

Rancher

Edge Analytics in Finance

Market Predictions

Rancher Prime by SUSE
K3s
SUSE Linux Enterprise Server

Navin Chandra, Google Summer of Code 2022 Contributor (openSUSE)

Edge Analytics in Finance

Market Predictions

Date: 2022-12-16

Summary

This document provides an introduction to the analytics edge ecosystem with Rancher by SUSE through the deployment of a machine learning model and Web-based user interface for generating and visualizing stock market predictions.

This project is part of the Google Summer of Code, a global online mentoring program focused on introducing new contributors to open source software development.

Disclaimer

Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Contents



1	Introduction	4
2	Prerequisites	5
3	Technical overview	6
4	ML model	7
5	Front-end	8
6	Build environment	10
7	Containerization	10
8	Deployment manifests	12
9	Deployment	14
10	Demonstration	16
11	Summary	19
12	Legal notice	21
13	GNU Free Documentation License	22

1 Introduction



1.1 Motivation

The finance industry relies on data analytics, including machine learning (ML), to describe, predict, and improve performance. Data can be anywhere, and is not always found in the core data center. Increasingly, the data needed by organizations is located at the edge - in branch sites and even in mobile and IoT devices. Moving data from edge to core introduces latency, costs, and security risks.

The modern container landscape, built on open source with Linux and Kubernetes, makes it possible to deploy powerful analytics and machine learning workloads at the edge and at scale. By reducing the need for data movement, organizations can improve security, reduce costs, and speed time to insight and time to value.

This project is part of the 2022 [Google Summer of Code](https://summerofcode.with-google.com/) (<https://summerofcode.with-google.com/>) , a global mentoring program focused on introducing new contributors to open source software development. Participating Google Summer of Code contributors are paired with mentors from open source organizations (such as [SUSE](https://suse.com/) (<https://suse.com/>) ) to gain exposure to real-world software development techniques. Contributors learn from experienced open source developers while writing code for real-world projects.

1.2 Scope


Edge analytics in the financial sector is simplified through the use of modern container technologies. This is illustrated by using [Rancher](https://www.rancher.com/products/rancher/) (<https://www.rancher.com/products/rancher/>) , by SUSE to deploy a machine learning model for stock price predictions and a Web-based front-end for user interaction and visualization into a lightweight, [K3s](https://k3s.io/) (<https://k3s.io/>)  Kubernetes cluster.

1.3 Audience

DevOps, MLOps, developers and data scientists looking to containerize their ML workloads and deploy them anywhere (including at the edge) will find useful guidance in this document. To get the most out of this guide, you should have basic familiarity with the Linux command line, Python, Docker, and Kubernetes.


2 Prerequisites

For this guide, you need:


- A Rancher (<https://www.rancher.com/products/rancher/>)  environment.
You can use Rancher or Rancher Prime by SUSE version 2.6 or later.
- A K3s Kubernetes cluster managed by Rancher.
K3s 1.21.3 + k3s1 or later should work without modification.



Note


K3s is an ideal Kubernetes distribution for edge deployments. With Rancher and Rancher Prime by SUSE, you can manage any [CNCF-certified Kubernetes distribution](https://www.cncf.io/certification/software-conformance/#logos) (<https://www.cncf.io/certification/software-conformance/#logos>) .

For this guide, your cluster should consist of at least one worker node with the following specifications:

- Processor: 2 vCPUs
 - Memory: 4 GB RAM
 - Storage: 10 GB (available)
 - Operating system: Linux (SLES 15 SP4 or later is recommended).
- The ability to install and execute [Python](https://www.python.org/) (<https://www.python.org/>)  programs.
Machine learning software is in rapid development, which can result in version incompatibilities. Python 3.10 or later should work with little to no modification.



Tip

See [Replicable Python Environments with JupyterLab and SUSE Linux Enterprise Server](https://documentation.suse.com/trd/jupyter/html/gs_sles_jupyter-jupyterlab/index.html) (https://documentation.suse.com/trd/jupyter/html/gs_sles_jupyter-jupyterlab/index.html)  for guidance on setting up your own Python environment.

- A Docker Hub (<https://hub.docker.com/>)  account.



Note

In this guide, you push your workload container image to Docker Hub and, later, pull this image for deployment onto your cluster. With appropriate command substitutions, any container image repository may be used.

3 Technical overview

In this guide, you use Rancher to deploy Fin-Dashboard, a stock prediction Web application, which you can obtain from the project's [code repository \(https://github.com/navin772/GSOC_S-stock_prediction/tree/master/fin_dashboard\)](https://github.com/navin772/GSOC_S-stock_prediction/tree/master/fin_dashboard). Fin-Dashboard is discussed in more detail later.

There can be many paths to creating and deploying analytics applications. A typical process involves the following steps:

1. Collect the required data and perform cleansing/preprocessing.

2. Create and test an ML model.

This can involve exploring different models in an interactive environment, like [JupyterLab \(https://jupyter.org/\)](https://jupyter.org/).

3. Enable a means of interacting with the model.

In this guide, you deploy a front-end Web application to accept user input and view predictions.

4. Containerize the application.










In this guide, you build a [Docker \(https://docs.docker.com/get-started/overview/\)](https://docs.docker.com/get-started/overview/) container.

5. Define a deployment mechanism.


This is done by creating manifest files.

6. Deploy the application.

Some additional components and tools used in this guide or by Fin-Dashboard include:

- Python packages, including [pandas](https://pandas.pydata.org/) (<https://pandas.pydata.org/>)  and [plotly](https://plotly.com/python/) (<https://plotly.com/python/>)  - for managing and visualizing numerical data.
- Jupyter Notebook (<https://jupyter.org/>)  - an interactive computational environment to explore and visualize data.
- Fbprophet (<https://facebook.github.io/prophet/>)  - to create the ML model to predict the stock/index price.
- Streamlit (<https://streamlit.io/>)  - to build a Web-based, interactive front-end for the ML model.
- VADER (<https://github.com/cjhutto/vaderSentiment>)  - to analyze market sentiment based on news content.
- Yahoo Finance (<https://pypi.org/project/yfinance/>)  - to fetch historical stock data.
- Kubectl (<https://kubernetes.io/docs/reference/kubectl/>)  - a command line tool for controlling the Kubernetes cluster manager.
- Kustomize (<https://kustomize.io/>)  - a template-free way to customize application configuration that simplifies the use of off-the-shelf applications.

4 ML model

In this section, you explore a subset of the Fin-Dashboard code by creating a [Jupyter Notebook](https://jupyterlab.readthedocs.io/en/stable/) (<https://jupyterlab.readthedocs.io/en/stable/>)  that illustrates how data is accumulated and processed, a model is created and used, and model predictions are visualized.

Open a new Jupyter Notebook and enter the code below.

1. Fetch the stock/index data from Yahoo Finance and store it as a dataframe.

```
import yfinance as yf
company = 'TSLA'
df = yf.download(tickers=company, period="7y", interval="1d")
```

2. Preprocess the data and rename columns for the Prophet ML model.

```
df = df.reset_index()
```

```
df = df[["Date", "Close"]]
df = df.rename(columns={"Date": "ds", "Close": "y"})
```

3. Initialize the ML model and make future predictions.

```
from fbprophet import Prophet
fp = Prophet(yearly_seasonality=True)
fp.fit(df)

future = fp.make_future_dataframe(periods=period)
forecast = fp.predict(future)
```

4. Visualize the predictions.

```
from prophet.plot import plot_plotly
fig = plot_plotly(fp, forecast)
fig.update_xaxes(title_text="Time")
fig.show()
```



5 Front-end

In addition to the ML analytics workload, Fin-Dashboard also provides a Web-based front-end to support user interaction.





5.1 Features

The front-end application features useful functionality for the user, including:

- Automatic detection of the user's country to show predictions for the top five, country-specific stocks
- Ability for users to select long- or short-term predictions and to specify the number of forecasted days
- Ability for users to choose other stocks and other stock exchanges
- For long term predictions, a comparison of the model's predicted price with the analyst's target price
- For short term predictions, a sentiment analysis of the stock news to enable visualization of its impact on stock price
- Correlation of a stock with the index

5.2 Components

Fin-Dashboard is composed of the following Python modules:

- `geo_location.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/geo_location.py) : Determines the country where the user is located to give country specific information.
- `index_correlation.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/index_correlation.py) : Correlates the stock with its parent index.
- `main.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/main.py) : Contains all the Streamlit configurations/page setup and uses the other files to create the application.
- `plot_charts.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/plot_charts.py) : Contains the modified and generalized form of the Jupyter Notebook code we created in the previous steps. Contains two python functions - one for long term and other for short term predictions.

- `sentiment.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/sentiment.py) : Collects stock news from Google News and uses the **VADER sentiment analysis** tool to determine the stock's trend.
- `stock_recomendation.py` (https://github.com/navin772/GSOC_Stock_prediction/blob/master/fin_dashboard/stock_recomendation.py) : Extracts the analyst's price target of a stock from Yahoo Finance.

6 Build environment


Create your build environment as follows:

1. Clone the Fin-Dashboard code from GitHub into an empty directory.

```
git clone https://github.com/navin772/GSOC_Stock_prediction
```

2. Elsewhere, create your working directory, called 'stock_frontend'.
3. Copy all the files listed in the previous section from the cloned 'fin_dashboard' directory into your working 'stock_frontend' directory.

7 Containerization

Container images are created by adding the application and its dependencies in one or more layers on top of a base container image. You define the base container image and the contents of the layers in a plain text file, called a **Dockerfile** (<https://docs.docker.com/engine/reference/builder/>) . When this is done, you simply call the `docker build` command to create the new container image.

1. Make sure you are in your 'stock_frontend' working directory.
2. Create the file 'requirements.txt'.

You use this file to specify the Python packages required for the application.

```
streamlit  
yfinance  
matplotlib  
plotly  
gnews
```

```
prophet
requests
nlTK
pandas
```

3. Create the file 'Dockerfile'.

```
FROM python:3.10-slim
COPY . /app
WORKDIR /app
EXPOSE 8501
RUN pip install -r requirements.txt
ENTRYPOINT ["streamlit", "run"]
CMD ["main.py"]
```

The Dockerfile tells the `docker build` command to:

- start with the 'python:3.10-slim' base image.
- copy the contents of the current directory (where the command is issued) to the '/app' directory in the image itself.
- use '/app' as the working directory for subsequent instructions.
- expose network port 8501, which is the default for the Streamlit service.
- install the Python packages per the 'requirements.txt' file created earlier.
- launch the Streamlit service.
- run the 'main.py' script.

4. Build the Docker image.

```
docker build -t fin_dashboard:latest .
```

5. Tag the image to your Docker Hub account.

```
docker tag fin_dashboard:latest <USERNAME>/fin_dashboard:latest
```

Replace `<USERNAME>` with your Docker Hub user name.


6. Push the image to your Docker Hub (<https://hub.docker.com/>) account.

This will make it easier to deploy to your Kubernetes cluster.

```
sudo docker push <USERNAME>/fin_dashboard:latest
```



Tip

Be sure you log in to your [Docker Hub \(https://hub.docker.com/\)](https://hub.docker.com/)  account before pushing the image. Use `docker login` on the command line to do this.


7. Test the docker container.

a. Run the container.

```
docker run -p 7000:8501 fin_dashboard:latest
```

b. Open your Web browser to `http://localhost:7000/` and verify that you can access the application.

8 Deployment manifests

A [manifest \(https://kubernetes.io/docs/concepts/overview/working-with-objects/kubernetes-objects/\)](https://kubernetes.io/docs/concepts/overview/working-with-objects/kubernetes-objects/)  specifies the desired state of an object that Kubernetes will maintain. Manifest files are typically in YAML format for human readability. For this application, you create three manifest files:

- 'namespace.yaml': to create a namespace where the application will be deployed.
- 'deployment.yaml': to define the deployment configuration.
- 'service.yaml': to expose the application so it can be accessed through a static port.

You also create a 'kustomization.yaml' file, which identifies the manifests to deploy.

1. In your 'stock-frontend' working directory, create the subdirectory, 'yaml_files', and enter it.
2. Create the 'namespace.yaml' manifest file with the contents listed below to define the 'finance' namespace.

```
apiVersion: v1
kind: Namespace
metadata:
  name: finance
```

3. Create the 'deployment.yaml' file with the listing below to define the deployment configuration.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app: stock-analysis
  name: stock-analysis
  namespace: finance
spec:
  replicas: 2 # Creating two PODs for our app
  selector:
    matchLabels:
      app: stock-analysis
  template:
    metadata:
      labels:
        app: stock-analysis
    spec:
      containers:
        - image: <USERNAME>/fin_dashboard:latest # Docker image name
          name: fin-markets # POD name
          ports:
            - containerPort: 8501
              protocol: TCP

```



Note

Here `image: <USERNAME>/fin_dashboard:latest` refers to the Docker image that you created and pushed to Docker Hub. Replace `<USERNAME>` with your Docker Hub user name.

4. Create the 'service.yaml' file below to expose the application so it can be accessed through a static port on the node.

```


apiVersion: v1
kind: Service
metadata:
  name: stock-analysis
  labels:
    run: stock-analysis
  namespace: finance
spec:
  type: NodePort
  ports:
    - port: 8501

```

```
targetPort: 8501
protocol: TCP
name: http
- port: 443
  protocol: TCP
  name: https
selector:
  app: stock-analysis
```



Note


Specifying a [NodePort](https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport) (<https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport>)  gives you the option to set up your own load balancing solution or expose one or more nodes' IP addresses directly.

5. Create the 'kustomization.yaml' file below to specify the manifests to deploy.

```
apiVersion: kustomize.config.k8s.io/v1beta1
kind: Kustomization
resources:
- namespace.yaml
- deployment.yaml
- service.yaml
```


9 Deployment

With your manifest files defining the deployment, you need to execute the deployment to your Kubernetes cluster.

1. Make sure you have [Kubectl installed and configured](https://documentation.suse.com/cloudnative/rancher-manager/latest/en/faq/install-and-configure-kubectl.html) (<https://documentation.suse.com/cloudnative/rancher-manager/latest/en/faq/install-and-configure-kubectl.html>)  to manage your Kubernetes cluster.



Note

Since version 1.14, Kubectl supports management of Kubernetes objects using a [Kustomize](https://kustomize.io/) (<https://kustomize.io/>)  kustomization file.

2. Deploy your application.

```
kubectl apply -k yaml_files
```



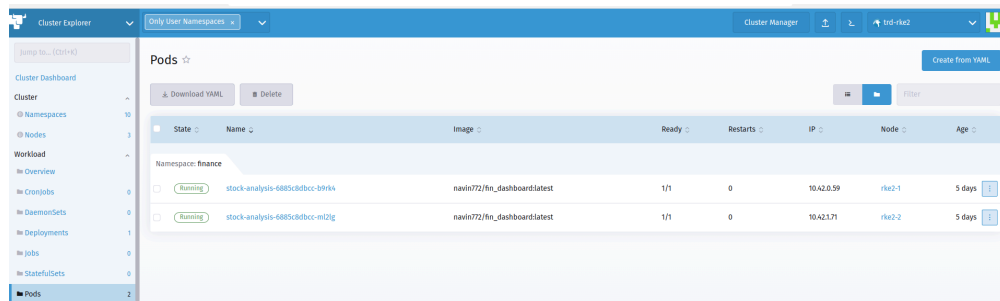
Note

The `-k` option tells `kubectl` to process the kustomization file in the directory.

3. Check whether all the pods are running in the 'finance' namespace.

```
kubectl get pods -n finance
```

NAME	READY	STATUS	RESTARTS	AGE
stock-analysis-6885c8dbcc-b9rk4	1/1	Running	0	2m25s
stock-analysis-6885c8dbcc-ml2lg	1/1	Running	0	2m25s



4. Get all the information of the 'finance' namespace.

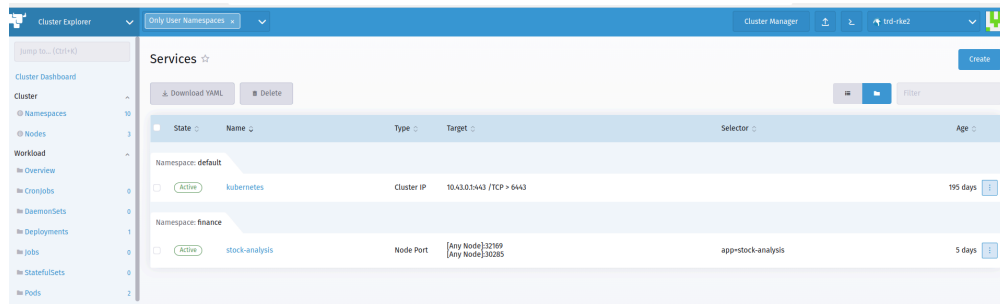
```
kubectl get all -n finance
```

NAME	READY	STATUS	RESTARTS	AGE
pod/stock-analysis-6885c8dbcc-b9rk4	1/1	Running	0	2m25s
pod/stock-analysis-6885c8dbcc-ml2lg	1/1	Running	0	2m25s

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
service/stock-analysis	NodePort	10.43.148.169	<none>	8501:32169/ TCP,443:30285/TCP

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
deployment.apps/stock-analysis	2/2	2	2	2m25s

NAME	DESIRED	CURRENT	READY	AGE
replicaset.apps/stock-analysis-6885c8dbcc	2	2	2	2m25s



Note

The deployment process might take some time depending on the resources of the cluster. Be patient!

- When the deployment is complete, note the 'CLUSTER-IP' and 'PORT(S)' associated with 'service/stock-analysis'.

In the example listing above, the service is available at 10.43.148.169 on port 32169.

10 Demonstration

Get started using Fin-Dashboard:

- Open your Web browser to the IP and port of the service (in our example, it would be <http://10.43.148.169:32169>).
- The Dashboard automatically detects your country, but you can select another from the drop-down box.


×

Select from below options

Select one

Dashboard

The Dashboard will show the predictions for the top 5 stocks for short/long term. You can also get predictions for any other stock using the next block.



Not your country? Select your country from the dropdown below

Change your country

India

See predictions for top stocks for your country India

Select your timeframe

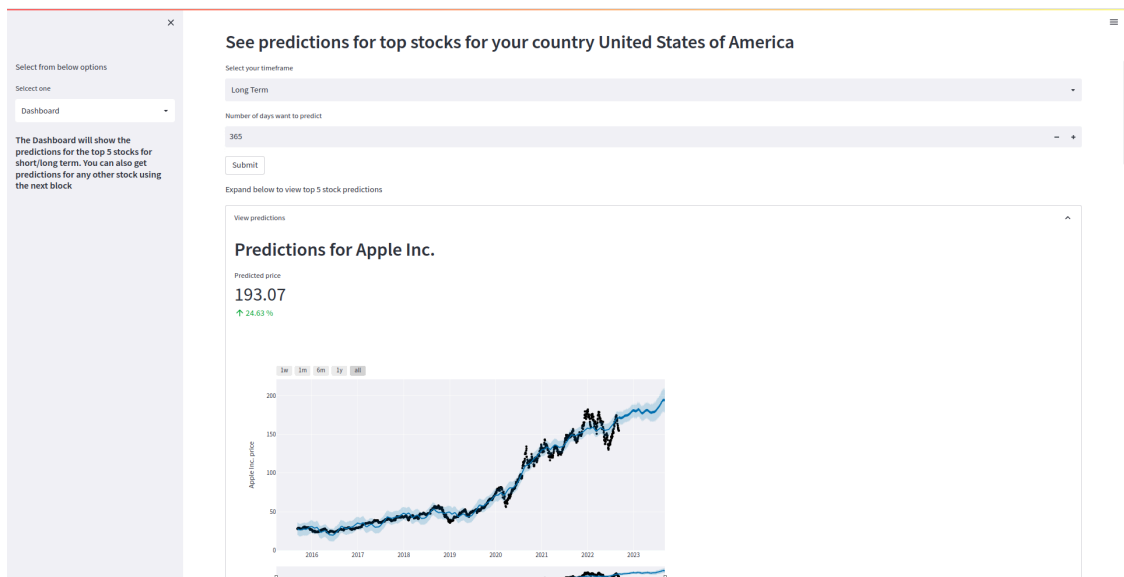
Short Term

Short Term

Long Term

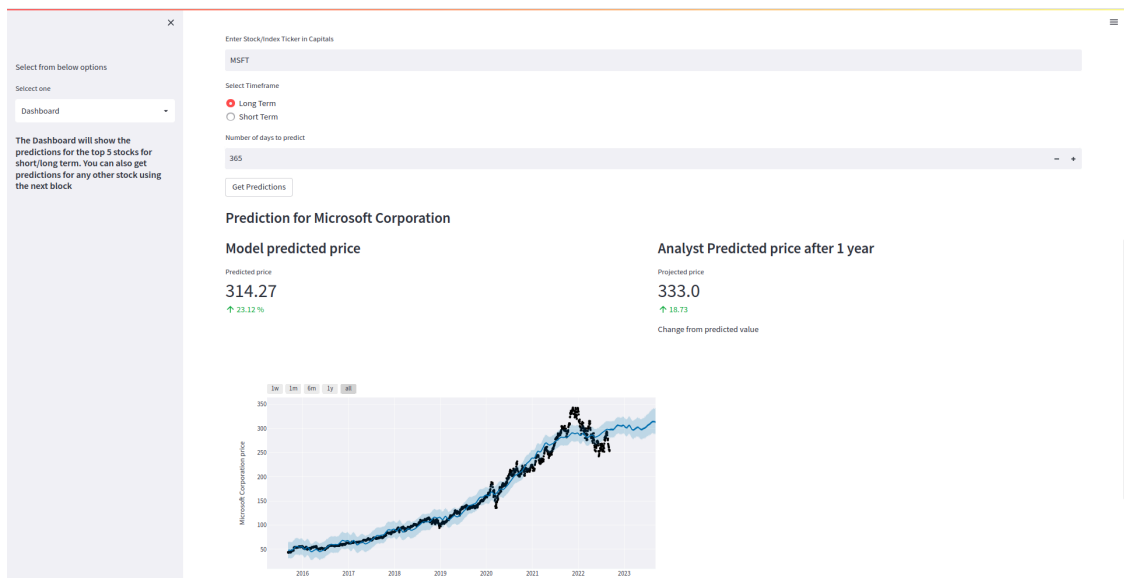
Submit

3. Select your timeframe, either 'Short Term' or 'Long Term'.
4. Enter the number of days to predict.
5. Click *Submit*.

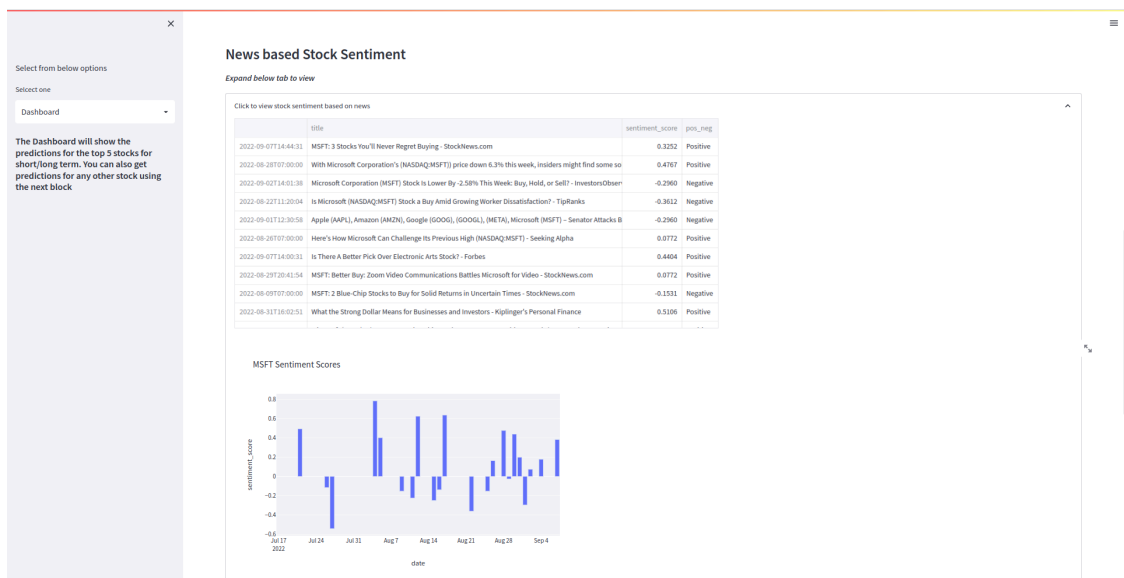


Explore the application features:

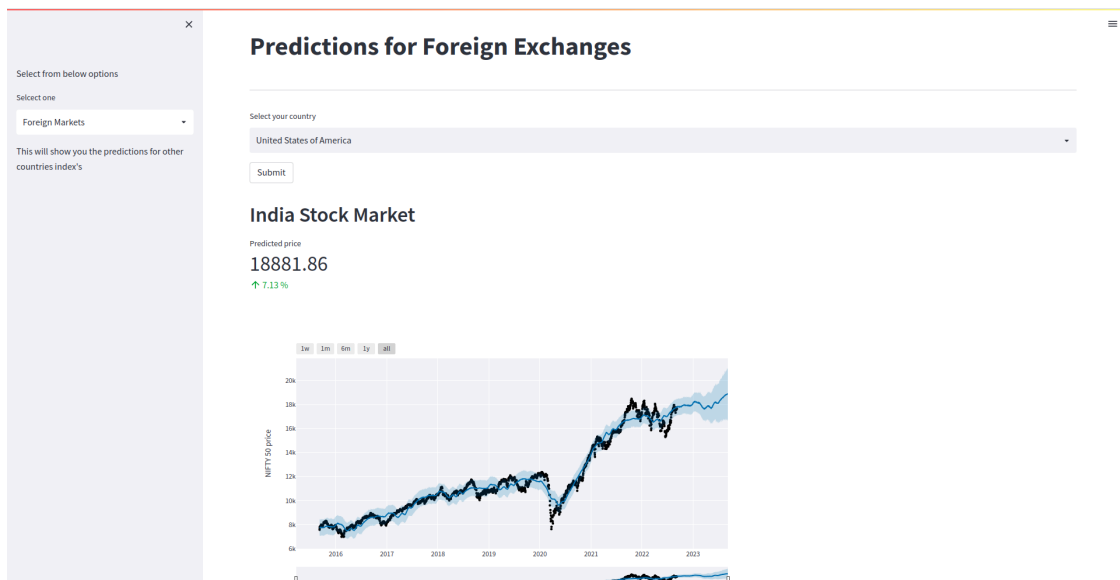
- For long term predictions, analysts' target price is also shown, which can be useful for comparing the model.




- For short term predictions, news sentiment analysis is shown to provide a sense of the market mood.



- You can view long term prediction charts for up to eight other countries.



- You can also view the correlation between the specified stock and the market index by selecting 'Index Correlation'.



11 Summary

Analytics at the edge is of growing importance in the financial space. Organizations realize they can achieve faster insights, maintain better security practices, and lower costs by deploying certain analytics workloads closer to or even at the edge. The scale and agility of cloud native architectures and tooling, like Kubernetes, make this possible.

In this guide, you learned how to containerize an ML application and deploy it into a Kubernetes environment managed by Rancher by SUSE.

Below are listed a few resources to help you continue your exploration of Rancher and K3s for edge analytics:

- [GitHub repository with the code \(https://github.com/navin772/GSOC_Stock_prediction\)](https://github.com/navin772/GSOC_Stock_prediction) ↗
- [SUSE Technical Reference Documentation \(https://documentation.suse.com/trd/suse/\)](https://documentation.suse.com/trd/suse/) ↗
- [Introductory Deployment of K3s \(https://documentation.suse.com/trd/suse/html/kubernetes_ri_k3s-slemicro/index.html\)](https://documentation.suse.com/trd/suse/html/kubernetes_ri_k3s-slemicro/index.html) ↗
- [SUSE & Rancher Community \(https://www.suse.com/community/\)](https://www.suse.com/community/) ↗

12 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> .

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

13 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.