The main objective of supervised contrastive learning can be described through its loss function. As show by Wang et al. in [24], let i $\in I \equiv \{1,...,N\}$ be the index of an arbitrary sample, for a set of N randomly sampled sample/label pairs. For a given one, with its representation z_i, and another sample from the same class with representation z_p, the supervised contrastive loss is given by:

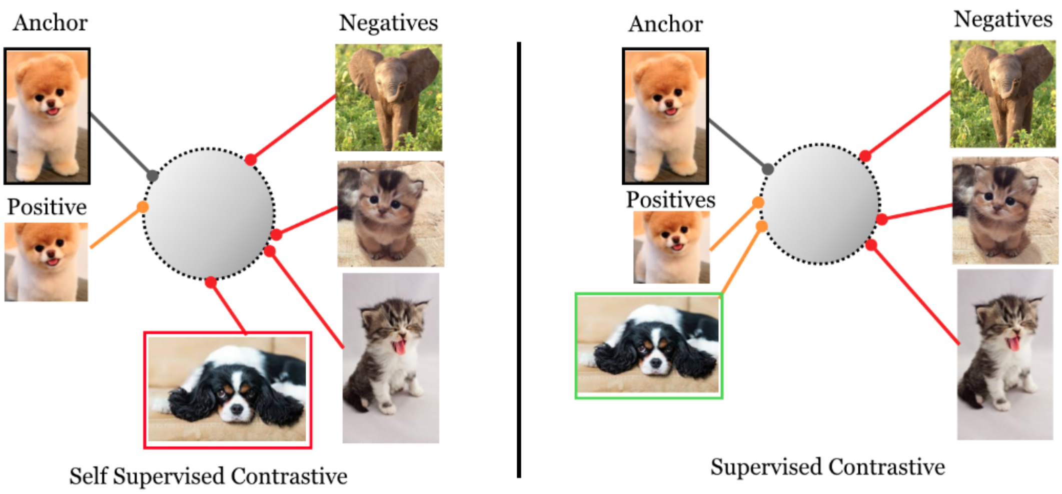

Figure 6.2: Supervised vs. self-supervised contrastive losses: The self-supervised contrastive loss (left) contrasts a single positive for each anchor (i.e., an augmented version of the same image) against a set of negatives consisting of the entire remainder of the batch. The supervised contrastive loss (right) considered in [24], however, contrasts the set of all samples from the same class as positives against the negatives from the remainder of the batch. As demonstrated by the photo of the black and white puppy, taking class label information into account results in an embedding space where elements of the same class are more closely aligned than in the self-supervised case. [24]

$$L_{SCN} = \sum_{i \in I} L_{SCN,i} =$$

$$= \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \in A(i)} \exp(z_i \cdot z_a / \tau)} \tag{6.2}$$

Here:

- z_i is the representation of an instance from a class y_i,
- $A(i) \equiv I \setminus \{i\}$, distinct from i,
- the $\cdot$ represents their inner product,
- $\tau \in R^+$ is a scalar temperature parameter,

- $P(i) \equiv \{p \in A(i) : y_p = y_i\}$ is the set of indices of all positives in the batch, and $|P(i)|$ is its cardinality.

The essence of this loss function is to maximize the similarity between positive pairs (samples from the same class) while minimizing their similarity with negative pairs. Supervised Contrastive Learning promotes the learning of more informative and discriminative features. This is achieved by ensuring that representations are not just class-specific, but also distinct from representations of other classes. This trains models to be more robust to variations within the class by emphasizing the importance of intra-class variations. That often leads to better generalization on unseen or noisy data. The representations learned via this loss are often more transferable to other related tasks. This makes supervised contrastive learning a powerful technique for pre-training models that can be fine-tuned for specific tasks. The way on how this loss function works, and how it differs from self-supervised techniques, is shown at high level in Fig.6.2.

In the context of this research, the principles of supervised contrastive learning have been seamlessly integrated with the ResNet architecture tailored for time series classification. The use of this loss during the training phase encourages the ResNet model to derive representations that are both discriminative and compact in the embedding space. This strategy has shown promise in enhancing classification accuracy across benchmark datasets of time series data.

6.3 Dataset Building and Labeling

The empirical analysis conducted in this study leverages a diverse array of datasets to ensure robustness and applicability across different domains. Primarily, the UEA and UCR Time Series Classification Repository serves as a benchmark [25], offering a wide range of datasets that span various fields. These repositories are renowned for their comprehensive collection of time series data, facilitating the evaluation of classification models across a spectrum of scenarios.

While the UEA and UCR datasets provide a valuable benchmark, the core of our analysis is deeply rooted in financial datasets. These datasets not only present unique challenges due to their volatility and noise but also offer significant opportunities for impactful insights into market behavior and dynamics. Therefore, our methodology prioritizes the construction and labeling of financial time series data to better understand and predict market movements.

6.3.1 Financial Datasets

The primary source for our financial dataset is the yfinance library [26], an accessible and reliable tool for extracting stock market data. From this library, we meticulously gathered information on 75 tickers encompassed within the NASDAQ 100 index, spanning an extensive 12-year period from 2010 to 2022. The NASDAQ 100, known for its heavy representation of technology companies, reflects the composition of the U.S. market, where tech industries play a predominant role, as showed in Fig.6.3. This deliberate selection of tickers is designed to capture a wide array of market behaviors and trends, offering a rich and diverse dataset for our analysis. By focusing on the NASDAQ 100, we not only align with the tech-centric nature of current market dynamics but also ensure our dataset embodies the complexity and innovation driving today's financial landscape.

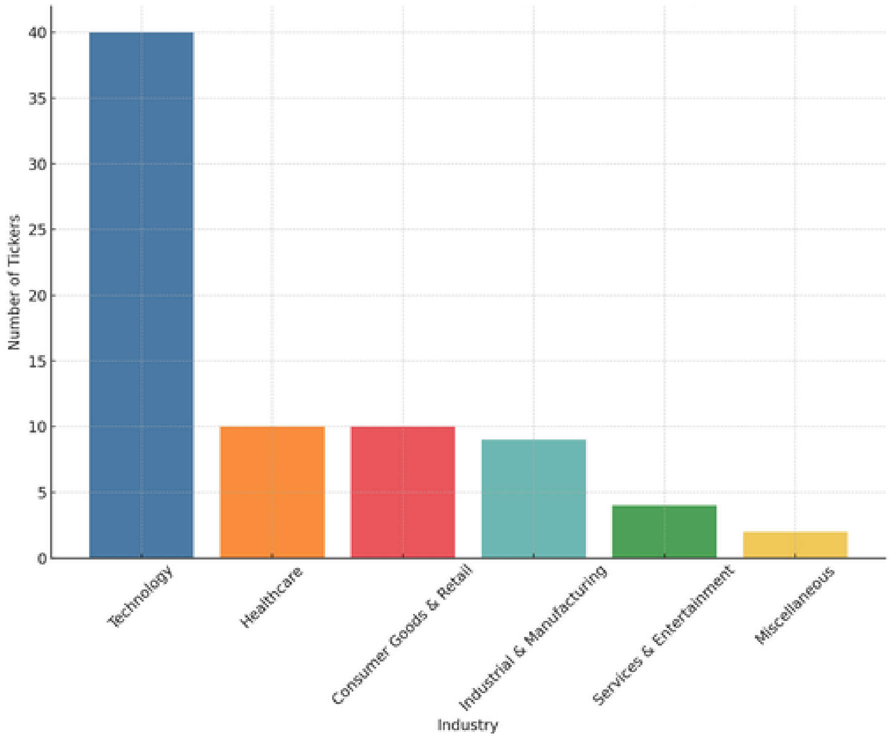

Figure 6.3: Selected NASDAQ tickers distribution among industries

Each ticker within our dataset represents daily stock activities, including *Open*, *High*, *Low*, *Close*, *Adj Close*, and *Volume*. This daily granularity allows for detailed examination of market fluctuations and trends. Furthermore, the diversity of industries represented by these tickers—ranging from technology and finance to manufacturing and healthcare—ensures a wide-ranging investigation into various market sectors.

6.3.2 Dataset Labeling

The dataset construction employs a sliding window approach, where each feature set corresponds to information captured within a ten-day window. This window slides forward by one day for each subsequent row of features, ensuring a comprehensive and continuous analysis of the time series data. To enrich the dataset with actionable insights, each ten-day window is assigned three distinct labels: *3 Days*, *5 Days*, and *7 Days*. These labels are determined based on the behavior of the closing prices in the days immediately following the ten-day window, providing a forward-looking perspective on market trends.

For the *3 Days* label, the methodology focuses on the closing price difference across the three days following the ten-day feature window (namely, the 11th, 12th, and 13th days). The assignment of label values—ranging from 0 to 2—is contingent upon the position of this closing price difference within the distribution of closing price differences for the respective ticker:

- *Label 0* is assigned if the closing price difference falls below the 33rd percentile of the distribution, indicating a bearish market trend.
- *Label 1* is assigned if the difference is between the 33rd and 66th percentiles, suggesting a neutral market condition.
- *Label 2* is given when the difference exceeds the 66th percentile, pointing to a bullish market trend.

The labeling approach for the 5 Days and 7 Days labels mirrors that of the 3 Days label, with the primary difference being the timeframe considered for the closing price difference. For the 5 Days label, the difference between the closing prices of the five days following the ten-day window is evaluated. Similarly, the 7 Days label is based on the closing price difference of the seven days post the ten-day window. The assignment of labels 0, 1, and 2 follows the same quantile-based criteria, facilitating a nuanced understanding of market behavior over different future intervals.

This labeling strategy aims to encapsulate the bearish or bullish behavior of the market, providing a foundation for subsequent analysis and modeling efforts. By categorizing the potential market trends in the days following the observed window, the dataset offers a valuable tool for forecasting market movements and developing trading strategies.

Chapter 7

Experiments and Results

7.1 Implementation Details

7.1.1 Datasets

This section outlines the datasets utilized in our experiments, which form the empirical basis for evaluating the performance of the Supervised Contrastive Learning (SCL) model in stock market time series analysis.

Financial Market Dataset

Our primary dataset consists of stock price movements among the ones available in the NASDAQ [3], providing a rich source of data for analyzing market behavior over time. This dataset includes daily trading information such as opening and closing prices, highest and lowest prices of the day, and trading volume, covering a comprehensive period that allows for a robust examination of stock market dynamics.

The choice of the NASDAQ 100 index was motivated by its representation of large-cap technology and non-technology sector companies, making it a valuable dataset for capturing a wide range of market sentiments and trends. This dataset's diversity and volume support our exploration into the effectiveness of SCL in discerning and classifying nuanced patterns within financial time series data.

To accurately conduct our experiments, we focused on a subset of features from the financial market dataset, specifically chosen for their relevance and impact on stock price movements. These features include *Open*, *High*, *Low*, *Close*, *Adj Close*, and *Volume*. Below is a detailed explanation of each feature:

- *Open*: The price at which a stock first trades upon the opening of an exchange on a trading day. It is a critical indicator of market sentiment and potential price movements throughout the day. High: The highest price at which a stock trades during the course of the trading day. This value is indicative of the peak buying interest for the stock.

- *Low*: The lowest price at which a stock trades during the course of the trading day, reflecting the minimum selling price that was accepted.
- *Close*: The price of a stock at the close of the trading day. It is used as a benchmark for financial reporting and by traders to assess market sentiment.
- *Adj Close*: The closing price after adjustments for all applicable splits and dividend distributions. This adjusted closing price gives a more accurate reflection of the stock's value.
- *Volume*: The number of shares or contracts traded in a security or an entire market during a given period. It is a measure of the total demand and supply for the stock and is often used to confirm trends and chart patterns.

These features were selected to provide a comprehensive understanding of market dynamics and to facilitate the modeling of stock price movements using the Supervised Contrastive Learning approach. The choice of features reflects an emphasis on both the quantitative aspects of the market (such as price and volume) and the adjustments made for external factors affecting stock valuations (reflected in the *Adj Close*). This selection is foundational to our experiments, enabling a nuanced analysis of the stock market's behavior through the lens of deep learning techniques.

UCR and UEA Time Series Classification Repositories

In addition to the financial market dataset, our experiments leverage datasets from the UCR (Univariate Time Series Classification Repository) and UEA (Multivariate Time Series Classification Repository) [25]. These repositories are well-regarded within the time series analysis community, offering a broad spectrum of time series datasets across various domains.

The inclusion of datasets from the UCR and UEA repositories allows us to benchmark the performance of our SCL model against a wide range of standard time series classification tasks. This comparative analysis helps establish the versatility and efficacy of SCL in handling different types of time series data, beyond the financial domain.

Together, these datasets provide a comprehensive platform for evaluating our SCL model. The subsequent sections will detail the experimental design, followed by an in-depth discussion of our findings.

Hardware

- Processor Intel(R) Xeon(R) CPU 2.20GHz
- System Memory 15.0 GB
- GPU NVIDIA Tesla T4 - Memory 15.0 GB

Experiments Execution Times

Model	**Train Time (s)**
SupCon	~ 550
TS2Vec	~ 750
Arima	~ 180
Random Forest	~ 0.35
GBoost	~ 0.38

Table 7.1: Models Execution Times per Dataset Train

Tabel 7.1 refers to the execution times for the training of the different models employed. SupCon and TS2Vec, with execution times of approximately 550 and 750 respectively, are more computationally intensive, likely due to their heavy deep learning frameworks, more suitable in extracting complex patterns from financial time series data, but also way more time consuming. ARIMA, on the other hand, shows a significantly lower execution time of around 180, reflecting its nature as a less computationally demanding statistical model suitable for linear trends and seasonality in time series. Random Forest and GBoost, with execution times of about 0.35 and 0.38 respectively, are the most efficient models listed. Their quick execution times can be attributed to their ensemble learning techniques, which, while efficient, may not capture deep temporal dependencies as effectively as the more time-consuming deep learning models.

7.2 Benchmark Results

This section presents the benchmark results of our Supervised Contrastive Learning (SupCon) model, evaluated across a selection of datasets from the UCR (Univariate) and UEA (Multivariate) repository, which are renowned for their comprehensive collection of time series datasets. The benchmarking process is crucial for assessing the performance of our model in comparison to other state-of-the-art time series classification methods. The datasets chosen span a wide range of domains, offering diverse challenges in time series classification. This diversity ensures a thorough evaluation of our model's adaptability and effectiveness across different types of time series data.

The results are presented in a detailed table, listing the accuracy scores achieved by our SupCon model alongside several other leading methods, including TS2Vec [20], T-Loss [23], TNC [22], TS-TCC [27], TST [28], and DTW [29].

7.2.1 Univariate Benchmark Dataset

As presented in Table 7.2, the evaluation of our Supervised Contrastive Learning (SupCon) model on the UCR univariate dataset showcases promising outcomes, affirming the model's robustness and versatility in handling time series data across a diverse set of domains.

Dataset	**SupCon**	TS2Vec	T-Loss	TNC	TS-TCC	TST	DTW
Coffee	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	0.821	**1.000**
Computers	**0.788**	0.660	0.664	0.684	0.704	0.696	0.700
DistalPhalanxOutlineCorrect	**0.783**	0.775	0.775	0.754	0.754	0.728	0.717
DistalPhalanxOutlineAgeGroup	0.720	0.727	0.727	0.741	0.755	0.741	**0.770**
DistalPhalanxTW	0.683	**0.698**	0.676	0.669	0.676	0.568	0.590
Earthquakes	**0.748**	**0.748**	**0.748**	**0.748**	**0.748**	**0.748**	0.719
ECG200	0.900	0.920	**0.940**	0.830	0.880	0.830	0.770
ECG5000	0.928	0.936	0.933	0.937	**0.941**	0.928	0.924
ElectricDevices	0.700	**0.721**	0.707	0.700	0.686	0.676	0.602
FaceAll	0.785	0.805	0.786	0.766	**0.813**	0.504	0.808
FordA	0.939	**0.948**	0.928	0.902	0.930	0.568	0.555
FordB	**0.826**	0.807	0.793	0.733	0.815	0.507	0.620
Ham	0.740	0.724	0.724	**0.752**	0.743	0.524	0.467
Herring	0.594	**0.641**	0.594	0.594	0.594	0.594	0.531
ItalyPowerDemand	0.930	**0.961**	0.954	0.928	0.955	0.845	0.950
PhalangesOutlinesCorrect	0.805	**0.823**	0.784	0.787	0.804	0.773	0.728
Plane	0.990	**1.000**	0.990	**1.000**	**1.000**	0.933	**1.000**
ProximalPhalanxOutlineCorrect	0.854	**0.900**	0.859	0.866	0.873	0.770	0.784
ProximalPhalanxOutlineAgeGroup	0.843	0.829	0.844	**0.854**	0.839	**0.854**	0.805
ProximalPhalanxTW	0.810	**0.824**	0.771	0.810	0.800	0.780	0.761
RefrigerationDevices	0.568	**0.589**	0.515	0.565	0.563	0.483	0.464
ScreenType	0.400	0.411	0.416	**0.509**	0.419	0.419	0.397
SmallKitchenAppliances	**0.755**	0.733	0.677	0.725	0.691	0.592	0.643
Strawberry	**0.967**	0.965	0.954	0.951	0.965	0.916	0.941
SwedishLeaf	0.820	**0.942**	0.914	0.880	0.923	0.738	0.792
SyntheticControl	**1.000**	0.993	0.987	**1.000**	0.990	0.490	0.993
ToeSegmentation1	0.830	**0.947**	0.939	0.864	0.930	0.807	0.772

Dataset	**SupCon**	TS2Vec	T-Loss	TNC	TS-TCC	TST	DTW
ToeSegmentation2	0.870	**0.915**	0.900	0.831	0.877	0.615	0.838
Trace	**1.000**	**1.000**	0.990	**1.000**	**1.000**	**1.000**	**1.000**
TwoPatterns	**1.000**	**1.000**	0.999	**1.000**	0.999	0.466	**1.000**
Worms	0.680	0.701	0.727	0.623	**0.753**	0.455	0.584
Chinatown	0.947	0.968	0.951	0.977	**0.983**	0.936	0.957
Crop	0.745	**0.756**	0.722	0.738	0.742	0.710	0.665
GunPointMaleVersusFemale	**1.000**	**1.000**	0.997	0.994	0.997	**1.000**	0.997
GunPointOldVersusYoung	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	0.838
InsectEPGRegularTrain	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	0.872
MixedShapesRegularTrain	0.855	**0.922**	0.905	0.911	0.855	0.879	0.842
PowerCons	0.890	**0.972**	0.900	0.933	0.961	0.911	0.878
SemgHandGenderCh2	0.880	**0.963**	0.890	0.882	0.837	0.725	0.802
SmoothSubspace	0.987	**0.993**	0.960	0.913	0.953	0.827	0.827
Average score per model:							
AVG	0.839	**0.855**	0.838	0.833	0.843	0.734	0.772
STD	0.147	0.148	0.153	0.141	0.148	0.190	0.177

Table 7.2: Accuracy scores of our method compared with those of other methods of unsupervised representation on 41 UCR datasets.

Our SupCon model achieved remarkable performance, demonstrating superior or competitive accuracy in numerous datasets, including achieving perfect scores in several categories like *Coffee*, indicating its exceptional ability to capture and leverage time series characteristics effectively.

Notably, the model outperformed traditional and some contemporary methods in datasets with challenging patterns, such as *Computers* and *Earthquakes*, highlighting its advanced feature extraction and representation learning capabilities. Even in highly competitive scenarios, SupCon showed its strength by closely matching or outperforming state-of-the-art methods, as seen in *ECG5000* and *FordB*, underscoring its potential for broad application in time series classification tasks.
The average accuracy score across all considered datasets positions our model as a strong contender, illustrating its consistency and efficiency in extracting meaningful patterns from complex time series data. These results underscore the effectiveness of Supervised Contrastive Learning in enhancing classification performance, offering valuable insights for future research and application in diverse time series analysis challenges.

7.2.2 Multivariate Benchmark Dataset

As shown in Table 7.3, results from the UEA datasets provide a comprehensive view of the performance of our Supervised Contrastive Learning (SupCon) model in comparison with other state-of-the-art time series classification methods.

Dataset	**SupCon**	TS2Vec	T-Loss	TNC	TS-TCC	TST	DTW
ArticularyWordRecognition	0.803	**0.987**	0.943	0.973	0.953	0.977	**0.987**
AtrialFibrillation	**0.333**	0.200	0.133	0.133	0.267	0.067	0.200
CharacterTrajectories	0.983	**0.995**	0.993	0.967	0.985	0.975	0.989
Cricket	0.862	0.972	0.972	0.958	0.917	**1.000**	**1.000**
Epilepsy	0.964	0.964	**0.971**	0.957	0.957	0.949	0.964
EthanolConcentration	0.315	0.308	0.205	0.297	0.285	0.262	**0.323**
FaceDetection	0.503	0.501	0.513	0.536	**0.544**	0.534	0.529
FingerMovements	0.490	0.480	**0.580**	0.470	0.460	0.560	0.530
HandMovementDirection	**0.405**	0.338	0.351	0.324	0.243	0.243	0.231
Heartbeat	0.722	0.683	0.741	0.746	**0.751**	0.746	0.717
JapaneseVowels	0.893	0.984	**0.989**	0.978	0.930	0.978	0.949
MotorImagery	0.500	0.510	0.580	0.500	**0.610**	0.500	0.500
NATOPS	**0.950**	0.928	0.917	0.911	0.822	0.850	0.883
PEMS-SF	**0.872**	0.682	0.676	0.699	0.734	0.740	0.711
PenDigits	**0.989**	0.989	0.981	0.979	0.974	0.560	0.977
PhonemeSpectra	0.150	0.233	0.222	0.207	**0.252**	0.085	0.151
RacketSports	**0.925**	0.855	**0.855**	0.776	0.816	0.809	0.803
SpokenArabicDigits	0.980	**0.988**	0.905	0.934	0.970	0.923	0.963
StandWalkJump	**0.467**	0.467	0.333	0.400	0.333	0.267	0.200
Average score per model:							
AVG	**0.689**	0.687	0.676	0.670	0.673	0.632	0.663
STD	0.277	0.295	0.305	0.298	0.286	0.321	0.317

Table 7.3: Accuracy scores of our method compared on multivariate UEA datasets.

Our SupCon model demonstrates competitive performance across a wide range of multivariate time series classification tasks, as evidenced by the accuracy scores on various UEA datasets. Notably, the SupCon model excels in certain datasets, indicating its robustness and adaptability to different types of time series data. The comparison with other models such as TS2Vec, T-Loss, TNC, TS-TCC, TST, and DTW reveals the strengths and limitations of our approach.

Performance

The SupCon model achieves the highest average score per model (0.689), marginally outperforming the TS2Vec model (0.687) and showing noticeable improvement over other models. This overall performance underscores the efficacy of the contrastive learning framework in capturing meaningful patterns in time series data.

The model shows substantial performance on datasets like NATOPS, PEMS-SF, PenDigits, and RacketSports, achieving the highest accuracy scores among the compared methods. These results highlight the model's capability in handling complex time series classification problems, especially where nuanced temporal

While the SupCon model leads in several datasets, there are instances where other models outperform it, such as in the ArticularyWordRecognition and CharacterTrajectories datasets. These outcomes suggest opportunities for further refining the model's architecture or training procedure to enhance its sensitivity to subtle temporal features.

The SupCon model's performance across a variety of domains within the UEA dataset collection demonstrates its versatility. This is particularly notable in its ability to handle both physical activity data (e.g., NATOPS, PEMS-SF) and sensor data (e.g., PenDigits), among others. The model's top performance in datasets with specific characteristics (e.g., HandMovementDirection, StandWalkJump) suggests that contrastive learning is particularly well-suited to scenarios where distinguishing between closely related time series is crucial.

7.3 Stock Prices Datasets Results

The task of predicting stock market movements represents a significant challenge, given the complexity and inherent unpredictability of financial markets. The Supervised Contrastive Learning (SupCon) model's evaluation on a comprehensive dataset, covering a span of 12 years with a division between a 10-year period for training and evaluation (2010 to 2020) and a 2-year period for testing (2021 to 2022), offers a nuanced insight into the model's capabilities and the intricacies of market prediction.
The performance metrics, including accuracy and F1 scores across different forecasting horizons (3-day, 5-day, and 7-day), provide a structured assessment of the model's predictive power. Despite the well-documented difficulties in financial forecasting, due to factors such as market volatility, economic indicators, and investor sentiment, the SupCon model demonstrates a promising ability to extract meaningful patterns from historical data on which further investigations can be performed.

7.3.1 Supervised Constrastive Learning Classification

The evaluation of the Supervised Contrastive Learning (SupCon) model across 3-day, 5-day, and 7-day forecasting intervals on the NASDAQ dataset revealed consistent accuracy levels, with a marginal improvement observed as the forecast horizon extends, as we can see from Table 7.4. Despite this, the model's F1 Score remains low, indicating challenges in harmonizing precision and recall. Another observation is the model's varied recall and precision across different labels, reflecting its nuanced ability to correctly identify specific market states.

	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
AVG	3D	0.44	0.20	0.61	0.04	0.34	0.26	0.02	0.17
AVG	5D	0.44	0.20	0.63	0.05	0.31	0.28	0.03	0.18
AVG	7D	0.45	0.20	0.64	0.07	0.27	0.28	0.03	0.14
STD	3D	5.40	0.02	0.48	0.19	0.47	0.20	0.08	0.23
STD	5D	6.24	0.02	0.47	0.24	0.44	0.20	0.10	0.21
STD	7D	5.85	0.03	0.47	0.20	0.45	0.21	0.14	0.25

Table 7.4: Average Scores of Supervised Contrastive Learning over NASDAQ dataset.

The standard deviation values highlighted considerable variability in accuracy across different evaluations, suggesting fluctuations in the model's performance.

This variability, combined with the consistent F1 Scores and the observed precision-recall dynamics, underscores the need for further model optimization to enhance predictive reliability and balance across different market states.

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
AAL	3D	36.10	0.28	0.61	0.43	0.00	0.37	0.35	0.00
AAL	5D	32.05	0.16	1.00	0.00	0.00	0.32	0.00	0.00
AAL	7D	34.48	0.17	1.00	0.00	0.00	0.34	0.00	0.00
AAPL	3D	47.06	0.21	1.00	0.00	0.00	0.47	0.00	0.00
AAPL	5D	45.84	0.21	1.00	0.00	0.00	0.46	0.00	0.00
AAPL	7D	45.03	0.21	1.00	0.00	0.00	0.45	0.00	0.00
ADBE	3D	45.23	0.21	1.00	0.00	0.00	0.45	0.00	0.00
ADBE	5D	44.22	0.20	1.00	0.00	0.00	0.44	0.00	0.00
ADBE	7D	42.19	0.20	1.00	0.00	0.00	0.42	0.00	0.00
ADI	3D	45.44	0.21	0.00	0.00	1.00	0.00	0.00	0.45
ADI	5D	39.96	0.26	0.75	0.00	0.18	0.42	0.00	0.33
ADI	7D	42.60	0.20	1.00	0.00	0.00	0.43	0.00	0.00
ADP	3D	37.73	0.18	1.00	0.00	0.00	0.38	0.00	0.00
ADP	5D	39.55	0.19	1.00	0.00	0.00	0.40	0.00	0.00
ADP	7D	22.52	0.20	0.01	0.73	0.31	0.04	0.09	0.65
ADSK	3D	47.47	0.21	1.00	0.00	0.00	0.47	0.00	0.00
ADSK	5D	47.06	0.21	1.00	0.00	0.00	0.47	0.00	0.00
ADSK	7D	46.86	0.21	1.00	0.00	0.00	0.47	0.00	0.00
AKAM	3D	38.54	0.19	1.00	0.00	0.00	0.39	0.00	0.00
AKAM	5D	37.93	0.18	0.00	0.00	1.00	0.00	0.00	0.38
AKAM	7D	35.90	0.19	0.02	0.00	0.94	0.25	0.00	0.36
ALGN	3D	48.07	0.22	1.00	0.00	0.00	0.48	0.00	0.00
ALGN	5D	40.37	0.19	0.00	0.00	1.00	0.00	0.00	0.40
ALGN	7D	52.33	0.23	1.00	0.00	0.00	0.52	0.00	0.00
AMAT	3D	47.87	0.22	0.00	0.00	1.00	0.00	0.00	0.48
AMAT	5D	44.62	0.21	1.00	0.00	0.00	0.45	0.00	0.00
AMAT	7D	46.05	0.21	1.00	0.00	0.00	0.46	0.00	0.00
AMD	3D	47.26	0.21	1.00	0.00	0.00	0.47	0.00	0.00
AMD	5D	47.47	0.21	1.00	0.00	0.00	0.47	0.00	0.00
AMD	7D	47.67	0.22	1.00	0.00	0.00	0.48	0.00	0.00
AMGN	3D	40.77	0.19	1.00	0.00	0.00	0.41	0.00	0.00
AMGN	5D	39.35	0.19	1.00	0.00	0.00	0.39	0.00	0.00
AMGN	7D	40.16	0.20	1.00	0.00	0.02	0.40	0.00	1.00
AMZN	3D	50.91	0.22	1.00	0.00	0.00	0.51	0.00	0.00
AMZN	5D	51.93	0.23	1.00	0.00	0.00	0.52	0.00	0.00

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
AMZN	7D	38.95	0.19	0.00	0.00	1.00	0.00	0.00	0.39
ASML	3D	46.65	0.21	1.00	0.00	0.00	0.47	0.00	0.00
ASML	5D	48.07	0.22	0.00	0.00	1.00	0.00	0.00	0.48
ASML	7D	48.48	0.22	0.00	0.00	1.00	0.00	0.00	0.48
ATVI	3D	37.32	0.18	0.00	0.00	1.00	0.00	0.00	0.37
ATVI	5D	46.86	0.21	1.00	0.00	0.00	0.47	0.00	0.00
ATVI	7D	33.87	0.17	0.00	0.00	1.00	0.00	0.00	0.34
AVGO	3D	45.23	0.21	1.00	0.00	0.00	0.45	0.00	0.00
AVGO	5D	42.19	0.20	1.00	0.00	0.00	0.42	0.00	0.00
AVGO	7D	48.68	0.22	0.00	0.00	1.00	0.00	0.00	0.49
BIDU	3D	39.55	0.19	1.00	0.00	0.00	0.40	0.00	0.00
BIDU	5D	42.80	0.20	1.00	0.00	0.00	0.43	0.00	0.00
BIDU	7D	41.99	0.20	1.00	0.00	0.00	0.42	0.00	0.00
BIIB	3D	41.38	0.20	1.00	0.00	0.00	0.41	0.00	0.00
BIIB	5D	42.80	0.20	1.00	0.00	0.00	0.43	0.00	0.00
BIIB	7D	40.16	0.19	1.00	0.00	0.00	0.40	0.00	0.00
BKNG	3D	42.80	0.20	0.00	0.00	1.00	0.00	0.00	0.43
BKNG	5D	39.15	0.19	1.00	0.00	0.00	0.39	0.00	0.00
BKNG	7D	43.41	0.20	0.00	0.00	1.00	0.00	0.00	0.43
BMRN	3D	36.71	0.18	1.00	0.00	0.00	0.37	0.00	0.00
BMRN	5D	37.12	0.18	1.00	0.00	0.00	0.37	0.00	0.00
BMRN	7D	39.76	0.21	0.98	0.00	0.04	0.40	0.00	0.46
CDNS	3D	50.71	0.22	0.00	0.00	1.00	0.00	0.00	0.51
CDNS	5D	43.00	0.20	1.00	0.00	0.00	0.43	0.00	0.00
CDNS	7D	44.22	0.20	1.00	0.00	0.00	0.44	0.00	0.00
CHKP	3D	41.58	0.20	1.00	0.00	0.00	0.42	0.00	0.00
CHKP	5D	38.95	0.19	0.00	0.00	1.00	0.00	0.00	0.39
CHKP	7D	41.78	0.20	1.00	0.01	0.00	0.42	0.17	0.00
CHRW	3D	41.38	0.23	0.00	0.06	0.98	0.00	0.27	0.42
CHRW	5D	38.13	0.18	1.00	0.00	0.00	0.38	0.00	0.00
CHRW	7D	39.55	0.19	1.00	0.01	0.00	0.39	1.00	0.00
CMCSA	3D	39.76	0.19	0.00	0.00	1.00	0.00	0.00	0.40
CMCSA	5D	42.19	0.20	1.00	0.00	0.00	0.42	0.00	0.00
CMCSA	7D	40.77	0.21	0.93	0.00	0.03	0.42	0.00	0.17
COST	3D	50.30	0.22	0.00	0.00	1.00	0.00	0.00	0.50
COST	5D	33.67	0.20	0.86	0.09	0.00	0.36	0.10	0.00
COST	7D	52.33	0.23	0.00	0.00	1.00	0.00	0.00	0.52
CSCO	3D	41.18	0.19	1.00	0.00	0.00	0.41	0.00	0.00
CSCO	5D	38.13	0.18	0.00	0.00	1.00	0.00	0.00	0.38

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
CSCO	7D	39.15	0.19	1.00	0.00	0.00	0.39	0.00	0.00
CSX	3D	27.99	0.25	0.00	0.82	0.36	0.00	0.17	0.60
CSX	5D	44.02	0.20	1.00	0.00	0.00	0.44	0.00	0.00
CSX	7D	42.19	0.20	0.00	0.00	1.00	0.00	0.00	0.42
CTSH	3D	40.37	0.19	1.00	0.00	0.00	0.40	0.00	0.00
CTSH	5D	29.61	0.23	0.00	0.84	0.23	0.00	0.28	0.33
CTSH	7D	42.60	0.20	1.00	0.00	0.00	0.43	0.00	0.00
DLTR	3D	39.55	0.19	1.00	0.00	0.00	0.40	0.00	0.00
DLTR	5D	40.37	0.19	1.00	0.00	0.00	0.40	0.00	0.00
DLTR	7D	30.43	0.26	0.55	0.15	0.14	0.36	0.11	0.42
EBAY	3D	40.16	0.19	0.00	0.00	1.00	0.00	0.00	0.40
EBAY	5D	48.88	0.22	1.00	0.00	0.00	0.49	0.00	0.00
EBAY	7D	47.67	0.22	1.00	0.00	0.00	0.48	0.00	0.00
EXPD	3D	27.59	0.23	0.00	0.92	0.28	0.00	0.19	0.62
EXPD	5D	42.19	0.20	1.00	0.00	0.00	0.42	0.00	0.00
EXPD	7D	42.60	0.20	0.00	0.00	1.00	0.00	0.00	0.43
FAST	3D	38.74	0.19	1.00	0.00	0.00	0.39	0.00	0.00
FAST	5D	21.09	0.12	0.00	1.00	0.00	0.00	0.21	0.00
FAST	7D	39.15	0.19	1.00	0.00	0.00	0.39	0.00	0.00
GILD	3D	34.48	0.17	1.00	0.00	0.00	0.34	0.00	0.00
GILD	5D	34.48	0.17	0.00	0.00	1.00	0.00	0.00	0.34
GILD	7D	31.24	0.16	1.00	0.00	0.00	0.31	0.00	0.00
GOOG	3D	45.44	0.21	0.00	0.00	1.00	0.00	0.00	0.45
GOOG	5D	46.05	0.21	0.00	0.00	1.00	0.00	0.00	0.46
GOOG	7D	44.62	0.21	1.00	0.00	0.00	0.45	0.00	0.00
GOOGL	3D	45.64	0.21	0.00	0.00	1.00	0.00	0.00	0.46
GOOGL	5D	44.42	0.21	1.00	0.00	0.00	0.44	0.00	0.00
GOOGL	7D	44.83	0.21	1.00	0.00	0.00	0.45	0.00	0.00
GRMN	3D	42.60	0.21	0.99	0.00	0.02	0.42	0.00	0.62
GRMN	5D	19.88	0.16	0.00	0.92	0.13	0.00	0.17	0.39
GRMN	7D	42.80	0.20	1.00	0.00	0.00	0.43	0.00	0.00
HSIC	3D	41.78	0.20	1.00	0.00	0.00	0.42	0.00	0.00
HSIC	5D	40.16	0.19	1.00	0.00	0.00	0.40	0.00	0.00
HSIC	7D	38.34	0.18	1.00	0.00	0.00	0.38	0.00	0.00
IDXX	3D	44.02	0.20	1.00	0.00	0.00	0.44	0.00	0.00
IDXX	5D	47.47	0.21	0.00	0.00	1.00	0.00	0.00	0.47
IDXX	7D	45.23	0.21	1.00	0.00	0.00	0.45	0.00	0.00
INCY	3D	35.70	0.18	1.00	0.00	0.00	0.36	0.00	0.00
INCY	5D	35.29	0.17	1.00	0.00	0.00	0.35	0.00	0.00

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
INCY	7D	35.90	0.18	1.00	0.00	0.00	0.36	0.00	0.00
INTC	3D	44.02	0.20	1.00	0.00	0.00	0.44	0.00	0.00
INTC	5D	45.64	0.21	1.00	0.00	0.00	0.46	0.00	0.00
INTC	7D	35.29	0.28	0.00	0.35	0.87	0.00	0.30	0.37
INTU	3D	48.07	0.22	0.00	0.00	1.00	0.00	0.00	0.48
INTU	5D	42.80	0.20	1.00	0.00	0.00	0.43	0.00	0.00
INTU	7D	51.93	0.23	0.00	0.00	1.00	0.00	0.00	0.52
ISRG	3D	42.39	0.20	1.00	0.00	0.00	0.42	0.00	0.00
ISRG	5D	49.29	0.22	0.00	0.00	1.00	0.00	0.00	0.49
ISRG	7D	44.02	0.20	1.00	0.00	0.00	0.44	0.00	0.00
LBTYA	3D	34.69	0.17	1.00	0.00	0.00	0.35	0.00	0.00
LBTYA	5D	36.31	0.27	0.83	0.25	0.00	0.33	0.48	0.00
LBTYA	7D	37.32	0.33	0.36	0.58	0.11	0.32	0.40	0.47
LBTYK	3D	33.87	0.17	0.00	0.01	0.99	0.00	0.33	0.34
LBTYK	5D	32.66	0.18	0.00	0.92	0.03	0.00	0.34	0.14
LBTYK	7D	34.28	0.17	1.00	0.00	0.00	0.34	0.00	0.00
LULU	3D	47.26	0.21	0.00	0.00	1.00	0.00	0.00	0.47
LULU	5D	43.61	0.20	1.00	0.00	0.00	0.44	0.00	0.00
LULU	7D	46.86	0.21	0.00	0.00	1.00	0.00	0.00	0.47
MAR	3D	43.81	0.20	0.00	0.00	1.00	0.00	0.00	0.44
MAR	5D	41.99	0.20	1.00	0.00	0.00	0.42	0.00	0.00
MAR	7D	43.41	0.23	0.98	0.00	0.05	0.42	0.00	0.79
MAT	3D	35.50	0.17	1.00	0.00	0.00	0.35	0.00	0.00
MAT	5D	36.51	0.18	1.00	0.00	0.00	0.36	0.00	0.00
MAT	7D	34.69	0.17	0.00	0.00	1.00	0.00	0.00	0.35
MDLZ	3D	36.51	0.18	1.00	0.00	0.00	0.36	0.00	0.00
MDLZ	5D	34.48	0.17	1.00	0.00	0.00	0.34	0.00	0.00
MDLZ	7D	35.09	0.17	1.00	0.00	0.00	0.35	0.00	0.00
MELI	3D	45.03	0.21	0.00	0.00	1.00	0.00	0.00	0.45
MELI	5D	46.05	0.21	0.00	0.00	1.00	0.00	0.00	0.46
MELI	7D	43.00	0.20	0.00	0.00	1.00	0.00	0.00	0.43
MNST	3D	42.80	0.20	0.00	0.00	1.00	0.00	0.00	0.43
MNST	5D	43.81	0.20	1.00	0.00	0.00	0.44	0.00	0.00
MNST	7D	42.60	0.20	1.00	0.00	0.00	0.43	0.00	0.00
MSFT	3D	44.22	0.20	1.00	0.00	0.00	0.44	0.00	0.00
MSFT	5D	43.00	0.20	1.00	0.00	0.00	0.43	0.00	0.00
MSFT	7D	44.42	0.21	1.00	0.00	0.00	0.44	0.00	0.00
MU	3D	41.18	0.19	0.00	0.00	1.00	0.00	0.00	0.41
MU	5D	42.19	0.20	0.00	0.00	1.00	0.00	0.00	0.42

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
MU	7D	45.64	0.21	1.00	0.00	0.00	0.46	0.00	0.00
NFLX	3D	46.25	0.21	1.00	0.00	0.00	0.46	0.00	0.00
NFLX	5D	43.41	0.20	0.00	0.00	1.00	0.00	0.00	0.43
NFLX	7D	45.23	0.21	0.00	0.00	1.00	0.00	0.00	0.45
NTES	3D	46.65	0.21	1.00	0.00	0.00	0.47	0.00	0.00
NTES	5D	48.07	0.22	1.00	0.00	0.00	0.48	0.00	0.00
NTES	7D	46.86	0.21	1.00	0.00	0.00	0.47	0.00	0.00
NVDA	3D	45.84	0.21	1.00	0.00	0.00	0.46	0.00	0.00
NVDA	5D	47.06	0.21	1.00	0.00	0.00	0.47	0.00	0.00
NVDA	7D	45.84	0.21	1.00	0.00	0.00	0.46	0.00	0.00
ORLY	3D	51.32	0.26	0.06	0.00	0.96	0.31	0.00	0.53
ORLY	5D	54.77	0.24	0.00	0.00	1.00	0.00	0.00	0.55
ORLY	7D	53.95	0.23	0.00	0.00	1.00	0.00	0.00	0.54
PAYX	3D	49.29	0.22	0.00	0.00	1.00	0.00	0.00	0.49
PAYX	5D	41.58	0.20	1.00	0.00	0.01	0.41	0.00	0.67
PAYX	7D	38.95	0.19	1.00	0.00	0.00	0.39	0.00	0.00
PCAR	3D	38.13	0.19	0.99	0.01	0.00	0.38	0.20	0.00
PCAR	5D	38.34	0.22	0.06	0.00	0.97	0.35	0.00	0.39
PCAR	7D	26.37	0.14	0.00	1.00	0.00	0.00	0.26	0.00
PEP	3D	34.89	0.22	0.85	0.00	0.11	0.35	0.00	0.37
PEP	5D	34.69	0.17	1.00	0.00	0.00	0.35	0.00	0.00
PEP	7D	43.81	0.20	0.00	0.00	1.00	0.00	0.00	0.44
QCOM	3D	44.83	0.21	1.00	0.00	0.00	0.45	0.00	0.00
QCOM	5D	46.05	0.21	1.00	0.00	0.00	0.46	0.00	0.00
QCOM	7D	46.65	0.21	1.00	0.00	0.00	0.47	0.00	0.00
REGN	3D	40.77	0.19	1.00	0.00	0.00	0.41	0.00	0.00
REGN	5D	41.78	0.20	1.00	0.00	0.00	0.42	0.00	0.00
REGN	7D	38.34	0.18	1.00	0.00	0.00	0.38	0.00	0.00
ROST	3D	46.86	0.21	0.00	0.00	0.99	0.00	0.00	0.47
ROST	5D	44.42	0.21	1.00	0.00	0.00	0.44	0.00	0.00
ROST	7D	41.99	0.20	0.01	0.00	1.00	1.00	0.00	0.42
SBUX	3D	42.80	0.20	0.00	0.00	1.00	0.00	0.00	0.43
SBUX	5D	44.22	0.20	1.00	0.00	0.00	0.44	0.00	0.00
SBUX	7D	44.42	0.21	1.00	0.00	0.00	0.44	0.00	0.00
SIRI	3D	37.32	0.18	1.00	0.00	0.00	0.37	0.00	0.00
SIRI	5D	41.18	0.19	1.00	0.00	0.00	0.41	0.00	0.00
SIRI	7D	39.55	0.19	1.00	0.00	0.00	0.40	0.00	0.00
SNPS	3D	43.81	0.20	1.00	0.00	0.00	0.44	0.00	0.00
SNPS	5D	49.90	0.22	0.00	0.00	1.00	0.00	0.00	0.50

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
SNPS	7D	43.20	0.20	1.00	0.00	0.00	0.43	0.00	0.00
SRCL	3D	34.08	0.17	1.00	0.00	0.00	0.34	0.00	0.00
SRCL	5D	34.89	0.17	1.00	0.00	0.00	0.35	0.00	0.00
SRCL	7D	34.89	0.17	0.00	1.00	0.00	0.00	0.35	0.00
STX	3D	44.62	0.21	1.00	0.00	0.00	0.45	0.00	0.00
STX	5D	34.08	0.22	0.76	0.14	0.00	0.41	0.11	0.00
STX	7D	42.19	0.20	0.98	0.00	0.00	0.43	0.00	0.00
SWKS	3D	46.05	0.21	1.00	0.00	0.00	0.46	0.00	0.00
SWKS	5D	46.05	0.29	0.83	0.01	0.16	0.49	0.05	0.39
SWKS	7D	35.70	0.18	0.00	0.00	1.00	0.00	0.00	0.36
TMUS	3D	43.61	0.20	1.00	0.00	0.00	0.44	0.00	0.00
TMUS	5D	40.37	0.19	0.00	0.00	1.00	0.00	0.00	0.40
TMUS	7D	39.55	0.19	0.00	0.00	1.00	0.00	0.00	0.40
TXN	3D	45.64	0.21	0.00	0.00	1.00	0.00	0.00	0.46
TXN	5D	43.00	0.20	0.00	0.00	1.00	0.00	0.00	0.43
TXN	7D	41.38	0.20	0.00	0.00	1.00	0.00	0.00	0.41
VOD	3D	27.99	0.19	0.83	0.00	0.11	0.31	0.00	0.16
VOD	5D	48.68	0.22	0.00	1.00	0.00	0.00	0.49	0.00
VOD	7D	30.83	0.16	1.00	0.00	0.00	0.31	0.00	0.00
VRSK	3D	41.78	0.20	1.00	0.00	0.00	0.42	0.00	0.00
VRSK	5D	48.07	0.22	0.00	0.00	1.00	0.00	0.00	0.48
VRSK	7D	48.68	0.22	0.00	0.00	1.00	0.00	0.00	0.49
VRTX	3D	40.97	0.19	0.00	0.00	1.00	0.00	0.00	0.41
VRTX	5D	37.73	0.18	1.00	0.00	0.00	0.38	0.00	0.00
VRTX	7D	37.32	0.18	1.00	0.00	0.00	0.37	0.00	0.00
WBA	3D	30.83	0.16	0.00	1.00	0.00	0.00	0.31	0.00
WBA	5D	29.21	0.21	0.00	0.15	0.80	0.00	0.32	0.29
WBA	7D	37.52	0.18	1.00	0.00	0.00	0.38	0.00	0.00
WDC	3D	37.52	0.18	1.00	0.00	0.00	0.38	0.00	0.00
WDC	5D	36.31	0.18	1.00	0.00	0.00	0.36	0.00	0.00
WDC	7D	37.73	0.18	1.00	0.00	0.00	0.38	0.00	0.00
XEL	3D	40.97	0.19	0.00	0.00	1.00	0.00	0.00	0.41
XEL	5D	43.20	0.20	0.00	0.00	1.00	0.00	0.00	0.43
Average results:									
AVG	3D	0.44	0.20	0.61	0.04	0.34	0.26	0.02	0.17
AVG	5D	0.44	0.20	0.63	0.05	0.31	0.28	0.03	0.18
AVG	7D	0.45	0.20	0.64	0.07	0.27	0.28	0.03	0.14
STD	3D	5.40	0.02	0.48	0.19	0.47	0.20	0.08	0.23

Ticker	label	Acc	F1 Score	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)
STD	5D	6.24	0.02	0.47	0.24	0.44	0.20	0.10	0.21
STD	7D	5.85	0.03	0.47	0.20	0.45	0.21	0.14	0.25

Table 7.5: Scores of Supervised Contrastive Learning over NASDAQ dataset.

A detailed look at the results, shown in Table 7.5 reveals variability in the model's performance across different stocks (tickers), with certain instances showcasing higher accuracy and F1 scores. This variability can be attributed to the diverse nature of companies and sectors represented in the dataset, each subject to unique market forces and investor behaviors. The model's ability to achieve notable accuracy, for the domain we are on, in specific cases highlights its potential for targeted financial analysis.

The endeavor to forecast stock market movements is fraught with challenges, primarily due to the dynamic and interconnected nature of global financial systems. The SupCon model's performance, while reflective of these challenges, also indicates the viability of using advanced machine learning techniques, such as contrastive learning, in deciphering the complex patterns underlying market trends.

The evaluation of the Supervised Contrastive Learning model on the financial dataset underscores both the potential and the challenges of applying deep learning techniques to stock market prediction. While the results highlight the difficulty of achieving high accuracy in this domain, they also suggest avenues for future research, particularly in refining models to better understand and predict the multifaceted dynamics of financial markets. The concept of supervised contrastive learning itself emerges as a powerful tool in the realm of deep learning for financial analysis. Its ability to leverage complex patterns from time series data points towards its potential for innovation in this field. Future explorations could include varying the underlying neural network architectures, optimizing hyperparameters, or integrating multimodal data sources to enhance predictive performance. This approach's flexibility and robustness make it a promising avenue for developing more accurate and reliable financial forecasting models, indicating a fertile ground for research and application in deciphering the intricacies of financial markets.

Configuration

A comprehensive grid of parameters was utilized to fine-tune the SupCon model. For each dataset, the best parameter combination was selected based on the evaluation phase, ensuring the most effective configuration for optimal model performance.

- batch size: Options included 16, and 32, allowing for variability in the number of samples processed before the model's internal parameters are updated.

- learning rate: Tested values were 0.1, and 0.01, to adjust the rate at which the model learns during the training phase.
- temperature: Experimented with values of 2, 2.8, and 1.8 to fine-tune the scaling factor applied to the output of the dot product in the contrastive loss calculation.
- epochs: The model was trained for 200 epochs, determining the number of complete passes through the entire training dataset.

7.3.2 Performance Comparison with Different Models

In assessing our Supervised Contrastive Learning (SupCon) model's performance on the financial dataset, it is crucial to compare its results with those of other models. This comparison includes TS2Vec, two shallow models (Gradient Boosting and Random Forest), and ARIMA. Each model brings a unique approach to time series forecasting:

- *TS2Vec*: A time series representation learning model that captures temporal dynamics and dependencies in data.
- *Gradient Boosting*: An ensemble technique that builds models sequentially to correct errors of the predecessors, using decision trees as the base learners.
- *Random Forest*: An ensemble learning method that operates by constructing multiple decision trees during training for more reliable and accurate predictions.
- *ARIMA (AutoRegressive Integrated Moving Average)*: A classic statistical model for analyzing and forecasting time series data, focusing on capturing different aspects of temporal patterns.

These models represent a spectrum of methodologies, from machine learning to traditional statistical approaches, offering a comprehensive view of current capabilities in financial market prediction.

TS2Vec

The analysis of TS2Vec results across the different label types (3 Days, 5 Days, and 7 Days) reveals a clear pattern of performance degradation as the prediction horizon extends. This trend is indicative of the increasing difficulty in forecasting market movements over longer periods using TS2Vec. Specifically, the model struggles with precision and recall, showing a notable decline in accuracy and F1 scores as the label type moves from 3 Days to 7 Days. This pattern underscores the challenges inherent in predicting stock market behavior, where longer-term predictions become progressively uncertain due to the growing influence of unforeseeable market variables and events.

The performance of the TS2Vec model on a 3-day prediction window, Table 7.6, evaluated through the NASDAQ dataset, presents a multifaceted view of its capabilities. With an average accuracy of 0.32, the model demonstrates a modest ability to predict market behaviors accurately. The precision, recall, and F1 scores across three labels (0, 1, and 2) reveal a nuanced performance: Label 1, associated with the highest recall of 0.57, indicates a better model sensitivity for this category, albeit with a lower precision of 0.16, suggesting potential overfitting or misclassification issues for other labels. The balance between precision and recall is further illustrated by the F1 scores, with Label 2 showing a relatively better balance compared to Labels 0 and 1.

	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AVG	0.32	0.27	0.25	0.24	0.16	0.57	0.18	0.27	0.36	0.29
STD	0.21	0.32	0.33	0.29	0.19	0.45	0.15	0.28	0.40	0.30

Table 7.6: Average Scores of TS2Vec on label 3D over NASDAQ dataset.

The standard deviation metrics highlight considerable variability in the model's performance across different runs, especially in precision and recall for all labels, which points to the model's sensitivity to the dataset's characteristics.

Proceeding with an in depth view for the 3 Days label, Table 7.7, TS2Vec exhibited moderate accuracy with certain stocks performing notably better than others, such as *AKAM* and *LULU*, which showcased a decent balance between recall and precision across classes. However, several stocks like *AAPL* and *ADBE* displayed extremely poor performance, indicating a struggle in capturing the short-term market movements accurately.

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AAL	0.55	0.68	0.43	0.53	0.55	0.45	0.50	0.48	0.81	0.60
AAPL	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
ADBE	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
ADI	0.13	0.00	0.00	0.00	0.13	1.00	0.23	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
ADP	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
ADSK	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
AKAM	0.64	0.64	0.79	0.71	0.89	0.07	0.13	0.63	0.83	0.71
ALGN	0.41	0.00	0.00	0.00	0.00	0.00	0.00	0.41	1.00	0.58
AMAT	0.45	0.38	0.01	0.03	0.00	0.00	0.00	0.45	0.98	0.62
AMD	0.52	0.81	0.18	0.30	0.00	0.00	0.00	0.48	0.96	0.64
AMGN	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
AMZN	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
ASML	0.06	0.00	0.00	0.00	0.06	1.00	0.12	0.00	0.00	0.00
ATVI	0.63	0.69	0.73	0.71	0.00	0.00	0.00	0.56	0.80	0.66
AVGO	0.49	0.00	0.00	0.00	0.00	0.00	0.00	0.49	1.00	0.65
BIDU	0.58	0.65	0.72	0.68	0.44	0.26	0.33	0.56	0.69	0.62
BIIB	0.52	0.53	0.76	0.63	0.00	0.00	0.00	0.50	0.67	0.57
BKNG	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
BMRN	0.57	0.56	0.79	0.66	0.00	0.00	0.00	0.59	0.76	0.66
CDNS	0.05	0.00	0.00	0.00	0.05	1.00	0.10	0.00	0.00	0.00
CHKP	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00
CHRW	0.52	0.69	0.43	0.53	0.27	0.06	0.10	0.47	0.89	0.62
CMCSA	0.19	0.00	0.00	0.00	0.19	1.00	0.33	0.00	0.00	0.00
COST	0.11	0.00	0.00	0.00	0.11	1.00	0.19	0.00	0.00	0.00
CSCO	0.57	0.54	0.87	0.67	0.31	0.04	0.07	0.64	0.57	0.61
CSX	0.58	0.53	0.94	0.67	0.00	0.00	0.00	0.79	0.39	0.53
CTSH	0.59	0.68	0.65	0.67	0.34	0.34	0.34	0.66	0.70	0.68
DLTR	0.49	0.45	0.64	0.53	0.00	0.00	0.00	0.55	0.58	0.56
EBAY	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
EXPD	0.51	0.57	0.43	0.49	0.25	0.46	0.32	0.68	0.61	0.64
FAST	0.20	0.00	0.00	0.00	0.20	1.00	0.34	0.00	0.00	0.00
GILD	0.40	0.36	0.98	0.52	0.00	0.00	0.00	0.67	0.25	0.37
GOOG	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
GOOGL	0.11	0.00	0.00	0.00	0.11	1.00	0.20	0.00	0.00	0.00
GRMN	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
HSIC	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
IDXX	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
INCY	0.53	0.54	0.80	0.64	0.00	0.00	0.00	0.53	0.72	0.61
INTC	0.50	0.66	0.43	0.53	0.28	0.34	0.31	0.50	0.70	0.59
INTU	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
ISRG	0.67	0.71	0.65	0.68	0.00	0.00	0.00	0.65	0.80	0.72
LBTYA	0.54	0.60	0.58	0.59	0.48	0.43	0.45	0.54	0.64	0.59
LBTYK	0.57	0.61	0.76	0.68	0.53	0.32	0.40	0.55	0.68	0.61
LULU	0.72	0.75	0.74	0.74	0.00	0.00	0.00	0.69	0.83	0.75
MAR	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
MAT	0.47	0.66	0.56	0.61	0.35	0.66	0.46	0.62	0.22	0.32
MDLZ	0.23	0.00	0.00	0.00	0.23	1.00	0.37	0.00	0.00	0.00
MELI	0.49	0.90	0.14	0.24	0.00	0.00	0.00	0.46	0.98	0.62
MNST	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
MSFT	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
MU	0.53	0.76	0.34	0.47	0.00	0.00	0.00	0.47	0.93	0.63
NFLX	0.61	0.75	0.46	0.57	0.00	0.00	0.00	0.55	0.89	0.68
NTES	0.65	0.78	0.58	0.66	0.00	0.00	0.00	0.58	0.87	0.70
NVDA	0.49	0.00	0.00	0.00	0.00	0.00	0.00	0.49	1.00	0.66
ORLY	0.11	0.00	0.00	0.00	0.11	1.00	0.19	0.00	0.00	0.00
PAYX	0.13	0.00	0.00	0.00	0.13	1.00	0.24	0.00	0.00	0.00
PCAR	0.52	0.60	0.59	0.59	0.29	0.18	0.23	0.54	0.71	0.61
PEP	0.19	0.00	0.00	0.00	0.19	1.00	0.31	0.00	0.00	0.00
QCOM	0.61	0.68	0.58	0.62	1.00	0.10	0.18	0.55	0.83	0.66
REGN	0.49	0.44	0.95	0.60	0.00	0.00	0.00	0.73	0.28	0.40
ROST	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
SBUX	0.15	0.00	0.00	0.00	0.15	1.00	0.25	0.00	0.00	0.00
SIRI	0.26	0.00	0.00	0.00	0.26	1.00	0.41	0.00	0.00	0.00
SNPS	0.08	0.00	0.00	0.00	0.08	1.00	0.14	0.00	0.00	0.00
SRCL	0.49	0.65	0.39	0.49	0.37	0.53	0.43	0.57	0.56	0.57
STX	0.16	0.57	0.08	0.13	0.10	0.56	0.18	0.33	0.10	0.16
SWKS	0.54	0.65	0.50	0.56	0.00	0.00	0.00	0.47	0.81	0.59
TMUS	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
TXN	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
VOD	0.52	0.51	0.64	0.57	0.56	0.42	0.48	0.47	0.56	0.51
VRSK	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
VRTX	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
WBA	0.29	0.00	0.00	0.00	0.00	0.00	0.00	0.29	1.00	0.45
WDC	0.31	0.00	0.00	0.00	0.00	0.00	0.00	0.31	1.00	0.48
XEL	0.16	0.00	0.00	0.00	0.16	1.00	0.27	0.00	0.00	0.00
Average results:										
AVG	0.32	0.27	0.25	0.24	0.16	0.57	0.18	0.27	0.36	0.29
STD	0.21	0.32	0.33	0.29	0.19	0.45	0.15	0.28	0.40	0.30

Table 7.7: Scores of TS2Vec on label 3D over NASDAQ dataset.

Looking at 5-day results, Table 7.8, the model's accuracy stands at a lower average of 0.19, suggesting challenges in general prediction capabilities over this timeframe. Precision and recall metrics highlight a pronounced discrepancy across labels, indicating the model's heightened sensitivity in identifying this particular state, yet this is coupled with a low precision of 0.11, reflecting a tendency for false positives. The F1 scores suggest difficulties in achieving a balanced precision-recall trade-off.

	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AVG	0.19	0.02	0.01	0.01	0.11	0.76	0.19	0.08	0.22	0.12
STD	0.11	0.08	0.09	0.07	0.08	0.43	0.13	0.16	0.41	0.21

Table 7.8: Average Scores of TS2Vec on label 5D over NASDAQ dataset.

The standard deviation values indicate a significant variability in the model's performance, particularly in recall and precision across labels, showcasing the model's fluctuating reliability. This variability, especially noted in the F1 scores, underscores the challenge of model consistency over the 5-day prediction horizon and highlights the importance of further optimization to improve prediction accuracy and stability.

5 Days label complete results, Table 7.9, were generally lower in performance, with many stocks showing a significant drop in accuracy and a tendency toward predicting a single class, often leading to a high recall but extremely low precision for that class. This suggests that as the prediction window widens, TS2Vec finds

it increasingly difficult to make accurate predictions, possibly due to the added complexity and variability in stock price movements over longer periods.

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AAL	0.29	0.00	0.00	0.00	0.00	0.00	0.00	0.29	1.00	0.44
AAPL	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
ADBE	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
ADI	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
ADP	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
ADSK	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
AKAM	0.23	0.00	0.00	0.00	0.23	1.00	0.38	0.00	0.00	0.00
ALGN	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.40	1.00	0.58
AMAT	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
AMD	0.05	0.00	0.00	0.00	0.05	1.00	0.10	0.00	0.00	0.00
AMGN	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
AMZN	0.06	0.00	0.00	0.00	0.06	1.00	0.11	0.00	0.00	0.00
ASML	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
ATVI	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.34	1.00	0.51
AVGO	0.48	0.00	0.00	0.00	0.00	0.00	0.00	0.48	1.00	0.65
BIDU	0.25	0.00	0.00	0.00	0.25	1.00	0.40	0.00	0.00	0.00
BIIB	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.32	1.00	0.49
BKNG	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00
BMRN	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.34	1.00	0.51
CDNS	0.06	0.00	0.00	0.00	0.06	1.00	0.12	0.00	0.00	0.00
CHKP	0.20	0.00	0.00	0.00	0.20	1.00	0.33	0.00	0.00	0.00
CHRW	0.24	0.00	0.00	0.00	0.24	1.00	0.39	0.00	0.00	0.00
CMCSA	0.19	0.00	0.00	0.00	0.19	1.00	0.31	0.00	0.00	0.00
COST	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
CSCO	0.23	0.00	0.00	0.00	0.23	1.00	0.37	0.00	0.00	0.00
CSX	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
CTSH	0.26	0.00	0.00	0.00	0.26	1.00	0.41	0.00	0.00	0.00
DLTR	0.15	0.00	0.00	0.00	0.15	1.00	0.27	0.00	0.00	0.00
EBAY	0.13	0.00	0.00	0.00	0.13	1.00	0.23	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
EXPD	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
FAST	0.21	0.00	0.00	0.00	0.21	1.00	0.35	0.00	0.00	0.00
GILD	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.34	1.00	0.51
GOOG	0.10	0.00	0.00	0.00	0.10	1.00	0.18	0.00	0.00	0.00
GOOGL	0.10	0.00	0.00	0.00	0.10	1.00	0.17	0.00	0.00	0.00
GRMN	0.16	0.00	0.00	0.00	0.16	1.00	0.27	0.00	0.00	0.00
HSIC	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00
IDXX	0.06	0.00	0.00	0.00	0.06	1.00	0.11	0.00	0.00	0.00
INCY	0.33	0.00	0.00	0.00	0.00	0.00	0.00	0.33	1.00	0.50
INTC	0.21	0.00	0.00	0.00	0.21	1.00	0.35	0.00	0.00	0.00
INTU	0.07	0.00	0.00	0.00	0.07	1.00	0.14	0.00	0.00	0.00
ISRG	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
LBTYA	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.30	1.00	0.46
LBTYK	0.31	0.00	0.00	0.00	0.00	0.00	0.00	0.31	1.00	0.48
LULU	0.11	0.00	0.00	0.00	0.11	1.00	0.19	0.00	0.00	0.00
MAR	0.13	0.00	0.00	0.00	0.13	1.00	0.23	0.00	0.00	0.00
MAT	0.28	0.00	0.00	0.00	0.28	1.00	0.43	0.00	0.00	0.00
MDLZ	0.25	0.00	0.00	0.00	0.25	1.00	0.39	0.00	0.00	0.00
MELI	0.05	0.00	0.00	0.00	0.05	1.00	0.10	0.00	0.00	0.00
MNST	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
MSFT	0.10	0.00	0.00	0.00	0.10	1.00	0.19	0.00	0.00	0.00
MU	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
NFLX	0.10	0.00	0.00	0.00	0.10	1.00	0.18	0.00	0.00	0.00
NTES	0.42	0.31	0.02	0.04	0.00	0.00	0.00	0.42	0.96	0.58
NVDA	0.36	0.00	0.00	0.00	0.00	0.00	0.00	0.50	0.72	0.59
ORLY	0.10	0.00	0.00	0.00	0.10	1.00	0.17	0.00	0.00	0.00
PAYX	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
PCAR	0.24	0.00	0.00	0.00	0.24	1.00	0.39	0.00	0.00	0.00
PEP	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
QCOM	0.16	0.00	0.00	0.00	0.16	1.00	0.27	0.00	0.00	0.00
REGN	0.44	0.43	0.81	0.56	0.00	0.00	0.00	0.49	0.24	0.32
ROST	0.13	0.00	0.00	0.00	0.13	1.00	0.22	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
SBUX	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
SIRI	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
SNPS	0.05	0.00	0.00	0.00	0.05	1.00	0.10	0.00	0.00	0.00
SRCL	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.28	1.00	0.44
STX	0.17	0.00	0.00	0.00	0.17	1.00	0.29	0.00	0.00	0.00
SWKS	0.40	0.54	0.12	0.20	0.00	0.00	0.00	0.38	0.91	0.54
TMUS	0.15	0.00	0.00	0.00	0.15	1.00	0.25	0.00	0.00	0.00
TXN	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
VOD	0.22	0.00	0.00	0.00	0.00	0.00	0.00	0.22	1.00	0.36
VRSK	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
VRTX	0.23	0.00	0.00	0.00	0.23	1.00	0.37	0.00	0.00	0.00
WBA	0.31	0.00	0.00	0.00	0.00	0.00	0.00	0.31	1.00	0.47
WDC	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.32	1.00	0.48
XEL	0.16	0.00	0.00	0.00	0.16	1.00	0.27	0.00	0.00	0.00
Average results:										
AVG	0.19	0.02	0.01	0.01	0.11	0.76	0.19	0.08	0.22	0.12
STD	0.11	0.08	0.09	0.07	0.08	0.43	0.13	0.16	0.41	0.21

Table 7.9: Scores of TS2Vec on label 5D over NASDAQ dataset.

The TS2Vec model's performance over a 7-day forecasting period, Table 7.10, displays a slightly improved accuracy of 0.20, offering a glimpse into its predictive nuances for longer-term forecasts. The model exhibits low precision and recall for Label 0 (Precision: 0.02, Recall: 0.01), indicating challenges in accurately identifying this particular market state, as reflected in the negligible F1 score (0.01). Contrastingly, Label 1 shows a stronger recall of 0.72, suggesting a relative sensitivity in detecting this label. Label 2 presents an improvement in both recall (0.26) and precision (0.09), achieving an F1 score of 0.14, hinting at a slightly better balance in prediction performance for this category.

	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AVG	0.20	0.02	0.01	0.01	0.11	0.72	0.18	0.09	0.26	0.14
STD	0.11	0.08	0.10	0.06	0.08	0.45	0.14	0.17	0.43	0.23

Table 7.10: Average Scores of TS2Vec on label 7D over NASDAQ dataset.

The standard deviation metrics highlight a considerable spread in performance across evaluations, especially notable in the recall and F1 scores for Labels 1 and 2, which indicates inconsistency in the model's predictive reliability across different instances of the dataset. This variability underscores the necessity for model adjustments and enhancements to bolster predictive consistency and accuracy in longer-duration forecasts.

The 7 Days label, shown in Table 7.11, analysis further confirmed the trend of declining performance with an extended forecast horizon. Despite a few exceptions, the overall ability of TS2Vec to generalize and accurately predict market movements significantly waned, with most stocks showing an inclination towards single-class predictions.

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
AAL	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.28	1.00	0.44
AAPL	0.12	0.00	0.00	0.00	0.12	1.00	0.21	0.00	0.00	0.00
ADBE	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
ADI	0.13	0.00	0.00	0.00	0.13	1.00	0.23	0.00	0.00	0.00
ADP	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
ADSK	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
AKAM	0.24	0.00	0.00	0.00	0.24	1.00	0.38	0.00	0.00	0.00
ALGN	0.41	0.00	0.00	0.00	0.00	0.00	0.00	0.41	1.00	0.58
AMAT	0.09	0.00	0.00	0.00	0.09	1.00	0.16	0.00	0.00	0.00
AMD	0.45	0.00	0.00	0.00	0.00	0.00	0.00	0.45	1.00	0.62
AMGN	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
AMZN	0.39	0.00	0.00	0.00	0.00	0.00	0.00	0.39	1.00	0.56
ASML	0.06	0.00	0.00	0.00	0.06	1.00	0.12	0.00	0.00	0.00
ATVI	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.34	1.00	0.51
AVGO	0.49	0.00	0.00	0.00	0.00	0.00	0.00	0.49	1.00	0.65
BIDU	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.30	1.00	0.46
BIIB	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.32	1.00	0.49
BKNG	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
BMRN	0.26	0.00	0.00	0.00	0.26	1.00	0.41	0.00	0.00	0.00
CDNS	0.05	0.00	0.00	0.00	0.05	1.00	0.10	0.00	0.00	0.00
CHKP	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
CHRW	0.24	0.00	0.00	0.00	0.24	1.00	0.38	0.00	0.00	0.00
CMCSA	0.19	0.00	0.00	0.00	0.19	1.00	0.33	0.00	0.00	0.00
COST	0.11	0.00	0.00	0.00	0.11	1.00	0.19	0.00	0.00	0.00
CSCO	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
CSX	0.13	0.00	0.00	0.00	0.13	1.00	0.24	0.00	0.00	0.00
CTSH	0.24	0.00	0.00	0.00	0.24	1.00	0.38	0.00	0.00	0.00
DLTR	0.19	0.00	0.00	0.00	0.19	1.00	0.32	0.00	0.00	0.00
EBAY	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
EXPD	0.17	0.00	0.00	0.00	0.17	1.00	0.29	0.00	0.00	0.00
FAST	0.20	0.00	0.00	0.00	0.20	1.00	0.34	0.00	0.00	0.00
GILD	0.35	0.19	0.04	0.06	0.00	0.00	0.00	0.37	0.93	0.53
GOOG	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
GOOGL	0.11	0.00	0.00	0.00	0.11	1.00	0.20	0.00	0.00	0.00
GRMN	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
HSIC	0.18	0.00	0.00	0.00	0.18	1.00	0.31	0.00	0.00	0.00
IDXX	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
INCY	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.34	1.00	0.51
INTC	0.19	0.00	0.00	0.00	0.19	1.00	0.32	0.00	0.00	0.00
INTU	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
ISRG	0.08	0.00	0.00	0.00	0.08	1.00	0.14	0.00	0.00	0.00
LBTYA	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.28	1.00	0.44
LBTYK	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.28	1.00	0.44
LULU	0.08	0.00	0.00	0.00	0.08	1.00	0.15	0.00	0.00	0.00
MAR	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
MAT	0.31	0.00	0.00	0.00	0.31	1.00	0.47	0.00	0.00	0.00
MDLZ	0.23	0.00	0.00	0.00	0.23	1.00	0.37	0.00	0.00	0.00
MELI	0.06	0.00	0.00	0.00	0.06	1.00	0.12	0.00	0.00	0.00
MNST	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
MSFT	0.07	0.00	0.00	0.00	0.07	1.00	0.13	0.00	0.00	0.00
MU	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.40	1.00	0.58
NFLX	0.10	0.00	0.00	0.00	0.10	1.00	0.18	0.00	0.00	0.00
NTES	0.43	0.25	0.02	0.04	0.00	0.00	0.00	0.44	0.95	0.60

Ticker	Acc	Prec(0)	Rec(0)	F1(0)	Prec(1)	Rec(1)	F1(1)	Prec(2)	Rec(2)	F1(2)
NVDA	0.49	0.00	0.00	0.00	0.00	0.00	0.00	0.49	1.00	0.66
ORLY	0.11	0.00	0.00	0.00	0.11	1.00	0.19	0.00	0.00	0.00
PAYX	0.13	0.00	0.00	0.00	0.13	1.00	0.24	0.00	0.00	0.00
PCAR	0.26	0.00	0.00	0.00	0.26	1.00	0.42	0.00	0.00	0.00
PEP	0.19	0.00	0.00	0.00	0.19	1.00	0.31	0.00	0.00	0.00
QCOM	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
REGN	0.40	0.39	0.85	0.53	0.00	0.00	0.00	0.48	0.18	0.26
ROST	0.09	0.00	0.00	0.00	0.09	1.00	0.17	0.00	0.00	0.00
SBUX	0.15	0.00	0.00	0.00	0.15	1.00	0.25	0.00	0.00	0.00
SIRI	0.26	0.00	0.00	0.00	0.26	1.00	0.41	0.00	0.00	0.00
SNPS	0.08	0.00	0.00	0.00	0.08	1.00	0.14	0.00	0.00	0.00
SRCL	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.28	1.00	0.44
STX	0.15	0.00	0.00	0.00	0.15	1.00	0.26	0.00	0.00	0.00
SWKS	0.38	0.53	0.12	0.19	0.00	0.00	0.00	0.36	0.90	0.52
TMUS	0.14	0.00	0.00	0.00	0.14	1.00	0.25	0.00	0.00	0.00
TXN	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
VOD	0.22	0.00	0.00	0.00	0.00	0.00	0.00	0.22	1.00	0.36
VRSK	0.12	0.00	0.00	0.00	0.12	1.00	0.22	0.00	0.00	0.00
VRTX	0.22	0.00	0.00	0.00	0.22	1.00	0.36	0.00	0.00	0.00
WBA	0.29	0.00	0.00	0.00	0.00	0.00	0.00	0.29	1.00	0.45
WDC	0.31	0.00	0.00	0.00	0.00	0.00	0.00	0.31	1.00	0.48
XEL	0.16	0.00	0.00	0.00	0.16	1.00	0.27	0.00	0.00	0.00
Average results:										
AVG	0.20	0.02	0.01	0.01	0.11	0.72	0.18	0.09	0.26	0.14
STD	0.11	0.08	0.10	0.06	0.08	0.45	0.14	0.17	0.43	0.23

Table 7.11: Scores of TS2Vec on label 7D over NASDAQ dataset.

Comparing TS2Vec with the Supervised Contrastive Learning (SupCon) approach, SupCon demonstrates a more robust performance across all label types. While both models face challenges in dealing with the unpredictable nature of stock market data, SupCon's method of leveraging contrastive learning appears to afford it a resilience in capturing complex patterns and relationships within the data. This is evidenced by its relatively higher accuracy and F1 scores. The comparison suggests that SupCon's methodology, which focuses on learning representations by contrasting positive and negative examples, may be better suited for the task of financial time series forecasting, particularly in handling the nuanced and often non-linear relationships that characterize market data.

This analysis underscores the value of exploring different deep learning architectures and methodologies in the quest for more effective and reliable stock market prediction models. The contrastive learning approach, as exemplified by SupCon, offers a promising avenue for future research, potentially paving the way for advancements in predictive accuracy and model robustness in the face of the stock market's inherent volatility and complexity.

For the TS2Vec model, a standardized configuration was adopted to maintain consistency across experiments:

- batch size: A fixed batch size of 8 was used, dictating the number of data points processed in a single training step.
- epochs: The model underwent training for 30 epochs, providing a sufficient number of iterations over the dataset to achieve substantial learning and adaptation.

Shallow Models

The summarized performance metrics for Random Forest (RF) and Gradient Boosting (GBoost) models, Table 7.12, over the NASDAQ dataset reveal insights into their predictive abilities across different forecasting periods. Both models exhibit comparable performance levels with accuracy hovering around the 0.40 mark, indicating a moderate ability to correctly predict market states.

For RF, accuracy remains consistent across all prediction windows, suggesting stability in its predictive capability. Precision and recall metrics across different labels (0, 1, 2) for RF indicate a balanced ability to identify each market state, with notable precision and recall for label 0 and label 2. However, both models struggle with label 1, showing low precision and recall, highlighting difficulties in predicting this specific market condition effectively.
GBoost shows a slight improvement in accuracy for the 5D predictions and maintains similar performance levels as RF in other aspects. It exhibits a slightly better or comparable precision and recall for most labels across different time frames, suggesting a nuanced edge in handling certain market conditions more effectively than RF.

	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
AVG	RF	3D	0.39	0.40	0.47	0.15	0.11	0.43	0.43
AVG	RF	5D	0.40	0.42	0.51	0.16	0.09	0.43	0.40
AVG	RF	7D	0.40	0.39	0.44	0.17	0.10	0.42	0.46
STD	RF	3D	0.07	0.09	0.27	0.13	0.15	0.08	0.26
STD	RF	5D	0.06	0.10	0.27	0.18	0.14	0.13	0.27
STD	RF	7D	0.06	0.12	0.29	0.19	0.15	0.12	0.29
Average results:									
AVG	GBoost	3D	0.40	0.42	0.46	0.16	0.11	0.43	0.44
AVG	GBoost	5D	0.41	0.43	0.49	0.14	0.08	0.42	0.43
AVG	GBoost	7D	0.40	0.43	0.41	0.16	0.11	0.43	0.49
STD	GBoost	3D	0.07	0.06	0.29	0.21	0.19	0.08	0.28
STD	GBoost	5D	0.07	0.10	0.29	0.18	0.17	0.13	0.29
STD	GBoost	7D	0.06	0.13	0.30	0.22	0.18	0.12	0.29

Table 7.12: Average Scores of Random Forest (RF) and Gradient Boosting (GBoost) over NASDAQ dataset.

The standard deviation values across metrics for both RF and GBoost indicate a degree of variability in their performance, with GBoost showing a tendency for higher variability in precision and recall for label 1 across prediction windows. This suggests potential areas for model refinement, particularly in improving consistency and reliability in predictions across different market states.

The examination of Random Forest (RF) and Gradient Boosting (GBoost) models' performance, Table 7.13, more in deep reveals distinct patterns and efficacy in stock market prediction. Both models showcased variability in accuracy, precision, recall across different labels, indicating the complexity of stock movement prediction and the sensitivity of models to different temporal scales.

For *Random Forest*, the average results show a moderate performance with slight variations across different prediction horizons. The 3-day label predictions exhibit a good response on average accuracy, precision, and recall, suggesting a balanced but not highly accurate prediction capability. The 5-day and 7-day labels show a slight increase in accuracy and precision, indicating a potentially better alignment of the model's predictive capability with slightly longer prediction horizons.

Gradient Boosting, on the other hand, displays a similar trend but with a generally higher performance in accuracy and precision across the prediction labels. This suggests that Gradient Boosting might be more adept at capturing the nuances of the stock market data for prediction purposes.

Comparing these results to the SupCon model, it's evident that SupCon provides a different approach to understanding and predicting stock market movements. While RF and GBoost rely on traditional feature-based learning, SupCon's contrastive learning framework might offer advantages in capturing complex, non-linear relationships in the data. The choice between these models could depend on specific use cases, data characteristics, and the trade-off between interpretability and predictive performance.

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
AAL	RF	3D	0.31	0.40	0.40	0.25	0.20	0.27	0.34
AAL	RF	5D	0.38	0.40	0.61	0.40	0.42	0.21	0.08
AAL	RF	7D	0.36	0.38	0.65	0.43	0.30	0.14	0.08
AAL	GBoost	3D	0.34	0.41	0.29	0.36	0.31	0.28	0.41
AAL	GBoost	5D	0.39	0.40	0.52	0.42	0.52	0.21	0.08
AAL	GBoost	7D	0.39	0.41	0.52	0.47	0.32	0.29	0.32
AAPL	RF	3D	0.15	0.39	0.17	0.04	0.55	0.67	0.08
AAPL	RF	5D	0.42	0.28	0.10	0.00	0.00	0.44	0.82
AAPL	RF	7D	0.39	0.28	0.09	0.00	0.00	0.41	0.81
AAPL	GBoost	3D	0.25	0.43	0.37	0.04	0.38	0.54	0.11
AAPL	GBoost	5D	0.43	0.36	0.14	0.00	0.00	0.45	0.80
AAPL	GBoost	7D	0.40	0.30	0.09	0.00	0.00	0.41	0.82
ADBE	RF	3D	0.49	0.51	0.31	0.25	0.03	0.48	0.72
ADBE	RF	5D	0.43	0.43	0.53	0.00	0.00	0.45	0.40
ADBE	RF	7D	0.41	0.44	0.24	0.07	0.13	0.49	0.60
ADBE	GBoost	3D	0.48	0.49	0.22	0.00	0.00	0.48	0.80
ADBE	GBoost	5D	0.46	0.45	0.50	0.00	0.00	0.47	0.48
ADBE	GBoost	7D	0.45	0.46	0.30	0.09	0.07	0.48	0.65
ADI	RF	3D	0.44	0.00	0.00	0.06	0.02	0.45	0.97
ADI	RF	5D	0.45	0.44	0.95	0.00	0.00	0.58	0.08
ADI	RF	7D	0.47	0.45	0.97	0.00	0.00	0.76	0.11
ADI	GBoost	3D	0.45	0.50	0.01	0.00	0.00	0.45	0.99
ADI	GBoost	5D	0.45	0.44	0.95	0.00	0.00	0.55	0.08
ADI	GBoost	7D	0.47	0.45	0.96	0.00	0.00	0.74	0.13
ADP	RF	3D	0.39	0.35	0.58	0.00	0.00	0.46	0.34
ADP	RF	5D	0.40	0.39	0.97	0.00	0.00	0.56	0.04
ADP	RF	7D	0.40	0.40	0.99	0.14	0.02	0.83	0.02
ADP	GBoost	3D	0.15	0.28	0.04	0.13	0.90	0.33	0.03
ADP	GBoost	5D	0.14	0.27	0.04	0.13	0.94	0.44	0.02
ADP	GBoost	7D	0.40	0.40	0.99	0.00	0.00	0.88	0.03
ADSK	RF	3D	0.44	0.51	0.41	0.08	0.10	0.46	0.53
ADSK	RF	5D	0.45	0.51	0.47	0.17	0.07	0.43	0.54
ADSK	RF	7D	0.47	0.51	0.48	0.06	0.04	0.48	0.54
ADSK	GBoost	3D	0.43	0.55	0.29	0.07	0.12	0.46	0.64

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
ADSK	GBoost	5D	0.45	0.50	0.41	0.00	0.00	0.41	0.62
ADSK	GBoost	7D	0.46	0.49	0.44	0.00	0.00	0.43	0.58
AKAM	RF	3D	0.37	0.40	0.53	0.16	0.08	0.40	0.39
AKAM	RF	5D	0.41	0.42	0.44	0.26	0.13	0.42	0.54
AKAM	RF	7D	0.39	0.45	0.48	0.17	0.05	0.38	0.52
AKAM	GBoost	3D	0.39	0.42	0.48	0.26	0.13	0.40	0.47
AKAM	GBoost	5D	0.43	0.47	0.48	0.22	0.02	0.41	0.64
AKAM	GBoost	7D	0.37	0.40	0.41	0.12	0.03	0.38	0.56
ALGN	RF	3D	0.47	0.48	0.62	0.00	0.00	0.44	0.36
ALGN	RF	5D	0.50	0.55	0.69	1.00	0.03	0.39	0.31
ALGN	RF	7D	0.47	0.65	0.19	0.00	0.00	0.44	0.89
ALGN	GBoost	3D	0.49	0.51	0.64	0.00	0.00	0.45	0.39
ALGN	GBoost	5D	0.51	0.57	0.69	0.00	0.00	0.40	0.36
ALGN	GBoost	7D	0.47	0.65	0.23	0.00	0.00	0.43	0.85
AMAT	RF	3D	0.49	0.48	0.95	0.00	0.00	0.57	0.07
AMAT	RF	5D	0.46	0.45	0.96	0.00	0.00	0.57	0.06
AMAT	RF	7D	0.44	0.23	0.01	0.00	0.00	0.44	0.96
AMAT	GBoost	3D	0.46	0.47	0.92	0.00	0.00	0.48	0.04
AMAT	GBoost	5D	0.46	0.45	0.96	0.00	0.00	0.60	0.06
AMAT	GBoost	7D	0.43	0.17	0.01	0.00	0.00	0.44	0.94
AMD	RF	3D	0.31	0.39	0.17	0.02	0.11	0.46	0.46
AMD	RF	5D	0.47	0.46	0.34	0.00	0.00	0.48	0.66
AMD	RF	7D	0.47	0.46	0.29	0.00	0.00	0.47	0.74
AMD	GBoost	3D	0.29	0.43	0.20	0.03	0.21	0.49	0.39
AMD	GBoost	5D	0.48	0.48	0.50	0.00	0.00	0.49	0.52
AMD	GBoost	7D	0.44	0.43	0.25	0.00	0.00	0.45	0.73
AMGN	RF	3D	0.43	0.44	0.48	0.00	0.00	0.44	0.60
AMGN	RF	5D	0.45	0.44	0.70	0.10	0.02	0.52	0.43
AMGN	RF	7D	0.42	0.42	0.60	0.29	0.02	0.42	0.45
AMGN	GBoost	3D	0.44	0.44	0.47	0.00	0.00	0.44	0.64
AMGN	GBoost	5D	0.45	0.44	0.65	0.00	0.00	0.46	0.49
AMGN	GBoost	7D	0.48	0.45	0.83	0.00	0.00	0.55	0.38
AMZN	RF	3D	0.50	0.53	0.64	0.00	0.00	0.44	0.40
AMZN	RF	5D	0.46	0.52	0.55	0.00	0.00	0.40	0.42
AMZN	RF	7D	0.45	0.55	0.43	0.50	0.06	0.37	0.54
AMZN	GBoost	3D	0.48	0.51	0.63	0.00	0.00	0.43	0.37
AMZN	GBoost	5D	0.41	0.49	0.44	0.00	0.00	0.36	0.43
AMZN	GBoost	7D	0.45	0.56	0.38	0.00	0.00	0.38	0.62
ASML	RF	3D	0.49	0.62	0.07	0.00	0.00	0.48	0.96
ASML	RF	5D	0.48	1.00	0.00	0.00	0.00	0.48	1.00
ASML	RF	7D	0.48	0.00	0.00	0.00	0.00	0.48	1.00
ASML	GBoost	3D	0.48	0.50	0.01	0.00	0.00	0.48	0.99
ASML	GBoost	5D	0.48	1.00	0.00	0.00	0.00	0.48	1.00

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
ASML	GBoost	7D	0.48	1.00	0.00	0.00	0.00	0.48	1.00
ATVI	RF	3D	0.38	0.40	0.27	0.33	0.01	0.37	0.68
ATVI	RF	5D	0.39	0.49	0.30	0.08	0.01	0.35	0.72
ATVI	RF	7D	0.36	0.46	0.19	0.00	0.00	0.34	0.80
ATVI	GBoost	3D	0.39	0.42	0.25	0.00	0.00	0.38	0.75
ATVI	GBoost	5D	0.39	0.47	0.32	0.00	0.00	0.36	0.69
ATVI	GBoost	7D	0.37	0.49	0.21	0.00	0.00	0.33	0.78
AVGO	RF	3D	0.48	0.00	0.00	0.00	0.00	0.48	1.00
AVGO	RF	5D	0.48	0.00	0.00	0.00	0.00	0.48	1.00
AVGO	RF	7D	0.49	0.00	0.00	0.00	0.00	0.49	1.00
AVGO	GBoost	3D	0.48	0.50	0.00	0.00	0.00	0.48	1.00
AVGO	GBoost	5D	0.48	0.00	0.00	0.00	0.00	0.48	1.00
AVGO	GBoost	7D	0.49	0.00	0.00	0.00	0.00	0.49	1.00
BIDU	RF	3D	0.32	0.33	0.21	0.30	0.32	0.32	0.44
BIDU	RF	5D	0.33	0.39	0.22	0.26	0.26	0.33	0.52
BIDU	RF	7D	0.32	0.44	0.25	0.24	0.22	0.31	0.52
BIDU	GBoost	3D	0.36	0.39	0.32	0.34	0.23	0.35	0.52
BIDU	GBoost	5D	0.33	0.44	0.26	0.30	0.29	0.30	0.46
BIDU	GBoost	7D	0.32	0.43	0.30	0.25	0.21	0.28	0.44
BIIB	RF	3D	0.40	0.48	0.37	0.25	0.17	0.39	0.58
BIIB	RF	5D	0.39	0.49	0.44	0.23	0.10	0.35	0.56
BIIB	RF	7D	0.38	0.44	0.43	0.22	0.06	0.35	0.59
BIIB	GBoost	3D	0.37	0.41	0.27	0.33	0.03	0.36	0.68
BIIB	GBoost	5D	0.39	0.47	0.39	0.29	0.03	0.35	0.68
BIIB	GBoost	7D	0.37	0.43	0.36	0.00	0.00	0.34	0.71
BKNG	RF	3D	0.40	0.41	0.68	0.17	0.06	0.42	0.27
BKNG	RF	5D	0.43	0.43	0.77	0.24	0.14	0.53	0.23
BKNG	RF	7D	0.45	0.45	0.77	0.16	0.08	0.57	0.28
BKNG	GBoost	3D	0.43	0.42	0.30	0.23	0.04	0.44	0.71
BKNG	GBoost	5D	0.43	0.43	0.78	0.23	0.10	0.50	0.24
BKNG	GBoost	7D	0.33	0.40	0.37	0.12	0.30	0.58	0.30
BMRN	RF	3D	0.35	0.38	0.46	0.25	0.15	0.36	0.41
BMRN	RF	5D	0.34	0.38	0.41	0.28	0.18	0.34	0.41
BMRN	RF	7D	0.37	0.40	0.48	0.21	0.08	0.36	0.45
BMRN	GBoost	3D	0.36	0.35	0.41	0.20	0.01	0.37	0.58
BMRN	GBoost	5D	0.37	0.39	0.33	0.20	0.01	0.37	0.73
BMRN	GBoost	7D	0.40	0.46	0.41	0.00	0.00	0.36	0.67
CDNS	RF	3D	0.46	0.27	0.06	0.00	0.00	0.48	0.87
CDNS	RF	5D	0.50	0.44	0.03	0.00	0.00	0.51	0.97
CDNS	RF	7D	0.48	0.46	0.96	0.00	0.00	0.69	0.11
CDNS	GBoost	3D	0.46	0.29	0.07	0.00	0.00	0.48	0.86
CDNS	GBoost	5D	0.50	0.47	0.31	0.00	0.00	0.52	0.72
CDNS	GBoost	7D	0.49	0.40	0.05	0.00	0.00	0.49	0.93

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
CHKP	RF	3D	0.39	0.44	0.53	0.25	0.20	0.38	0.34
CHKP	RF	5D	0.39	0.41	0.63	0.16	0.07	0.41	0.30
CHKP	RF	7D	0.38	0.43	0.58	0.21	0.23	0.40	0.23
CHKP	GBoost	3D	0.36	0.39	0.57	0.19	0.05	0.34	0.30
CHKP	GBoost	5D	0.36	0.39	0.59	0.12	0.04	0.35	0.28
CHKP	GBoost	7D	0.37	0.41	0.62	0.16	0.16	0.44	0.19
CHRW	RF	3D	0.33	0.40	0.43	0.17	0.26	0.36	0.25
CHRW	RF	5D	0.39	0.40	0.61	0.42	0.17	0.37	0.32
CHRW	RF	7D	0.39	0.46	0.44	0.26	0.24	0.39	0.42
CHRW	GBoost	3D	0.35	0.38	0.44	0.19	0.19	0.41	0.34
CHRW	GBoost	5D	0.37	0.39	0.58	0.29	0.04	0.35	0.36
CHRW	GBoost	7D	0.39	0.49	0.32	0.29	0.24	0.37	0.55
CMCSA	RF	3D	0.39	0.43	0.41	0.19	0.10	0.40	0.50
CMCSA	RF	5D	0.27	0.37	0.26	0.16	0.36	0.35	0.25
CMCSA	RF	7D	0.32	0.33	0.20	0.09	0.06	0.36	0.59
CMCSA	GBoost	3D	0.29	0.32	0.10	0.18	0.36	0.37	0.45
CMCSA	GBoost	5D	0.31	0.45	0.28	0.17	0.37	0.39	0.32
CMCSA	GBoost	7D	0.33	0.35	0.18	0.15	0.14	0.38	0.59
COST	RF	3D	0.41	0.40	0.93	0.00	0.00	0.53	0.08
COST	RF	5D	0.43	0.40	0.97	0.33	0.02	0.82	0.10
COST	RF	7D	0.41	0.39	0.95	0.00	0.00	0.68	0.11
COST	GBoost	3D	0.40	0.40	0.90	0.00	0.00	0.44	0.09
COST	GBoost	5D	0.43	0.41	0.98	0.00	0.00	0.84	0.10
COST	GBoost	7D	0.42	0.39	0.95	0.00	0.00	0.66	0.12
CSCO	RF	3D	0.43	0.47	0.50	0.20	0.08	0.43	0.52
CSCO	RF	5D	0.39	0.40	0.51	0.29	0.14	0.40	0.41
CSCO	RF	7D	0.40	0.41	0.66	0.26	0.16	0.46	0.28
CSCO	GBoost	3D	0.45	0.47	0.44	0.31	0.13	0.46	0.61
CSCO	GBoost	5D	0.39	0.38	0.49	0.30	0.14	0.43	0.45
CSCO	GBoost	7D	0.41	0.40	0.68	0.47	0.06	0.41	0.33
CSX	RF	3D	0.38	0.43	0.61	0.17	0.29	0.53	0.17
CSX	RF	5D	0.46	0.49	0.43	0.00	0.00	0.44	0.65
CSX	RF	7D	0.39	0.46	0.46	0.00	0.00	0.37	0.45
CSX	GBoost	3D	0.43	0.44	0.90	0.21	0.08	0.59	0.09
CSX	GBoost	5D	0.52	0.49	0.87	0.00	0.00	0.62	0.33
CSX	GBoost	7D	0.44	0.47	0.55	0.17	0.06	0.42	0.43
CTSH	RF	3D	0.37	0.44	0.52	0.18	0.16	0.38	0.32
CTSH	RF	5D	0.33	0.40	0.38	0.26	0.31	0.32	0.29
CTSH	RF	7D	0.32	0.41	0.29	0.25	0.25	0.29	0.40
CTSH	GBoost	3D	0.38	0.44	0.53	0.18	0.10	0.35	0.37
CTSH	GBoost	5D	0.33	0.40	0.41	0.25	0.31	0.32	0.25
CTSH	GBoost	7D	0.36	0.43	0.38	0.33	0.24	0.31	0.42
DLTR	RF	3D	0.38	0.38	0.71	0.25	0.06	0.40	0.20

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
DLTR	RF	5D	0.39	0.40	0.78	0.10	0.01	0.36	0.16
DLTR	RF	7D	0.40	0.39	0.78	0.19	0.04	0.46	0.20
DLTR	GBoost	3D	0.39	0.40	0.75	0.00	0.00	0.37	0.21
DLTR	GBoost	5D	0.39	0.40	0.81	0.17	0.01	0.34	0.13
DLTR	GBoost	7D	0.41	0.41	0.79	0.56	0.05	0.43	0.22
EBAY	RF	3D	0.42	0.43	0.24	0.35	0.15	0.42	0.72
EBAY	RF	5D	0.43	0.49	0.70	0.11	0.03	0.33	0.23
EBAY	RF	7D	0.36	0.39	0.22	0.18	0.04	0.36	0.67
EBAY	GBoost	3D	0.43	0.48	0.22	0.27	0.07	0.43	0.80
EBAY	GBoost	5D	0.45	0.50	0.72	0.17	0.01	0.34	0.27
EBAY	GBoost	7D	0.41	0.47	0.54	0.33	0.01	0.35	0.42
EXPD	RF	3D	0.42	0.41	0.93	0.29	0.02	0.51	0.10
EXPD	RF	5D	0.41	0.41	0.65	0.25	0.01	0.39	0.29
EXPD	RF	7D	0.40	0.40	0.81	0.25	0.01	0.42	0.16
EXPD	GBoost	3D	0.41	0.40	0.95	0.00	0.00	0.56	0.09
EXPD	GBoost	5D	0.39	0.39	0.63	0.00	0.00	0.39	0.29
EXPD	GBoost	7D	0.41	0.41	0.96	0.00	0.00	0.44	0.04
FAST	RF	3D	0.39	0.41	0.90	0.10	0.03	0.41	0.07
FAST	RF	5D	0.39	0.41	0.85	0.27	0.26	0.75	0.04
FAST	RF	7D	0.39	0.42	0.84	0.11	0.04	0.35	0.11
FAST	GBoost	3D	0.40	0.41	0.87	0.32	0.10	0.43	0.11
FAST	GBoost	5D	0.39	0.41	0.87	0.27	0.24	0.55	0.03
FAST	GBoost	7D	0.40	0.42	0.86	0.11	0.03	0.41	0.13
GILD	RF	3D	0.33	0.34	0.70	0.25	0.08	0.31	0.17
GILD	RF	5D	0.30	0.31	0.80	0.38	0.05	0.20	0.08
GILD	RF	7D	0.30	0.31	0.72	0.25	0.13	0.33	0.09
GILD	GBoost	3D	0.35	0.35	0.73	0.41	0.08	0.35	0.22
GILD	GBoost	5D	0.32	0.33	0.96	0.39	0.04	0.00	0.00
GILD	GBoost	7D	0.30	0.31	0.76	0.27	0.11	0.26	0.07
GOOG	RF	3D	0.47	0.47	1.00	0.00	0.00	0.57	0.02
GOOG	RF	5D	0.32	0.50	0.60	0.10	0.45	0.80	0.02
GOOG	RF	7D	0.47	0.28	0.02	1.00	0.04	0.47	0.98
GOOG	GBoost	3D	0.46	0.46	0.99	0.00	0.00	0.25	0.01
GOOG	GBoost	5D	0.32	0.47	0.63	0.08	0.33	0.40	0.02
GOOG	GBoost	7D	0.48	0.31	0.02	0.80	0.08	0.48	0.99
GOOGL	RF	3D	0.46	0.46	0.98	0.00	0.00	0.50	0.03
GOOGL	RF	5D	0.45	0.45	0.98	0.00	0.00	0.35	0.03
GOOGL	RF	7D	0.44	0.09	0.00	0.00	0.00	0.45	0.98
GOOGL	GBoost	3D	0.46	0.46	0.99	0.00	0.00	0.60	0.03
GOOGL	GBoost	5D	0.45	0.45	0.99	0.50	0.02	0.43	0.03
GOOGL	GBoost	7D	0.44	0.22	0.02	1.00	0.02	0.45	0.97
GRMN	RF	3D	0.44	0.50	0.20	0.29	0.09	0.44	0.82
GRMN	RF	5D	0.25	0.45	0.20	0.14	0.51	0.33	0.20

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
GRMN	RF	7D	0.25	0.52	0.20	0.12	0.46	0.35	0.22
GRMN	GBoost	3D	0.44	0.48	0.19	0.00	0.00	0.43	0.84
GRMN	GBoost	5D	0.24	0.45	0.20	0.14	0.53	0.30	0.17
GRMN	GBoost	7D	0.21	0.36	0.12	0.10	0.36	0.30	0.24
HSIC	RF	3D	0.30	0.41	0.52	0.15	0.34	0.26	0.04
HSIC	RF	5D	0.36	0.37	0.48	0.04	0.01	0.38	0.40
HSIC	RF	7D	0.39	0.39	0.76	0.27	0.04	0.39	0.20
HSIC	GBoost	3D	0.42	0.42	0.90	0.11	0.01	0.45	0.09
HSIC	GBoost	5D	0.41	0.41	0.93	0.10	0.01	0.50	0.07
HSIC	GBoost	7D	0.40	0.40	0.96	0.14	0.01	0.48	0.05
IDXX	RF	3D	0.46	0.44	0.76	0.00	0.00	0.50	0.24
IDXX	RF	5D	0.49	0.50	0.27	0.00	0.00	0.48	0.76
IDXX	RF	7D	0.42	0.46	0.57	0.04	0.06	0.47	0.33
IDXX	GBoost	3D	0.44	0.44	0.62	0.00	0.00	0.46	0.34
IDXX	GBoost	5D	0.47	0.47	0.13	0.00	0.00	0.47	0.87
IDXX	GBoost	7D	0.40	0.48	0.50	0.04	0.11	0.48	0.35
INCY	RF	3D	0.34	0.36	0.38	0.29	0.16	0.35	0.44
INCY	RF	5D	0.36	0.38	0.37	0.20	0.05	0.36	0.62
INCY	RF	7D	0.38	0.40	0.40	0.27	0.08	0.38	0.62
INCY	GBoost	3D	0.40	0.45	0.23	0.17	0.01	0.39	0.85
INCY	GBoost	5D	0.38	0.49	0.29	0.00	0.00	0.35	0.83
INCY	GBoost	7D	0.35	0.39	0.23	0.00	0.00	0.34	0.80
INTC	RF	3D	0.37	0.41	0.24	0.29	0.29	0.39	0.58
INTC	RF	5D	0.35	0.46	0.26	0.25	0.14	0.32	0.59
INTC	RF	7D	0.37	0.55	0.29	0.30	0.22	0.31	0.57
INTC	GBoost	3D	0.40	0.47	0.21	0.30	0.31	0.41	0.67
INTC	GBoost	5D	0.37	0.50	0.26	0.30	0.17	0.34	0.66
INTC	GBoost	7D	0.32	0.46	0.22	0.22	0.15	0.29	0.57
INTU	RF	3D	0.12	0.50	0.05	0.06	0.79	0.56	0.10
INTU	RF	5D	0.49	0.44	0.06	0.00	0.00	0.49	0.94
INTU	RF	7D	0.53	0.51	0.32	0.00	0.00	0.55	0.78
INTU	GBoost	3D	0.12	0.50	0.06	0.06	0.82	0.57	0.09
INTU	GBoost	5D	0.49	0.43	0.06	0.00	0.00	0.49	0.94
INTU	GBoost	7D	0.53	0.49	0.36	0.00	0.00	0.55	0.75
ISRG	RF	3D	0.44	0.41	0.46	0.00	0.00	0.48	0.50
ISRG	RF	5D	0.39	0.43	0.59	0.07	0.13	0.50	0.26
ISRG	RF	7D	0.43	0.43	0.70	0.08	0.02	0.47	0.25
ISRG	GBoost	3D	0.46	0.37	0.17	0.00	0.00	0.48	0.78
ISRG	GBoost	5D	0.41	0.40	0.65	0.21	0.09	0.45	0.27
ISRG	GBoost	7D	0.43	0.43	0.64	0.00	0.00	0.43	0.31
LBTYA	RF	3D	0.34	0.36	0.33	0.34	0.37	0.32	0.31
LBTYA	RF	5D	0.32	0.33	0.30	0.33	0.31	0.31	0.35
LBTYA	RF	7D	0.33	0.38	0.31	0.35	0.30	0.28	0.40

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
LBTYA	GBoost	3D	0.35	0.35	0.25	0.34	0.45	0.35	0.35
LBTYA	GBoost	5D	0.29	0.31	0.21	0.33	0.32	0.26	0.35
LBTYA	GBoost	7D	0.33	0.40	0.27	0.34	0.33	0.27	0.40
LBTYK	RF	3D	0.34	0.37	0.37	0.33	0.33	0.31	0.31
LBTYK	RF	5D	0.33	0.35	0.35	0.32	0.30	0.33	0.35
LBTYK	RF	7D	0.33	0.36	0.36	0.33	0.30	0.30	0.34
LBTYK	GBoost	3D	0.31	0.27	0.14	0.32	0.43	0.32	0.37
LBTYK	GBoost	5D	0.32	0.26	0.20	0.35	0.34	0.33	0.44
LBTYK	GBoost	7D	0.32	0.34	0.33	0.35	0.20	0.29	0.47
LULU	RF	3D	0.49	0.50	0.53	0.25	0.03	0.48	0.51
LULU	RF	5D	0.43	0.44	0.35	0.50	0.02	0.43	0.62
LULU	RF	7D	0.50	0.53	0.29	0.12	0.02	0.49	0.78
LULU	GBoost	3D	0.49	0.50	0.51	0.00	0.00	0.49	0.54
LULU	GBoost	5D	0.45	0.47	0.29	0.00	0.00	0.45	0.71
LULU	GBoost	7D	0.46	0.47	0.22	0.00	0.00	0.46	0.77
MAR	RF	3D	0.45	0.43	0.35	0.25	0.01	0.46	0.69
MAR	RF	5D	0.45	0.44	0.86	0.20	0.03	0.54	0.18
MAR	RF	7D	0.44	0.44	0.90	0.08	0.01	0.52	0.16
MAR	GBoost	3D	0.46	0.44	0.33	0.00	0.00	0.46	0.73
MAR	GBoost	5D	0.45	0.43	0.87	1.00	0.03	0.52	0.17
MAR	GBoost	7D	0.47	0.45	0.89	0.00	0.00	0.59	0.23
MAT	RF	3D	0.35	0.35	0.43	0.33	0.37	0.39	0.27
MAT	RF	5D	0.36	0.39	0.40	0.29	0.26	0.36	0.39
MAT	RF	7D	0.33	0.37	0.24	0.29	0.41	0.36	0.34
MAT	GBoost	3D	0.35	0.38	0.51	0.32	0.43	0.30	0.11
MAT	GBoost	5D	0.31	0.35	0.27	0.20	0.16	0.34	0.48
MAT	GBoost	7D	0.37	0.44	0.42	0.34	0.43	0.35	0.28
MDLZ	RF	3D	0.37	0.31	0.18	0.25	0.01	0.39	0.76
MDLZ	RF	5D	0.34	0.34	0.75	0.24	0.18	0.60	0.07
MDLZ	RF	7D	0.38	0.38	0.91	0.27	0.11	0.54	0.07
MDLZ	GBoost	3D	0.36	0.27	0.11	0.42	0.04	0.38	0.79
MDLZ	GBoost	5D	0.36	0.35	0.91	0.29	0.05	0.67	0.08
MDLZ	GBoost	7D	0.26	0.49	0.10	0.23	0.89	0.68	0.06
MELI	RF	3D	0.46	0.50	0.23	0.00	0.00	0.45	0.78
MELI	RF	5D	0.48	0.54	0.28	0.00	0.00	0.46	0.75
MELI	RF	7D	0.43	0.47	0.25	0.50	0.12	0.41	0.68
MELI	GBoost	3D	0.48	0.55	0.19	0.00	0.00	0.46	0.85
MELI	GBoost	5D	0.49	0.56	0.24	0.00	0.00	0.47	0.80
MELI	GBoost	7D	0.45	0.51	0.21	0.00	0.00	0.43	0.79
MNST	RF	3D	0.47	0.45	0.81	0.14	0.01	0.55	0.30
MNST	RF	5D	0.48	0.47	0.70	0.00	0.00	0.53	0.42
MNST	RF	7D	0.46	0.48	0.63	0.11	0.10	0.55	0.41
MNST	GBoost	3D	0.47	0.44	0.65	1.00	0.01	0.50	0.44

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
MNST	GBoost	5D	0.48	0.49	0.42	0.00	0.00	0.48	0.71
MNST	GBoost	7D	0.38	0.45	0.33	0.12	0.25	0.50	0.47
MSFT	RF	3D	0.44	0.44	0.99	0.00	0.00	0.50	0.02
MSFT	RF	5D	0.44	0.43	1.00	0.00	0.00	0.91	0.04
MSFT	RF	7D	0.48	0.23	0.01	0.00	0.00	0.48	0.96
MSFT	GBoost	3D	0.44	0.44	0.99	0.00	0.00	0.50	0.02
MSFT	GBoost	5D	0.44	0.43	1.00	0.00	0.00	0.89	0.03
MSFT	GBoost	7D	0.48	0.23	0.01	0.00	0.00	0.49	0.96
MU	RF	3D	0.40	0.42	0.16	0.19	0.05	0.41	0.79
MU	RF	5D	0.41	0.43	0.12	0.20	0.05	0.41	0.81
MU	RF	7D	0.41	0.45	0.48	0.17	0.06	0.39	0.45
MU	GBoost	3D	0.42	0.46	0.19	1.00	0.02	0.41	0.80
MU	GBoost	5D	0.41	0.38	0.09	0.20	0.02	0.42	0.88
MU	GBoost	7D	0.43	0.47	0.63	0.00	0.00	0.37	0.35
NFLX	RF	3D	0.50	0.50	0.70	0.00	0.00	0.51	0.38
NFLX	RF	5D	0.45	0.50	0.61	0.08	0.08	0.46	0.35
NFLX	RF	7D	0.41	0.39	0.32	0.11	0.02	0.42	0.58
NFLX	GBoost	3D	0.48	0.47	0.61	0.00	0.00	0.49	0.43
NFLX	GBoost	5D	0.45	0.50	0.51	0.03	0.02	0.46	0.49
NFLX	GBoost	7D	0.43	0.42	0.34	0.00	0.00	0.45	0.62
NTES	RF	3D	0.42	0.42	0.22	0.00	0.00	0.42	0.71
NTES	RF	5D	0.42	0.45	0.24	0.00	0.00	0.44	0.70
NTES	RF	7D	0.45	0.47	0.30	0.00	0.00	0.46	0.69
NTES	GBoost	3D	0.44	0.46	0.58	0.00	0.00	0.41	0.39
NTES	GBoost	5D	0.43	0.46	0.21	0.00	0.00	0.43	0.76
NTES	GBoost	7D	0.44	0.46	0.28	0.06	0.02	0.45	0.68
NVDA	RF	3D	0.45	0.45	0.83	0.00	0.00	0.49	0.14
NVDA	RF	5D	0.52	0.49	0.83	0.00	0.00	0.61	0.26
NVDA	RF	7D	0.45	0.36	0.10	0.00	0.00	0.47	0.82
NVDA	GBoost	3D	0.49	0.47	0.76	0.00	0.00	0.57	0.28
NVDA	GBoost	5D	0.48	0.36	0.07	0.00	0.00	0.49	0.89
NVDA	GBoost	7D	0.47	0.40	0.11	0.00	0.00	0.48	0.85
ORLY	RF	3D	0.48	0.39	0.23	0.00	0.00	0.51	0.77
ORLY	RF	5D	0.31	0.37	0.59	0.14	0.54	0.58	0.08
ORLY	RF	7D	0.33	0.36	0.79	0.09	0.13	0.56	0.06
ORLY	GBoost	3D	0.45	0.41	0.69	0.00	0.00	0.54	0.36
ORLY	GBoost	5D	0.51	0.36	0.12	0.00	0.00	0.54	0.87
ORLY	GBoost	7D	0.33	0.36	0.81	0.09	0.13	0.64	0.05
PAYX	RF	3D	0.44	0.39	0.47	0.00	0.00	0.49	0.52
PAYX	RF	5D	0.40	0.40	0.84	0.00	0.00	0.40	0.12
PAYX	RF	7D	0.38	0.39	0.98	0.00	0.00	0.40	0.01
PAYX	GBoost	3D	0.39	0.38	0.95	0.00	0.00	0.47	0.04
PAYX	GBoost	5D	0.41	0.41	0.95	0.00	0.00	0.39	0.04

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
PAYX	GBoost	7D	0.38	0.38	0.99	0.00	0.00	0.25	0.00
PCAR	RF	3D	0.33	0.37	0.17	0.21	0.17	0.35	0.58
PCAR	RF	5D	0.37	0.40	0.57	0.20	0.12	0.39	0.31
PCAR	RF	7D	0.33	0.40	0.17	0.28	0.65	0.42	0.27
PCAR	GBoost	3D	0.34	0.36	0.22	0.21	0.12	0.37	0.61
PCAR	GBoost	5D	0.38	0.39	0.59	0.28	0.09	0.40	0.36
PCAR	GBoost	7D	0.34	0.43	0.14	0.27	0.59	0.42	0.36
PEP	RF	3D	0.37	0.37	0.85	0.34	0.14	0.43	0.08
PEP	RF	5D	0.34	0.34	0.91	0.00	0.00	0.37	0.06
PEP	RF	7D	0.39	0.38	0.89	0.26	0.05	0.55	0.10
PEP	GBoost	3D	0.36	0.35	0.87	0.33	0.01	0.46	0.12
PEP	GBoost	5D	0.34	0.34	0.92	0.00	0.00	0.29	0.04
PEP	GBoost	7D	0.37	0.39	0.82	0.18	0.11	0.47	0.10
QCOM	RF	3D	0.39	0.39	0.25	0.12	0.06	0.41	0.64
QCOM	RF	5D	0.39	0.44	0.13	0.00	0.00	0.38	0.85
QCOM	RF	7D	0.38	0.46	0.11	0.12	0.01	0.38	0.86
QCOM	GBoost	3D	0.41	0.43	0.27	0.17	0.01	0.41	0.69
QCOM	GBoost	5D	0.40	0.48	0.12	0.00	0.00	0.38	0.89
QCOM	GBoost	7D	0.36	0.25	0.03	0.12	0.01	0.37	0.90
REGN	RF	3D	0.42	0.42	0.66	0.17	0.06	0.47	0.33
REGN	RF	5D	0.42	0.43	0.71	0.10	0.03	0.45	0.28
REGN	RF	7D	0.41	0.41	0.77	0.14	0.03	0.47	0.23
REGN	GBoost	3D	0.43	0.42	0.78	0.14	0.01	0.49	0.26
REGN	GBoost	5D	0.44	0.43	0.72	0.00	0.00	0.46	0.32
REGN	GBoost	7D	0.42	0.40	0.79	0.20	0.01	0.49	0.25
ROST	RF	3D	0.46	0.47	0.58	0.09	0.03	0.46	0.41
ROST	RF	5D	0.46	0.47	0.59	0.00	0.00	0.45	0.46
ROST	RF	7D	0.45	0.49	0.57	0.00	0.00	0.43	0.41
ROST	GBoost	3D	0.45	0.44	0.52	0.00	0.00	0.46	0.45
ROST	GBoost	5D	0.44	0.45	0.51	0.00	0.00	0.43	0.49
ROST	GBoost	7D	0.44	0.47	0.50	0.00	0.00	0.42	0.47
SBUX	RF	3D	0.42	0.46	0.24	0.00	0.00	0.42	0.72
SBUX	RF	5D	0.40	0.38	0.17	0.50	0.01	0.41	0.80
SBUX	RF	7D	0.40	0.44	0.14	0.00	0.00	0.40	0.83
SBUX	GBoost	3D	0.43	0.45	0.20	0.00	0.00	0.42	0.78
SBUX	GBoost	5D	0.42	0.45	0.24	0.00	0.00	0.41	0.75
SBUX	GBoost	7D	0.41	0.43	0.10	0.00	0.00	0.41	0.90
SIRI	RF	3D	0.39	0.40	0.77	0.28	0.04	0.37	0.24
SIRI	RF	5D	0.37	0.40	0.64	0.24	0.08	0.32	0.23
SIRI	RF	7D	0.34	0.38	0.48	0.21	0.10	0.32	0.36
SIRI	GBoost	3D	0.41	0.41	0.53	0.58	0.06	0.41	0.52
SIRI	GBoost	5D	0.36	0.37	0.52	0.20	0.03	0.36	0.38
SIRI	GBoost	7D	0.33	0.36	0.51	0.22	0.04	0.29	0.33

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
SNPS	RF	3D	0.49	0.33	0.03	0.00	0.00	0.49	0.95
SNPS	RF	5D	0.51	0.53	0.16	0.00	0.00	0.51	0.89
SNPS	RF	7D	0.49	0.00	0.00	1.00	0.03	0.49	0.99
SNPS	GBoost	3D	0.46	0.40	0.28	0.00	0.00	0.48	0.67
SNPS	GBoost	5D	0.50	0.48	0.47	0.00	0.00	0.51	0.58
SNPS	GBoost	7D	0.51	0.74	0.08	0.00	0.00	0.50	0.98
SRCL	RF	3D	0.31	0.30	0.21	0.35	0.37	0.29	0.37
SRCL	RF	5D	0.32	0.31	0.25	0.38	0.39	0.26	0.32
SRCL	RF	7D	0.34	0.42	0.17	0.36	0.44	0.30	0.44
SRCL	GBoost	3D	0.33	0.33	0.23	0.34	0.47	0.31	0.27
SRCL	GBoost	5D	0.30	0.36	0.11	0.34	0.39	0.25	0.42
SRCL	GBoost	7D	0.35	0.50	0.09	0.37	0.59	0.31	0.40
STX	RF	3D	0.44	0.46	0.82	0.24	0.07	0.39	0.16
STX	RF	5D	0.43	0.43	0.89	0.33	0.02	0.40	0.13
STX	RF	7D	0.41	0.45	0.88	0.12	0.03	0.27	0.08
STX	GBoost	3D	0.44	0.46	0.85	0.22	0.03	0.36	0.14
STX	GBoost	5D	0.41	0.43	0.87	0.25	0.02	0.32	0.11
STX	GBoost	7D	0.42	0.45	0.89	0.33	0.05	0.24	0.08
SWKS	RF	3D	0.43	0.46	0.69	0.08	0.01	0.38	0.27
SWKS	RF	5D	0.43	0.47	0.66	0.00	0.00	0.36	0.31
SWKS	RF	7D	0.39	0.48	0.50	0.00	0.00	0.30	0.40
SWKS	GBoost	3D	0.39	0.51	0.48	0.15	0.25	0.39	0.33
SWKS	GBoost	5D	0.43	0.47	0.63	0.20	0.01	0.37	0.34
SWKS	GBoost	7D	0.44	0.51	0.65	1.00	0.01	0.33	0.33
TMUS	RF	3D	0.39	0.51	0.26	0.13	0.21	0.44	0.57
TMUS	RF	5D	0.46	0.55	0.34	0.25	0.01	0.42	0.74
TMUS	RF	7D	0.40	0.51	0.19	0.08	0.01	0.39	0.80
TMUS	GBoost	3D	0.36	0.43	0.24	0.17	0.37	0.45	0.48
TMUS	GBoost	5D	0.45	0.53	0.36	0.00	0.00	0.41	0.70
TMUS	GBoost	7D	0.41	0.51	0.17	0.00	0.00	0.39	0.83
TXN	RF	3D	0.45	0.45	0.82	0.00	0.00	0.51	0.18
TXN	RF	5D	0.49	0.48	0.91	0.25	0.02	0.62	0.19
TXN	RF	7D	0.33	0.40	0.01	0.13	0.34	0.43	0.69
TXN	GBoost	3D	0.46	0.45	0.86	0.00	0.00	0.50	0.15
TXN	GBoost	5D	0.50	0.48	0.89	0.00	0.00	0.61	0.22
TXN	GBoost	7D	0.29	0.00	0.00	0.13	0.48	0.45	0.56
VOD	RF	3D	0.28	0.23	0.20	0.37	0.23	0.25	0.47
VOD	RF	5D	0.33	0.26	0.23	0.48	0.38	0.21	0.35
VOD	RF	7D	0.40	0.27	0.15	0.46	0.69	0.20	0.11
VOD	GBoost	3D	0.32	0.25	0.24	0.40	0.42	0.25	0.24
VOD	GBoost	5D	0.35	0.27	0.25	0.43	0.49	0.23	0.19
VOD	GBoost	7D	0.36	0.26	0.20	0.47	0.57	0.17	0.14
VRSK	RF	3D	0.47	0.47	0.62	0.28	0.08	0.48	0.43

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
VRSK	RF	5D	0.47	0.44	0.43	0.00	0.00	0.50	0.63
VRSK	RF	7D	0.52	0.49	0.70	0.00	0.00	0.58	0.52
VRSK	GBoost	3D	0.47	0.46	0.48	0.50	0.02	0.47	0.58
VRSK	GBoost	5D	0.46	0.42	0.39	0.00	0.00	0.48	0.63
VRSK	GBoost	7D	0.55	0.51	0.68	0.00	0.00	0.59	0.59
VRTX	RF	3D	0.35	0.35	0.42	0.25	0.08	0.36	0.42
VRTX	RF	5D	0.33	0.35	0.50	0.14	0.04	0.34	0.33
VRTX	RF	7D	0.31	0.31	0.45	0.13	0.05	0.34	0.32
VRTX	GBoost	3D	0.40	0.39	0.41	0.17	0.02	0.41	0.58
VRTX	GBoost	5D	0.35	0.37	0.54	0.09	0.01	0.34	0.36
VRTX	GBoost	7D	0.35	0.33	0.47	0.00	0.00	0.38	0.42
WBA	RF	3D	0.33	0.37	0.42	0.28	0.29	0.33	0.26
WBA	RF	5D	0.30	0.32	0.48	0.22	0.06	0.27	0.32
WBA	RF	7D	0.30	0.34	0.51	0.22	0.07	0.26	0.30
WBA	GBoost	3D	0.35	0.37	0.45	0.31	0.28	0.34	0.28
WBA	GBoost	5D	0.32	0.35	0.37	0.22	0.09	0.32	0.50
WBA	GBoost	7D	0.35	0.38	0.39	0.38	0.23	0.32	0.45
WDC	RF	3D	0.36	0.41	0.48	0.27	0.18	0.34	0.39
WDC	RF	5D	0.38	0.38	0.41	0.45	0.23	0.34	0.48
WDC	RF	7D	0.35	0.39	0.46	0.31	0.24	0.32	0.32
WDC	GBoost	3D	0.37	0.41	0.57	0.28	0.16	0.34	0.33
WDC	GBoost	5D	0.39	0.40	0.66	0.60	0.16	0.32	0.33
WDC	GBoost	7D	0.35	0.39	0.56	0.39	0.12	0.30	0.35
XEL	RF	3D	0.42	0.46	0.51	0.09	0.02	0.41	0.48
XEL	RF	5D	0.42	0.43	0.49	0.33	0.01	0.40	0.50
XEL	RF	7D	0.44	0.44	0.62	0.19	0.04	0.47	0.43
XEL	GBoost	3D	0.45	0.48	0.46	0.00	0.00	0.43	0.62
XEL	GBoost	5D	0.41	0.44	0.46	0.21	0.09	0.41	0.48
XEL	GBoost	7D	0.42	0.43	0.59	0.14	0.06	0.46	0.39
Average results:									
AVG	RF	3D	0.39	0.40	0.47	0.15	0.11	0.43	0.43
AVG	RF	5D	0.40	0.42	0.51	0.16	0.09	0.43	0.40
AVG	RF	7D	0.40	0.39	0.44	0.17	0.10	0.42	0.46
STD	RF	3D	0.07	0.09	0.27	0.13	0.15	0.08	0.26
STD	RF	5D	0.06	0.10	0.27	0.18	0.14	0.13	0.27
STD	RF	7D	0.06	0.12	0.29	0.19	0.15	0.12	0.29
Average results:									
AVG	GBoost	3D	0.40	0.42	0.46	0.16	0.11	0.43	0.44
AVG	GBoost	5D	0.41	0.43	0.49	0.14	0.08	0.42	0.43
AVG	GBoost	7D	0.40	0.43	0.41	0.16	0.11	0.43	0.49

Ticker	Model	label	Acc	Prec(0)	Rec(0)	Prec(1)	Rec(1)	Prec(2)	Rec(2)
STD	GBoost	3D	0.07	0.06	0.29	0.21	0.19	0.08	0.28
STD	GBoost	5D	0.07	0.10	0.29	0.18	0.17	0.13	0.29
STD	GBoost	7D	0.06	0.13	0.30	0.22	0.18	0.12	0.29

Table 7.13: Scores of Random Forest (RF) and Gradient Boosting (GBoost) over NASDAQ dataset.

ARIMA

Analyzing the performance of the ARIMA model across 3-day (3D), 5-day (5D), and 7-day (7D) forecast horizons, the summarized metrics, Table 7.14, highlight its distinct performance characteristics. The average accuracy gradually increases from 0.15 for 3D predictions to 0.18 for 7D forecasts, suggesting a modest improvement in the model's overall predictive ability with longer forecast horizons.

Remarkably, the ARIMA model exhibits a high recall for label 1 across all periods, peaking at 0.99 for both 3D and 5D forecasts, which indicates its strong capacity to identify this specific market condition. However, its ability to detect label 0 and label 2 conditions remains significantly low, as reflected by the minimal recall values. Precision scores vary across labels and time frames, with a noticeable trend of improvement in the precision for label 2 (0.34 for 3D to 0.16 for 7D), albeit starting from a higher baseline.

The F1 Score, a harmonic mean of precision and recall, remains low across all forecasts, pointing to challenges in achieving a balanced predictive performance for all market states. This suggests that while ARIMA is adept at recognizing certain conditions, it struggles to maintain consistency across all market behaviors.

	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
AVG	3D	0.15	3-e4	0.99	2-e02	0.07	0.15	0.34	0.08
AVG	5D	0.17	3-e2	0.99	4-e02	0.19	0.17	0.18	0.09
AVG	7D	0.18	6-e2	0.97	0.02	0.19	0.17	0.16	0.10
STD	3D	0.09	0.00	0.00	0.00	0.25	0.09	0.47	0.04
STD	5D	0.09	0.01	0.01	0.01	0.33	0.10	0.30	0.04
STD	7D	0.10	0.01	0.12	0.12	0.33	0.09	0.26	0.04

Table 7.14: Average Scores of ARIMA over NASDAQ dataset.

Standard deviation metrics indicate variability in the model's performance, especially in precision across different labels, highlighting potential instability in the ARIMA model's predictions across various runs. This variability, combined with the overall performance metrics, underscores the need for cautious application and potential model adjustments or integration with other models to enhance reliability and predictive accuracy in financial time series forecasting.

The analysis of the complete ARIMA model results, Table 7.15, over the NASDAQ dataset reveals a distinct performance pattern compared to the SupCon approach. ARIMA's average accuracy across different forecasting horizons (3D, 5D, and 7D) is notably lower. The precision and recall metrics indicate that ARIMA struggles to correctly predict class 0 and class 2, with significantly higher performance for class 1 predictions across all horizons. This suggests ARIMA's tendency to favor one class significantly over others, likely due to its nature of capturing linear trends and seasonality, which may not adequately represent the complex, non-linear patterns present in stock price movements.

Comparatively, the SupCon model demonstrated a more balanced performance across classes, benefiting from its ability to learn complex, high-dimensional representations of data through contrastive learning. This suggests that for the task of stock market prediction, where data inherently contains non-linear patterns and is influenced by a myriad of factors beyond historical prices, models that can capture these complex relationships, like SupCon, may offer better performance than traditional time series models like ARIMA.

In summary, while ARIMA provides a baseline for time series forecasting, its limitations in handling the complexities of stock market data are evident. The contrastive learning approach, represented by SupCon, shows promise in overcoming some of these challenges by leveraging a richer representation of the data, ultimately leading to improved prediction accuracy.

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
AAL	7D	0.37	0.00	0.99	0.00	0.00	0.37	0.00	0.18
AAL	3D	0.33	0.00	1.00	0.00	0.00	0.33	0.00	0.16
AAL	5D	0.39	0.00	1.00	0.00	0.00	0.39	0.00	0.19
AAPL	7D	0.13	0.03	1.00	0.00	0.60	0.12	0.00	0.09
AAPL	3D	0.06	0.00	1.00	0.00	0.00	0.06	1.00	0.04
AAPL	5D	0.08	0.00	1.00	0.00	0.25	0.08	0.00	0.05
ADBE	7D	0.10	0.01	1.00	0.00	0.27	0.09	0.00	0.06
ADBE	3D	0.07	0.00	1.00	0.00	0.00	0.07	0.00	0.05
ADBE	5D	0.08	0.02	1.00	0.00	0.67	0.07	0.00	0.06
ADI	3D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
ADI	5D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
ADI	7D	0.14	0.00	0.97	0.01	0.00	0.13	0.27	0.09

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
ADP	5D	0.13	0.00	1.00	0.00	0.00	0.13	0.00	0.08
ADP	7D	0.09	0.01	0.93	0.00	0.08	0.09	0.00	0.06
ADP	3D	0.14	0.01	1.00	0.00	1.00	0.13	0.00	0.08
ADSK	7D	0.10	0.00	0.96	0.02	0.00	0.10	0.40	0.07
ADSK	3D	0.08	0.00	1.00	0.00	0.00	0.08	0.00	0.05
ADSK	5D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
AKAM	5D	0.23	0.00	1.00	0.00	0.00	0.23	0.00	0.13
AKAM	3D	0.23	0.00	1.00	0.00	0.00	0.23	0.00	0.12
AKAM	7D	0.23	0.00	0.99	0.00	0.00	0.24	0.00	0.13
ALGN	3D	0.06	0.00	1.00	0.00	0.00	0.06	0.00	0.04
ALGN	5D	0.07	0.00	1.00	0.02	0.00	0.06	0.56	0.06
ALGN	7D	0.08	0.00	1.00	0.04	0.00	0.07	0.60	0.07
AMAT	7D	0.11	0.00	1.00	0.05	0.00	0.09	0.58	0.08
AMAT	5D	0.08	0.00	1.00	0.01	0.00	0.08	0.20	0.06
AMAT	3D	0.06	0.00	1.00	0.00	1.00	0.05	0.00	0.04
AMD	7D	0.09	0.00	1.00	0.03	0.00	0.08	0.33	0.07
AMD	3D	0.06	0.00	1.00	0.00	0.00	0.06	1.00	0.04
AMD	5D	0.06	0.00	0.96	0.02	0.00	0.05	0.50	0.05
AMGN	5D	0.22	0.00	1.00	0.01	0.00	0.22	1.00	0.13
AMGN	7D	0.22	0.00	1.00	0.01	0.00	0.22	0.17	0.12
AMGN	3D	0.21	0.00	1.00	0.01	0.00	0.21	1.00	0.12
AMZN	7D	0.08	0.02	1.00	0.00	0.45	0.07	0.00	0.06
AMZN	5D	0.07	0.01	1.00	0.00	0.60	0.06	0.00	0.05
AMZN	3D	0.06	0.00	1.00	0.00	0.00	0.05	1.00	0.04
ASML	5D	0.09	0.00	0.98	0.01	0.00	0.08	0.33	0.06
ASML	3D	0.06	0.00	1.00	0.00	0.00	0.06	1.00	0.04
ASML	7D	0.08	0.00	1.00	0.02	0.00	0.07	0.45	0.06
ATVI	5D	0.19	0.02	0.99	0.00	0.80	0.19	0.00	0.12
ATVI	7D	0.18	0.02	0.99	0.00	0.67	0.18	0.00	0.11
ATVI	3D	0.17	0.00	1.00	0.01	0.00	0.17	1.00	0.10
AVGO	5D	0.10	0.00	1.00	0.01	0.00	0.10	0.67	0.06
AVGO	7D	0.09	0.00	0.96	0.00	0.00	0.09	0.17	0.06
AVGO	3D	0.07	0.00	1.00	0.00	0.00	0.07	1.00	0.05
BIDU	5D	0.25	0.00	1.00	0.00	0.50	0.25	0.00	0.14
BIDU	7D	0.28	0.00	1.00	0.00	0.00	0.28	0.00	0.15
BIDU	3D	0.27	0.00	0.99	0.00	0.00	0.27	0.00	0.14
BIIB	7D	0.28	0.00	0.98	0.01	0.00	0.28	0.29	0.15
BIIB	3D	0.22	0.00	1.00	0.00	0.00	0.22	0.00	0.12
BIIB	5D	0.25	0.00	0.99	0.00	0.00	0.25	0.00	0.13
BKNG	7D	0.15	0.00	1.00	0.00	0.00	0.15	0.00	0.09
BKNG	5D	0.18	0.00	1.00	0.00	0.00	0.18	0.00	0.10
BKNG	3D	0.17	0.00	1.00	0.00	0.00	0.17	0.00	0.10
BMRN	7D	0.27	0.03	1.00	0.00	0.86	0.26	0.00	0.16

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
BMRN	3D	0.29	0.00	1.00	0.00	0.00	0.29	0.00	0.15
BMRN	5D	0.29	0.00	1.00	0.01	0.00	0.29	0.33	0.15
CDNS	7D	0.07	0.02	1.00	0.00	0.83	0.06	0.00	0.05
CDNS	5D	0.06	0.00	1.00	0.00	0.00	0.06	0.00	0.04
CDNS	3D	0.05	0.00	1.00	0.00	0.00	0.05	1.00	0.04
CHKP	7D	0.18	0.00	1.00	0.00	0.00	0.18	0.00	0.10
CHKP	3D	0.20	0.00	1.00	0.00	1.00	0.20	0.00	0.11
CHKP	5D	0.20	0.00	1.00	0.00	0.20	0.20	0.00	0.12
CHRW	5D	0.24	0.00	1.00	0.01	0.00	0.24	0.50	0.13
CHRW	7D	0.23	0.00	1.00	0.00	0.00	0.24	0.00	0.13
CHRW	3D	0.19	0.00	1.00	0.00	0.00	0.20	0.00	0.11
CMCSA	7D	0.20	0.00	1.00	0.00	0.25	0.20	0.00	0.11
CMCSA	5D	0.19	0.00	0.99	0.00	0.20	0.19	0.00	0.11
CMCSA	3D	0.19	0.00	1.00	0.00	0.00	0.19	0.00	0.11
COST	5D	0.10	0.01	1.00	0.00	0.50	0.09	0.00	0.06
COST	3D	0.10	0.00	1.00	0.00	0.00	0.10	0.00	0.06
COST	7D	0.12	0.03	0.96	0.00	0.67	0.11	0.00	0.08
CSCO	3D	0.19	0.00	1.00	0.00	0.00	0.19	0.00	0.11
CSCO	5D	0.23	0.00	1.00	0.00	0.00	0.23	0.00	0.13
CSCO	7D	0.22	0.00	1.00	0.01	0.00	0.22	1.00	0.12
CSX	7D	0.15	0.00	1.00	0.03	0.00	0.14	0.50	0.10
CSX	3D	0.16	0.00	1.00	0.00	0.00	0.16	0.00	0.09
CSX	5D	0.14	0.00	1.00	0.00	0.00	0.14	0.00	0.08
CTSH	3D	0.22	0.00	1.00	0.00	0.00	0.22	0.00	0.12
CTSH	7D	0.24	0.02	1.00	0.00	1.00	0.24	0.00	0.14
CTSH	5D	0.26	0.00	0.99	0.00	0.50	0.26	0.00	0.14
DLTR	3D	0.17	0.00	1.00	0.00	0.00	0.17	0.00	0.10
DLTR	5D	0.15	0.00	1.00	0.00	0.00	0.15	0.00	0.09
DLTR	7D	0.19	0.00	1.00	0.00	0.00	0.19	0.14	0.11
EBAY	7D	0.16	0.00	0.96	0.03	0.00	0.15	0.60	0.11
EBAY	5D	0.13	0.00	0.97	0.01	0.00	0.13	0.40	0.08
EBAY	3D	0.16	0.00	1.00	0.00	0.00	0.16	0.00	0.09
EXPD	7D	0.17	0.01	0.99	0.00	0.67	0.17	0.00	0.10
EXPD	3D	0.17	0.00	1.00	0.00	0.00	0.17	0.00	0.10
EXPD	5D	0.14	0.00	0.99	0.00	0.00	0.14	0.00	0.08
FAST	7D	0.20	0.00	1.00	0.00	0.00	0.21	0.00	0.11
FAST	5D	0.21	0.00	1.00	0.00	0.00	0.21	0.00	0.12
FAST	3D	0.18	0.00	1.00	0.00	0.00	0.18	0.00	0.10
GILD	3D	0.31	0.00	1.00	0.00	0.00	0.31	0.00	0.16
GILD	5D	0.34	0.00	1.00	0.01	0.00	0.34	0.25	0.17
GILD	7D	0.32	0.00	1.00	0.00	0.00	0.32	0.00	0.16
GOOG	5D	0.11	0.00	1.00	0.01	0.00	0.10	0.50	0.07
GOOG	3D	0.08	0.00	1.00	0.00	0.00	0.08	1.00	0.05

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
GOOG	7D	0.10	0.00	1.00	0.00	0.00	0.10	0.50	0.06
GOOGL	5D	0.10	0.00	1.00	0.01	0.00	0.10	1.00	0.06
GOOGL	3D	0.09	0.00	1.00	0.00	0.00	0.08	1.00	0.05
GOOGL	7D	0.11	0.00	1.00	0.00	0.00	0.11	0.00	0.07
GRMN	3D	0.16	0.00	1.00	0.00	0.00	0.16	0.00	0.09
GRMN	5D	0.16	0.00	1.00	0.00	0.00	0.16	0.33	0.09
GRMN	7D	0.14	0.00	1.00	0.00	0.00	0.14	0.00	0.08
HSIC	3D	0.17	0.00	1.00	0.00	0.00	0.18	0.00	0.10
HSIC	5D	0.18	0.00	0.99	0.00	0.00	0.18	0.00	0.10
HSIC	7D	0.18	0.00	1.00	0.01	0.00	0.18	0.33	0.11
IDXX	7D	0.08	0.00	0.97	0.02	0.00	0.07	0.50	0.05
IDXX	5D	0.06	0.01	1.00	0.00	0.67	0.06	0.00	0.04
IDXX	3D	0.07	0.00	1.00	0.00	0.00	0.06	1.00	0.04
INCY	7D	0.30	0.00	1.00	0.00	0.00	0.30	0.00	0.15
INCY	5D	0.32	0.00	1.00	0.01	0.00	0.31	0.67	0.17
INCY	3D	0.28	0.00	0.99	0.00	0.00	0.28	0.00	0.14
INTC	3D	0.19	0.00	1.00	0.01	0.00	0.19	1.00	0.11
INTC	5D	0.21	0.00	0.99	0.00	0.00	0.21	0.00	0.12
INTC	7D	0.19	0.00	0.99	0.02	0.00	0.19	0.43	0.12
INTU	3D	0.07	0.00	1.00	0.00	1.00	0.07	0.00	0.05
INTU	7D	0.09	0.02	1.00	0.00	0.57	0.08	0.00	0.06
INTU	5D	0.09	0.03	1.00	0.00	1.00	0.07	0.00	0.06
ISRG	7D	0.09	0.01	1.00	0.00	1.00	0.08	0.00	0.06
ISRG	3D	0.08	0.00	1.00	0.00	0.00	0.08	1.00	0.05
ISRG	5D	0.10	0.01	1.00	0.00	1.00	0.09	0.00	0.06
LBTYA	3D	0.32	0.00	0.99	0.00	0.00	0.32	0.00	0.16
LBTYA	5D	0.38	0.00	0.99	0.01	0.00	0.38	0.67	0.19
LBTYA	7D	0.38	0.00	0.99	0.00	0.00	0.38	0.00	0.18
LBTYK	3D	0.31	0.00	0.99	0.00	0.00	0.31	0.00	0.16
LBTYK	5D	0.35	0.00	0.99	0.02	0.00	0.34	0.75	0.18
LBTYK	7D	0.37	0.00	0.99	0.01	0.00	0.37	0.25	0.19
LULU	5D	0.11	0.00	1.00	0.01	0.00	0.11	0.38	0.07
LULU	3D	0.07	0.00	1.00	0.00	0.00	0.07	1.00	0.05
LULU	7D	0.09	0.00	0.98	0.02	0.00	0.08	0.45	0.06
MAR	3D	0.13	0.00	0.99	0.00	0.00	0.13	0.00	0.08
MAR	5D	0.14	0.01	1.00	0.00	1.00	0.13	0.00	0.08
MAR	7D	0.16	0.02	1.00	0.00	0.67	0.15	0.00	0.10
MAT	3D	0.28	0.00	1.00	0.00	0.00	0.28	0.00	0.15
MAT	7D	0.31	0.00	0.98	0.00	0.00	0.31	0.00	0.16
MAT	5D	0.28	0.00	1.00	0.00	0.00	0.28	0.00	0.15
MDLZ	7D	0.23	0.03	0.99	0.00	0.83	0.23	0.00	0.14
MDLZ	3D	0.23	0.00	1.00	0.00	0.00	0.23	0.00	0.13
MDLZ	5D	0.25	0.01	1.00	0.00	0.50	0.25	0.00	0.14

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
MELI	7D	0.09	0.00	1.00	0.05	0.00	0.07	0.50	0.07
MELI	3D	0.05	0.00	1.00	0.00	0.00	0.05	1.00	0.04
MELI	5D	0.06	0.02	1.00	0.00	0.83	0.05	0.00	0.05
MNST	5D	0.14	0.00	1.00	0.00	0.00	0.14	0.00	0.08
MNST	7D	0.13	0.00	1.00	0.00	0.00	0.13	0.00	0.07
MNST	3D	0.15	0.00	1.00	0.00	0.00	0.15	0.00	0.09
MSFT	7D	0.07	0.00	1.00	0.00	0.00	0.07	0.00	0.05
MSFT	5D	0.10	0.00	1.00	0.00	0.00	0.10	0.00	0.06
MSFT	3D	0.09	0.00	1.00	0.00	0.00	0.09	0.00	0.05
MU	3D	0.13	0.00	1.00	0.00	0.00	0.13	0.00	0.08
MU	7D	0.15	0.00	0.97	0.01	0.00	0.14	0.33	0.09
MU	5D	0.12	0.00	1.00	0.00	0.00	0.12	0.17	0.07
NFLX	3D	0.08	0.00	1.00	0.00	0.00	0.08	0.00	0.05
NFLX	5D	0.11	0.02	0.98	0.00	0.62	0.10	0.00	0.07
NFLX	7D	0.11	0.02	1.00	0.00	0.38	0.10	0.00	0.08
NTES	5D	0.11	0.00	1.00	0.04	0.00	0.10	0.80	0.08
NTES	7D	0.12	0.00	1.00	0.06	0.00	0.09	0.87	0.09
NTES	3D	0.10	0.00	1.00	0.00	0.00	0.09	1.00	0.06
NVDA	5D	0.05	0.00	1.00	0.02	0.00	0.03	0.60	0.04
NVDA	3D	0.04	0.00	1.00	0.00	0.00	0.03	1.00	0.02
NVDA	7D	0.49	0.00	0.00	1.00	0.00	0.00	0.49	0.22
ORLY	3D	0.10	0.00	1.00	0.00	0.00	0.10	0.00	0.06
ORLY	5D	0.10	0.00	1.00	0.00	0.00	0.10	0.00	0.06
ORLY	7D	0.12	0.00	1.00	0.02	0.00	0.11	0.42	0.08
PAYX	7D	0.14	0.01	0.97	0.00	0.14	0.14	0.00	0.08
PAYX	5D	0.12	0.01	0.98	0.00	0.50	0.12	0.00	0.08
PAYX	3D	0.13	0.01	1.00	0.00	1.00	0.13	0.00	0.08
PCAR	7D	0.27	0.00	1.00	0.01	0.00	0.26	1.00	0.15
PCAR	5D	0.24	0.00	1.00	0.01	0.00	0.24	1.00	0.14
PCAR	3D	0.24	0.00	1.00	0.01	0.00	0.24	1.00	0.13
PEP	3D	0.19	0.00	1.00	0.00	0.00	0.19	0.00	0.11
PEP	5D	0.22	0.02	0.98	0.00	0.60	0.22	0.00	0.13
PEP	7D	0.20	0.03	1.00	0.00	0.60	0.19	0.00	0.13
QCOM	3D	0.14	0.00	1.00	0.00	0.00	0.13	1.00	0.08
QCOM	5D	0.16	0.00	0.99	0.00	0.00	0.16	0.00	0.09
QCOM	7D	0.14	0.00	0.97	0.00	0.00	0.14	0.00	0.08
REGN	7D	0.18	0.00	1.00	0.00	0.00	0.18	0.00	0.10
REGN	5D	0.16	0.00	1.00	0.00	0.00	0.16	0.00	0.09
REGN	3D	0.16	0.00	1.00	0.00	0.00	0.16	0.00	0.09
ROST	3D	0.07	0.00	1.00	0.00	0.00	0.07	0.00	0.05
ROST	5D	0.13	0.01	1.00	0.00	1.00	0.13	0.00	0.08
ROST	7D	0.10	0.02	1.00	0.00	0.67	0.09	0.00	0.07
SBUX	5D	0.16	0.02	1.00	0.00	0.83	0.15	0.00	0.10

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
SBUX	3D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
SBUX	7D	0.16	0.03	1.00	0.00	0.86	0.15	0.00	0.10
SIRI	3D	0.24	0.00	1.00	0.01	0.00	0.24	1.00	0.13
SIRI	5D	0.22	0.00	1.00	0.00	0.00	0.22	0.00	0.12
SIRI	7D	0.27	0.03	1.00	0.00	0.50	0.26	0.00	0.16
SNPS	3D	0.07	0.00	1.00	0.00	0.00	0.07	1.00	0.04
SNPS	7D	0.08	0.00	1.00	0.00	0.00	0.08	0.00	0.05
SNPS	5D	0.05	0.00	1.00	0.00	0.00	0.05	0.00	0.03
SRCL	5D	0.37	0.00	1.00	0.01	0.00	0.37	0.50	0.18
SRCL	3D	0.35	0.00	1.00	0.01	0.00	0.35	1.00	0.18
SRCL	7D	0.35	0.00	1.00	0.00	0.00	0.35	0.00	0.17
STX	7D	0.15	0.00	1.00	0.00	0.14	0.15	0.00	0.09
STX	5D	0.17	0.00	0.99	0.00	0.00	0.17	0.00	0.10
STX	3D	0.14	0.00	1.00	0.00	0.00	0.14	0.00	0.08
SWKS	7D	0.16	0.00	1.00	0.04	0.00	0.15	0.70	0.11
SWKS	5D	0.15	0.00	0.99	0.02	0.00	0.15	0.50	0.10
SWKS	3D	0.14	0.00	1.00	0.00	0.00	0.13	1.00	0.08
TMUS	7D	0.15	0.02	1.00	0.00	1.00	0.14	0.00	0.10
TMUS	3D	0.14	0.00	1.00	0.00	0.00	0.14	0.00	0.08
TMUS	5D	0.15	0.01	1.00	0.00	0.67	0.15	0.00	0.09
TXN	3D	0.10	0.00	1.00	0.00	0.00	0.10	1.00	0.06
TXN	5D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
TXN	7D	0.13	0.00	0.97	0.01	0.00	0.13	0.33	0.08
VOD	7D	0.47	0.00	0.99	0.00	0.00	0.47	0.00	0.21
VOD	5D	0.49	0.00	1.00	0.02	0.00	0.49	0.67	0.23
VOD	3D	0.44	0.00	1.00	0.00	0.00	0.44	0.00	0.20
VRSK	7D	0.13	0.03	0.97	0.00	0.42	0.13	0.00	0.09
VRSK	5D	0.13	0.02	1.00	0.00	1.00	0.12	0.00	0.09
VRSK	3D	0.12	0.00	1.00	0.00	0.00	0.12	0.00	0.07
VRTX	7D	0.22	0.00	1.00	0.00	0.00	0.22	0.00	0.12
VRTX	5D	0.22	0.00	0.99	0.00	0.00	0.23	0.00	0.12
VRTX	3D	0.21	0.00	1.00	0.00	0.00	0.21	0.00	0.12
WBA	3D	0.30	0.00	1.00	0.00	0.00	0.30	0.00	0.16
WBA	5D	0.33	0.00	0.99	0.03	0.00	0.32	0.80	0.18
WBA	7D	0.34	0.00	0.98	0.03	0.00	0.33	0.44	0.18
WDC	7D	0.30	0.00	0.99	0.00	0.00	0.31	0.00	0.16
WDC	5D	0.32	0.00	1.00	0.00	0.00	0.32	0.00	0.16
WDC	3D	0.30	0.00	1.00	0.01	0.00	0.29	1.00	0.16
XEL	7D	0.17	0.03	0.99	0.00	0.88	0.16	0.00	0.11
XEL	3D	0.17	0.00	1.00	0.00	0.00	0.17	0.00	0.09
XEL	5D	0.16	0.01	0.99	0.00	0.50	0.16	0.00	0.10
Average results:									
AVG	3D	0.15	3-e4	0.99	2-e02	0.07	0.15	0.34	0.08

Ticker	label	Acc	Rec(0)	Rec(1)	Rec(2)	Prec(0)	Prec(1)	Prec(2)	F1 Score
AVG	5D	0.17	3-e2	0.99	4-e02	0.19	0.17	0.18	0.09
AVG	7D	0.18	6-e2	0.97	0.02	0.19	0.17	0.16	0.10
STD	3D	0.09	0.00	0.00	0.00	0.25	0.09	0.47	0.04
STD	5D	0.09	0.01	0.01	0.01	0.33	0.10	0.30	0.04
STD	7D	0.10	0.01	0.12	0.12	0.33	0.09	0.26	0.04

Table 7.15: Scores of ARIMA over NASDAQ dataset.

7.3.3 Qualitative Analysis

In the domain of financial markets, the pursuit of analytical precision transcends mere quantitative metrics, ushering in the indispensable realm of qualitative analysis. This realm, far from being ancillary, complements and enriches the numerical rigor of quantitative assessments, providing a multi-dimensional view of the intricate schema of market dynamics. This comprehensive approach is particularly pivotal when delving into the complex interplay of market forces, where numerical data alone may not fully capture the subtleties of market sentiment, investor behavior, and the undercurrents shaping market trends.

Qualitative analysis, therefore, stands not as a secondary or optional endeavor but as an essential counterpart to quantitative methodologies. It allows for a nuanced exploration of the factors that quantitative metrics may overlook, such as the qualitative aspects of company management, brand value, market sentiment, and emerging trends that might not yet be reflected in the numbers. In this context, the qualitative analysis of t-SNE visualizations of a sample of stock market data, namely *ADP - Automatic Data Processing Inc.*, *AMGN - AMGEN Inc.*, and *MSFT - Microsoft Corporation* tickers, obtained through Supervised Contrastive Learning (SupCon), and compared with the one extracted from TS2Vec exemplifies the synthesis of numerical precision and qualitative insight.

The t-SNE (t-Distributed Stochastic Neighbor Embedding) visualizations serve as a powerful tool for distilling the essence of complex, high-dimensional data into comprehensible, two-dimensional maps. This transformation, while preserving the relative proximities of data points, unveils patterns, clusters, and anomalies that might remain obscured in the high-dimensional space. In the realm of financial data analysis, such visualizations offer a window into the learned representations of market states, beyond what traditional metrics can convey. Furthermore, the juxtaposition of t-SNE visualizations with other financial indicators and charts, such as candlestick plots and Bollinger Bands, enriches the qualitative analysis. It allows for a holistic view of how the model's representations align with or deviate from traditional technical analysis indicators. This alignment or deviation, in turn, can offer insights into the model's sensitivity to various market signals and its potential for uncovering novel patterns or trends not captured by conventional analysis.

ADP Ticker Analysis

Our visual inspection commences with the separability of classes in the t-SNE plots. The distinctiveness of class clusters serves as an indicator of the model's ability to discriminate between different classes. The observed visualizations depict a scarse degree of class separability, with notable overlap between Class 0 and Class 1, as shown in Fig. 7.1, Fig. 7.2, and Fig. 7.3. Such overlap is indicative of the model's challenges in learning discriminative features that could facilitate a clear delineation of classes.

The t-SNE projections demonstrate consistency in cluster shapes across various dimensions, suggesting a degree of stability in the learned feature space. Despite this, the varying spread of points within clusters points to the model's inconsistent representation of features, possibly hinting at its sensitivity to the dimensionality of input data. The data point density within clusters varied significantly. Areas of high density suggest higher model confidence, whereas the more dispersed arrangements imply lower certainty. This variance in density highlights the potential inconsistencies in the model's confidence across the dataset, which could impact the model's predictive accuracy.

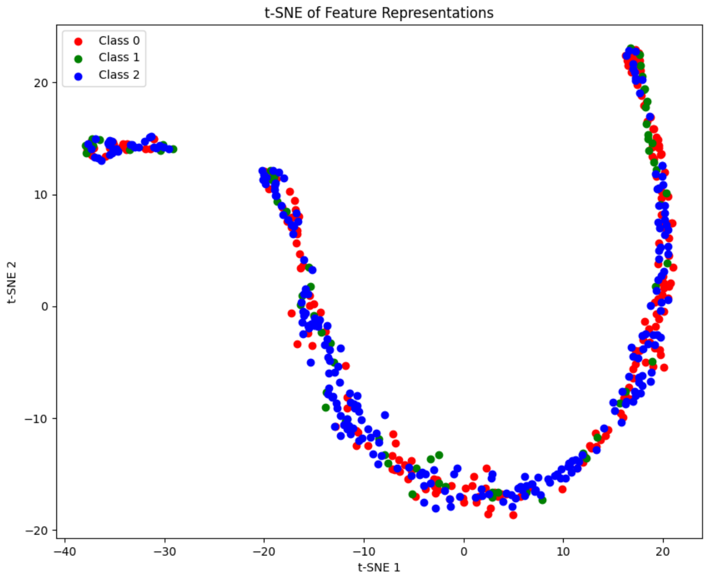

Figure 7.1: SupCon t-SNE visualization of ADP embeddings for 3D label type

The presence of outliers, as evident in the visualizations, raises concerns regarding anomalies and noise within the data. Such discrepancies could stem from extraordinary market events or artifacts introduced during the feature extraction process, underscoring the need for robust feature engineering and noise reduction techniques. Considering the temporal span of the dataset (2021-2022), a period

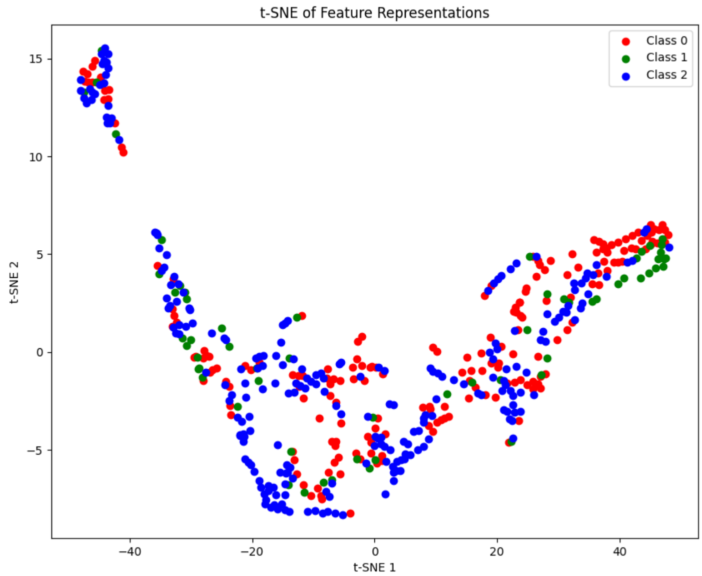

Figure 7.2: SupCon t-SNE visualization of ADP embeddings for 5D label type

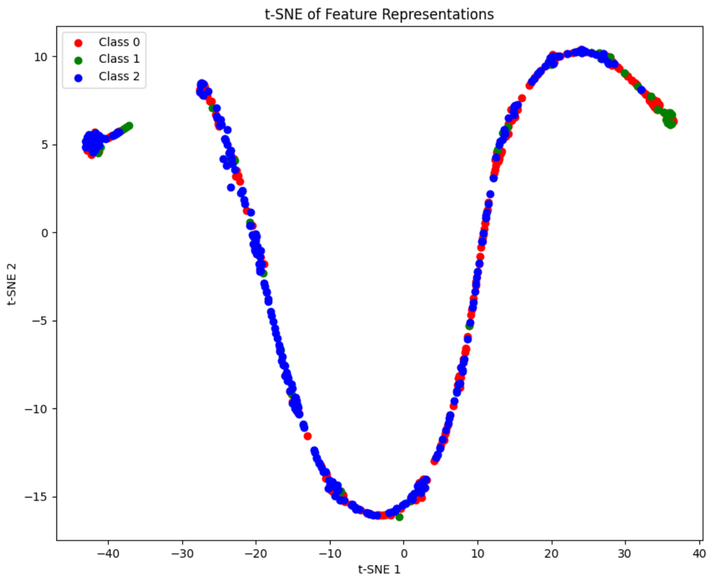

Figure 7.3: SupCon t-SNE visualization of ADP embeddings for 7D label type

characterized by significant market volatility, the overlap in class representations may reflect the temporal instability inherent in financial time series. The model's poor performance could be attributed to the non-stationary nature of stock prices, where patterns evolve and the discriminative power of features diminishes over time.

Building upon the previous qualitative analysis of feature representations obtained from supervised contrastive learning, we extend our examination to include the t-SNE visualization of the dataset used for feature extraction. Looking at Figures 7.4, 7.5, and 7.6, comparative analysis aims to discern patterns and correlations between feature representations and the actual data distribution, offering deeper insights into the model's learning behavior.

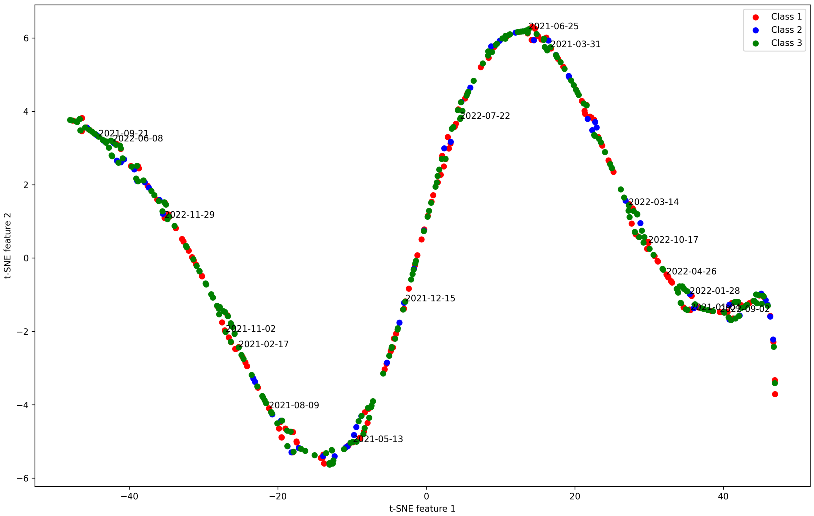

Figure 7.4: t-SNE visualization of ADP dataset for 3D label type

The additional t-SNE visualizations annotate points with corresponding dates, revealing the temporal progression within the feature space. By comparing these plots with the earlier visualizations, we can examine how well the temporal dimension is captured within the feature representations. It appears that data points that are temporally closer tend to cluster together, suggesting the model is capturing some aspects of temporal continuity. When observing the dispersion of points from the same class across different times, a pattern emerges where certain classes drift over the feature space as time progresses. This drift could be symptomatic of the model's adaptation to evolving market conditions, reflected in the time series data.

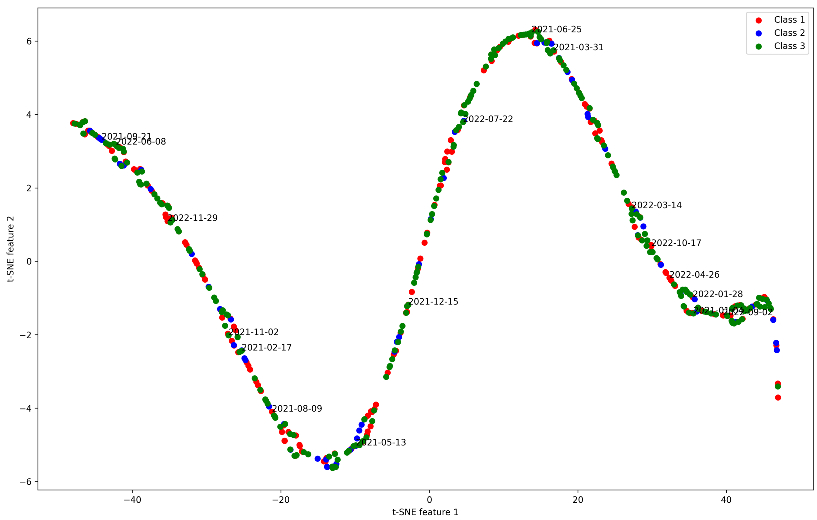

Figure 7.5: t-SNE visualization of ADP dataset for 5D label type

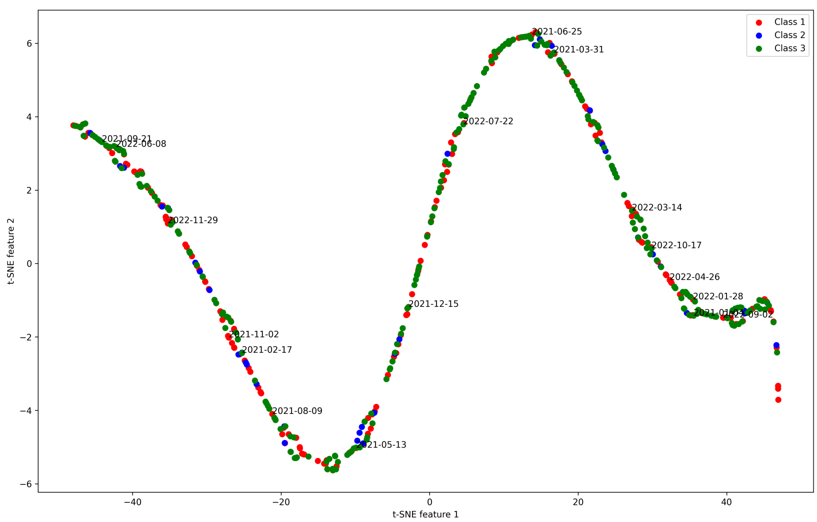

Figure 7.6: t-SNE visualization of ADP dataset for 7D label type

Such a phenomenon underscores the non-stationary nature of the financial time series data and the challenges it poses for feature extraction and subsequent learning.

A direct comparison of dataset and feature representation visualizations showcases discrepancies in the model's ability to maintain temporal coherence. Certain dates, such as periods of market volatility, seem to result in more scattered representations, indicating that the model's performance is susceptible to external market dynamics.

To deepen the qualitative analysis of the supervised contrastive learning model for stock price prediction, we also integrate technical analysis components, including Bollinger Bands, volume data, and candlestick patterns, as shown in Figures 7.8, 7.9, and 7.7. These elements are traditionally used in stock market analysis to understand market sentiment, volatility, and price trends. We correlate these technical indicators with the t-SNE feature representations to provide a holistic understanding of the model's performance and the underlying data characteristics.

Figure 7.7: Candlestick visualization of ADP dataset

Bollinger Bands, a measure of volatility and price levels relative to moving averages, are depicted in one of the provided charts. The t-SNE visualizations, when compared alongside Bollinger Bands, show that feature representations corresponding to periods of high volatility (where the bands widen) might be contributing to the more dispersed clusters in the t-SNE plots. This suggests that the model might be sensitive to volatility, capturing the dramatic shifts in price movement within the feature space. *Volume* charts reflect trading activity intensity. By examining the volume spikes and their corresponding dates, we can identify periods of high market

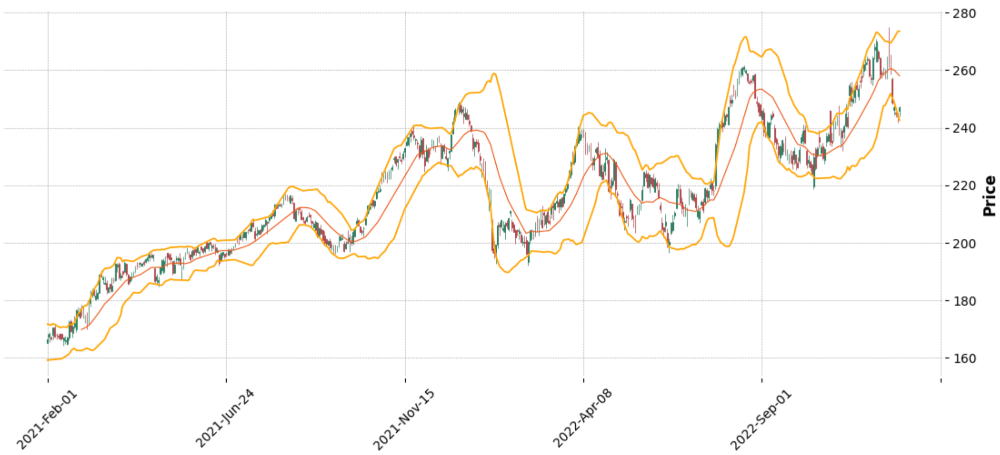

Figure 7.8: Bollinger Bands visualization of ADP dataset

Figure 7.9: Volumes visualization of ADP dataset

activity that may correlate with anomalies or outliers in the t-SNE plots. Such a correlation might imply that significant trading volumes impact the model's ability to form consistent feature representations, leading to potential misclassifications or poor performance in stock price prediction.

The *candlestick* chart provides insights into the price action. Analyzing the candlestick patterns alongside the t-SNE feature representations, one can discern if specific price movements, like bullish periods or bearish reversals, correspond to the clustering of points in the t-SNE visualizations. If certain patterns consistently correspond to misclassified points or cluster boundaries, this would indicate the model's response to short-term price movements, as opposed to capturing more long-term trends effectively.

To enhance our analysis, we now consider also a comparison between the features extracted from the Supervised Contrastive Learning model and, as shown in Figures 7.10, 7.11, and 7.12, Time Series to Vectors (TS2Vec). By visualizing and comparing their resulting t-SNE plots, we aim to understand the differences in feature space representations, which may offer insights into each method's strengths and potential improvements for stock price prediction models.

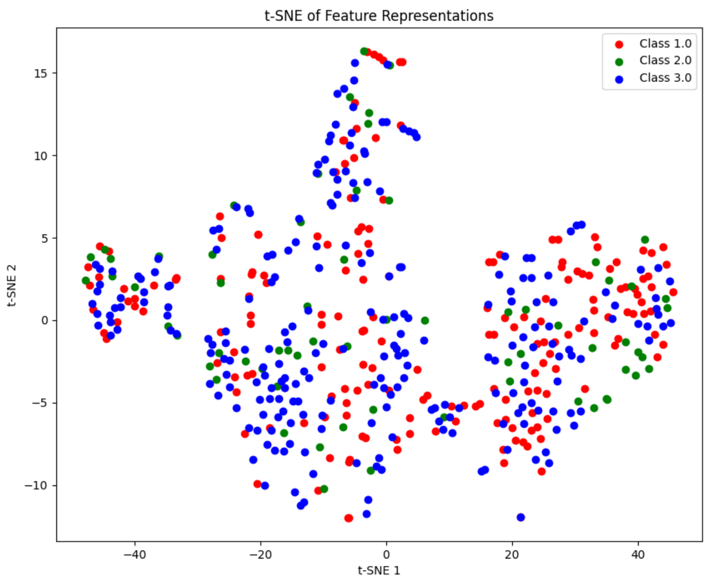

Figure 7.10: TS2Vec t-SNE visualization of ADP embeddings for 3D label type

The t-SNE plots of the features extracted through Supervised Contrastive Learning depict clusters that are somewhat entangled, with overlaps between classes. This might suggest that while the model is capable of learning a nuanced representation of the data, it may struggle with clearly demarcating the boundaries necessary for high-precision classification tasks. The t-SNE visualizations of features extracted via TS2Vec show different clustering patterns. These plots appear to have distinct clusters with less overlap between classes than the Supervised Contrastive

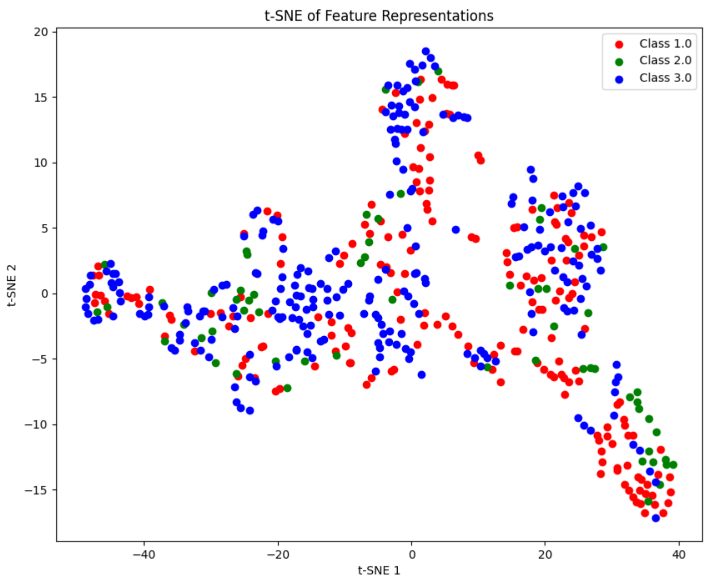

Figure 7.11: TS2Vec t-SNE visualization of ADP embeddings for 5D label type

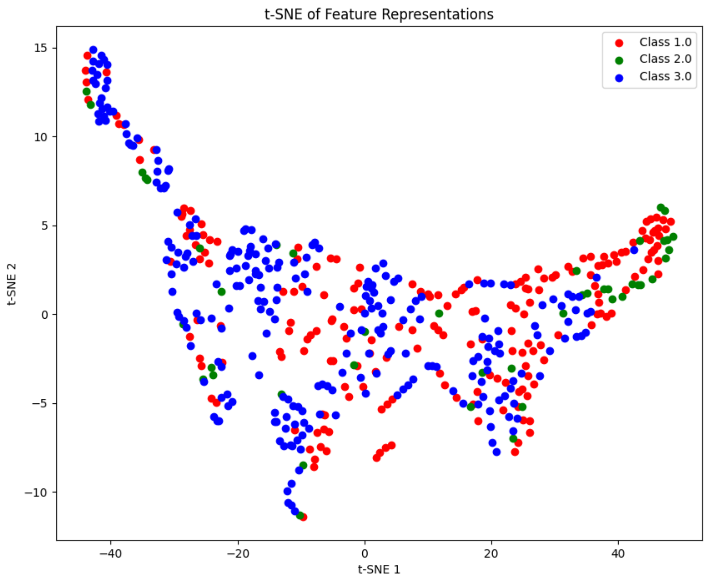

Figure 7.12: TS2Vec t-SNE visualization of ADP embeddings for 7D label type

Learning features. The more defined separation could indicate that TS2Vec is capturing more discriminative features conducive to distinguishing between different classes in the dataset, but less marked patterns in the information extracted. This could imply that TS2Vec extracts features less sensitive to the noise and variance within the time series data. In contrast, the Supervised Contrastive Learning representations, while indicating some degree of separation, suggest a feature space where the boundaries between classes are less clear. This may result from the contrastive loss function prioritizing relative distances over absolute positioning, which can sometimes lead to a trade-off between class compactness and separability.

AMGN Ticker Analysis

The t-SNE visualizations for the AMGN ticker, based on features extracted from the Supervised Contrastive Learning (SupCon) model, present an informative view of the feature space's extracted by our main model.

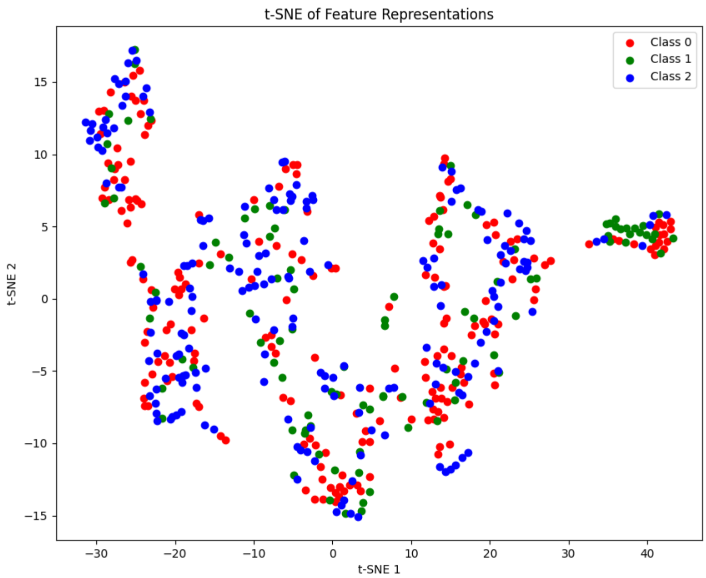

Figure 7.13: SupCon t-SNE visualization of AMGN embeddings for 3D label type

Considering the plot for 3Days type label, Fig. 7.13, we observe a moderate degree of class overlap, particularly between Classes 0 and 2, suggesting potential challenges in delineating these classes. Class 1, however, appears relatively more distinct, indicating some level of effective feature separation. Proceeding with plot

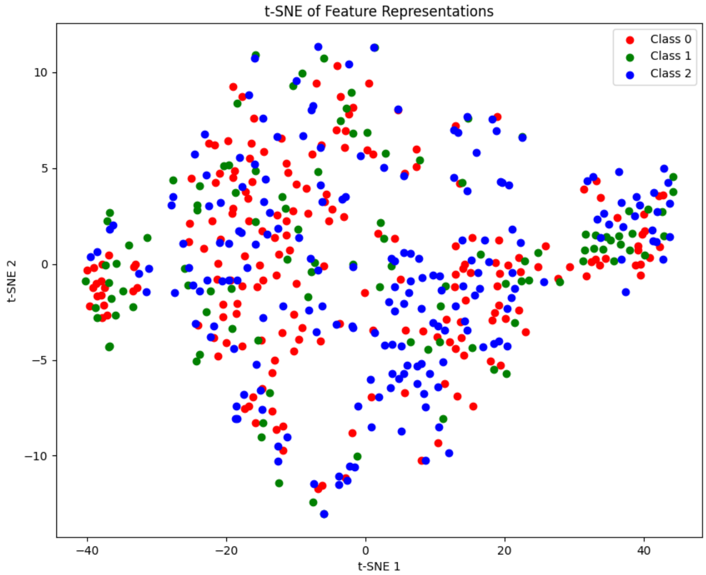

Figure 7.14: SupCon t-SNE visualization of AMGN embeddings for 5D label type

in Fig. 7.14, the dispersion of data points suggests a higher level of feature space complexity, due to also the highest complexity in trying to predict fluctuations in a larger time span. The clusters for each class are somewhat scattered, with significant intermixing, which could imply a lack of distinct boundaries necessary for high-precision classification tasks. Eventually, Fig. 7.15 visualization further highlights the blending of Class 0 and Class 2, with Class 1 maintaining a semblance of separation. The tendency of points to form elongated clusters may reflect certain temporal patterns within the data.

The patterns observed suggest that the SupCon model might be capturing meaningful temporal dynamics, yet may struggle to consistently differentiate between classes in higher-dimensional feature spaces.

As presented for ADP tickers, the t-SNE visualizations of the dataset annotated with dates offer an additional dimension of analysis by incorporating the temporal sequence of the data. This provides a perspective on how the feature space adapts over time for each dimensionality considered.

Refering to figure 7.16, this visualization indicates that the raw data possesses an inherent temporal pattern, with clusters formed around certain dates potentially corresponding to specific market behaviors or events. As we move to a 5D

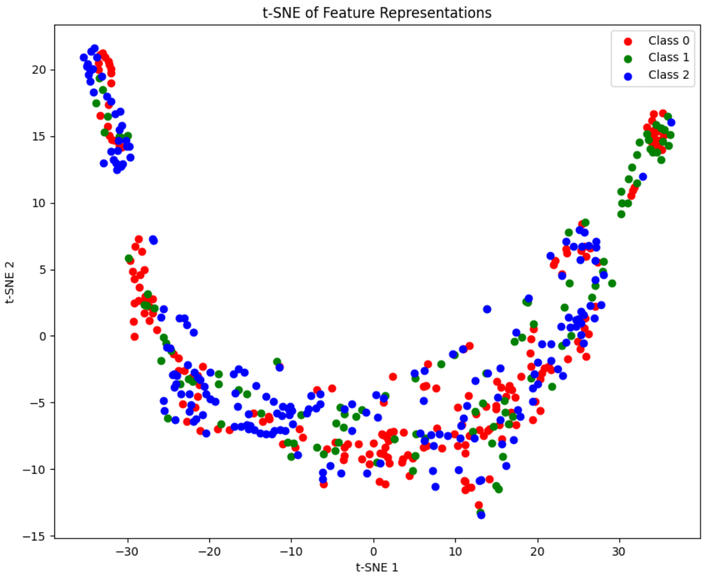

Figure 7.15: SupCon t-SNE visualization of AMGN embeddings for 7D label type

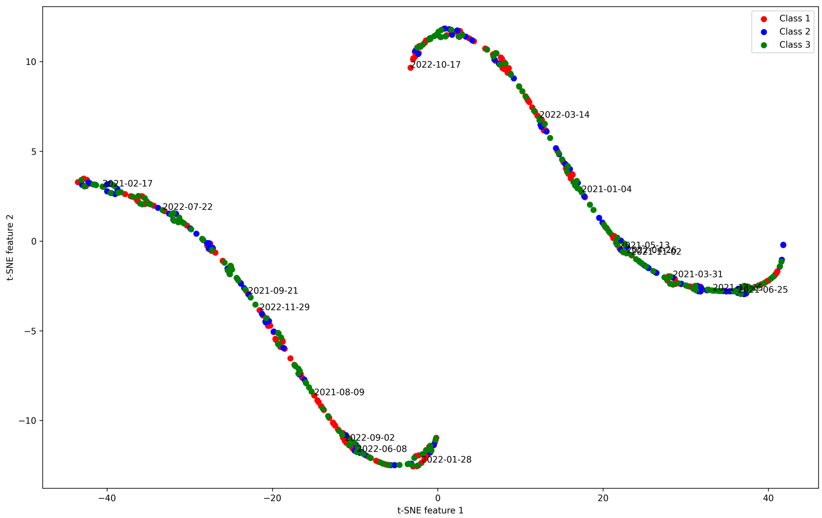

Figure 7.16: t-SNE visualization of AMGN dataset for 3D label type

representation, the temporal labels reveal overlapping time periods within clusters. This overlap may suggest that the data for these dates share similar characteristics or that the market conditions during these periods result in similar stock behaviors, as reflected in the dataset. The 7D t-SNE plot further emphasizes the complex interplay between time and data structure. The spread of temporal points within clusters is indicative of the multifaceted nature of stock data, capturing more nuanced time-related patterns and variances in the stock's behavior.

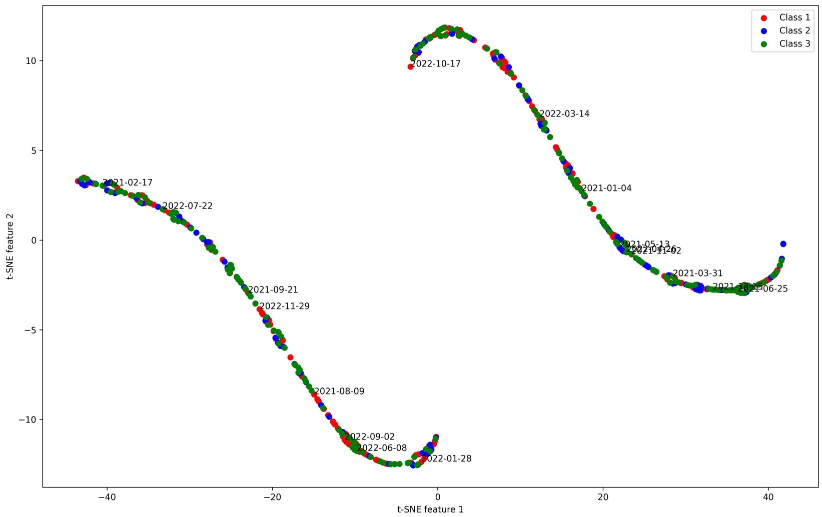

Figure 7.17: t-SNE visualization of AMGN dataset for 5D label type

As done for the previous ticker, we shall now extend our qualitative analysis by examining traditional technical indicators such as Bollinger Bands, volume data, and candlestick patterns alongside the t-SNE visualizations of the AMGN dataset.

By observing the periods where the bands widen, we can infer increased market volatility, which could correspond to regions of greater dispersion or overlapping clusters in the t-SNE plots. This indicates that the data during volatile periods might be more challenging to cluster distinctly, a factor that should be considered when interpreting the t-SNE visualizations.

The volume charts reflect the intensity of trading activity, with spikes often coinciding with significant price movements or market events. When examining the candlestick patterns instead, in relation to the t-SNE clusters, it might be possible to identify how certain patterns of price movements are represented within the feature space, potentially as outliers or boundary points between clusters.

Integrating technical indicators with t-SNE visualizations for the AMGN features

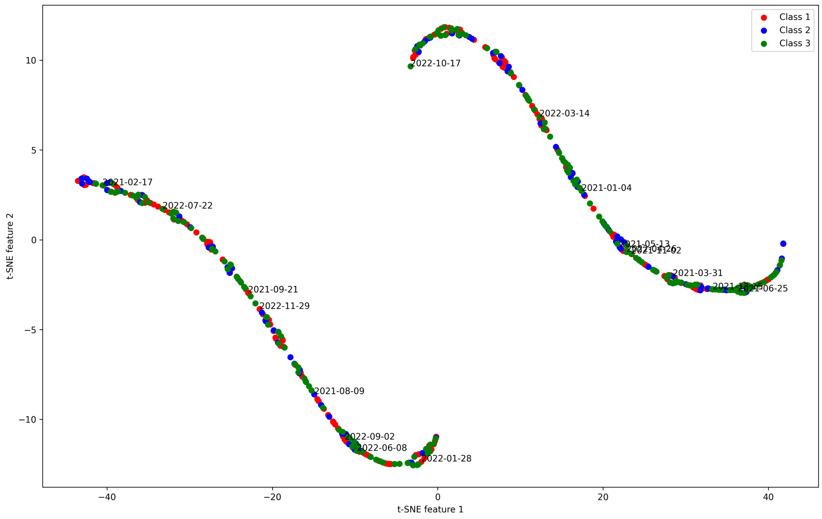

Figure 7.18: t-SNE visualization of AMGN dataset for 7D label type

underscores again the intricate relationship between market dynamics and the data's representational structure.

Figure 7.19: Candlestick visualization of AMGN dataset

Figure 7.20: Bollinger Bands visualization of AMGN dataset

Figure 7.21: Volumes visualization of AMGN dataset

Eventually, the t-SNE visualizations based on TS2Vec features, Figures 7.22, 7.23, and 7.24, appear to show a different pattern, with clusters being more defined and less inter-class overlap. This could indicate that TS2Vec is capturing unique aspects of the data that allow for better class separation, potentially making it more suitable for classification tasks that require clear discrimination between different states of the stock. However, neither in this case the separation between classes is evident, and both models are poorly performing the target task.

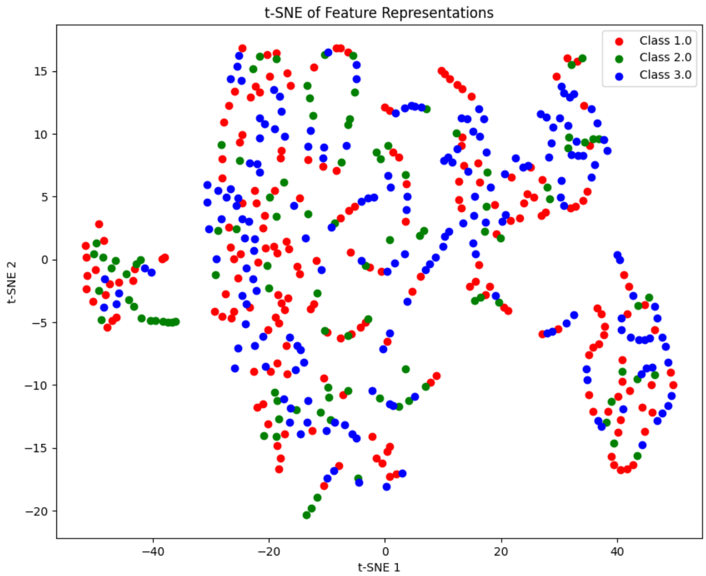

Figure 7.22: TS2Vec t-SNE visualization of AMGN embeddings for 3D label type

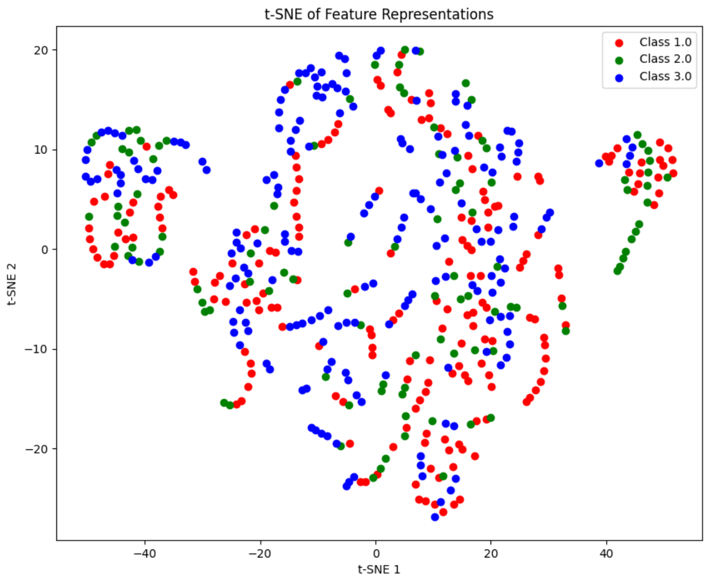

Figure 7.23: TS2Vec t-SNE visualization of AMGN embeddings for 5D label type

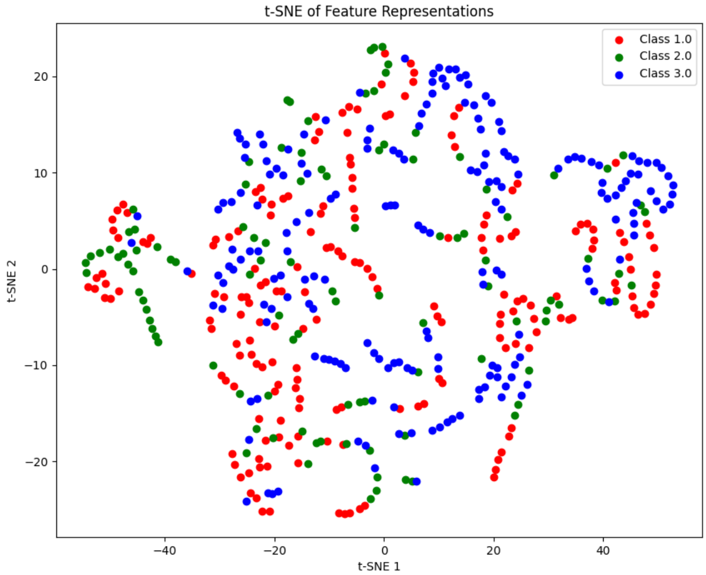

Figure 7.24: TS2Vec t-SNE visualization of AMGN embeddings for 7D label type

MSFT Ticker Analysis

We eventually conduct an analysis of the MSFT stock dataset by examining t-SNE visualizations of features extracted through the Supervised Contrastive Learning (SupCon) model.

Exploring the feature representations of the MSFT dataset obtained from Supervised Contrastive Learning, we've found a pattern that aligns with observations from other stock tickers we've analyzed. Specifically, in Figures 7.25, 7.26, and 7.27, the displayed class overlaps and the spread of the clusters show that, although there is some discernible structure in the data, it's not quite enough for a clear-cut distinction among the classes.

The t-SNE plots reveal that classes aren't as neatly separable as we would like, suggesting that the model might not be capturing all the nuances necessary to distinguish one class from another accurately. This problem is compounded by an apparent preference for clustering the data into one or two classes, which may point to a tendency of the model to 'favor' certain data patterns. Such favoritism can lead to skewed representations, potentially causing some classes to be underrepresented in the feature space.

As we delve into the t-SNE representations for the MSFT dataset, it becomes

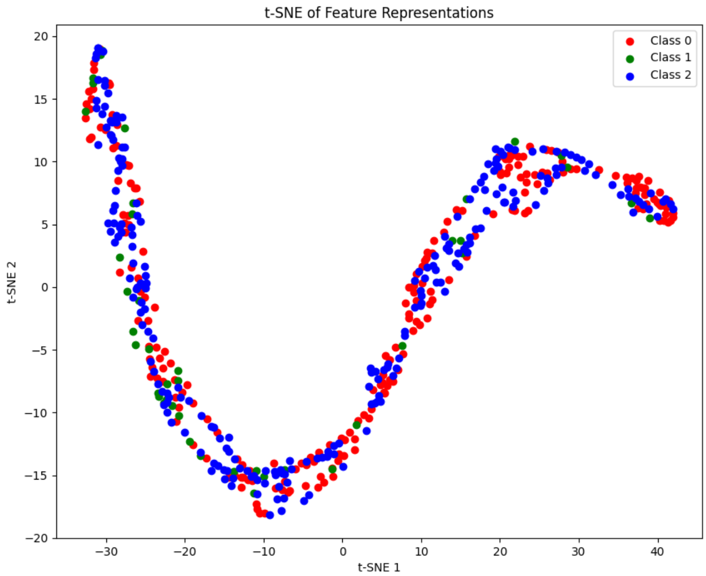

Figure 7.25: SupCon t-SNE visualization of MSFT embeddings for 3D label type

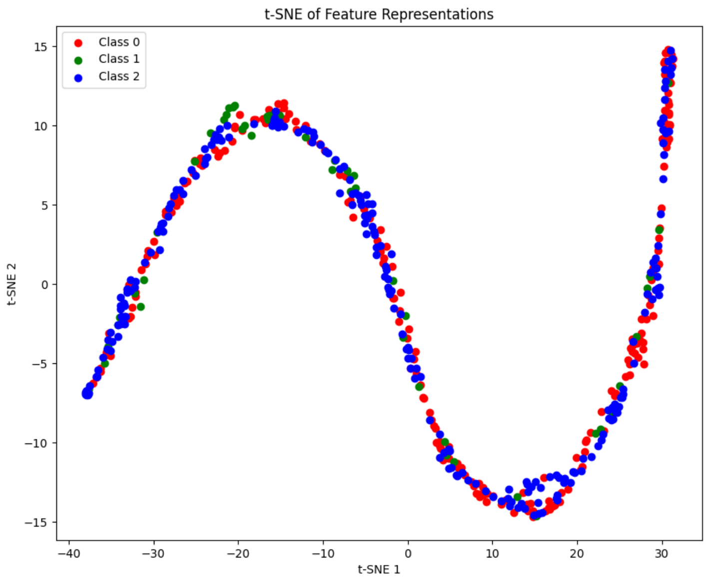

Figure 7.26: SupCon t-SNE visualization of MSFT embeddings for 5D label type

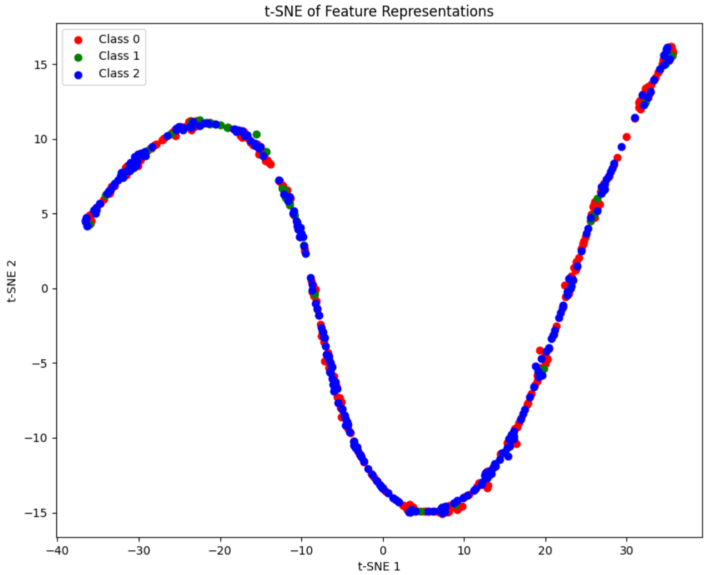

Figure 7.27: SupCon t-SNE visualization of MSFT embeddings for 7D label type

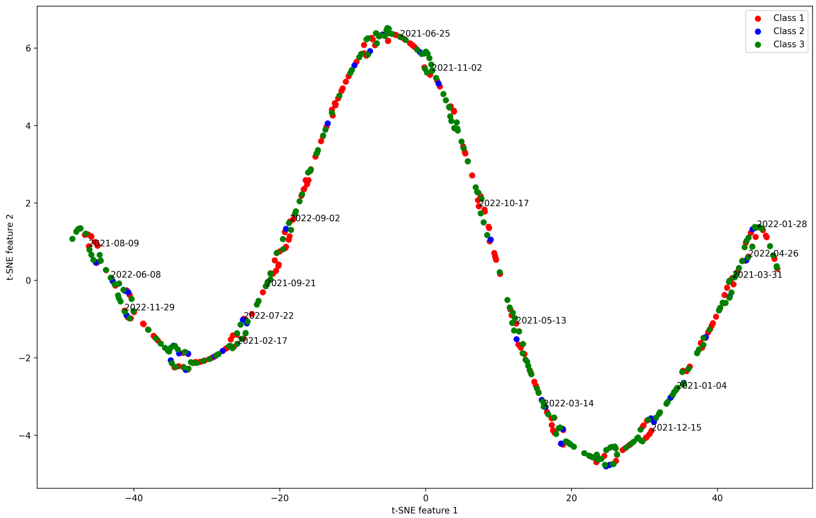

Figure 7.28: t-SNE visualization of MSFT dataset for 3D label type

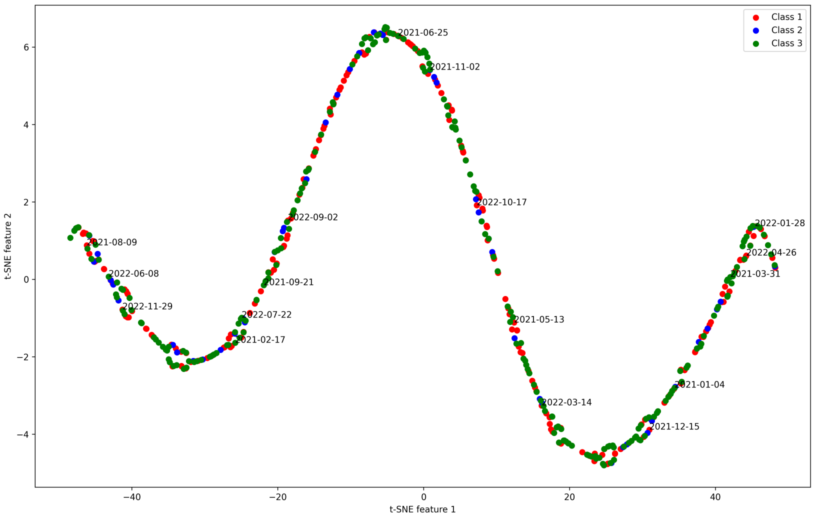

Figure 7.29: t-SNE visualization of MSFT dataset for 5D label type

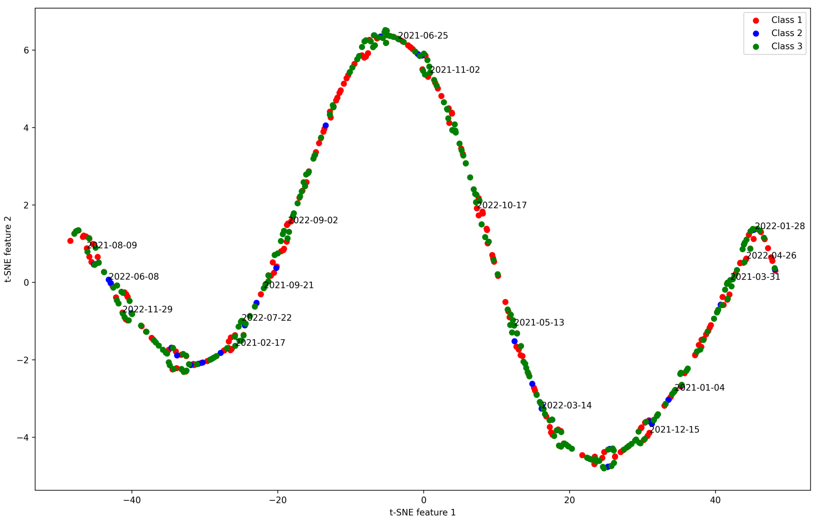

Figure 7.30: t-SNE visualization of MSFT dataset for 7D label type

evident that there are patterns present that echo the findings from the embeddings analysis. Notably, the overlap of classes in the t-SNE plots could be indicative of significant market activities, which is an observation that becomes particularly interesting when juxtaposed with additional financial metrics. By integrating this insight with other economic indicators represented in Figures 7.32, 7.31, and 7.33, we're able to discern a pattern of daily high trading activity. Such vigorous trading can contribute to the challenge of distinguishing between different classes, as it introduces a layer of complexity to the stock's behavior.

Figure 7.31: Candlestick visualization of MSFT dataset

Moreover, the Bollinger Bands provide a visual narrative of high volatility, particularly during the winter of 2021 and the fall of 2022. This volatility is not just a simple fluctuation; it adds substantial noise to the dataset, which could be one of the factors that make class differentiation more challenging. In wrapping up this analysis, it's clear that while the t-SNE visualization aids in identifying data patterns, the inherent market activity and volatility reflected in the economic indicators further complicate the separation of classes.

Examining the t-SNE feature representations obtained from the Time Series to Vectors (TS2Vec) model on the MSFT dataset gives us an insightful perspective into the extracted temporal patterns. The TS2Vec model, known for its capability to learn deep representations of time series data, provides us with a different angle to observe the clustering behavior of the stock market data.

The t-SNE visualizations derived from the TS2Vec model exhibit certain traits of clustering, which might reflect the underlying temporal dynamics captured from

Figure 7.32: Bollinger Bands visualization of MSFT dataset

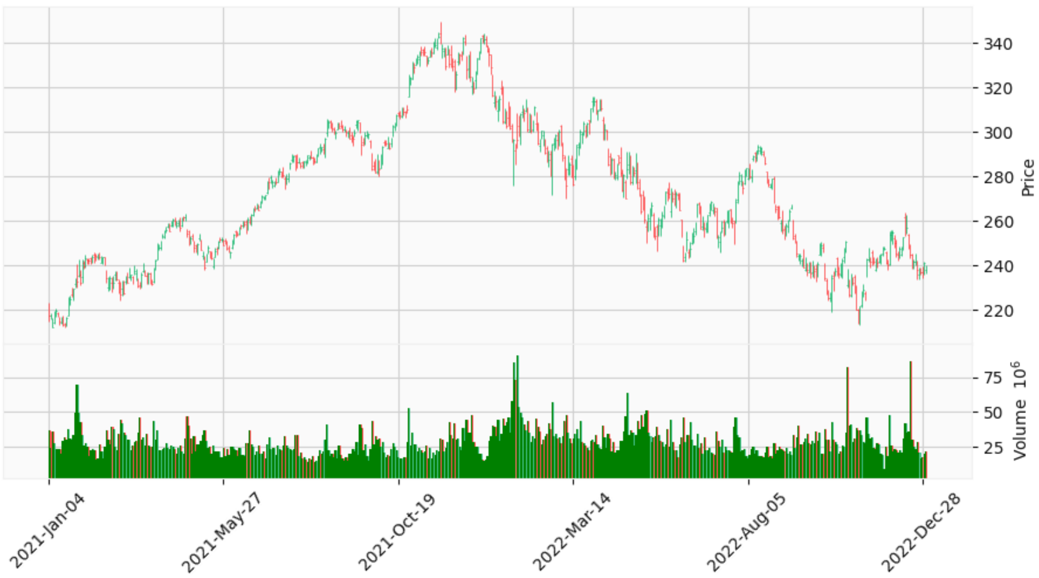

Figure 7.33: Volumes visualization of MSFT dataset

the MSFT time series. Unlike the Supervised Contrastive Learning model, the TS2Vec appears to organize the data into more outlined clusters.

Across the 3D, 5D, and 7D t-SNE visualizations, there is a visible tendency for the classes to form distinct groupings, although some overlap persists. This clustering behavior suggests that TS2Vec is learning somehow some sort of representations that encapsulate significant features of the time series, potentially

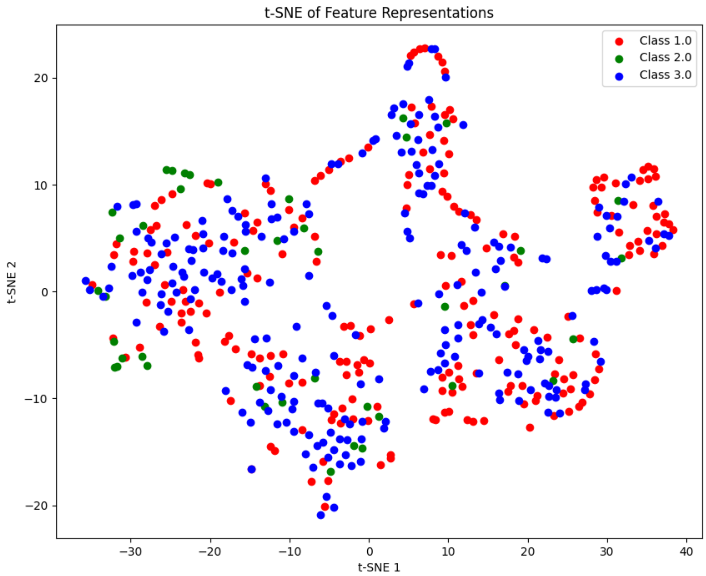

Figure 7.34: TS2Vec t-SNE visualization of MSFT embeddings for 3D label type

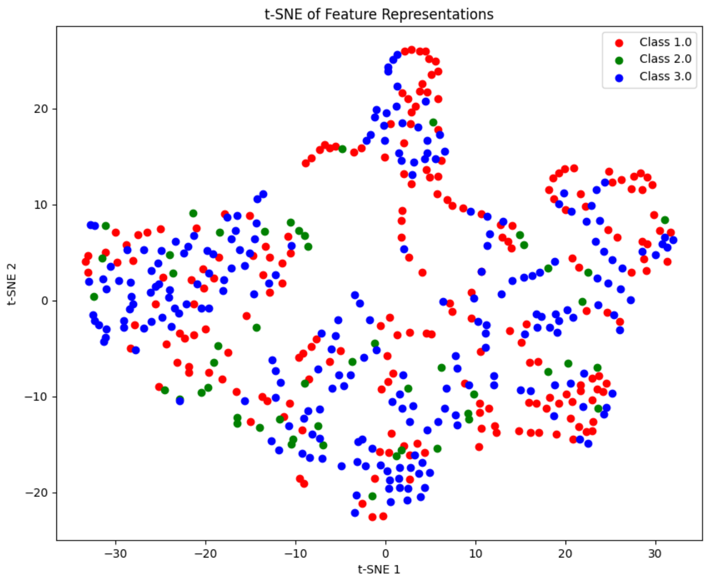

Figure 7.35: TS2Vec t-SNE visualization of MSFT embeddings for 5D label type

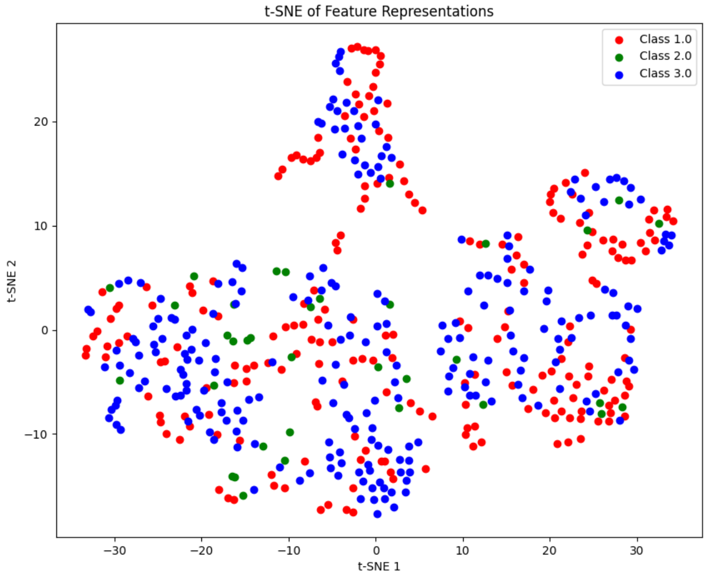

Figure 7.36: TS2Vec t-SNE visualization of MSFT embeddings for 7D label type

improving the discriminability between different periods or market conditions.

It's crucial to note that the presence of class overlaps may still pose challenges for any downstream tasks, such as forecasting or anomaly detection. Hence, there may be a need for further model tuning or combining TS2Vec with additional feature engineering techniques to enhance class separability.

7.3.4 Cross-Analysis Summary

Throughout the comprehensive analysis of feature representations from the Supervised Contrastive Learning and Time Series to Vectors (TS2Vec) models on three distinct stock datasets — ADP, AMGN, and MSFT — we have gleaned valuable insights into the intricate dynamics of financial time series data. Our qualitative analysis, underpinned by t-SNE visualizations, revealed recurring themes of class overlap and varying degrees of cluster density across all tickers, highlighting the challenges in achieving clear class separability in high-dimensional feature spaces.

These qualitative observations are further contextualized by the quantitative analysis conducted in previous discussions. The numerical assessments complemented our visual findings, quantifying the degree of class dispersion and the models' discriminative power. Where t-SNE plots suggested class biases or feature space complexities, quantitative metrics provided objective evidence of the models' performance, including their accuracy in classifying and forecasting market trends. In combining both qualitative and quantitative analyses, we arrive at a multidimensional understanding of feature extraction models in financial time series.

The datasets examined posed incredibly challenging scenarios for feature extraction. The complex nature of financial data, characterized by non-linear relationships and influenced by a multitude of external factors, was mirrored in the subtleties these models sought to unravel. Particularly, the Supervised Contrastive Learning model, though falling short in performance metrics, brought to light the demanding nature of financial datasets. Its struggle to distinctly separate classes within the feature space was telling of the inherent difficulty presented by such sophisticated data. Despite the SupCon model's limitations in performance, it has laid a foundation for further investigation. Its qualitative outcomes have provided us with directional cues for deeper inquiry and have prompted a rethinking of approach and methodology.

Chapter 8

Conclusions and Future Works

This thesis embarked on an exploratory journey into the innovative realm of Supervised Contrastive Learning (SCL) for the classification of market stock series. It set out with the ambition to delve into the intricacies of financial time series analysis and leverage the potential of SCL as a novel approach in this domain. Through a rigorous methodology, encompassing minimal data preprocessing and the deployment of a deep Residual Network (ResNet), we aimed to train a model capable of discerning between different market states with high accuracy. Despite the challenges encountered, particularly the model's performance on the financial dataset, this research journey has been both enlightening and enriching.

The investigation provided valuable insights into the application of deep learning techniques, particularly SCL, in the analysis of financial time series data. The approach demonstrated promise in theoretical discussions and initial benchmarks, illustrating the potential of SCL to enhance classification tasks by fostering robust and discriminative feature learning. Although the expected performance uplift in real-world financial data classification was not realized, the exploration shed light on the complexities and challenges inherent in financial time series analysis. It underscored the nuanced nature of market data and the intricate patterns that dictate market behaviors.

This thesis contributes to the broader academic and practical discourse on financial technology innovation, highlighting the importance of continuous exploration and experimentation in developing analytical tools for financial markets. The journey through this research has reaffirmed the belief in the potential of deep learning and contrastive learning techniques to transform financial analysis, despite the challenges encountered.

The path of research is never-ending, and every conclusion opens new paths for exploration. The journey undertaken in this thesis lays the groundwork for future investigations into the application of Supervised Contrastive Learning in time series classification, particularly within the financial domain. Several directions are identified for future works:

- *Exploring Alternative Base Models*: The current research utilized a Residual Network as the base model. Future work could explore alternative architectures, such as Transformer models or more advanced CNNs, to evaluate their efficacy in conjunction with SCL for financial time series analysis.
- *Diversifying Domains*: While this thesis focused on the financial market, SCL's application could be tested across various domains where time series data is critical, such as healthcare, energy consumption, or climatology. Such exploration could uncover domain-specific insights and model adaptability.
- *In-Depth Data Augmentation and Preprocessing*: This thesis adopted minimal data preprocessing to maintain the raw essence of financial time series. Future studies could delve into sophisticated data augmentation and preprocessing techniques to uncover potentially hidden patterns within the data, enhancing model performance.
- *Hybrid Models and Ensemble Techniques*: Combining SCL with other learning paradigms, such as unsupervised learning or reinforcement learning, could offer innovative approaches to time series classification. Ensemble methods that integrate multiple models could also be explored to improve prediction accuracy.

In conclusion, while the model did not meet all the anticipated outcomes in classifying financial time series data, the research underscores the complexity of financial markets and the potential of advanced machine learning techniques in navigating them. The exploration of Supervised Contrastive Learning in this context opens up new horizons for innovation in financial analysis, promising a future where machine learning not only aids but also enhances decision-making in the financial industry. The journey continues, and the quest for knowledge and improvement remains relentless.

Bibliography

[1] Periklis Gogas and Theophilos Papadimitriou. «Machine Learning in Economics and Finance». In: *Computational Economics* 57.1 (2021), pp. 1–4. DOI: `https://doi.org/10.1007/s10614-021-10094-w` (cit. on p. 1).

[2] John W Goodell, Satish Kumar, Weng Marc Lim, and Debidutta Pattnaik. «Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis». In: *Journal of Behavioral and Experimental Finance* 32 (2021), p. 100577 (cit. on p. 2).

[3] NASDAQ. *NASDAQ 100 (NDX) Latest Quotes, Charts, Data & News.* Accessed: 2023-10-07. 2023. URL: `https://www.nasdaq.com/market-activity/index/ndx` (cit. on pp. 3, 44).

[4] John McCarthy. «WHAT IS ARTIFICIAL INTELLIGENCE?» In: (2007). http://www-formal.stanford.edu/jmc/ (cit. on p. 5).

[5] Director of Chappuis Halder Co. Sébastien Meunier. *International Banker: The Impacts and Challenges of Artificial Intelligence in Finance.* 2018. URL: `https://internationalbanker.com/finance/the-impacts-and-challenges-of-artificial-intelligence-in-finance/` (cit. on pp. 5, 6).

[6] NetBase Quid. *Data on global corporate investment in AI from 2013 to 2022.* Tech. rep. The AI Index is an independent initiative at the Stanford University Institute for Human-Centered Artificial Intelligence. Their flagship output is the annual AI Index Report, which has been published since 2017. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2023. URL: `https://drive.google.com/drive/folders/1ma9WZJzKreS8f2It1rMy_KkkbX6XwDOK` (cit. on p. 6).

[7] IBM. *What is machine learning?* Accessed: 2023-10. URL: `https://www.ibm.com/topics/machine-learning` (cit. on p. 8).

[8] MathWorks. *Reinforcement Learning Explained.* `https://www.mathworks.com/discovery/reinforcement-learning.html`. Accessed: 2024-03-16. 2023 (cit. on p. 8).

[9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. `http://www.deeplearningbook.org`. MIT Press, 2016 (cit. on p. 11).

[10] IBM. *What is artificial intelligence?* Accessed: 2023-10. URL: `https://www.ibm.com/topics/artificial-intelligence` (cit. on p. 12).

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. *Deep Residual Learning for Image Recognition*. 2015. arXiv: `1512.03385 [cs.CV]` (cit. on p. 12).

[12] R. Hadsell, S. Chopra, and Y. LeCun. «Dimensionality Reduction by Learning an Invariant Mapping». In: *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*. Vol. 2. 2006, pp. 1735–1742. DOI: `10.1109/CVPR.2006.100` (cit. on pp. 14, 15).

[13] Florian Schroff, Dmitry Kalenichenko, and James Philbin. «FaceNet: A Unified Embedding for Face Recognition and Clustering». In: *CoRR* abs/1503.03832 (2015). arXiv: `1503.03832`. URL: `http://arxiv.org/abs/1503.03832` (cit. on pp. 14, 15).

[14] V7 Labs. *The Beginner's Guide to Contrastive Learning*. 2022. URL: `https://www.v7labs.com/blog/contrastive-learning-guide` (cit. on p. 16).

[15] Gaudenz Boesch. *Image Data Augmentation for Computer Vision (2023 Guide)*. 2023. URL: `https://viso.ai/computer-vision/image-data-augmentation-for-computer-vision/` (cit. on p. 18).

[16] John J Murphy. *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. Penguin, 1999 (cit. on pp. 20, 22).

[17] Investopedia. *Fundamental vs. Technical Analysis: What's the Difference?* Accessed: 2023-10. URL: `https://www.investopedia.com/ask/answers/difference-between-fundamental-and-technical-analysis/` (cit. on p. 23).

[18] Patara Trirat, Yooju Shin, Junhyeok Kang, Youngeun Nam, Jihye Na, Minyoung Bae, Joeun Kim, Byunghyun Kim, and Jae-Gil Lee. *Universal Time-Series Representation Learning: A Survey*. 2024. arXiv: `2401.03717 [cs.LG]` (cit. on p. 30).

[19] Qianli Ma, Zhen Liu, Zhenjing Zheng, Ziyang Huang, Siying Zhu, Zhongzhong Yu, and James T. Kwok. *A Survey on Time-Series Pre-Trained Models*. 2023. arXiv: `2305.10716 [cs.LG]` (cit. on p. 30).

[20] Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. *TS2Vec: Towards Universal Representation of Time Series*. 2022. arXiv: `2106.10466 [cs.LG]` (cit. on pp. 32, 38, 47).

[21] Zhiguang Wang, Weizhong Yan, and Tim Oates. *Time Series Classification from Scratch with Deep Neural Networks: A Strong Baseline*. 2016. arXiv: 1611.06455 [cs.LG] (cit. on pp. 33, 37, 38).

[22] Sana Tonekaboni, Danny Eytan, and Anna Goldenberg. *Unsupervised Representation Learning for Time Series with Temporal Neighborhood Coding*. 2021. arXiv: 2106.00750 [cs.LG] (cit. on pp. 34, 47).

[23] Jean-Yves Franceschi, Aymeric Dieuleveut, and Martin Jaggi. *Unsupervised Scalable Representation Learning for Multivariate Time Series*. 2020. arXiv: 1901.10738 [cs.LG] (cit. on pp. 34, 47).

[24] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. *Supervised Contrastive Learning*. 2021. arXiv: 2004.11362 [cs.LG] (cit. on pp. 39, 40).

[25] Eamonn Keogh and Anthony Bagnall. *The UCR/UEA Time Series Classification Repository*. https://www.cs.ucr.edu/~eamonn/time_series_data_2018/. Accessed: date-of-access. 2018 (cit. on pp. 41, 45).

[26] Ran Aroussi. *yfinance: Download market data from Yahoo! Finance API*. Python package available from PyPi. 2023. URL: https://pypi.org/project/yfinance/ (cit. on p. 42).

[27] Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Chee Keong Kwoh, Xiaoli Li, and Cuntai Guan. *Time-Series Representation Learning via Temporal and Contextual Contrasting*. 2021. arXiv: 2106.14112 [cs.LG] (cit. on p. 47).

[28] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. *A Transformer-based Framework for Multivariate Time Series Representation Learning*. 2020. arXiv: 2010.02803 [cs.LG] (cit. on p. 47).

[29] Yanping Chen, Bing Hu, Eamonn Keogh, and Gustavo EAPA Batista. «DTW-D: time series semi-supervised learning from a single example». In: (2013), pp. 383–391 (cit. on p. 47).